

THE IMPACT OF DATA PRESENTATION PRACTICES ON PERCEPTIONS OF POLITICAL POLARISATION

JonRobert James Tartaglione

Homerton College

July 2024



UNIVERSITY OF
CAMBRIDGE

This dissertation is submitted for the degree of Doctor of Philosophy

DECLARATION

This thesis is the result of my own work and includes nothing which is the outcome of work done in collaboration except as declared in the preface and specified in the text. It is not substantially the same as any work that has already been submitted, or is being concurrently submitted, for any degree, diploma, or other qualification at the University of Cambridge or any other University or similar institution except as declared in the preface and specified in the text. It does not exceed the prescribed word limit for the relevant Degree Committee.

ABSTRACT

By multiple metrics, the United States is a highly polarised sociopolitical environment. However, some researchers – while acknowledging the presence of some degree of polarisation – have deemed it fair to question whether citizens are as divided as they *believe themselves to be*. “False polarisation” is a phenomenon wherein true levels of polarisation are consistently overestimated by political partisans, leading to pervasive interparty misperceptions and arguably precipitating a self-fulfilling prophecy which ultimately accelerates *actual* polarisation. While false polarisation is purported to be driven by well-documented socio-cognitive propensities such as oversimplification and categorisation, the present research questions whether it may also be a result of the manner in which we choose to *present data*. Chapter 1 tested the impact of data visualisation format on participants political perceptions and cross-party behaviour. A series of four inter-related studies yielded a number of key findings, including 1) type of visualisation plays a highly-significant role in perceptions of polarisation, 2) more “balanced” visualisation formats (e.g., those which depict both group differences *and* similarities or those which do not visually *accentuate* group differences) promote more accurate perceptions of intergroup agreement, and 3) providing participants with certain types of “sub-optimal” data visualisations may be superior to providing them with no visualisation at all. Chapter 2 examines the results from a large-scale replication of our work in Chapter 1, which broadened both the scale and scope (i.e., two separate geopolitical contexts) of our original set of four studies. Chapter 3 investigated the longitudinal impact of repeated exposure to different formats of data summarisation on partisan political perceptions over time. Although results broadly failed to identify significant differences *between* summarisation techniques, a pattern materialised which seemed to indicate that *any* summarisation technique promoted more positive intergroup perceptions than a no data baseline. Chapter 4 sought to understand whether different visualisation choices and summarisation techniques might impact the way in which an audience A) engages with content or B) perceives sources that present content in a particular manner. Results suggested that while different data presentation choices seem to have little impact on engagement with content, they exert significant influence on metrics related to source perception (e.g., level of trust in source, beliefs about intent, etc.), especially when the relative advantages and disadvantages of presentation methods were made explicit. Finally, Chapter 5 tested whether it was possible to “inoculate” individuals against perceiving exaggerated levels of polarisation when presented with interparty data depicted via truncated bar charts. Results indicated that a short inoculation intervention made individuals significantly more capable of correctly differentiating between misleading and non-misleading visualisations and ultimately significantly less likely to succumb to truncation-induced polarisation effects.

ACKNOWLEDGEMENTS

Overt sentimentality – at least the type that is not ultimately undercut by some sarcastic follow-up – is not really my thing. However, as this is likely the culminating chapter of my academic career, I want to take just a moment to acknowledge a few individuals who have played a defining role in shaping my path to this point.

First and foremost, to my parents, who have instilled in me a belief in the value of education and who have always nurtured my ambitions and passions, whatever they may be.

To Dr. Ellen Langer, whose kind response to an email from a young man with no psychology experience to speak of meant more than she will likely ever know.

To Dr. Lee de-Wit, whose knowledge, patience, and astute guidance made him an outstanding supervisor, and whose character and integrity made him a leader worth following.

To the members of the Political Psychology Lab, past and present, whose brilliance, collaborative spirit, humour, and humility created an ideal environment in which to propose, test, and improve ideas.

And finally, to my wife, Sumedha, whose support has never wavered, and whose pride in me fuels me endlessly.

Thank you.

With love and gratitude,

JonRobert “Tat” Tartaglione

P.S. Oh, and to the American and British people, whose refusal to spell words in the same way has made for a *really* seamless writing experience for an American attempting to obtain a Ph.D. in the UK.

P.P.S. Sorry, I couldn't keep myself from the sarcastic follow-up.

"All people are born alike – except Republicans and Democrats."

-Comedian Groucho Marx

TABLE OF CONTENTS

| | |
|----------------------------------------------------------------------------------------------------------------------------------------|----|
| Introduction: Political Polarisation, Pervasive Socio-Political Misperceptions, and The Possible Role of Data Presentation..... | 12 |
| The Current State of Political Polarisation..... | 12 |
| The Case for False (or Misperceived) Polarisation..... | 15 |
| <i>Evidence of the Overestimation of Polarisation and the Presence of More General Inter-Party Misperceptions.....</i> | 16 |
| <i>Socio-Cognitive and Affective Processes Theorised to Contribute to Partisan Misperception and False Polarisation.....</i> | 20 |
| <i>False Polarisation to True Polarisation: The Case for a Self-Fulfilling Prophecy.....</i> | 21 |
| Combating False Polarisation: The Value and Promise of Misperception Correction Interventions..... | 22 |
| The Double-Edged Sword of Data Presentation: A Tool to Inform, A Weapon to Manipulate..... | 23 |
| <i>Data Visualisations and Summary Statistics as a Form of Framing.....</i> | 31 |
| The Overarching Observations Motivating Our Research..... | 33 |
| <i>Observation 1: In the Context of Political Perception, Choice of Data Presentation Format Is A Significant Decision.....</i> | 33 |
| <i>Observation 2: Some Traditional Formats of Data Presentation May Reinforce Political Misperceptions.....</i> | 36 |
| <i>Observation 3: Certain Non-Traditional Data Presentation Formats May Mitigate Misperceptions of Political Polarisation.....</i> | 44 |
| A Brief Overview of Our Research Approach, Trajectory, and Objectives..... | 49 |
| Chapter 1: How Data Visualisation Choices Affect and Correct Political (Mis)Perceptions..... | 51 |
| Abstract..... | 51 |
| Research Approach..... | 55 |

| | |
|------------------------------------------------------------------------------------------------------------------------------------------|----|
| Study 1: Investigating the Effects of Mode of Presentation for an Unspecified Issue with Political and Non-Political Groups..... | 55 |
| Method..... | 57 |
| <i>Participants</i> | 57 |
| <i>Design</i> | 58 |
| <i>Procedure</i> | 58 |
| <i>Measures</i> | 61 |
| <i>Hypotheses</i> | 61 |
| Results..... | 62 |
| Discussion..... | 64 |
| Study 2: Investigating the Effects of Mode of Presentation for a Specified Issue Attributed to Elite and Non-Elite Political Groups..... | 66 |
| Method..... | 68 |
| <i>Participants</i> | 68 |
| <i>Design</i> | 69 |
| <i>Procedure</i> | 70 |
| <i>Measures</i> | 71 |
| <i>Hypotheses</i> | 72 |
| Results..... | 72 |
| Discussion..... | 77 |
| Study 3: Investigating the Intersection of Mode of Presentation and Source Effects..... | 79 |
| Method..... | 81 |
| <i>Participants</i> | 81 |
| <i>Design</i> | 82 |
| <i>Procedure</i> | 83 |
| <i>Measures</i> | 84 |

| | |
|------------------------------------------------------------------------------------------------------------------------------------------|------------|
| <i>Hypotheses</i> | 85 |
| Results..... | 86 |
| Discussion..... | 90 |
| Study 4: Investigating the Immediate and Longitudinal Effects of Mode of Presentation Using Real Partisan Data on a Polarised Issue..... | 92 |
| Method..... | 94 |
| <i>Participants</i> | 94 |
| <i>Design</i> | 95 |
| <i>Procedure</i> | 96 |
| <i>Measures</i> | 98 |
| <i>Hypotheses</i> | 99 |
| Results..... | 99 |
| Discussion..... | 105 |
| General Discussion..... | 108 |
| Chapter 2: Investigating the Replicability of Mode of Presentation Effects Across Geopolitical Contexts..... | 111 |
| Abstract..... | 111 |
| Method..... | 114 |
| <i>Participants</i> | 114 |
| <i>Procedure</i> | 115 |
| <i>Measures</i> | 116 |
| <i>Hypotheses</i> | 118 |
| Results..... | 118 |
| Discussion..... | 124 |
| Conclusion..... | 127 |
| Chapter 3: Investigating Whether Data Summarisation Choices Shift Views of Polarisation and Interparty Dynamics Over Time..... | 129 |

| | |
|----------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| Abstract..... | 129 |
| Method..... | 133 |
| <i>Participants</i> | 133 |
| <i>Design</i> | 134 |
| <i>Procedure</i> | 135 |
| <i>Measures</i> | 136 |
| <i>Hypotheses</i> | 139 |
| Results..... | 140 |
| Discussion..... | 149 |
| Shifting Focus: From Identifying Challenges to Investigating Potential Solutions | 153 |
| Chapter 4: Exploring A Systemically-Focused Intervention: Elucidating the Impact of Data Presentation Choices on Media Reputation and Trust | 155 |
| Abstract..... | 155 |
| Method..... | 160 |
| <i>Participants</i> | 160 |
| <i>Design</i> | 161 |
| <i>Procedure</i> | 162 |
| <i>Measures</i> | 165 |
| <i>Hypotheses</i> | 166 |
| Results..... | 168 |
| <i>Visualisation Study</i> | 168 |
| <i>Summarisation Study</i> | 178 |
| Discussion..... | 185 |
| <i>Limitations</i> | 187 |
| Conclusion..... | 188 |

| | |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------|------------|
| Chapter 5: Exploring An Individual-Focused Intervention: Testing the Prospect of Inoculating Against Truncation-Induced Polarisation Effects..... | 190 |
| Abstract..... | 190 |
| Method..... | 193 |
| <i>Participants.....</i> | <i>193</i> |
| <i>Design.....</i> | <i>194</i> |
| <i>Procedure.....</i> | <i>195</i> |
| <i>Measures.....</i> | <i>196</i> |
| <i>Hypotheses.....</i> | <i>198</i> |
| Results..... | 198 |
| Discussion..... | 202 |
| General Discussion: The Impact of Data Presentation on Political Perception and Intergroup Dynamics..... | 207 |
| A Summary of Research Objectives..... | 207 |
| Synthesising Key Findings..... | 208 |
| <i>Visualisation Choice Exerts a Significant Impact on Perceptions of Ideological Polarisation.....</i> | <i>208</i> |
| <i>More “Balanced” Visualisations Catalyse More Accurate Perceptions – But Still Significantly Underestimate Actual Levels of Intergroup Agreement.....</i> | <i>210</i> |
| <i>In the Context of Reducing Polarised Perceptions, Exposure to Any Form of Data Often Seems Preferable to Exposure to No Data at All.....</i> | <i>211</i> |
| <i>Data Presentation Choices Impact How Individuals Perceive and Intend to Interact with Sources.....</i> | <i>213</i> |
| <i>Individuals Can Be Inoculated Against Truncation-Induced Polarisation Effects.....</i> | <i>214</i> |
| <i>Shifting Perceptions of Ideological Polarisation Might Not Exert Meaningful Influence on Levels of Affective Polarisation or Interparty Behaviour.....</i> | <i>215</i> |
| Directions for Future Research..... | 217 |

| | |
|--------------------------------------------------------------------------------------------------------------------------|------------|
| <i>Does Shifting Perceptions on Specific Types of Interparty Data Reduce Affective Polarisation?</i> | 217 |
| <i>Does Shifting Perceptions of Ideological Polarisation Lead to Reliable Shifts in Any Interparty Behaviours?</i> | 218 |
| <i>Will Similar Effects Emerge in Dissimilar Samples?</i> | 219 |
| Concluding Thoughts..... | 220 |
| Appendices | 223 |
| Appendix A: The Construction and Usage of Icon Array Histograms (Introduction, Chapter 1, and Chapter 2)..... | 223 |
| Appendix B: The Appropriateness of Different Measures of Similarity for Different Data Visualisations (Chapter 1)..... | 224 |
| Appendix C: Chapter 1, Study 1 Additional Hypotheses and Results..... | 225 |
| Appendix D: Chapter 1, Study 2 Additional Hypotheses and Results..... | 227 |
| Appendix E: Chapter 1, Study 3 Additional Hypotheses and Results..... | 228 |
| Appendix F: Chapter 1, Study 4 Additional Hypotheses and Results..... | 231 |
| Appendix G: Rationale for Selecting Immigration as the Issue Presented to Participants (Chapter 2)..... | 234 |
| Appendix H: The Interaction Between Education and Perceived Issue Polarisation (Chapter 2)..... | 234 |
| Appendix I: Stimuli Accompanying Summary Statistics for Participants in Visualisation Conditions (Chapter 3)..... | 234 |
| Appendix J: Pairwise T-Tests between Experimental Summary Statistic Conditions (Chapter 3) | 237 |
| Appendix K: Additional Analyses, Results, and Discussion (Chapter 4)..... | 239 |
| Appendix L: Full Script of the Inoculation Intervention (Chapter 5)..... | 257 |
| References | 262 |

INTRODUCTION

POLITICAL POLARISATION, PERVASIVE SOCIO-POLITICAL MISPERCEPTIONS, AND THE POSSIBLE ROLE OF DATA PRESENTATION

The Current State of Political Polarisation

Challenging though it may be to believe given the vexing incompatibility that seems to characterise the parties of the modern American political system, for the majority of the 20th century, many of America's foremost political thinkers bemoaned the *dearth* of political polarisation within the country (Fiorina, 1980), and thus political science organisations in the United States sought to find ways to *create* greater degrees of political differentiation. In 1950, the American Political Science Association (APSA) released a 98-page report titled "Towards a More Responsible Two-Party System" (American Political Science Association, 1950) which asserted that the composition of the current political system was too ill-defined to be reliably navigable by the average citizen. More specifically, the authors – among them some of the most eminent political scholars of the era – contended that parties contained too much intraparty diversity of opinion and were too willing to pursue interparty cooperation and compromise. The authors argued that the heterogeneity within political parties, coupled with their lack of coherent policy preferences or clear proposals for governmental programs, rendered the parties devoid of discernible political identities. Since parties are a political instrument designed to provide the electorate with a range of options related to different courses of action (American Political Science Association, 1950), political scientists believed that ill-defined party identities compromised the public's ability to make informed choices.

While such concerns, at the time, were likely not misplaced, less than a century later, the notion of *intentionally* manufacturing political polarisation sounds satirical, especially as a formidable body of literature has documented the surging volatility and persistent acceleration of political polarisation which has afflicted the United States over the past several decades (e.g., Park, 2018; McCarty, Poole, & Rosenthal, 2016; Abramowitz & Saunders, 2008).

Broadly, scholars tend to categorise political polarisation in one of two ways: ideological or affective. Ideological polarisation refers to policy-based divisions, wherein individuals disagree on the correct way to handle economic and social issues (e.g., Iyengar, Sood, & Lelkes, 2012). This form of polarisation can be measured via mechanisms such as opinion polls, ideological self-placement surveys, or an examination of

voting patterns. Conversely, affective polarisation refers to a general dislike, distrust, or reticence to engage with members of an opposing political party (e.g., Druckman & Levendusky, 2019). This form of polarisation has been traditionally measured by apparatuses such as “feeling thermometers” (e.g., Tyler & Iyengar, 2023) and social distance items (e.g., “How comfortable would you be having close friends from the opposing political party?”; Druckman & Levendusky, 2019). By several metrics, America has experienced a precipitous rise in levels of both affective and ideological polarisation for at least the past twenty years (Boxell, Gentzkow, & Shapiro, 2022; McCarty, Poole, & Rosenthal, 2016; Pew Research Center, 2017; Druckman & Levy, 2022; see Figure 1).

Figure 1

Progression of Intergroup Affective Polarisation Over Time (from Druckman & Levy, 2022)

The consequences of political polarisation are both robust and increasingly diffuse. Polarisation endangers a multitude of democracy-dependent functions, including facilitating reductions in cross-party engagement (Baldassarri & Page, 2021; Fiorina & Abrams, 2008), erosions of legislative productivity (McCarty, 2016; Mayhew, 2005), and the inhibition of political compromises which might otherwise prove mutually-beneficial (Levendusky & Malhotra, 2016a; Whitt et al., 2021). Rising levels of polarisation can also spur dangerous shifts in interpersonal dynamics among partisans, such as heightened levels of interparty hostility (Iyengar, Lelkes, Levendusky, Malhotra, & Westwood, 2019) which can, at times, engender cross-party dehumanisation (Cassese, 2021). Concerningly, not only does political polarisation (predictably) lead to domain-specific issues, but its effects do not appear to be confined to the realm of expressly political activity, as evidence has shown that polarisation is also capable of contributing to a host of downstream (ostensibly) apolitical issues such as diminished concern for climate change (Linde, 2020) and suboptimal responses to collective risk scenarios (Milosh, Painter, Sonin, Van Dijcke, & Wright, 2021).

Some measures suggest that prejudice towards political opponents (along with other deleterious consequences of political sectarianism) has escalated to its highest point in four decades within the United States (Finkel et al., 2020). For example, “ticket splitting,” a once-prevalent trend among American voters where a voter would, for instance, cast their vote for a Democratic Senator but also for a Republican President, has virtually evaporated from the behavioural repertoire of the American voter (Klein, 2020). Whereas approximately 1 in 4 voters split their ticket in elections throughout the 1970s and 1980s, in recent elections, the percentage of voters casting a split-ticket ballot has decreased to approximately 10% (American National Election Studies, 2021).

During a similar timeframe, levels of affective polarisation have steadily climbed. Ratings of the opposing party on a 0-100 “feeling thermometer” have declined from an average of 45 in 1980 to a paltry 29 by 2016 (Abramowitz & Webster, 2016). Hetherington and Weiler (2018) characterise any feeling thermometer rating of 20 or less as “hatred,” and note how in the two decades prior to 2000, the number of partisans whose ratings fell below that threshold never exceeded 20 percentage points; by 2016, that number had risen to 48 percent for Democrats and 50 percent for Republicans.

Furthermore, the motivations which underlie political allegiance might be shifting, as well. *Negative partisanship* describes a phenomenon by which an individual declares their political allegiances based not on *who they are*, but instead based upon who they *are not* (e.g., Abramowitz & Webster, 2018). More specifically, with negative partisanship, the motivation underlying the selection of parties to support or policies to back is driven more by a *rejection* of the alternatives (as opposed to an affirmation of merits of the selection itself; Abramowitz & Webster, 2016). For example, a 2016 poll found that independents who tend to “lean” towards voting for one party over another now do so predominantly because the “other” party endorses policies which they believe will hurt the country (Pew Research Center, 2016).

Lilliana Mason (2018) argues that the widening divide, particularly between Republicans and Democrats, is not sufficiently explained by merely examining disagreements about policy. On the contrary, she states that increases in the levels of bias, anger, and political activity on each side have far outstripped the much more moderate pace by which policy disagreements have widened. Instead, she contends that the divergence has to be understood through the lens of the strengthening and crystallisation of partisan identities, leading to a type of tribalistic intolerance.

When polarisation escalates beyond mere political disagreement into territory scholars have referred to as “pernicious polarisation” (a state characterised by “the division of society into mutually distrustful ‘Us versus Them’ camps in which political identity becomes a social identity”; Somer, McCoy, & Luke, 2023, p. 61), consequences affect not only the quality of intergroup relations, but also begin to erode the integrity of democratic systems. A state of pernicious polarisation can ultimately create perverse incentives, giving

license to parties to attain and retain control via the use of unscrupulous means, often resulting in the instigation of anti-democratic behaviour when it serves the needs of political actors grappling for power (Somer, McCoy, & Luke, 2023).

In fact, some commentators have gone as far as to suggest that America, while still a single geographic unit, has effectively become a nation housing two fundamentally incompatible political *blocs*, each operating as a separate sociocultural entity (Brownstein, 2022). One such commentator, political strategist Michael Podhorzer, contends that thinking of America as a single nation is an “essential error”; rather, “[America is] more like a federated republic of two nations: Blue Nation [Democrats] and Red Nation [Republicans]. This is not a metaphor; it is a geographic and historical reality” (Brownstein, 2022).

In their 2022 chapter on affective polarisation in the “Handbook on Politics and Political Opinion,” James Druckman and Jeremy Levy write “it has become a point of consensus that *something* concerning is going on” (p. 257). However, what exactly that “something” is remains a point of contention within the academic community.

The Case for False (or Misperceived) Polarisation

While growing levels of polarisation is undoubtedly an issue worth monitoring, some researchers have taken to questioning the accuracy of the prevailing wisdom that the circumstances unfolding in the United States reflects the dynamics of genuine, runaway polarisation (such as suggested by teams like Abramowitz & McCoy, 2019 or McCoy, Rahman, & Somer, 2018). For example, some have taken issue with the manner in which levels of polarisation have traditionally been *measured* (e.g., Levendusky & Pope, 2011; Mehlhaff, 2023), arguing that classic analytical methodologies may ignore relevant metrics such as *intragroup* heterogeneity. Others contend that while democratic contests may have the *appearance* of intense polarisation, this need not necessarily be the result of an escalating ideological cleavage, but rather a dearth of moderate *candidates* (e.g., Binder, 2004) who may more appropriately represent the actual views of the less extreme political majority on both sides. Levendusky and Pope (2011), for instance, argue that it is “misleading to discuss a polarization of beliefs in the presence of so much ideological overlap” (p. 228); instead, a more apt way to think about large-scale polarisation is through the lens of “a polarization of *choices* available to the electorate” (p. 228).

Perhaps most intriguingly, certain cohorts of researchers have highlighted what appears to be a pervasive disconnect between *actual* levels of polarisation and *perceived* levels of polarisation, with the latter being consistently greater than the former – a phenomenon known as “false polarisation” (Fernbach & Van Boven, 2022). Researchers investigating false polarisation (or, alternatively, *misperceived polarisation*; e.g., Lees & Cikara, 2021) do not contend that polarisation itself is not occurring. Instead, they forward a series of less

sensationalised claims, including: 1) that levels of polarisation are habitually overestimated by members of the general public, 2) that this pattern of overestimation and general inter-party misperception can, at least in part, be explained by a myriad of natural socio-cognitive and affective proclivities, and 3) that such misperceptions, however exaggerated, are capable of influencing partisan attitudes, emotions, and behaviour, which can promote cycles of reciprocal escalation of political tension and, consequently, have the power to precipitate *actual* increases in polarisation via the process of a self-fulfilling prophecy.

Evidence of the Overestimation of Polarisation and the Presence of More General Inter-Party Misperceptions

A collection of research has concluded that political misperceptions have important consequences for interparty relations (Ahler & Sood, 2018; Moore-Berg, Ankori-Karlinsky, Hameiri, & Bruneau, 2020), and empirical evidence documents a persistent pattern of overestimation of levels of polarisation as well as more general misperceptions about “political others.”

Some of the earlier documentations of misperceptions across ideological divides came from Robinson, Keltner, Ward, and Ross (1995), who found that partisans greatly exaggerated the differences between themselves and their political adversaries. For instance, when asked to read scenarios about women undergoing procedures for abortion, pro-choice participants rated themselves as “very sympathetic” and expected pro-life participants to express dramatically lower levels of sympathy, when in reality, the two groups were nearly equally-sympathetic. A similarly important finding from the study was that *non-partisans also overestimated* the differences between liberals and conservatives, which implies that false polarisation may not be exclusively driven by misperceptions from (and of) the fringes of the political spectrum, but may, in fact, also be supported by misperceptions from the so-called “middle.”

Related research has shown that partisans struggle to identify the actual position of their adversaries on a host of key issues. Chambers, Baron, and Inman (2006) documented the attempts of partisans to estimate the levels of political disagreement across several “value issues” (e.g., eliminating poverty, the presence of a strong military, etc.) and found that, on some issues, perceived disagreement was more than *three times* as great as levels of actual disagreement. Reporting similar patterns, Moore-Berg et al. (2020) found that, while Republicans and Democrats do undoubtedly harbour some degree of ill-will towards one another, their estimations of the degree of prejudice and dehumanisation held by their political opponents were approximately twice as high as the actual levels.

Ahler (2014) likens the habitual tendency of partisans on both sides of the aisle to overestimate the extremity of ideological conviction of political others as being “a phenomenon akin to pluralistic ignorance of public moderateness” (Ahler, 2014, p. 607). Essentially, while a majority of the population likely harbours moderate, “relatively centrist” social and political views, their belief that they are in the minority in these views can result in the stifling of public support and defences of such moderation. In the pluralistic ignorance

literature, this phenomenon – wherein a misapprehension of other’s beliefs leads to a reticence to state one’s true feelings – is referred to as the “spiral of silence” (e.g., Taylor, 1982). Politically, such a spiral may result in the wide-scale misperceptions of political extremity as moderates, despite their superior numbers, fail to speak up and, correspondingly, the voices of extremists – emboldened by the overestimation of their prevalence in the electorate – grow louder.

Misperception research has brought to light what scholars have termed “perception gaps,” or large differences between where partisans actually stand on issues and where their political rivals *believe* they stand. A 2019 publication released by the global, non-partisan political organisation *More in Common* sought to examine these perception gaps, and thereby identify the most misleadingly distorted cross-party beliefs held by Republicans and Democrats (Yudkin, Hawkins, & Dixon, 2019). Republicans, for instance, vastly underestimated the percentage of Democrats who would disagree with the statement “*Most police are bad people,*” submitting an average guess of 48% when, in reality, 85% of Democrats indicated disagreement with this statement – a difference of 37 percentage points. Similarly, Democrats substantially underestimated the percentage of Republicans who would agree with the statement “*Properly controlled immigration can be good for America,*” submitting an average guess of 52% when the true number of Republicans endorsing such a statement was actually 85% – a difference of 33 percentage points.

Ahler (2014) argues that misperceptions may, in part, be driven by a persistent media bias towards framing societal problems as necessarily arising from polarisation. The media’s arguable overreliance on a polarisation narrative to rationalise a host of societal ills (Fiorina, Abrams, & Pope, 2011) likely permeates the public conscience. Mutz (1998) notes that while there may not exist sufficient evidence to suggest that mass media successfully influences *what people think*, she emphasises that these agencies are “tremendously influential in telling people what others are *thinking about* and experiencing [and these] perceptions, in turn, have important consequences for the political behaviour of mass publics and political elites as well” (p. 5; italics added).

Alternatively, Ahler and Sood (2018) note that misperceptions may also be caused by stereotypic thinking about the parties. They forward the idea that parties have, in essence, become “sociopolitical brands” (p. 964) whose images are inextricably linked to the groups with whom they have come to be associated. Campbell, Green, and Layman (2007) suggest that, when assessing political parties, many citizens tend to intuitively ask themselves questions such as “What kinds of social groups come to mind as I think about [these parties]?” Such an exercise is likely to generate prototypes of a “typical” Republican or Democrat, who – by the very nature of prototypical thinking – are unlikely to capture the true degree of diversity and nuance which can be found within any ideological cohort. Consequently, the *salience* of group prototypes might result in an overestimation of prototype-congruent party composition, as Ahler and Sood (2018) note

that “Americans believe that 32% of Democrats are gay, lesbian, or bisexual (only 6.3% are in reality), and that 38% of Republicans earn over \$250,000 per year (just 2.2% do in reality)” (p. 965).

More recent research has extended beyond mere inter-party perceptions to examine the field of group *meta-perceptions*, or “how we believe others perceive us” (Lees & Cikara, 2020, p. 279). Work on meta-perceptions has documented similar inaccuracies to those seen in the aforementioned, “first-order” perception literature, with scholars observing a persistent “negativity bias” in how groups believe they are seen by others (Lees & Cikara, 2020). Moreover, this phenomenon – which can lead to a host of negative out-group attributions – appears to also replicate across a number of diverse international samples (Ruggeri et al., 2021).

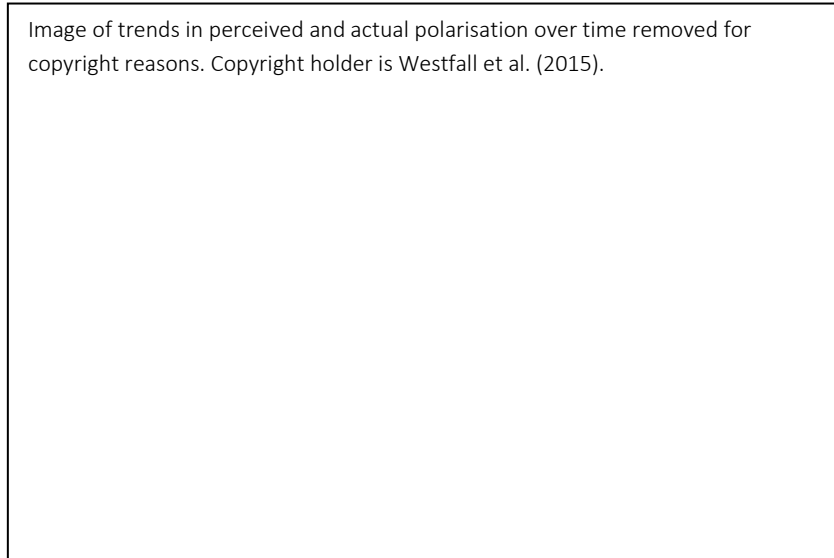
Misperceptions have also been shown to permeate the troubling domain of dehumanisation, with researchers documenting inaccuracies in partisan “meta-dehumanisation” (i.e., the perception of the degree to which you believe your group is dehumanised by an outside group; Landry, Ihm, Kwit, & Schooler, 2021). While partisans have indeed been shown to dehumanise members of rival political parties, Moore-Berg et al. (2020) found that the *actual* degree of dehumanisation was mild compared to the degree to which partisans *believed* their group would be dehumanised by political outgroups. Specifically, while the average partisan evaluated members of the outgroup to be approximately 21 points¹ “less human” than members of their own party, partisans anticipated being judged to be approximately 60 points less human by their rivals.

Finally, other researchers have opted to examine (mis)perceptions of polarisation in the aggregate by reviewing sets of archival data. In an impressive series of retrospective longitudinal studies, Westfall, Van Boven, Chambers, and Judd (2015) demonstrated a consistent model of incongruity between perceived and actual polarisation in the United States in a pattern which can be traced over the span of more than three decades (Figure 2).

¹ Using the “Ascent of Man” dehumanisation scale as seen in Kteily, Bruneau, Waytz, and Cotterill (2015), which asks participants to rank targets from 0 (“Least Evolved”) to 100 (“Most Evolved”).

Figure 2

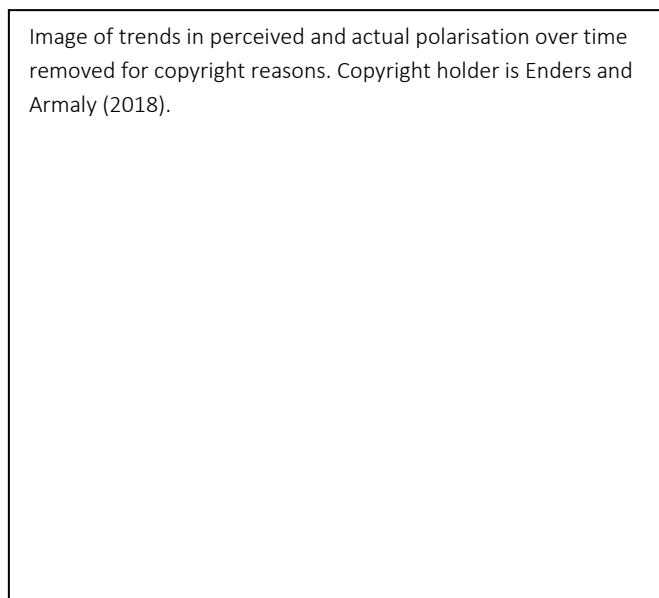
Perceived and Actual Polarisation Between Republicans and Democrats Over Time (from Westfall et al., 2015)



Similar work by Enders and Armaly (2019) also shows a historical disconnect between actual and perceived levels of polarisation (see Figure 3). However, their visualisation – derived from cumulative file data from the American National Election Studies between the years of 1972 and 2012 – seems to capture a *rapidly accelerating divergence* between the two types of polarisation beginning in the early 1980s.

Figure 3

Widening Gap between Actual and Perceived Polarisation from 1972 until 2012 (from Enders & Armaly, 2018)



Socio-Cognitive and Affective Processes Theorised to Contribute to Partisan Misperception and False Polarisation

False polarisation is caused by partisan propensities to overestimate the scale of interparty disagreement relative to true political attitudes. These pervasive inaccuracies in partisan perception are theorised to be driven by a variety of phenomena, such as unrealistic beliefs about ideological consistency (Fernbach & Van Boven, 2022), moral stereotyping of political others (Graham, Nosek, & Haidt, 2012), and exaggerated interparty differences (Keltner & Robinson, 1993).

The possibility of false polarisation is buoyed by a plethora of earlier research in social and cognitive psychology which has consistently shown that humans are prone to categorical thinking (Rosch, 1973). Humans effortlessly (and perhaps instinctively) categorise groups as “us” versus “them” (even within so-called “minimal group paradigms”; Tajfel, 1969; Tajfel, 1981; Tajfel & Turner, 1986). While such cognitive processes are arguably adaptive insofar as they enable us to better traverse a complex social world, categorical thinking often has the unintended consequence of *oversimplifying* our realities (Sloman, Sloman, & Fernbach, 2018), producing biased representations of those who fall within and outside of each social category (with obvious repercussions for polarised political environments). We project greater degrees of ideological coherence within social groups than actually exist (De Langhe & Fernbach, 2019), and simultaneously perceive outgroups as having less variability and uniqueness among their members (i.e., the *outgroup homogeneity effect*; Ostrom & Sedikides, 1992; Judd, Ryan, & Park, 1991), leading to the facilitation of *intergroup accentuation* processes (e.g., Rothermund & Deutsch, 2024). Moreover, we tend to exaggerate the ideological differences of ingroups and outgroups (Keltner & Robinson, 1993; Westfall, Van Boven, Chambers, & Judd, 2015), leading to incorrect beliefs about the extremity of the views of the “typical” political opponent relative to our own (Graham, Nosek, & Haidt, 2012; Ahler & Sood, 2018). What’s more, research has shown that these cognitive proclivities can be exacerbated via the experience of anger (Fernbach & Van Boven, 2022; Huber, Van Boven, Park, and Pizzi, 2015), a distressing realisation considering the heightened level of partisan animus and the ubiquity of combative rhetoric (Gervais, 2019) that currently characterises the American political arena.

Essentially, our penchant for categorical thinking facilitates the creation of crisp distinctions where there is actually noisy data, ultimately dichotomising more nuanced sociopolitical continuums. Moreover, we tend to overestimate both the magnitude of disagreement within the political sphere as well as the ideological consistency of the political actors inhabiting it. As Fernbach and Van Boven (2022, p. 3) note:

“We amplify differences across category boundaries, assuming that individuals from different categories are more different than they are, and we compress within categories, assuming that individuals from the same category are more similar than they are.”

These combinations of socio-cognitive and affective propensities unfortunately produce an ideal set of circumstances to accelerate political divisions wherever they might form, and even *create* political division where it may not currently exist.

False Polarisation to True Polarisation: The Case for a Self-Fulfilling Prophecy

In summarising the seminal work of Allport, Clark, and Pettigrew (1954), Westfall et al. (2015) note how “one of psychological science’s fundamental insights is that the perception of social reality has potent—if not prepotent—influence on cognitions, emotions, and behaviors” (p. 145). In light of such an insight, it should come as no surprise that misperceived polarisation, in spite of its detachment from political reality, can still contribute considerably to material political discord. In fact, the degree of *actual* polarisation currently observed might, in part, have *begun* as mere (mis)perception and is now ultimately the result of a self-fulfilling prophecy fuelled by inaccurate cross-party assumptions, pervasive intergroup misattributions, and flawed inferences about political “others.”

Scholars such as Huxley (1978) have contended that the escalation of group conflict is frequently attributable to fundamental social misperceptions. Echoing these sentiments, more recent work has suggested that polarisation can be cyclical in nature (Lees & Cikara, 2021), with upward shifts in domains such as affective polarisation producing elevated *perceptions* of division within the domain of ideological polarisation (Armaly & Enders, 2021; Levendusky & Malhotra, 2016b), eventually engendering a feedback loop that can catalyse a rapid acceleration in both areas. In essence, though changes in the *perceptions* of political polarisation do not necessarily entail any corresponding changes in the actual levels of polarisation, they can nevertheless contribute to real shifts through a process of mutual reinforcement (Lees & Cikara, 2021).

Ahler (2014) posits a similar rationale, arguing that political extremism may, at least in part, be attributable to pernicious feedback loops. He notes that participants who were exposed to the *true* distribution of attitudes on political issues tended to express greater moderation (i.e., reduced extremism) in their own views than participants who either A) did not have their perceptions corrected or B) has their misperceptions *enforced*, corroborating how allowing misperceptions to persist can bolster extremist attitudes. Landry et al. (2021) discussed how misperceptions of meta-dehumanisation threaten one’s social identity via the erosion of group-derived self-esteem (Branscombe, Ellemers, Spears, and Doosje, 1999). They conclude that misperceptions of meta-dehumanisation can promote an escalation of *actual* dehumanisation due to the fact that social identity threat promotes intergroup hostility via “a desire to restore the ingroup’s status by reciprocally denigrating the offending group (Branscombe & Wann, 1994)” (p. 52). Relatedly, Moore-Berg et al. (2020) found that support for policies that would harm the country and undercut democratic norms was directly attributable to inaccurate meta-perceptions concerning levels of prejudice and dehumanisation directed toward one’s political group by opposing political factions. Most recently, D’Amore, van Zomeren,

and Koudenberg (2024) found evidence to suggest that it was *perceptions* of greater polarisation that drove attitudinal moralisation among political partisans during the 2020 US election cycle.

In summarising the very real consequences of partisan misperception, Enders and Armaly (2019) conducted research which concluded that *perceived polarisation* is “more strongly related to negative affective evaluations of out-parties and out-party candidates, voting, [political] participation, trust, and efficacy than is actual polarisation” (p. 815).

Thus, false polarisation, instigated by misperceptions, perpetuated by a suite of socio-cognitive predispositions, and capable of ultimately facilitating the emergence of genuine polarisation, poses a significant challenge to those who wish to improve the prospects of cross-party cooperation and democratic civility. And so a critical question becomes: how might we combat such a phenomenon?

Combating False Polarisation: The Value and Promise of Misperception Correction Interventions

While research shows that false polarisation is fuelled by misperception, complementary research also suggests that it might be extinguished (or at least mitigated) via the *correction* of these misperceptions.

Misperception correction interventions have proven to be a simple and effective tactic for reducing misperceptions across a myriad of social and political domains. Conceptually-similar to norm-based interventions for producing attitude and behaviour change (e.g., Miller & Prentice, 2016), misperception correction paradigms generally feature a fairly straightforward design: participants, after disclosing their own beliefs or perceptions (e.g., about the ideological stances or views of a political other), are then exposed to *actual* data pertaining to the stances or views of the target political group. Assuming there is incongruity² between their own views and the reality reflected in the data, exposure can often cause a corresponding adjustment in perception inaccuracy.

Several of the studies referenced in the preceding section, while documenting political misperceptions, also sought to examine whether one might be able to *mitigate* these misperceptions via the mechanism of misperception correction. Ahler (2014), via an informational intervention, was able to reduce attitudinal extremity (by roughly 8-11 percentage points) in participants (and, thus, effectively bolster political moderation) simply by exposing them to true levels of out-group policy extremity (which was often far less inflated than participants had expected). Lees and Cikara (2020) found that correcting exaggerated group

² The underlying assumption here is that the incongruity would be in the *desired* direction (e.g., a participant believing an out-group member is more extreme in their beliefs than they actually are). While this is often the case, the risk of such interventions lies in the possibility that some participants may *underestimate* the degree of political disagreement, leading to so-called “boomerang effects” (e.g., Schultz, Nolan, Cialdini, Goldstein, & Griskevicius, 2007) wherein exposure to actual data prompts one to adjust their attitude or behaviour in the opposite direction of the intervention’s intention. In the political context, if certain subgroups underestimate the degree to which, for instance, Republicans and Democrats disagree, perception-correction interventions have the potential to backfire, creating greater degrees of polarisation, attitudinal extremity, etc. than may have existed prior to the exposure.

meta-perceptions effectively reduced negative outgroup attributions concerning obstructionism. And within the aforementioned domain of inaccurate perceptions of meta-dehumanisation (i.e., Moore-Berg et al., 2020), research demonstrated that correcting such erroneous assumptions (by exposing them to true levels of out-party dehumanisation) was not only successful in moderating dehumanisation in the short-term, but also proved to be a robust mitigator even one week post-intervention (Landry, Schooler, Willer, & Seli, 2023). In addition, recent undertakings by groups like Mernyk, Pink, Druckman, and Willer (2022) have investigated the troubling trend of support for partisan violence. They determined that partisan beliefs about their rivals' support for (and willingness to engage in) political violence were consistently overestimated by levels occasionally eclipsing 400%. Crucially, however, a brief informational intervention aimed at correcting these mistaken meta-perceptions attenuated both support for violence as well as one's willingness to engage in politically-motivated violence, with effects of the simple intervention persisting for more than one month post-administration.

Lees and Cikara (2021) state that, relative to alternative strategies such as political fact-checking (whose efficacy as an intervention is mixed; see Walter, Cohen, Holbert, & Morag, 2020), "addressing first-order misperceptions regarding polarisation appears to be promising" (p. 3). Supporting this perspective, Moore-Berg et al. (2020) note that "interventions that correct *intergroup* meta-perceptions may be even stronger than those that target *ideological* meta-perceptions" (p. 14870). Considering the promise that such misperception-correction paradigms hold, it's fair to wonder whether there might be any overlooked mechanisms or mediums through which practitioners could deploy these tactics at scale.

The Double-Edged Sword of Data Presentation: A Tool to Inform, A Weapon to Manipulate

"If I can't picture it, I can't understand it."

-Albert Einstein (Horgan, 1991)

Friendly and Denis (2001) posit that data visualisation subsumes two primary domains: statistical graphics and thematic cartography. While the advent of cartography, and thus the first evidence for data visualisation, can be traced back to at least 6,200 BC (Friendly & Denis, 2001),³ advances to facilitate the emergence of statistical graphics, such as proto-bar graphs (produced by French philosopher Nicole Oresme to portray a theoretical function), did not emerge until the 14th century (Friendly, 2008).⁴

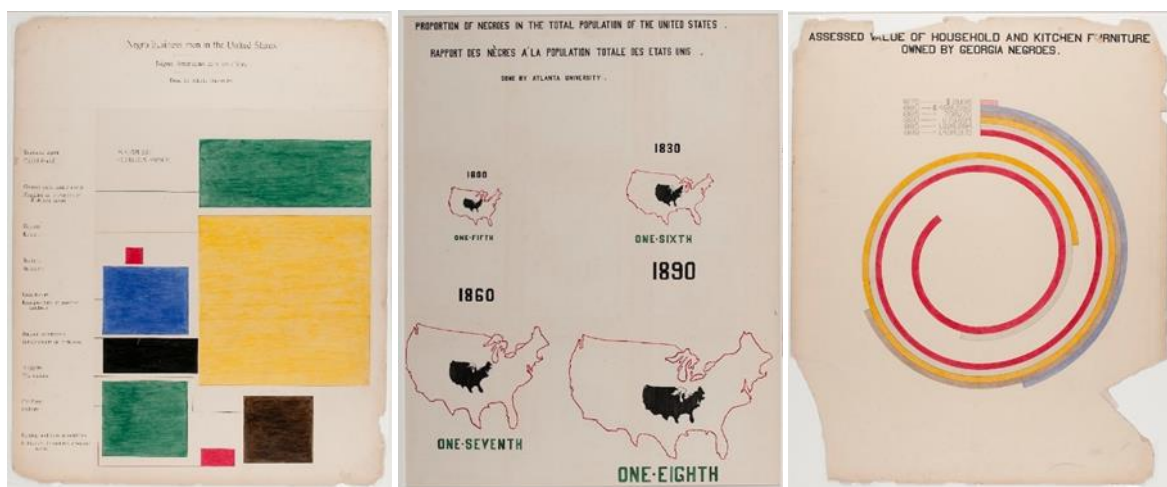
³ A drawing of the town of Catalhöyük, near present-day Konya, Turkey, with what appears to be an erupting volcano in the background (Friendly, 2008).

⁴ Tufte (2001) argues that the advent of some of the more modern statistical graphics (e.g., time-series, scatterplots, multivariate displays, etc.) can be traced back roughly to the period between 1750 and 1800, with many of the most significant advancements attributable to William Playfair, who sought a substitute for traditional mathematical tables as a means of visually depicting his "linear arithmetic" (Tufte, 2001). Spence (2006) corroborates and defends the magnitude of Playfair's contributions, noting that while some scholars point to "isolated instances of earlier diagrams that might be considered to be statistical...the publication of statistical graphs began with William Playfair" (p. 2426).

Since their emergence, some of the most important figures and revolutionary thinkers throughout history have recognised the potential power well-crafted visualisations have to inform and persuade. For example, in 1900, W. E. B. Du Bois and a team of African American sociologists debuted an exhibit called “The American Negro” at the *Exposition Universelle* (i.e., the Paris World Fair; Mansky, 2018), which featured a unique series of data visualisations and infographics (see Figure 4) that powerfully captured both the current state as well as the retrospective trajectory of African-American life in the US in an accessible, intuitive manner.

Figure 4


A Selection of Visualisations Created by William Edward Burghardt Du Bois (better known as W. E. B. Du Bois) and his Team of African American Sociologists for the 1900 Paris World Fair



Others have used data visualisations as a way to reliably track, analyse, and thus predict patterns in a chaotic world. For example, during World War II, Winston Churchill’s military and domestic strategies were informed, in part, by collection of data visualisations (see Figure 5) which closely monitored war-relevant data, including the weight of bombs dropped on enemy and occupied territory, the estimated production of oil each month in Axis Europe, and the manufacture of different fighter aircraft within the United Kingdom.

Figure 5

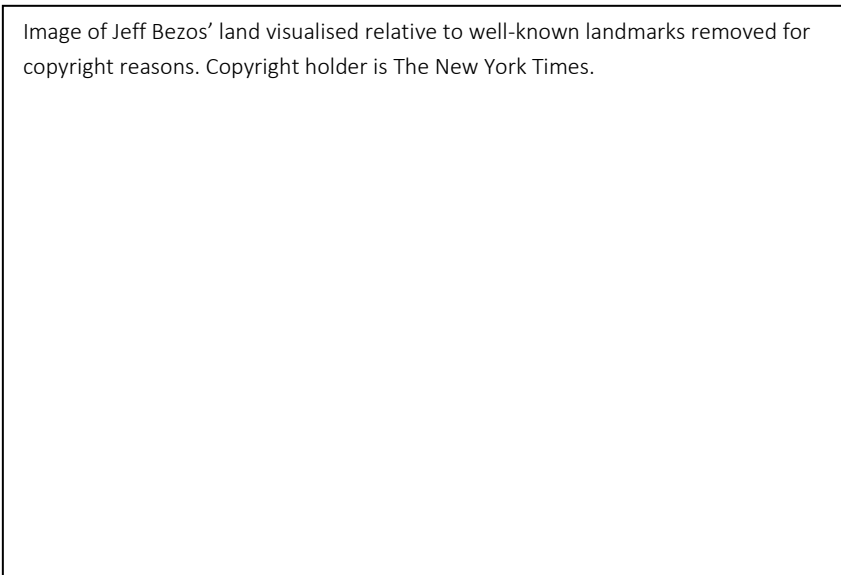
Collection of Data Visualizations Tracking War-Relevant Metrics in the Churchill War Rooms in London



Others, such as Mona Chalabi, have used data visualisations as a means of *shifting perspectives*. Chalabi, a contributor to *The New York Times*, won the 2023 Pulitzer Prize for Illustrated Reporting and Commentary for her work which utilised clever and thought-provoking visual depictions to help readers comprehend the sheer magnitude of Jeff Bezos’ wealth (see Figure 6 for an example).

Figure 6

A Visualisation of Jeff Bezos’ Land from Mona Chalabi’s Award-Winning work



For decades, scholars have mused about the relationship between one’s ability to visualise data (either internally or externally) and critical thinking. Walkup (1965) believed that the ability of one to “visualize in

manipulatable images” (p. 39) was a talent which bestowed a great advantage to those capable of harnessing it, and – not coincidentally – one which connected countless great thinkers across disparate domains.⁵

Supporting Walkup’s claims concerning the value of “visual thinking,” writer and perceptual psychologist Rudolf Arnheim (1980) contended that not only is the notion that this form of internal perception is inferior to thinking in its contribution to cognition misguided, but also that the mere notion that the two phenomena can be neatly disentangled and thus categorised as mutually-independent is itself flawed. To demonstrate, he borrows a thought experiment (described as a “mental puzzle”) from Walkup (1965) in which a reader is asked to imagine a large cube made up of twenty-seven smaller cubes (arranged in the same manner as that of a Rubik’s Cube⁶) whose entire outer edge is then painted red: how many of the twenty-seven smaller cubes (which form the larger whole) will have precisely three sides that have now been painted red? He continues on to elucidate the likely thought process of the reader as they attempt to solve the problem, noting how the key to success is one’s ability to form and perceive a “visual conception” (i.e., mental image; p. 491) of the object as described. After walking through how one might arrive at an answer via the aid of this visual conception, Arnheim (p. 492) muses:

“Was it seeing or was it thinking that solved the problem? Obviously, the distinction is absurd. In order to see we had to think; and we had nothing to think about if we were not looking. But our claim goes farther. We assert not only that perceptual problems can be solved by perceptual operations but that productive thinking solves any kind of problem in the perceptual realm because there exists no other arena in which true thinking can take place.”

Data visualisation is a medium by which one can visually represent quantitative data, but, perhaps more importantly, one via which, through the simple exercise of visualising it, they can gain a *better understanding* of this data (i.e., data visualisation as a means to conduct “exploration and discovery”; Friendly & Denis, 2001).

One of the most brilliant illustrations of how data visualisations can serve as a valuable tool for not only portraying data, but for identifying useful patterns and drawing meaningful conclusions, comes from Dr. John Snow, who, in the midst of a devastating cholera outbreak besieging central London in September of 1854, opted to plot the location of cholera deaths and, critically, the location of water pumps (Figure 7; Snow, 1855). After examining the map, Snow noticed a particularly high occurrence of deaths in the vicinity of the Broad Street pump and, conjecturing that the outbreak might be originating from a water

⁵ For example, he cited how Michael Faraday was able to “see” the complex composition of electromagnetic fields; how August Kekulé’s discovery of the structure of benzene was derived from a vision he had in a dream of a snake eating its’ own tail; how prolific mathematician Jacques Hadamard had described how his thinking consisted exclusively of visual representations of the elements of the problems and theories upon which he worked.

⁶ The Rubik’s Cube was not invented until 1974, and thus this description was not used in Walkup’s original paper (nor in the 1980 reproduction by Arnheim), but provides a useful reference for modern readers.

contamination there, petitioned to have the pump handle removed. Snow's insight, derived from a uniquely effective data visualisation, halted an epidemic which had taken more than 500 lives (Tufté & Robins, 1997).

Figure 7

Dr. John Snow's Map of Central London's 1854 Cholera Outbreak

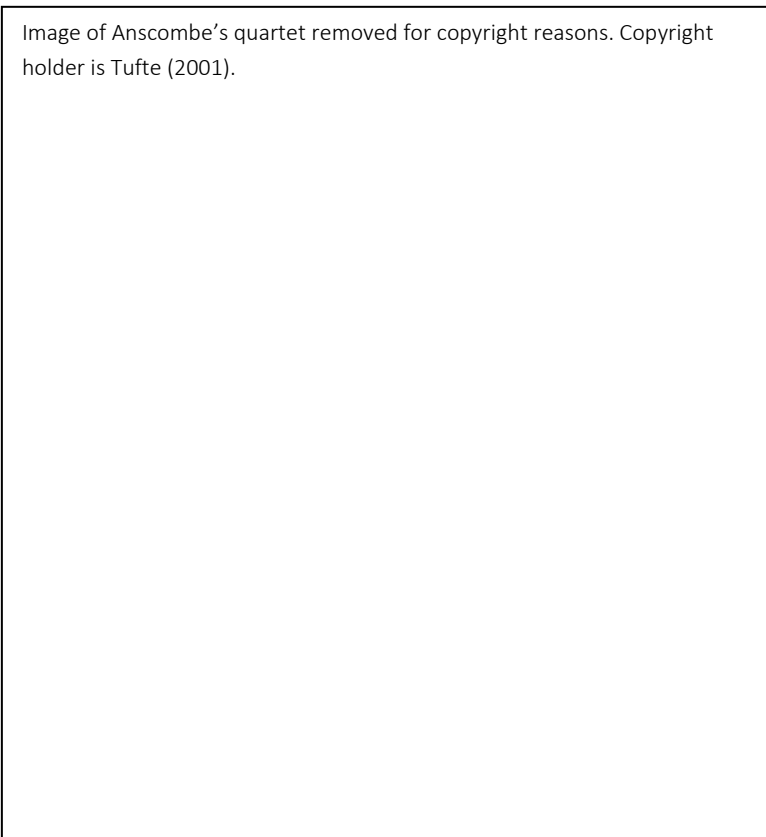


Note. Snow represented deaths from the disease using a black dot and neighbourhood water pumps (due to the disease's ability to be transmissible through water) using a cross and, after examining the visualisation, noticed that an inordinate number of deaths seemed to be occurring near the Broad Street pump. The realisation (and consequent removal of the handle from the Broad Street pump) likely saved many lives.

Moreover, as opposed to merely being *complementary* to conventional statistical computations, there are certain circumstances under which data visualisations are a *superior* (and perhaps *necessary*) means to truly understanding the composition of data sets. For example, Tufté (2001) cites a quirky but illustrative mathematical exemplar known as Anscombe's quartet (see Anscombe, 1973), which features four datasets that have identical descriptive statistics (i.e., the same mean, standard deviation, and regression line used to capture the data). In such a circumstance, devoid of any visual depiction, one might be inclined to think that the arrangement of the data which comprises the four data sets would be indistinguishable from one another, but a glimpse at the visualisations shows that their respective configurations are qualitatively quite different (see Figure 8).

Figure 8

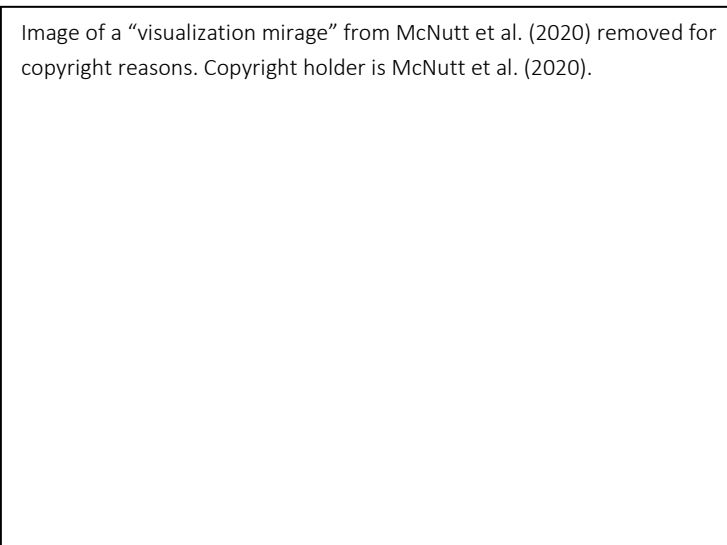
Illustration of Anscombe's Quartet



As Anscombe's quartet neatly illustrates, data visualisations, far from merely an auxiliary source of analysis, are often a necessary component of truly capturing the form and substance of data as they permit an audience to see underlying structures and patterns often not discernible in written or tabular formats. While some seemingly-identical data can, once visualised, be shown to be very different, conversely, some very different data can be cast as identical with inappropriate or deliberately manipulative visualisation efforts. Particularly germane to the present research agenda, McNutt, Kindlmann, and Correll (2020) point out how multiple different data sets, each with unique patterns of distributions (which, presumably, would promote different inferences and conclusions about the included groups) will nevertheless all be depicted as a single, identical bar chart (Figure 9), underscoring the problematic nature of bar charts as they relate to concealing the underlying structure of responses and oversimplifying complex data (see Observation 2 later in this chapter).

Figure 9

Example of a “Visualization Mirage” (from McNutt et al., 2020)



In light of the rapidly evolving dynamics of information dissemination, data graphics is becoming an important new medium of persuasion and expression. With the exponential growth of the Internet over the past several decades and the advent of popular information sharing platforms like social media, we are operating in an era where data visualisation will likely enjoy unprecedented relevance (Kennedy & Engebretsen, 2020). The growing importance of data visualisation within the modern media environment was captured succinctly by Weber (2020, p. 307) who writes:

“...the nature of journalistic storytelling is changing enormously, and data visualization is shaping this change. This change will affect how we shape facts, communicate news, and share knowledge in society in future.”

Corroborating this contention in the context of knowledge acquisition and decision-making, Kennedy, Hill, Aiello, and Allen (2016) note that “data are becoming increasingly valued and relied upon, as they come to play an ever more important role in decision-making and knowledge about the world” (p. 715).


Unfortunately, not all depictions of data can be trusted equally. While, at their best, data visualisations create a persuasive and accessible narrative (Pandey, Manivannan, Nov, Satterthwaite, & Bertini, 2014) via the dissemination of accurate and easily comprehensible depictions of trends, relationships, and key takeaways within a particular dataset, at their worst, data visualisations can be used to deliberately manipulate an audience, skewing facts and figures in such a manner as to artificially engineer patterns that don’t exist or purposefully obscure those that do.

Manipulating data in the hopes of fitting a particular political agenda or supporting a predetermined narrative is far from a modern issue. Darrell Huff discussed the problematic nature of communicating with

data in his infamous 1954 publication “How to Lie with Statistics.” But in an age where misinformation pervades popular media outlets and disinformation campaigns run rampant (Jerit & Zhao, 2020; Guess & Lyons, 2020), there is no shortage of outright manipulation from sources engaging in data distortion via unscrupulous visualisation practices such as the utilisation of disproportionate comparison sizes, the covert usage of non-linear scales, and the reorientation (and, in extreme cases, complete inversion) of y-axes (Pandey, Rall, Satterthwaite, Nov, & Bertini, 2015).

One of the most infamous examples in recent memory was a graphic depicting the number of murders committed by firearm each year in Florida (Figure 10). The graphic explicitly draws the viewers’ attention to the data point depicting the year 2005 – the year when Florida enacted its’ controversial “Stand Your Ground” law. The law made the use of deadly force in self-defence legally justifiable, and notably removed the legal burden concerning “duty to retreat,” which stipulated that an individual is not permitted to use deadly force if it is possible to remove oneself from the danger entirely by the act of retreating. While a cursory view of the overall trendline within the graphic might lead a casual viewer to conclude that gun deaths *declined* following the enactment of the new law (as is ostensibly evidenced by the pronounced “drop” in the data in the years following the year 2005), a more astute viewer will notice that the y-axis is actually *inverted*, meaning data points lower on the y-axis indicate *higher* numbers of deaths.

Figure 10
Graphic Depicting Yearly Murders via Firearms in Florida using an Inverted Y-Axis



Data visualisations are a double-edged sword: simultaneously a tool to inform and enlighten (when yielded by a capable and responsible actor) as well as a weapon to distort and mislead (when employed by incompetent or unscrupulous agents of influence). And while the decision to use data presentation as a

vehicle to either educate or manipulate is undoubtedly an important determinant of the way in which it will influence an audience's perception, we argue that similar impacts may also be achieved – whether appreciated by the decision-maker or not – by the mere *choice* of which data presentation formats to use.

Data Visualisations and Summary Statistics as a Form of Framing

Ultimately, data visualisation is a *choice* (or, more precisely, a series of choices). The first choice is to opt to visualise data in the first place. However, once that decision has been made, one then needs to determine what kind of visualisation(s) to use and how to construct those visualisations.

Some of these choices necessarily come at the expense of alternative options. If one can only include one graphic for an article or paper, and they opt for a bar chart, they have also simultaneously opted to *not* use a line graph, a pie chart, etc. Now, if the underlying data used to compose each of these types of visualisations is the same, it's fair to wonder whether such a choice is truly consequential. Should the manner in which identical data is displayed elicit meaningfully different responses from an audience? To respond to such an inquiry, it's perhaps appropriate to turn to a domain which poses parallel questions: the field of message framing.

Message framing (and the resultant phenomenon of *framing effects*; Druckman, 2001a) refers to the practice of making small, often seemingly insignificant changes to the way in which an issue or event is presented and thereby yielding noticeable changes in opinion or preference (Chong & Druckman, 2007). While a number of novel framing tactics have emerged in the past few decades (e.g., metaphoric framing and “moral values” framing; Thibodeau & Boroditsky, 2011 and Feinberg & Willer, 2015, respectively), theorists tend to broadly categorise framing efforts in two distinct manners: as either *emphasis* framing or as *equivalence* framing (Schaffner & Sellers, 2009). Emphasis framing (e.g., Nelson, Clawson, & Oxley, 1997) refers to the manner in which one might emphasise two distinct aspects, arguments, or other “alternative consideration[s]” (Druckman, 2011, p. 282) from a single issue or event (e.g., how those on one side of the abortion debate might emphasise an unborn child's right to life, while those on the other side might emphasise a woman's right to choose; McCaffrey & Keys, 2000). Meanwhile, equivalence framing (Druckman, 2001a) is more narrowly defined as the competing manners in which one might present identical data or information; what Druckman (2011, p. 282) refers to as “semantically distinct but logically equivalent portrayals [of information].” For example, in describing the outcome of a coin flip, one might say it “landed on heads” or “did not land on tails.” If we are to assume a coin flip only has two possible outcomes, then each of the descriptions offers the same “truth value” (i.e., each captures an identical state of the world; Kahneman, 2011).

Framing effects have received extensive attention in the literature in part because they seem to violate assumptions of rational choice theory and the so-called "rational actor model," upon which much of classical economic treatise is based. These models maintain that choices are driven by considerations of optimisation and utility maximisation and thus should not be susceptible to the superficial manipulations relating to the presentation of choice (e.g., Tversky & Kahneman, 1989⁷). However, a multitude of experiments have demonstrated that humans are unexpectedly vulnerable to framing effects. Classic examples include Tversky and Kahneman (1974) demonstrating that presenting the same policy proposal as either a *gain* or a *loss* can dramatically shift individual perceptions of (and preferences for) the policy, as well as Tversky and Kahneman (1989) showing that surgical outcomes, framed either via their *survival* probability or *mortality* probability, can significantly influence one's preference to accept or decline surgical intervention. More recently, framing has been shown capable of exerting significant influence on areas as disparate as consumer attitudes (e.g., Mitchell, Bakewell, Jackson, & Heslin, 2015), patient responses to medical advice (e.g., Glare, Fridman, & Ashton-James, 2018), the motivation to pursue health behaviours (e.g., Gallagher & Updegraff, 2012), the intention to exhibit pro-environmental preferences (e.g., Homar & Cvelbar, 2021), and the reduction of reckless and anti-social behaviour (e.g., Gerend & Cullen, 2008).

Most framing effects are engendered due to the (seemingly insignificant) *choices* a communicator makes: whether to describe something as a loss or a gain; whether to display a figure as a frequency or a probability (Lotto, Tasso, Gavaruzzi, Carnaghi, & Rumiati, 2014); whether to format a statistic in relative or absolute terms (Hux & Naylor, 1995). Similarly, we contend that it is not inconceivable to believe that the choices one makes when selecting and customising data visualisations – insignificant as they may seem – may also be capable of triggering downstream consequences for comprehension, perception, and attitude crystallisation (or change). Just as a survival frame and a mortality frame can produce disparate judgments despite conveying the same underlying data, so, too, perhaps, could a bar chart and an alternative mode of visualisation (e.g., histogram) – derived from an identical data set – produce divergent perceptions about the state of political polarisation.

Moreover, an analogous phenomenon could emerge with regards to the use of *summary statistics*. Just as how data visualisations attempt to visually capture and convey the overall patterns of a data set, summary statistics are used as a way to succinctly summarise the primary values of interest within a data set. These values often include (but are not limited to) measures of central tendency (e.g., mean, median, mode) and measures of spread or dispersion (e.g., range and standard deviation). And similar to how the choice of

⁷ It should be noted, however, that in a series of rebuttals of Kahneman and Tversky's claims that framing effects expose the bounds of human rationality, researchers like Gerd Gigerenzer have forwarded counter-claims that certain instances of supposed "irrationality" are instead, in fact, examples of *ecological rationality*, wherein actors consider the context and circumstances under which the decision takes place, which can sometimes generate rational choices which may be mislabelled as irrational (Todd & Gigerenzer, 2012).

visualisation could feasibly impact perceptions of political polarisation, so, too, could the choice of summary statistic used.

When only one summary statistic is used to summarise a data set, one could argue it's tantamount to emphasis framing, as one aspect of the data is necessarily highlighted at the expense of other aspects (which are omitted). For example, a hypothetical data set containing Republican and Democrat opinions on abortion is most commonly summarised via the group means (e.g., "the average response for Republicans on a 1 to 7 scale was 3.1, while the average response for Democrats was 5.7"). However, while such a summary is not inaccurate, it A) is incomplete, and B) accentuates certain elements of the data while ignoring others. In regards to the former, it does not provide any information on the *variance* of responses within groups, which at best leads one to concede they know nothing about the level of intragroup attitude heterogeneity and at worst permits them to rely on biases which *presume* high levels of homogeneity for out-groups (e.g., Judd et al., 1991). As for the latter, employing group means as the sole summary statistic may necessarily accentuate the degree of group *difference*, whereas alternative summary statistics (e.g., *percentage of common scores*; Hanel et al., 2019) may instead permit the emphasis of the degree of group *similarity*.

Ultimately, we liken the choices one makes during the process of preparing data presentations (i.e., both visualisation and summarisation) to the choices made which ultimately produce framing effects, which routinely yield meaningfully different responses to the same source material in spite of their purportedly trivial nature.

The Overarching Observations Motivating Our Research

The research contained herein was motivated by a trio of broad observations situated at the intersection of data presentation, perception modulation, and polarisation.

Observation 1: In the Context of Political Perception, Choice of Data Presentation Format Is A Significant Decision

Cairo (2020) forwards the perspective that data visualisations represent a new *technology*, and like any technological innovation, they have the potential to fundamentally shift how we "see and relate to reality" (p. 17). The notion of how individuals see and relate to reality as a consequence of data presentation choice represents the crux of this investigation, and thus it is important to understand *if* and *how* certain data presentation formats may alter audience processing, perception, and inference.

While it is likely self-evident that outright forgery of data or blatant manipulation of visualisations (such as an inverted y-axis; Figure 10) will inevitably skew the perception of the data for an audience, over the past few decades, there have been a number of multidisciplinary efforts aimed at investigating the impact of data

visualisation components and techniques on perception⁸, with a special effort to identify best (and worst; e.g., Wainer, 1984) practices. As a result of these efforts, scholars have gleaned a number of important insights concerning how even minor shifts in layout and design can impact the experience and interpretation of a viewer.

Lauer and O'Brien (2020) suggest that it was the work of Pandey et al. (2014) that was "first to empirically show that participants are likely to be misled in their interpretations of data visualisations that employed deceptive tactics such as message reversal (e.g., an inverted axis) and message exaggeration (e.g., a truncated y-axis)" (p. 329). In the referenced research, Pandey and colleagues (2014) found that while the potency of visualisation-based persuasion was dependent on attitudinal priors, there emerged "consistent results that charts lead to higher persuasion when participants do not possess a strong initial attitude about the topic" (p. 2219). In related research, Pandey et al. (2015) showed that common deceptive tactics, such as axis or area distortion of charts, can result in overestimations between 58.5% and 129.5% greater than the corresponding control conditions. Lauer and O'Brien (2020) themselves found that several common deceptive visualisation tactics (including y-axis truncation) engendered exaggerated perceptions of group differences across several graph types (i.e., bar charts, line graphs, pie charts, and bubble graphs) relative to the control version of each graph.

The number of presentation decisions and design components that can contribute to altering the manner in which visual data is perceived is so vast that a number of scholars have made ambitious attempts to wrangle and meaningfully categorise them. Lo et al. (2022) sought to synthesise the extant literature on misleading visualisations and correspondingly devised a taxonomy (via open-coding of more than a thousand real-world visualisations reported as deceptive) of misleading visualisation elements. Among the 74 reported issues include the "inappropriate" use of varying chart types (e.g., bar, pie, line, etc.), "hidden distribution," and, most prominently, "truncated y-axis" (which was flagged 155 times in the bar chart category, by far the most frequently cited issue). Dimara, Franconeri, Plaisant, Bezerianos, and Dragicevic (2018) compiled a similar taxonomical account of visualisation issues, theirs focusing more on the cognitive biases that may corrupt accurate estimation, attribution, and recall of visualised data. Work by Bresciani and Eppler (2009) attempted to document a comprehensive list of the "disadvantages and risks" of visualisations. Overall, they list nearly 50 types of risk, each related to one of three types of effects: cognitive, social, and emotional (à la Bürgi & Roos, 2003) – all of which were further sub-divided into "designer-induced" and "user-induced" categories. Effects identified included "misleading" (e.g., Van Wijk, 2006), "hiding/obscuring" (e.g., Wainer, 1984), and "over-simplification" (e.g., Nicolini, 2007). Additionally, Nguyen, Jung, and Gupta (2021) couch a

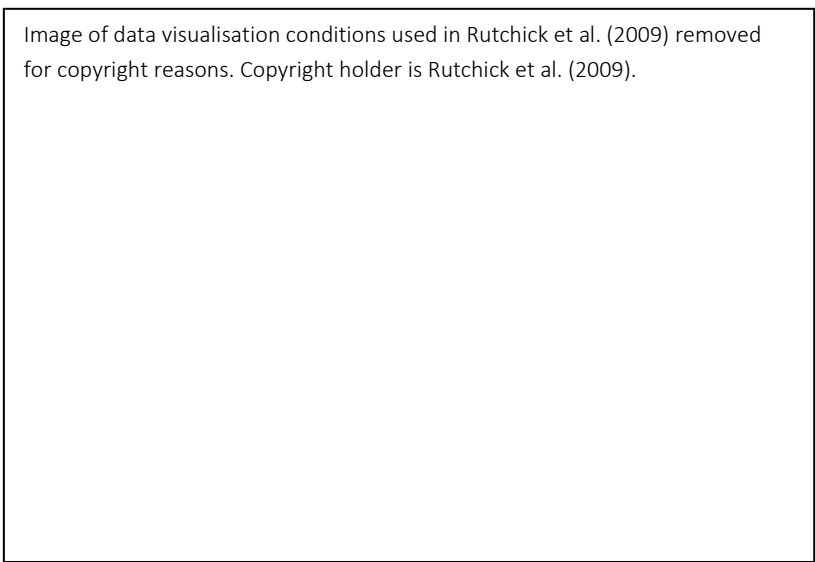
⁸ Although, it is worth noting that, even as late as the mid-2000s, scholars noted that there existed few *systematic, empirically-rigorous* investigations into the perception of quantitative information (Feldman-Stewart, Brundage, & Zotov, 2007). William S. Cleveland and Robert McGill were cited as notable exceptions, and credited with pioneering many of the more rigorous initial forays into the field.

discussion of the most common visualisation “pitfalls” in the context of visual perception research, noting how data visualisations can mislead or misinform due to myriad subtle variations in colour, size, shape, and orientation.

Although there is little reason to suspect that the impact of minute changes in *apolitical* data presentation would fail to generalise to portrayals of *political* information, such an assumption requires empirical verification. While other researchers have investigated the circumstances and contexts in which data visualisations might prove capable of changing attitudes (e.g., Markant, Rogha, Karduni, Wesslen, & Dou, 2022; Pandey et al., 2014), few have directly focused on how certain depictions of data may affect beliefs about the *extent* of political polarisation (Alieva, 2023; Santos et al., 2017). However, this sub-domain is not without precedent. For example, Rutchick, Smyth, and Konrath (2009) found that electoral maps depicting state-level support for a presidential candidate with a *shade* of purple along a red and blue continuum (in proportion to the degree of support for the Republican or Democratic candidate) reduced stereotyping as well as perceptions of political division compared to the traditional visualisation technique which uses a *binary* categorisation of “red and blue” states based solely on which party carried the electoral college votes (Figure 11).

Figure 11

Four Visualisation Conditions featured in Rutchick et al. (2009)



Such an effect was predicted several years earlier by Seyle and Newman (2006), who feared that the popularisation of the term “red and blue states” created a metaphoric vehicle by which members of the general public could neatly categorise (and necessarily oversimplify) “us versus them,” precipitating a host of negative social identity processes related to stereotyping and intergroup conflict. The authors made particular note of the simplified red and blue maps which sought to capture the metaphoric landscape,

warning that such a dichotomous visual distinction was likely to produce misperceptions due to its misleading implication of ideological uniformity within the states.

Overall, considering the way in which minute, seemingly inconsequential tweaks in data presentation can influence perception – including within the political domain – we believe it’s fair to assume that such effects may be capable of exerting sway over beliefs about polarisation, as well.

Observation 2: Some Traditional Formats of Data Presentation May Reinforce Political Misperceptions


A large portion of our research was inspired by the work of Hanel, Maio, and Manstead (2019), who, while conceding that rival social, cultural, and political factions may undoubtedly be different in notable ways (e.g., their worldviews, values, cognitive proclivities, policy preferences, etc.; e.g., Deppe et al., 2015; Graham, Haidt, & Nosek, 2009), question whether the *manner* in which we have chosen to historically *display* these differences might contribute to how *fundamentally incompatible* we now perceive “others” to be (e.g., Iyengar et al., 2012).

Broadly, Hanel and colleagues argue that the field of psychology tends to place a disproportionate focus on identifying and publicising statistically significant *differences* between groups (Fanelli, 2010; Open Science Collaboration, 2015), with studies that “fail” to find such differences often being discarded, contributing to the proverbial “file-drawer” problem (Rosenthal, 1979). Relatedly, when experimental results are displayed, the predominant method of doing so within academic journals is via bar graphs depicting *mean group differences* (Lane & Sándor, 2009). While not necessarily misleading, per se, a habitual and unilateral focus on differences necessarily comes at the expense of the examination of arguably equally-relevant similarities (see Hanel & Wolf, 2020). Critically, Hanel et al. (2019) found that simply changing the *type* of visualisation used to depict identical data had a significant impact across an array of intergroup assessments, with modes that more effectively highlighted similarity information (e.g., superimposed normal distributions) resulting in both more accurate intergroup perceptions as well as more positive outgroup appraisals. In light of this, we believe it is fair to A) continue to interrogate the potentially negative impact that an over-reliance on group mean information (either via bar charts or summary statistics) might have on political perception (Observation 2) and B) consider whether alternative means of data visualisation and summarisation strategies may generate more positive outcomes (Observation 3).

The development of the first known bar chart is credited to Scottish political economist William Playfair, who remarkably also pioneered several graphical standards which remain ubiquitous in data visualisation to this day, such as the time-series line graph and the pie chart (Spence, 2006). The bar chart, which depicted the exports and imports of Scotland for the period between Christmas 1780 and Christmas 1781, was featured in Playfair’s *Commercial and Political Atlas* (1786) and can be seen below (Figure 12):

Figure 12

The First Known Bar Chart, from William Playfair's "Commercial and Political Atlas" (1786)



However, Playfair himself bemoaned the inferiority of the bar chart relative to time-series plots (which comprised 43 of the 44 visualisations in the first edition of the *Commercial and Political Atlas*; Tufte, 2001), noting that the development and inclusion of the bar chart was merely a product of necessity (as he did not have complete year-to-year data to construct a time-series plot). In speaking of the bar chart, he wrote:

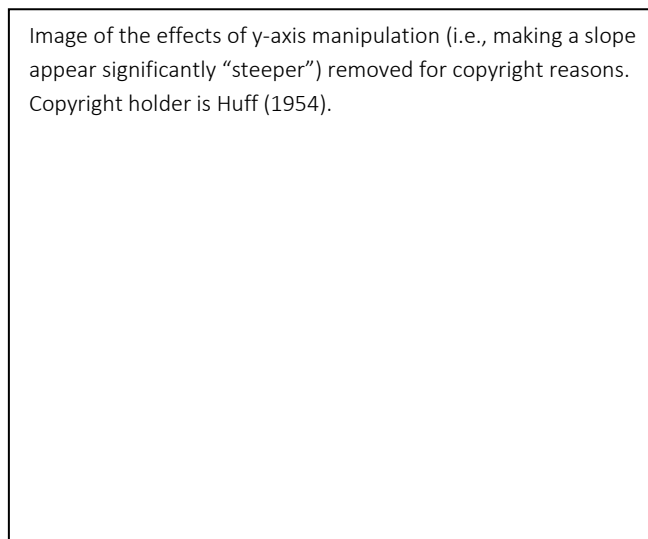
“...it does not comprehend any portion of time, and it is much inferior in utility to those that do; for though it gives the extent of the different branches of trade, it does not compare the same branch of commerce with itself at different periods; nor does it imprint upon the mind that distinct idea, in doing which, the chief advantage of Charts consists: for as it wants the dimension that is formed by duration, there is no shape given to the quantities” (Playfair, 1786, p. 101).

Beyond Playfair’s commentary, a larger concern associated with this format in the context of data visualisation is how easily information depicted via bar chart can be manipulated. A common method of manipulation is the practice of *truncating* (or shortening) the y-axis. By reducing the scale length of a graph, one is effectively “zooming in” on the data, which has the effect of magnifying otherwise minute differences or gradual shifts in trends. As an example of the effect truncation can have on the appearance of data (and, consequently, the “story” data tells), Huff (1954) cites an incident from a 1938 publication called “Dun’s Review” wherein an editorial writer reproduced a chart purporting to show a dramatic increase in government payrolls, with the (seemingly) image-appropriate headline “Govt. Pay Rolls Up!” The Dun’s

Review staff member then altered the scale of the y-axis and juxtaposed the new chart with the old one, complete with a revised (and *still* image-appropriate) headline “Govt. Pay Rolls Stable!” (see Figure 13).

Figure 13

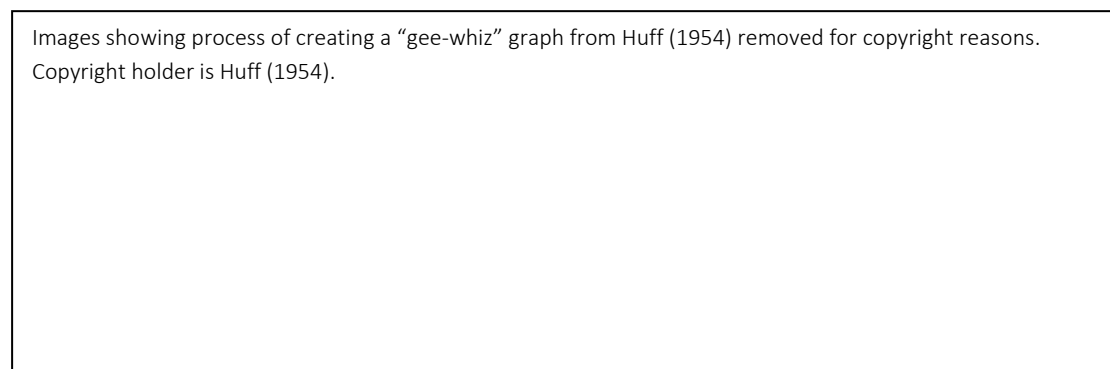
Example of the Effects of Y-Axis Manipulation (featured in Huff’s “How to Lie with Statistics”)



Huff (1954) facetiously refers to these visualisation efforts as creating “Gee-Whiz!” graphs, or graphs distorted so as to make a trend, data point, or difference seem more impressive than it may actually be (Figure 14).

Figure 14

Huff’s (1954) Three-Step Process for Creating a “Gee-Whiz” Graph




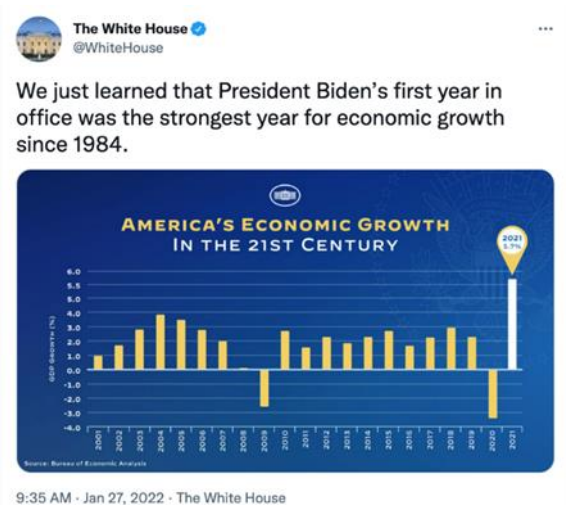
Note. Huff explains that all one must do is begin with a standard line graph (left), “chop off the bottom” (centre), and finally “change the proportion between the ordinate and the abscissa” (e.g., “stretch” the chart in such a manner that it covers the same space as the original, but now the range of the y-axis measures just 2.0 as opposed to the original range of 24; right).

Naturally, y-axis manipulation has become a widely-used tactic by political entities as a means of forwarding their preferred narratives or conclusions. To underscore the bipartisan nature of this practice, below are two

examples from opposite sides of the political aisle. The left image (in Figure 15) is from the (Biden) White House’s official Twitter account, wherein the creators of the graphic have inserted an additional scale point (i.e., 5.5), effectively making the distance between 5.0 and 6.0 double the size of any other 1-point interval in an effort to “stretch” the bar representing 2021’s level of economic growth. The right image (in Figure 15) is from the FOX News program *Tucker Carlson Tonight*, wherein the creators have truncated a 100-point response scale to one-fifth of its original range, making a 12 percentage point decline in self-professed Christian identity look like a much more substantial drop.

Figure 15

Examples of Y-Axis Manipulation by Different Sources

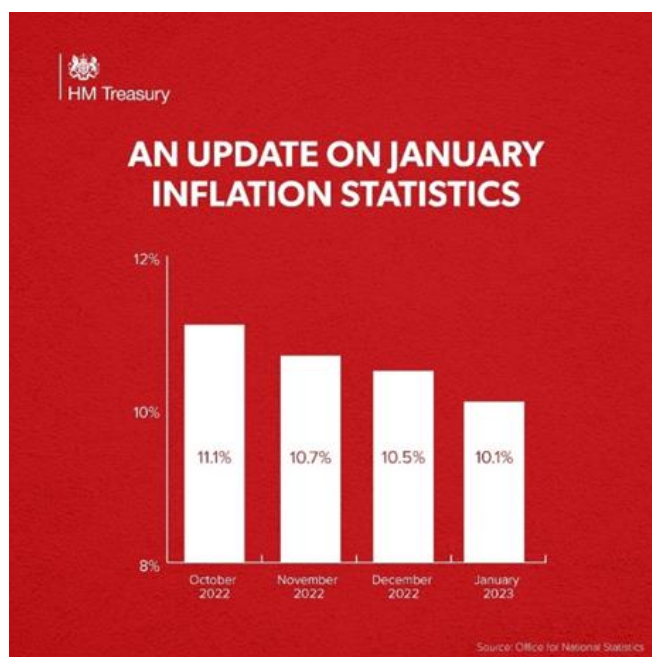


Once again, these questionable visualisation practices – while heavily featured within the US – are not constrained to US politics. A similar example of arguably inappropriate truncation (Figure 16) was originally posted by the official Twitter account of HM Treasury (i.e., @HMTreasury), the UK Government’s Economic and Finance Ministry, on February 15th, 2023.⁹ As the title of the graph indicates, it depicts the country’s monthly inflation statistics, from October 2022 through January 2023. While the bar chart shows a steady, almost linear decline – a reassuring sight for those whose economy, like so many across the world throughout 2022, had been battered for months by persistent, runaway inflation – some took issue with the manner in which the visualisation was crafted. One individual in particular, Ed Humpherson, Director General of the UK Statistics Authority, authored a letter (published on the Office for Statistics Regulation website) to David Pares, Head of Communication at HM Treasury, formally condemning the graphic (Humpherson, 2023). In his letter, he states that “the graph gives a misleading impression of the scale of the deceleration in inflation, due to the y-axis beginning at 8%.”

⁹ <https://twitter.com/hmtreasury/status/1625826967688806400>

Figure 16

An Example of Y-Axis Truncation from the Official Twitter Account of HM Treasury (the UK Government's Economic and Finance Ministry)



Although excessive truncation (or, arguably, *any* truncation) can be argued to be misleading relative to an unrestricted depiction,¹⁰ troublingly, rules about the degree of y-axis truncation are virtually non-existent for many media sources,¹¹ and many academic journals also take a relatively permissive stance, with minimal prohibitions regulating the practice. With few official guidelines, the degree to which y-axes are restricted is left largely up to the discretion of those who publish the data. This allows for considerable flexibility in graph construction – a potentially problematic reality for those concerned with preserving the integrity of data.¹²

However, it's not just manipulations of the y-axis that can cause data presentation formats to (intentionally or otherwise) become misleading. Though not necessarily *manipulative*, per se, certain types of statistics may be classified as more or less *appropriate* for particular purposes based upon the information they provide and, critically, the information they may conceal. Means (or averages) are a good example of this type of potentially misleading statistic.

¹⁰ Huff (1954), for instance, when commenting on the practice of truncation, writes that “a truncated bar chart has, and deserves, exactly the same [bad] reputation as the truncated line graph.”

¹¹ Although Humpherson notes, in the inflation example above, that the UK's Government Analysis Function provides clear guidelines of data visualisation best practices, which includes ensuring that the y-axis “always start at zero” when utilising a bar chart.

¹² Even when visualisations *are* regulated, Tufte (2001) concedes that there is inevitably a degree of relativism embedded in the assessment of graphical distortion, musing how it is impossible to determine whether distortion should be measured strictly based on the physical features (e.g., dimensions, relative sizes, etc.) of the graphic *as they are objectively composed* (i.e., “physically measured on the surface of the graphic”; Tufte, 2001, p. 55) or on how they are *subjectively perceived* (i.e., “the perceived visual effect”; Tufte, 2001, p. 55). Though it imperfectly addresses this philosophical quandary, he proposes a measure referred to as the “Lie Factor,” which is calculated as the size of the effect as shown in the graphic divided by the actual size of the effect in the data, noting that if the resulting calculation falls 0.05 below or above 1, the graphic can be categorized as “[substantially] distorted” (Tufte, 2001, p. 57).

Means are one of the most widely used statistics. This can likely be explained by the fact that means are simple to understand (i.e., even individuals possessing limited statistical literacy can grasp what a mean conveys) and they provide a useful “broad-stroke” summary of large swaths of data. However, means – especially when serving as the *sole* summary statistic for a set of observations (i.e., not complemented by alternate measures of central tendency or accompanied by other statistical measures such as standard deviation or range) – can provide an incomplete (and potentially misleading) picture of a dataset. For example, a phenomenon sometimes colloquially referred to as the “Bill Gates effect” (Curbal, 2016) describes how the exorbitant wealth of an ultra-high net worth individual can radically skew the mean net worth of any group of which they happen to find themselves a part. Imagine a bar filled with ten patrons, each of whom possesses a net worth of exactly \$50,000. In such a circumstance, naturally, the average of all ten of the bar’s patrons would also be \$50,000. But should one patron leave and be replaced by Bill Gates – worth roughly \$113.4 billion at the time of writing – suddenly the average net worth of the pub patrons skyrockets to more than \$11 billion. In such a situation, few would maintain that the average provides an accurate reflection of the *actual* summary of net worth for most patrons present in the bar (here, a measure such as the mode or even the median would serve as a more appropriate summary statistic), but yet this use of the group mean is, though perhaps misleading, not *explicitly* manipulative (merely, one could argue, *inappropriate*).

A similar issue with means can be observed in the manner in which they fail to capture the *range* or distribution of data. In expounding on their problematic nature, Huff (1954) notes how the average temperature in Oklahoma City is a pleasant 60.2 degrees Fahrenheit – a figure which might entice campers looking for a comfortable location to spend the day. However, what the average fails to disclose is that, depending on the time of year, Oklahoma City temperatures can get as high as 113 degrees Fahrenheit and as low as -17 degree Fahrenheit; a range of 130 degrees which was otherwise hidden by the publication of only the pleasant average (see Figure 17).

Figure 17

Depiction of How Averages Can Conceal Critical Information Such as Range of Outcomes

Image of how average temperatures can conceal important information about temperature ranges. Copyright holder is Huff (1954).

A critical component of the observation is the role that rapid *impressions* seem to play in perception. William Playfair broached the topic of impressions in his claim regarding why graphical representations of data were superior to their written or tabular counterparts. He was specifically concerned with the inability of tabular data to forge a lasting impression, and contended that this is where graphical depictions may exert the greatest comparative advantage, noting:

“...a man who has carefully investigated a printed table finds, when done, that he has only a very faint and partial idea of what he has read; and that like a figure imprinted on sand, is soon totally erased...[data] are capable of being as easily represented in drawing...though, till now, it has not been attempted. Upon that principle these Charts were made; and, while they give a simple and distinct idea, they are as near perfect accuracy as is any way useful. On inspecting any one of these Charts attentively, a sufficiently distinct impression will be made, to remain unimpaired for a considerable time, and the idea which does remain will be simple and complete” (Playfair, 1786, p. 3-4).

The concept of impressions is one which should be given serious consideration in the domain of perception as research has demonstrated that individuals generate near-instantaneous impressions of social phenomena (e.g., Willis & Todorov, 2006). The reflex to form rapid impressions is a critical piece of our social cognition repertoire, and is likely an evolutionary adaptation which operates largely under subconscious

control (Fiske, Cuddy, & Glick, 2007; Uleman & Kressel, 2013). Critically, research has shown that such impressions also extend into the political domain, wherein a single second of exposure to candidate pictures is enough to generate immediate judgments of candidate competence which, in turn, can predict the outcomes of US congressional elections at a rate significantly better than chance (Todorov, Mandisodza, Goren, & Hall, 2005).

Huff (1954) points out that the danger of practices like y-axis truncation is that, technically speaking, “nothing has been falsified – *except the impression that it gives*” (p. 61; italics added). In light of this, communicators need to take great care in considering the impressions certain forms of data presentation might create. We believe, for instance, that bar charts – and especially truncated bar charts – are likely to promote greater impressions of *difference*¹³; that is, upon exposure, a viewer’s first reaction is likely to take note of the degree to which the groups diverge.

We argue that impressions derived from bar charts (and, relatedly, from using group means as standalone summary statistics) promote an oversimplification of group opinions which necessarily encourages a disproportionate focus on group difference. Such an oversimplification of cross-party beliefs coupled with the highlighting of group difference may risk facilitating *essentialised* perceptions (see Haslam & Whelan, 2008). Essentialising could be conceptualised as either a cause or effect of categorisation (Gelman & Rhodes, 2012), and refers to the tendency to perceive members of a particular category as having a “fixed, underlying nature” (Haslam, Bastian, Bain, & Kashima, 2006); as comprising of distinct *essences* compared to those in other categories, making their category membership (and even the categories themselves) a natural phenomenon as opposed to a human construction (Gelman & Rhodes, 2012; Brick, Hood, Ekroll, & de-Wit, 2020). Beliefs in “natural kinds” have been found to amplify group differences (Rothbart & Taylor, 1992) and the attribution of “essences” to outgroups has been linked to classic conceptions of prejudice (Allport, Clark, & Pettigrew, 1954). Moreover, research has suggested that essentialist beliefs (Haslam, Bastian, & Bissett, 2004) appear to be a necessary component for certain forms of dehumanisation (Haslam, 2006).

Evergreen (2019) cautions readers that certain visualisation formats, while not intentionally manipulative, may not be *appropriate* for certain pieces of data. Ultimately, we believe that bar charts and group means (as standalone summary statistics), although not necessarily *inherently* misleading or manipulative in isolation, may be *inappropriate* ways to present intergroup data in a highly-polarised political ecosystem due to their inability to capture any information pertaining to intergroup *similarity* or intragroup response diversity. Such presentation formats – due to their disproportionate focus on group differences, misleading

¹³ In fact, with the exception of chart elements like standard error bars (which report some semblance of response range and variation), one could argue that bar charts *only* show difference.

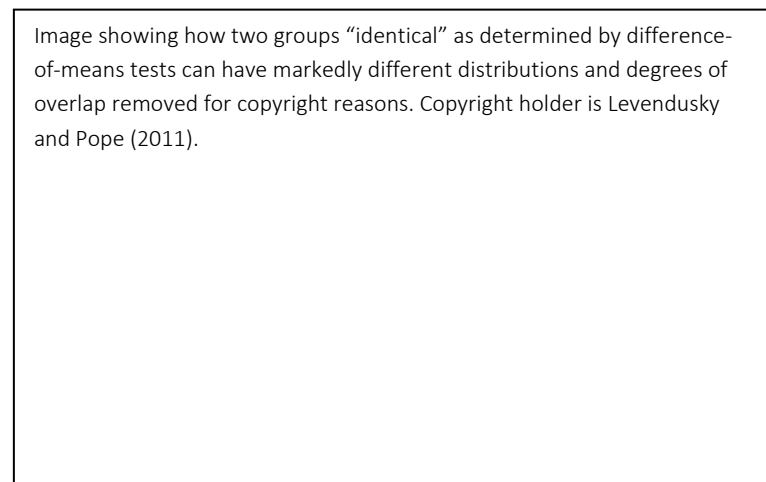
oversimplification of group opinion homogeneity, and exaggeration of group distinctiveness¹⁴ – may serve to reinforce (and perhaps *exacerbate*) perceptions of polarisation.

Observation 3: Certain Non-Traditional Data Presentation Formats May Mitigate Misperceptions of Political Polarisation

In their aptly titled paper “Red and Blue States: Going Beyond the Mean,” Matthew Levendusky and Jeremy Pope (2011) make a compelling case that a significant portion of polarisation research hinges on the analysis of *difference-of-mean tests*. Consequently, if the average opinion of members of one group (e.g., Republicans) differs meaningfully from the average opinion of members of another group (e.g., Democrats), the groups may then safely be categorised as “polarised.” However, such an approach suffers from the fact that A) cohorts of ultra-extreme voters can dramatically skew the group’s mean and B) it fails to capture the “common ground” of the two groups by ignoring the (sometimes strikingly overlapping) distribution of responses (see Figure 18).

Figure 18

Graphic Depicting the Shortcomings of Solely Using Difference-of-Means to Define Intergroup Polarisation (from Levendusky & Pope, 2011)



Note. Unilateral reliance on difference-of-mean tests would conclude that the two pairs above are identically different, but a cursory examination of the distributions makes it clear that the pair on the right has significantly more “common ground” than the pair on the left.

Similarly, Mehlhaff (2023), in an attempt to develop a more accurate model to predict and interpret polarisation, stresses the need for analyses to take into account not just the degree of intergroup heterogeneity, but importantly the degree of *intragroup homogeneity*, as well. This is because, despite the existence of pronounced intergroup heterogeneity, substantial *intragroup* heterogeneity in opinions may

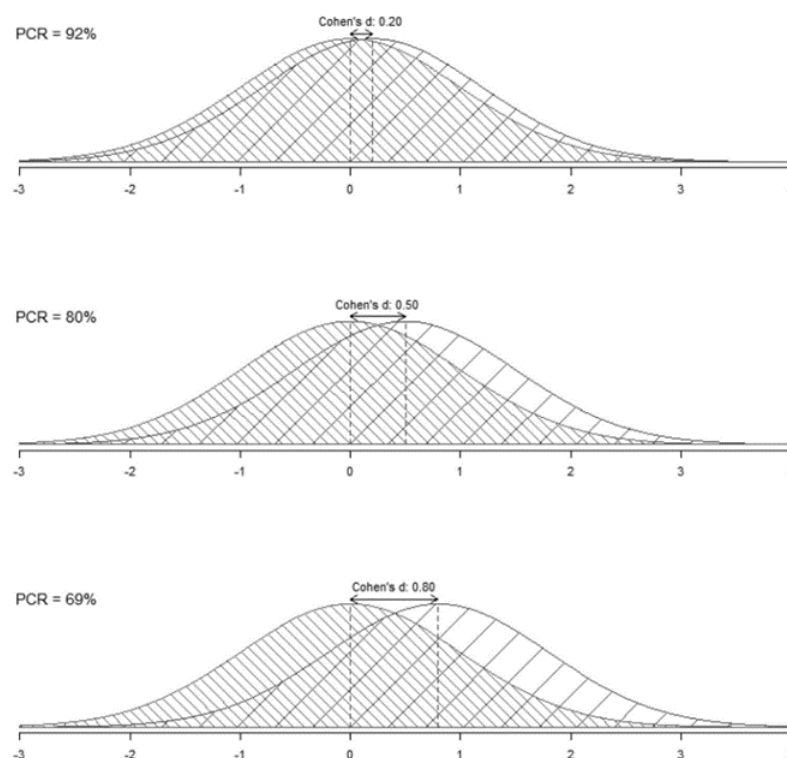
¹⁴ Coincidentally, many of the proposed contributors of false polarisation (Fernbach & Van Boven, 2022).

also be present. Unfortunately, data presentation formats which rely heavily on group mean information may conceal this and, subsequently, neglect the possibility of intergroup opinion overlap.¹⁵ McNutt et al. (2020) employ the phrase “*visualisation mirages*” to refer to visualisations which “superficially convey a particular message” due to their “[hiding] of important details” or their “[failure] to represent uncertainty” (p. 1).

Corroborating the critique offered by Levendusky and Pope (2011), Hanel et al. (2019) describe a statistical measure they refer to as the “percentage of common responses” (PCR), which is derived by taking the overlap coefficient (OVL) – an estimation of the percentage of overlap between two normal distributions derived from a transformation of Cohen’s *d* (Inman & Bradley, 1989) – and multiplying it by 100. They show that even in circumstances where a statistical analysis might reveal a “large” effect size between two groups (e.g., a Cohen’s *d* of 0.80), the PCR is still approximately 69 percent (see Figure 19).

Figure 19

A Visual Representation of the Percentage of Common Responses for Normally Distributed Groups at Different Effect Size Thresholds (from Hanel et al., 2019)



This complex interplay between, on the one hand, statistically-significant *differences in group means* and, on the other hand, substantial levels of *intergroup similarity*, was empirically reproduced by Hanel & Wolf

¹⁵ Or, perhaps more appropriately, the *inevitability* of intergroup opinion overlap, as Levendusky and Pope (2011) assert “If there is a good deal of overlap between red- and blue-state citizens, a claim of polarisation needs serious qualification. And, as will be seen below, the question of overlap very quickly becomes one of degree. There is no case of complete heterogeneity; there are only cases of more and less overlap” (p. 231).

(2020), who showed that despite “Leavers” and “Remainers” of the Brexit referendum achieving statistically-significant group differences across a variety of personality, belief, and behavioural measures (e.g., Becker, Fetzer, & Novy, 2017; Harper & Hogue, 2019; Hobolt, 2016; Swami, Barron, Weis, & Furnham, 2018), when examining the full *distribution* of responses, the two groups still had an average response overlap of 90 percent across measures which examined factors such as prejudice, human values (via the *Portrait Values Questionnaire*; Schwartz et al., 2001), and ingroup identity.

As mentioned, traditional means of data presentation fail to highlight the often substantial levels of intergroup agreement and similarity – an especially unfortunate reality as recent research has found that simply by presenting American political partisans with data that emphasises levels of psychological similarity between the two groups, one can engender more positive feelings towards the outgroup as well as an enhanced belief that common ground could be reached on pressing social issues (Syropoulos & Leidner, 2023). However, alternative options of data presentation exist – ones which may better capture intergroup *similarity* or intragroup opinion *heterogeneity*. For example, in the realm of summary statistics, there are options such as the aforementioned PCR or, relatedly, PCS (i.e., percentage of common scores; Hanel et al., 2019), both of which describe levels of intergroup opinion similarity.¹⁶ In terms of alternative data visualisations, we might examine options such as the aforementioned superimposed normal distributions or icon arrays.

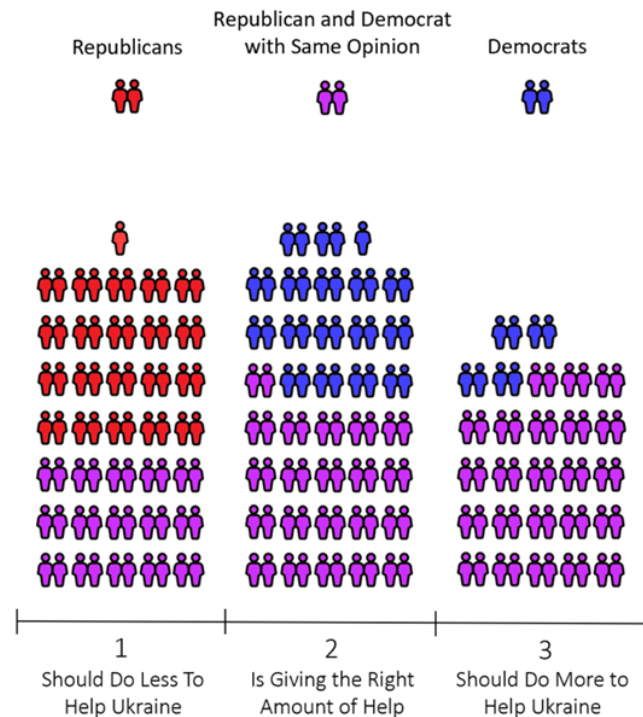
Icon arrays, or sets of icons used to visually depict proportions or the likelihood of incidents within a population (e.g., contracting a disease, experiencing a side effect, etc.), feature prominently in the risk communication literature and have shown promise across a multitude of comprehension and retention metrics. Icon arrays have proven capable of improving the accuracy of risk perception (even in low numeracy populations; Galesic, Garcia-Retamero, & Gigerenzer, 2009), mitigating the phenomenon of “denominator neglect” (see Garcia-Retamero, Galesic, & Dhami, 2013) across multiple age groups (Garcia-Retamero, Galesic, & Gigerenzer, 2010), and, more recently, reducing concern over rare COVID-19 side effects which, consequently, resulted in a diminished aversion to vaccination (Fansher et al., 2022). Moreover, concerning the importance of impression as discussed in Observation 2, Feldman-Stewart, Brundage, and Zotov (2007) argued that icons seemed to provide the optimal balance of accurate processing for both high-level “gist information” (i.e., “giving the gist”; see Spiegelhalter, 2017) as well as more granular, detail-level information.

In light of the promising research on icon arrays, for the purposes of this research (and with a hope of producing a visualisation technique that would accurately capture intragroup opinion *heterogeneity as well as* intergroup agreement), we devised what we’ve termed *icon array histograms*. Icon array histograms are essentially two superimposed histograms (each representing one group) which are constructed by using

¹⁶ And, by logical extension, *dissimilarity*, as well, making them potentially less one-sided than summary statistics which *only* describe difference.

icons as the building blocks (see Figure 20 for an example, and Appendix A for a more detailed explanation of their construction). To our knowledge, they represent a unique contribution to the visualisation toolkit.

Figure 20
Example of Icon Array Histograms



Note. The visualisation above displays actual data from YouGov (US) on partisan opinions concerning whether the US government was supplying an appropriate level of help to the Ukraine in their ongoing conflict with Russia (see Chapter 3). Icon array histograms attempt to display a more comprehensive picture of the range of group responses while simultaneously highlighting “opinion overlap” (depicted by the purple icons, each of which represents a Republican and Democrat who have given the same response).

Famed statistician John Tukey (1977) suggests that the greatest value a visualisation can have is to force us to notice something unexpected. We believe icon array histograms – unlike bar charts – force an audience to confront levels of intergroup agreement which they may not have expected; complicating partisan divides as opposed to concretising them.

While icon array histograms are undoubtedly a slightly more complex form of data visualisation than bar charts, we believe *oversimplified* visualisations may be equally problematic, and contend that a certain degree of complexity in these visualisation formats is both necessary and useful. Not only has recent research suggested that equipping parties with more complex sets of information may counteract individuals’ proclivity to oversimplify issues and promote increased perceptions of “resolution tractability” (Kugler &

Coleman, 2020), but Tufte (2001) suggests that good visualisations should “display an accessible complexity of detail” and concludes the second edition of *The Visual Display of Quantitative Information* with the following excerpt:

“What is to be sought in designs for the display of information is the clear portrayal of complexity. Not the complication of the simple; rather the task of the designer is to give visual access to the subtle and the difficult – that is, the revelation of the complex” (Tufte, 2001, p. 191).

We argue that data visualisations such as icon array histograms, by depicting data in a manner that is capable of simultaneously showing opinion divergence *as well as* agreement, “[gives] visual access to the subtle and the difficult.” Moreover, in expounding upon the tenets of good graphical design, Tufte (2001) suggests that visualisations should “have a narrative quality, a story to tell about the data” (p. 192). Our contention is while bar charts (especially truncated bar charts) often tell a one-sided story of disagreement, icon array histograms may be capable of bringing a degree of nuance to the story the data tells, at once acknowledging group differences while also challenging preconceptions about the degree of outgroup homogeneity and attitudinal extremity.

In speaking of the nature of data visualisations, Lauer and O’Brien (2020, p. 338) warn that “graphs are not objective vehicles, delivering visualised data in a rhetorical and straightforward manner.” Relatedly, Weber (2020) argues that, in the context of journalism, data visualisations are a medium through which storytelling can happen and narrativity often emerges. In light of this perspective, it’s fair to question if (and, if so, *how*) different visualisation modalities (and data summarisation choices) may be impacting the stories and narratives an audience absorbs.

Broadly, we contend that, far from being incidental or inconsequential details, the manner in which we *choose* to present data has the ability to fundamentally affect the way we perceive it. Following this, modes of data visualisation that *only depict* group means (e.g., bar charts) as well as summary statistics which *only discuss* group means are likely to confirm or worsen misperceptions about the degree of political polarisation. Specifically, we hypothesise that these formats of presentation, and their simplified reliance on group differences, are likely to exacerbate an individual’s socio-cognitive and affective propensities to rely on processes such as categorical thinking, oversimplification, and intergroup accentuation – all of which are theorised to underlie false polarisation.

On the contrary, data visualisations which also depict group *distributions* (and, critically, any overlap occurring across these intergroup distributions) or summary statistics which highlight corresponding levels of both group disagreement *and* group agreement are likely to implicitly *challenge* beliefs about attitudinal extremity and ideological coherence as well as disrupt inclinations to oversimplify or neatly categorise, ultimately moderating exaggerated perceptions of polarisation. In this way, we may conceptualise non-traditional data presentation formats as *implicit* perception correction vehicles. Although they do not *explicitly* identify an inaccuracy, by simply presenting data in a comparatively novel and, arguably, richer manner, they are more likely than their traditional counterparts to engender surprise, deeper engagement, and, perhaps, a re-examination of preconceived notions about polarisation.

A Brief Overview of Our Research Approach, Trajectory, and Objectives

The current research sought to explore the intersection of polarisation and data presentation, broadly asking how the manner in which we present data – either via one’s choice of data visualisation or their choice of summary statistic(s) – may impact perceptions of polarisation, inter-party sentiment, and cross-party behavioural dynamics.

The research can be broadly categorised into three distinct parts, each of which aims to make a unique (but interconnected) contribution to the whole:

Chapter 1 systematically investigates the impact mode of presentation has on intergroup perceptions (primarily related to polarisation and level of agreement) and dynamics (e.g., trust, perceived willingness to compromise, etc.). The research comprising this chapter consists of four separate yet sequential studies. This cluster of studies was designed to serve as a proof-of-concept test concerning if and how mode of visualisation might impact partisan attitudes and beliefs. We began in Study 1 by testing whether different modes of data visualisation would shift perceptions of polarisation (and other intergroup dynamics) in an apolitical, “sterile” data environment (i.e., using clean, neatly curated artificial data sets) and, with each subsequent study, we progress through different levels of experimental abstraction, permitting us the ability to systematically test how incrementally introducing levels of ecological validity may impact the differential efficacy of the modes.

Chapter 2 examines a large-scale, international replication of our preceding work (i.e., the studies covered in Chapter 1). Made possible by funding awarded via a Cambridge Impact Grant, we partnered with global market research firm Ipsos MORI to attempt to investigate the reproducibility of our results from Chapter 1 in two large, representative samples in separate geopolitical contexts (i.e., the United States and the United Kingdom).

Chapter 3 examines how repeated exposure to certain summary statistics impacts political perception and behaviour over an extended period. More specifically, participants were exposed to actual data on recent, hot-button political issues where the positions of political partisans were summarised using either group means, levels of opinion overlap (and non-overlap; i.e., percentage of common scores), or both to determine if such repeated exposure to disparate summaries of the data might meaningfully impact perceptions of interparty dynamics over time.

Finally – in light of the findings from the first three chapters – Chapter 4 and Chapter 5 consider what interventions might be employed to mitigate undesirable mode of presentation effects at both an institutional and individual level. In regards to the former, Chapter 4 examines how data presentation choices affect perceived media trust and credibility, with the (admittedly idealistic) underlying assumption that, if certain presentation tactics are found to reduce these reputation-relevant perceptions, it might organically spur institutional change and adaptation to more appropriate presentation formats. In regards to the latter, Chapter 5 reviews the possibility of “inoculating” individuals against the pernicious polarisation effects produced by one of the most widely-used data visualisation manipulation tactics: y-axis truncation.

We conclude with a general discussion of findings, insights, and suggestions for further research.

CHAPTER 1

HOW DATA VISUALISATION CHOICES AFFECT AND CORRECT POLITICAL (MIS)PERCEPTIONS

Abstract

Can a simple change in the way data is presented help to improve the accuracy of intergroup perceptions and ultimately reduce polarisation? In a series of four studies (N = 3,509) using politically-balanced US samples, we aimed to determine the role “mode of presentation” – or the manner in which data is visualised – plays in partisan intergroup perceptions, particularly those related to levels of agreement, polarisation, and cooperation. Findings indicate that mode of presentation plays a significant role across a myriad of intergroup perception metrics, and the effects even persist on polarised issues for which individuals would likely hold strong prior beliefs. Importantly, mode of presentation not only proved capable of reducing perceptions of issue-specific polarisation, but also exhibited mitigation effects that extended to perceptions of overall ideological polarisation, as well. However, mode of presentation only appears to reliably impact perceptions in the short term, and does not appear to significantly impact partisan behaviour.

For years, affective and ideological polarisation have been increasing across the United States as well as in certain other western democracies (Garzia, Ferreira da Silva, & Maye, 2023; Druckman & Levy, 2022; McCarty, Poole, & Rosenthal, 2016; Pew Research Center, 2017). Rising levels of polarisation threatens effective legislative functioning (McCarty, 2016), reduces the willingness of partisans to engage with those outside their political ingroup (Fiorina & Abrams, 2008; Baldassarri & Page, 2021), and has the potential to exacerbate prejudice toward and dehumanisation of political opponents (Cassese, 2021; Iyengar, Sood, & Lelkes, 2012). Moreover, rising levels of polarisation may reduce support for or otherwise inhibit mutually-beneficial compromises between political rivals (Levendusky & Malhotra, 2016b; Whitt et al., 2021). However, while polarisation is undoubtedly increasing in several metrics across the US, a growing number of scholars have begun to address the phenomenon known as *false polarisation* (e.g., Fernbach & Van Boven, 2022), wherein individuals overestimate both the magnitude of disagreement within the political sphere as well as the ideological consistency of the actors occupying it.

Research has shown that partisans struggle to identify the actual position of their adversaries on key issues (Chambers, Baron, & Inman, 2006) and often base their assumptions on stereotypic thinking which serves

to exaggerate the differences between one's ingroup and outgroup (Keltner & Robinson, 1993; Graham, Nosek, & Haidt, 2012). These propensities can yield profound "perception gaps" (Yudkin, Hawkins, & Dixon, 2019), wherein distorted beliefs about where the other party stands on issues differs dramatically from the actual positions they hold. More broadly, Westfall, Van Boven, Chambers, and Judd (2015) cite evidence showing that Americans have exhibited a pattern of overestimation of polarisation which has spanned more than three decades. Consistent with such a finding, recent research has begun to focus on group meta-perceptions, where similar inaccuracies continue to emerge (Lees & Cikara, 2020), both within and outside of the United States (Ruggeri et al., 2021).

While political misperceptions have important consequences for interparty relations (Ahler & Sood, 2018; Moore-Berg, Ankori-Karlinsky, Hameiri, & Bruneau, 2020), the prevalence of erroneous beliefs regarding the stances of political others is especially troubling as scholars such as Huxley (1978) have contended that the escalation of group conflict is frequently attributable to fundamental social misperceptions. Echoing these sentiments, more recent work has suggested that polarisation can be cyclical in nature (Lees & Cikara, 2021), with upward shifts in domains such as affective polarisation producing elevated perceptions within the domain of ideological polarisation (Armaly & Enders, 2021; Levendusky & Malhotra, 2016b), ultimately engendering a feedback loop that can catalyse a rapid acceleration in both areas. In essence, though changes in the perceptions of political polarisation do not necessarily entail any corresponding changes in the actual levels of polarisation, they can nevertheless contribute to actual shifts through a process of mutual reinforcement (Lees & Cikara, 2021). Thus, it's important that such misperceptions are corrected whenever possible.

Interventions which simply reveal the *true* positions of partisans (often relative to the *believed* positions) have shown promise in accomplishing this, proving capable of reducing negative outgroup attributions (Lees & Cikara, 2020) and attenuating beliefs about attitudinal extremity, which can ultimately engender greater political moderation (Ahler, 2014). However, Druckman and Levy (2022) note that individual interventions, though promising, might still struggle to overcome persistent misperceptions on a large scale due to the fact that many misperceptions are institutionally driven and sustained (see Wilson, Parker, & Feinberg, 2020). More specifically, because exposure to (often exaggerated) polarisation narratives is near-ubiquitous in today's modern media landscape, sometimes being promoted by agents with immense scope for influence (e.g., partisan media outlets, political elites, etc.), interventions seeking to effectively counter such influence should have the capacity to be deployed on an equally-large scale, ideally with minimal "training" or requirements of sustained attention on the part of the target audience.

We believe an underexplored route of engendering large-scale misperception corrections may be as simple as changing the way in which we visualise data. As data visualisations are already an omnipresent feature of

print, television, and social media, should certain visualisations demonstrate the ability to reduce political misperceptions, simply shifting current data presentation methodologies to these superior visualisation formats would represent a promising avenue for catalysing widespread change with minimal effort. Such an intervention would also have an excellent cost-benefit profile and superior scalability to individually-tailored interventions, permitting depolarisation agents to easily embed perception-correction mechanisms across multiple domains and mediums in which visualisations already exist.

There are reasons to believe altering visualisations may be a viable mechanism through which one could shift perceptions of groups. Scholars like Hanel, Maio, and Manstead (2019) have argued that traditional means of presenting data about groups – such as bar charts with truncated y-axes – may (purposely or otherwise) accentuate group *differences* while ignoring equally-important similarities (e.g., Hanel & Wolf, 2020). In discussing psychology’s possible over-reliance on group differences, they note that not only has past research found that more than 90% of published data reports statistically significant differences (e.g., Fanelli, 2010), but they call attention to the troubling reality that authors who fail to show such differences often have their research relegated to the proverbial “file drawer” (Hanel et al., 2019).¹⁷ While some have suggested that certain degrees of y-axis truncation may promote appropriate levels of sensitivity to effect sizes (Witt, 2019), we suspect that truncated bar charts – which emphasise mean differences while providing minimal distributional information about the samples – may inadvertently encourage categorical thinking and oversimplification, cognitive predispositions implicated as supporting factors in false polarisation (Fernbach & Van Boven, 2022). Moreover, an oversimplification of group beliefs coupled with exaggerated depictions of group differences may risk potentially facilitating more “essentialised” perceptions of the depicted groups (see Haslam & Whelan, 2008), wherein individuals conceptualise social categories – and, consequently, the delineation between different groups – as both natural and fixed. Beliefs in “natural kinds” have been found to amplify group differences (Rothbart & Taylor, 1992) and the attribution of “essences” to outgroups has been linked to classic conceptions of prejudice (Allport, Clark, & Pettigrew, 1954).

While the modes that have traditionally been used to visualise intergroup data may unintentionally magnify perceived group differences and lead to exaggerated beliefs about political polarisation, Hanel et al. (2019) found that simply changing the type of visualisation used to depict identical data had a significant impact across an array of intergroup assessments, with modes that more effectively highlight similarity information (e.g., superimposed normal distributions) resulting in both more accurate intergroup perceptions as well as more positive outgroup appraisals. While such findings might seem modest, we believe they have the potential to be particularly impactful in the context of political polarisation. Researchers have consistently

¹⁷ Further corroborating the discipline’s preoccupation with identifying differences is the fact that most inferential statistics and some of the most common analysis techniques are only appropriate for assessing difference but cannot render definitive conclusions concerning degrees of intergroup similarity.

shown that minute, sometimes seemingly inconsequential changes to the way in which data is visualised can profoundly alter how that data is perceived and interpreted (e.g., see Nguyen, Jung, & Gupta, 2021; Lo et al., 2022). Considering this, we believe it's fair to question whether the manner in which we've *traditionally* visualised data about political groups may contribute to political misperceptions and false polarisation. In addition to expounding upon the disadvantages of traditional modes of presentation and the possible dangers of y-axis truncation in depicting group data, the results reported by Hanel et al. (2019) also provide preliminary evidence for the potential power of using comparatively novel visualisation techniques (i.e., those which feature aspects such as full ranges and distributions of responses, levels of intergroup agreement as opposed to simply disagreement, etc.) to shift intergroup perceptions. We believe this should motivate researchers to consider how expanding their data depiction methodologies beyond the traditional (e.g., bar charts) might similarly alter *political* perceptions. For example, modes such as icon arrays, which feature prominently in the risk communication literature, have shown promise in areas such as increasing the accuracy of risk perception (Galesic, Garcia-Retamero, & Gigerenzer, 2009) and mitigating phenomena such as denominator neglect (Garcia-Retamero, Galesic, & Gigerenzer, 2010),¹⁸ but have yet to be extensively explored in the context of polarisation.

While other researchers have investigated the circumstances and contexts in which data visualisations might prove capable of changing attitudes (e.g., Markant, Rogha, Karduni, Wesslen, & Dou, 2022; Pandey, Manivannan, Nov, Satterthwaite, & Bertini, 2014), few have directly focused on how certain depictions of data may affect beliefs about the extent of political polarisation (Alieva, 2023; Santos et al., 2017). However, there is some precedence for believing that “better” visualisations might constitute useful interventions to reduce political polarisation. For example, Rutchick, Smyth, and Konrath (2009) found that maps depicting state-level support for a presidential candidate with a shade of purple along a red and blue continuum (in proportion to the degree of support for the Republican or Democratic candidate) reduced stereotyping as well as perceptions of political division compared to the traditional binary red and blue state-by-state visualisation.

Building upon the conceptual groundwork laid by Hanel et al. (2019), the present research will seek to further elucidate the impact that simple shifts in visualisation format may have on political perceptions. Broadly, we endeavour to determine A) whether (and, if so, to what degree) y-axis truncation of bar charts that depict intergroup political information shifts perceptions relative to bar charts with a full-range y-axis, B) whether “novel” modes of data visualisation (e.g., such as those involving icon arrays) – which depict aspects of intergroup data such as response distributions and intergroup agreement – promote different political perceptions than more “traditional” modes of presentation, and C) whether exposure to visualised

¹⁸ More recently, icon arrays have even shown themselves capable of assisting in pressing, real-world issues such as a reticence to get vaccinated against COVID-19 by reducing the concern over exceedingly rare side effects of vaccines (Fansher et al., 2022).

data leads participants to hold political perceptions which differ significantly from those who are provided with no such visualisations. Should the results indicate that different modes of presentation – despite depicting identical underlying data – yield significant differences across polarisation metrics, such findings would have important implications for the process by which individuals and organisations go about making visualisation choices, and ultimately may provide a subtle yet highly-scalable addition to the depolarisation repertoire.

Research Approach

The four studies contained herein were conceptualised as a proof-of-concept series, designed to incrementally assess the viability of mode of presentation interventions in the political sphere at different levels of abstraction and realism. Study 1 begins by testing mode of presentation effects in a relatively sterile environment (e.g., using a pair of artificial datasets with no information provided about the specific issue being depicted by the data) to determine how the modes perform devoid of the noise introduced by actual political issues. As the studies progress, layers of realism and complexity are added to the original design. For example, Study 2 attributes the data to an *actual* political issue, and then Study 3 begins integrating different sources through which the visualisations are disseminated. This incremental approach enables us to tease apart the effect of particular variables on the efficacy of mode while also systematically investigating the potential utility of the intervention and boundary conditions of the effects.

Study 1: Investigating the Effects of Mode of Presentation for an Unspecified Issue with Political and Non-Political Groups¹⁹

The findings of Hanel et al. (2019) were revelatory in the sense that they provided evidence to suggest that an alteration in the way identical data is visually depicted can have substantial repercussions on the way that data is perceived. Specifically, Hanel et al. (2019) found two particular effects of “mode of presentation” which we sought to replicate in Study 1.

The first concerned the significant impact “mode of presentation” played on perceptions of intergroup similarity. Several studies within their work (e.g., Study 2, Study 5) demonstrated the ability of visualisation format to significantly impact perceptions of similarity. Consequently, our first confirmatory hypothesis concerned the successful replication of these effects, positing that main effects of mode of presentation would be observed on perceptions of intergroup similarity. Moreover, Hanel et al. (2019) found that certain modes seemed to promote higher estimates of intergroup similarity than others. For example, Study 3 found that, when participants were asked to evaluate how similar the values of British people and Polish people were (after viewing a collection of ten value visualisations, depicted either via superimposed normal

¹⁹ University of Cambridge ethics identification number 2651.105

distributions, an “unrestricted” (i.e., full-range) bar chart, and a “restricted” (i.e., truncated) bar chart), participants exposed to either superimposed normal distributions or unrestricted bar charts came to view the two groups as significantly more similar than participants exposed to the data as restricted bar charts. Consequently, we hypothesised that a similar pattern would emerge in the present investigation.

The second effect, as evidenced by the results of Study 2 from Hanel et al. (2019), demonstrated that certain modes, beyond promoting different *perceptions* of similarity, are also capable of generating more *accurate estimates* of group similarities and differences. Within Study 2, Hanel et al. (2019) observed a pattern of results that seemed to indicate that data presented via superimposed normal distributions promoted the most accurate estimates of group similarity, significantly more accurate across almost every interval of Cohen’s *d* depiction (with the exception of a Cohen’s *d* of 2.0, which showed no significant difference in estimates between any of the visualisation formats) than full-range bar charts. Consequently, we hypothesised that superimposed normal distributions would encourage accurate estimates of intergroup similarity. However, beyond this confirmatory hypothesis, we also forwarded an additional hypothesis – one which arguably runs counter to the data discussed in Hanel et al. (2019). Specifically, we hypothesised that the icon array histograms would also encourage accurate estimates of intergroup similarity. Such a hypothesis is not necessarily supported by the data, as Hanel et al. (2019) found that superimposed histograms – the mode which is arguably most conceptually similar to the icon array histograms – typically elicited *less* accurate perceptions than the superimposed normal distributions of unrestricted bar charts. However, given the fact that the extant literature has shown icon arrays themselves to be useful at promoting accurate perceptions (e.g., Galesic, Garcia-Retamero, & Gigerenzer, 2009), we hypothesised that such arrays, serving as the foundation of the icon array histograms (see *Appendix A*), would aid in generating accurate estimates among participants.

Beyond these two primary effects, we also sought to determine whether perceptions would differ simply as a function of the groups being depicted as political or non-political in nature. While Hanel et al. (2019) did describe a brief foray into the political (or at least politically-*adjacent*) domain with their fifth study (which, among other things, found that when the moral foundations of liberals and conservatives were presented as superimposed normal distributions, participants believed the two groups to be more similar than if the same data was presented via restricted barplots), the majority of their findings were obtained with non-political groups. As our ultimate goal is to determine whether manipulating the manner in which data is presented can act as a viable intervention to combat political polarisation, we have a particular interest in determining how altering the political nature of the groups being depicted might impact the overall effects. Broadly, we hypothesised that, due to the pervasive misperceptions that permeate the American political landscape (e.g., Westfall et al., 2015), participants would have a tendency to *underestimate* similarities

between political groups (e.g., Republicans and Democrats) relative to identical data attributed to non-political groups.

Method

Participants

A sample of American participants were recruited via the Prolific Academic platform and paid an approximate rate of £7.50/hour. Following exclusions for non-consent, attention check failure²⁰, or incomplete data, the final sample was 792.²¹ Table 1 summarises the demographic composition of the sample.

Table 1

Composition of Sample for Study 1

| Identity Characteristic | Sample Composition |
|-------------------------|-----------------------------------------------------------------------------------------------------------|
| Sex | Female: 50.88% ($n = 403$), Male: 47.85% ($n = 379$), Other: 1.26% ($n = 10$) |
| Age | $M = 35.95$, Range: 18-81 |
| Education | 48.74% ($n = 386$) had attained a bachelor's degree or higher |
| Political Affiliation | Democrat: 32.45% ($n = 257$), Republican: 34.34% ($n = 272$), Independent/Other: 33.21% ($n = 263$) |

Among participants who openly identified as Democrat or Republican, 54.47% ($n = 140$) categorised themselves as “strong” Democrats and 45.96% ($n = 125$) categorised themselves as “strong” Republicans. Among those who openly identified as Independent/Other, when forced to choose which of the two major US political parties they considered themselves closest to, 60.84% ($n = 160$) of this group chose Democrat and 39.16% ($n = 103$) chose Republican. Ideologically, the sample skewed slightly liberal ($M = 46.39$ on a 0 to 100 ideological self-placement scale).

²⁰ The attention check simply asked participants to specify what the data visualisation format to which they had been exposed depicted – a detail which had been reiterated a minimum of three times prior to the check being performed. There were three available response options from which participants could choose: the full range of responses for each group, the average responses for each group, or the most extreme responses for each group. 76.60% of participants gave the correct response to the attention check.

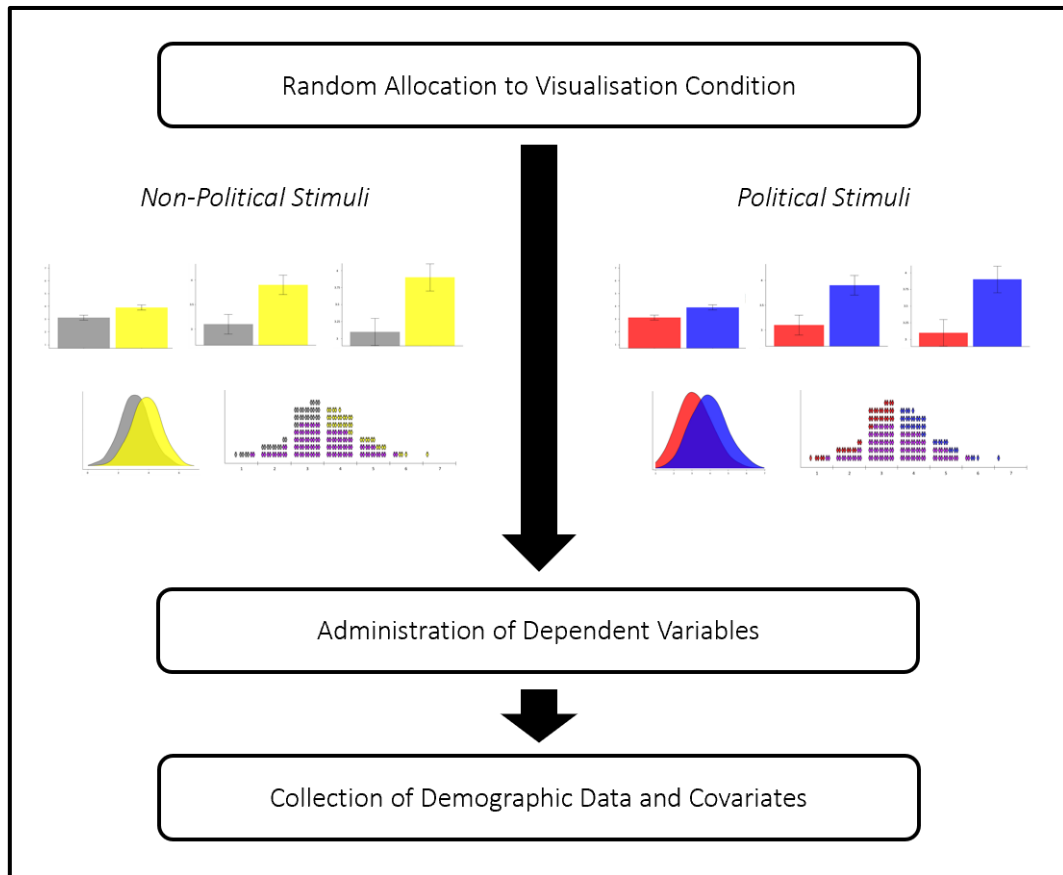
²¹ A priori power analyses were conducted to determine the sample size required to achieve sufficient statistical power to detect a medium effect size (i.e., Cohen's f of 0.25) with a power level of 0.90 and an alpha threshold of 0.05. The analysis indicated a sample size of approximately 33 participants per group would be required. The effect size was chosen based on the work of Hanel et al. (2019), who routinely found large effect sizes, but whose work was novel and, due to the lack of corroborating evidence within the extant literature, we wished to conduct a more conservative power analysis. Ultimately, given our budget for the project, we recruited more than was required A) to account for potential dropout and B) in the hopes of detecting more subtle effects within the analyses. Additionally, prior to exclusions for non-consent or incomplete data, the total sample was 1,034. For those who provided data, the sample composition prior to exclusions had a mean age of 36.22, was comprised of 489 females, 477 males, and 14 other (with 54 non-responses), and had 315 self-reported Democrats, 338 self-reported Republicans, and 327 self-reported other (with 54 non-responses).

Design

The study utilized a 5x2 (mode of presentation by group label), between-subjects design (Figure 1).

Figure 1

Design Outline for Study 1



Procedure

To ensure precision and maximise experimental control, two artificial datasets were created and the underlying data points were used to construct the visualisations for each mode of presentation. Each of the artificial datasets was normally-distributed and consisted of 100 observations that ranged from 1 to 7 (similar to the range of responses one might encounter were participants to respond to a survey via a 7-point Likert scale). The final datasets featured means of 3.1 and 3.9, and each had a standard deviation of 1. Critically, the intergroup effect size of the two datasets was a Cohen's d of exactly 0.80, the common threshold for characterising an effect as "large," thus ensuring the visualisations would constitute an accurate simulation of the type of response patterns one might encounter when examining an issue for which there is substantial disagreement between two groups.

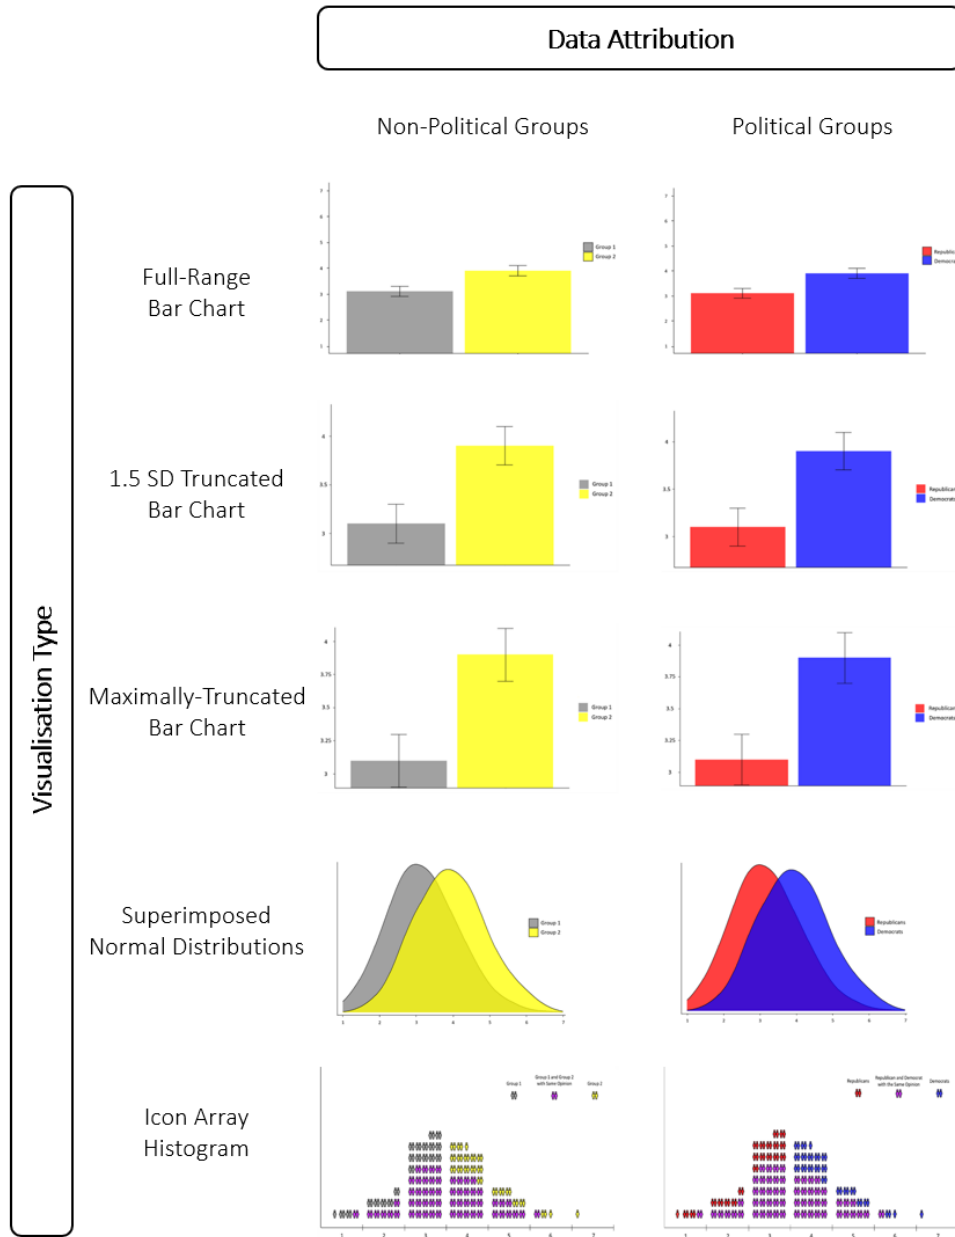
Following an affirmative response to the consent questionnaire, participants were randomly allocated into one of five modes of presentation (see Figure 2):

1. A bar chart with a full-range y-axis depicting mean group differences with a 95% confidence interval (henceforth simply “full-range bar chart”)
2. A bar chart with a 1.5 SD y-axis depicting mean group differences with a 95% confidence interval (henceforth simply “1.5 SD bar chart”)
3. A bar chart with a maximally-truncated²² y-axis depicting mean group differences with a 95% confidence interval (henceforth simply “maximally-truncated bar chart”)
4. Superimposed normal distributions
5. Overlapping histograms (using icon arrays to depict data for each bin; henceforth “icon array histogram”)²³

²² For the purposes of this study, “maximally-truncated” is defined as the greatest degree of y-axis truncation that does not restrict any of the chart’s core features (e.g., group means, the upper or lower bounds of the CI bars, etc.).

²³ See Appendix A for more detail on their construction.

Figure 2
Condition Stimuli for Study 1



Each of the five modes were further sub-divided into two separate conditions (only one of which a participant would see): a condition which depicted the data using non-political group labels (and colour schemes) and a condition which depicted the data using political group labels (and colour schemes). For the former, responses were attributed to groups specified only as “Group 1” and “Group 2,” and depicted using grey and yellow – colours not readily associated with any political parties in the American political context. For the latter, responses were attributed to Republicans and Democrats, and depicted using colours for which each party has come to adopt as their own (i.e., red and blue, respectively).

Depending on the random allocation, participants were exposed to one of the ten data visualisations shown above, each of which contained the following statement²⁴ below it:

“[Members of two groups OR Republicans and Democrats] were asked, on a scale from 1 (Strongly Disagree) to 7 (Strongly Agree), how they felt about a particular issue. The graph above depicts the [average OR full range of]²⁵ responses from each group.”

Finally, after exposure to the data via one of ten unique visualisation conditions, participants were presented with a set of dependent variables, randomised so as to mitigate possible order effects, which aimed to assess perceived group similarity, estimates of intergroup opinion overlap, perceived issue polarisation, perceived affective polarisation on the issue, beliefs about intergroup willingness to cooperate, optimism about intergroup compromise, and levels of trust concerning out-party fairness on the issue (see *Measures* for details of item wording and scoring).

Measures

The study featured two primary dependent variables which would serve as the crux of the replication:

Perceived Group Similarity: Participants were asked to indicate their response (via a 101-point scale that ranged from “Extremely Different” and “Extremely Similar”) to the item “Based on the graph, how similar or different do you perceive the groups to be on this issue?”

Estimates of Intergroup Opinion Overlap: Participants were asked to indicate their response (via a 101-point scale) to the item “Based on the graph, what percentage of the two groups share the same opinion on the issue?”

Hypotheses

Study 1 featured several hypotheses²⁶ (see Table 2):

²⁴ This statement was not true, as the visualisations were merely depicting the artificial data created solely for the purposes of this study, but the use of deception here – designed to ensure participant reactions similar to those observed when encountering actual intergroup data depictions – was cleared by the University of Cambridge Ethics Committee.

²⁵ Depending on whether the mode was a bar chart (in which case “average” was used) or either an icon array histogram or superimposed normal distribution (in which case “full range of” was used).


²⁶ For additional hypotheses (and their accompanying results), see Appendix C.

Table 2*Hypotheses for Study 1*

| Label | Hypothesis |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H1 | Main effects of mode of presentation will be observed on measures of perceived intergroup similarity. |
| H1 _A | The maximally-truncated bar chart will elicit lower perceptions of intergroup similarity than the bar chart with the full-range y-axis. |
| H1 _B | The maximally-truncated bar chart will elicit lower perceptions of intergroup similarity than the superimposed normal group distributions. |
| H2 _A | The superimposed normal distributions will elicit accurate estimates of intergroup opinion overlap. |
| H2 _B | The icon array histograms will elicit accurate estimates of intergroup opinion overlap. |
| H3 | Main effects of group label will be observed on measures of perceived intergroup similarity (with perceptions of intergroup similarity being lower in the conditions with political group labels than the non-political group labels) |

Results

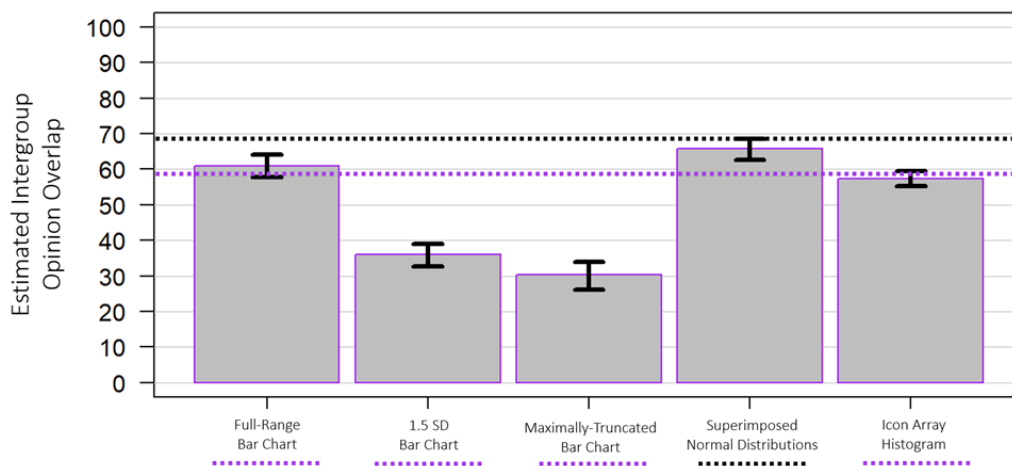
To assess the main effects of mode of presentation on assessments of intergroup similarity, a 5x2 (mode by group label) ANOVA was conducted. Results indicated a significant main effect of mode of presentation on intergroup similarity assessments [H1] ($F(4, 782) = 61.020, p < .001, \eta_p^2 = 0.238$; see Figure 3). Notably, Tukey's Honest Significant Difference (HSD) post-hoc tests confirmed highly-significant differences (i.e., $p < .001$) between the maximally-truncated bar chart ($M = 38.64, SD = 26.18$) and each of the full-range bar chart ($M = 62.47, SD = 20.02$), superimposed normal distributions ($M = 68.39, SD = 17.05$), and icon array histogram ($M = 57.57, SD = 17.52$) conditions. T-tests confirmed highly-significant differences between the maximally-truncated bar chart and the full-range bar chart [H1_A] ($t(252.42) = 8.691, p < .001, d = 1.03$), the superimposed normal distributions [H1_B] ($t(232.28) = 11.219, p < .001, d = 1.35$), and the icon array histogram ($t(216.71) = 7.425, p < .001, d = 0.88$). Exploratory t-tests determined that there was no significant difference between the maximally-truncated bar chart and the 1.5 SD bar chart ($M = 40.15, SD = 23.02$) on this measure ($t(271.64) = 0.5166, p = .606, d = 0.06$).

Figure 3*Perceptions of Intergroup Similarity by Mode*

Main effects of mode also emerged from a 5x2 (mode by group label) ANOVA for estimates of intergroup opinion overlap ($F(4, 782) = 101.451, p < .001, \eta_p^2 = 0.342$ (see Figure 4). Beyond main effects (which only tell us whether the modes differ significantly *from one another* on this metric), we also used two accuracy benchmarks (created and utilised by Hanel et al., 2019) to determine *how accurate* participants exposed to each mode were at estimating intergroup opinion overlap. For every mode except the superimposed normal distributions, *percentage of common scores* (PCS) was used; for superimposed normal distributions, *percentage of common responses* (PCR) was used.²⁷ For the artificial datasets featured in the study, the PCS was 59% while the PCR was 69%. While the superimposed normal distributions ($M = 65.80, SD = 17.17$) more closely approached their PCR benchmark ($t(142) = 2.2313, p = .027, d = 0.19$) than either the 1.5 SD bar chart ($M = 35.97, SD = 19.75$) or maximally-truncated bar chart ($M = 30.26, SD = 23.44$) approached their PCS benchmark ($t(147) = 14.185, p < .001, d = 1.17$ and $t(136) = 14.347, p < .001, d = 1.23$, respectively), only the full-range bar chart ($M = 61.01, SD = 20.79$) and icon array histogram ($M = 57.38, SD = 15.10$) managed to catalyse estimates that did not differ significantly from their accuracy benchmark [H2_A, H2_B] (i.e., a PCS of 59%; $t(158) = 1.2209, p = .224, d = 0.10$ and $t(204) = 1.5355, p = .126, d = 0.11$, respectively).

Figure 4

Estimates of Intergroup Opinion Overlap by Mode



Note. The purple and black horizontal lines refer to the PCS and PCR measures, respectively – the accuracy benchmarks used in the analysis.


Finally, the original 5x2 ANOVA also revealed a significant main effect of group label (i.e., political versus non-political) on perceptions of intergroup similarity [H3], with non-political group labels ($M = 55.66, SD = 23.88$) leading to higher perceptions of similarity ($F(1, 782) = 6.892, p = .009, \eta_p^2 = 0.009$; see Figure 5) than

²⁷ While PCS is an appropriate accuracy benchmark for all other modes, PCR is the more appropriate benchmark for the superimposed normal distributions. This mode effectively “smooths” over of the original, discrete response bins, and while PCS and PCR are generally well-correlated (Hanel et al., 2019), with this particular combination of datasets, the superimposed normal distributions depict an overlap that is 69%, approximately 10 percentage points higher than the actual PCS. See Appendix B for a more detailed explanation.

political group labels ($M = 52.24$, $SD = 23.33$). T-tests confirmed a significant difference in the pairwise comparison ($t(789.96) = 2.0345$, $p = .042$, $d = 0.14$). No significant interaction between mode and group label was detected for this metric ($F(4, 782) = 0.595$, $p = .666$, $\eta_p^2 = 0.003$).

Figure 5

Perceptions of Intergroup Similarity by Group Label Condition



* $p < .05$, ** $p < .01$, *** $p < .001$

Discussion

Study 1 was able to broadly corroborate the findings originally obtained by Hanel et al. (2019) concerning “mode of presentation” and its effects on perceptions of intergroup similarity. Critically, we observed a highly-significant main effect of “mode of presentation” – or the manner in which one chooses to visually depict data – on perceptions of intergroup similarity. Communicated differently, this means that if identical data is depicted via different visualisation formats, one can expect significantly different perceptions of similarity to emerge depending on the format chosen. Such a finding is critical as it suggests that choices concerning *how* to visualise data can ultimately prove consequential to the ultimate interpretations and inferences that the data might elicit. For example, while two publications might opt to report on the same poll, the one which presents the data as a truncated bar chart might (inadvertently or otherwise) promote perceptions of intergroup similarity which will be meaningfully lower than the publication which chooses to present the data as a full-range bar chart. Such an insight is important as it suggests that data visualisation format is not a merely an inconsequential vehicle through which data is passively presented, but actually plays an active role in how that data will be perceived.

However, findings from Study 1 differed slightly from the results obtained by Hanel et al. (2019) in regards to which types of visualisations lead to more or less *accurate* perceptions. While Hanel et al. (2019) found superimposed normal distributions to (generally) promote more accurate estimates of intergroup similarity than (full-range) bar charts (which, generally, promoted more accurate estimates than overlapping histograms; see Study 2), we determined that only the full-range bar chart and the icon array histogram

catalysed estimates that did not significantly differ from their accuracy benchmark. The remaining modes (i.e., superimposed normal distributions, 1.5 SD bar charts, and maximally-truncated bar charts) all promoted estimates which significantly *underestimated* the true degree of intergroup opinion similarity. Importantly, however, both forms of truncated bar charts yielded particularly pronounced levels of underestimation, with participants within these conditions believing group attitudes to be more than 20 percentage points *less similar* than they actually were.

Finally, in an extension of Hanel et al.'s (2019) work, we found that simply specifying the *types* of groups being depicted can exert a significant influence on perceptions of similarity. More specifically, the study found that group label (i.e., whether identical group data are attributed to political or non-political groups) has a significant impact on perceptions of group similarity independent of that of mode, with data said to be depict opinions of political rivals (i.e., Republicans and Democrats) being perceived as having significantly *less* intergroup opinion similarity than identical data attributed to the opinions of non-political groups (e.g., Group 1 and Group 2). Such a finding is important as it implies that the mere fact that groups being depicted are political adversaries might inherently skew perception, influencing individuals to see the groups as “more divided” than if identical data had been attributed to non-political groups.

While the study permitted the successful replication of a number of Hanel et al.'s (2019) findings (in addition to some modest novel contributions), it is not without its flaws. Of particular note is the fact that the design was deliberately kept as “sterile” as possible (i.e., using perfectly-normal, artificial datasets and not specifying an issue to which to attribute the data). While such a choice permitted a higher degree of experimental control, it necessarily raises the question as to whether responses to such “clean” datasets amidst such minimal context will provide insights that carry any predictive validity in a messier, real world situation. Can individuals respond meaningfully to a chart with such information scarcity? Will response patterns continue to look similar as ecological validity is increased? While we cannot confidently answer the first question, our next studies will seek to shed light on the latter.

The current study is valuable in that it provides corroborating evidence that the “mode of presentation” effects observed by Hanel et al. (2019) can be replicated in a similar experimental context. While we find slightly different results concerning which modes promote the most *accurate* perceptions, it's worth noting that we utilised two different accuracy benchmarks (i.e., PCS and PCR) while Hanel et al. (2019) only use one (i.e., PCR, which could have contributed to the asymmetric findings). Overall, however, the pattern of findings suggest that the manner in which one chooses to visually depict identical data (as well as the *types* of groups being depicted) can have tangible impacts on perceptions of similarity.

Study 2: Investigating the Effects of Mode of Presentation for a Specified Issue Attributed to Elite and Non-Elite Political Groups²⁸

While Study 1 confirmed that “mode of presentation” effects, as originally observed in Hanel et al. (2019), could be replicated for intergroup perceptions of *similarity*, it did not provide any insight as to whether such effects would continue to persist for expressly *political* perceptions (e.g., polarisation), particularly when the data – rather than remaining unspecified – is attributed to a divisive political issue. Consequently, Study 2 sought to shift focus from investigating how mode of presentation impacts perceptions of intergroup similarity to investigating how it impacts perceptions of politically-relevant items, such as ideological and affective polarisation. As scholars have documented a persistent pattern of overestimation of political polarisation among the American public (Westfall et al., 2015), determining whether the manner in which data is visualised might be capable of attenuating (or exacerbating) these effects is a critical step in evaluating its value as an effective intervention strategy within the political sphere.

To test for mode of presentation effects in slightly more ecologically-valid manner, Study 2 built upon Study 1 by replacing the previously (intentionally) unspecified issue with a specified one: immigration. Immigration was chosen as a suitable issue for two primary reasons. Firstly, we found that immigration items, such as those found in recent waves of the American National Election Studies, routinely generated effect sizes between Republican and Democratic respondents around a Cohen’s *d* of 0.80, which would enable us to continue using the visualisations derived from the original artificial datasets in Study 1 (which depict two groups with a Cohen’s *d* of *precisely* 0.80) without concern that the depictions would be unnecessarily deceptive. Secondly, immigration – beyond being a perennial hot-button issue within US politics – also appears to be an issue for which there are profound misperceptions on *both* sides of the political aisle. For example, a report by *More in Common* (Yudkin, Hawkins, & Dixon, 2019) found that immigration was one of the only issues for which both Republicans *and* Democrats substantially misjudged the other side’s stance.

The presence of an actual issue attributed to the data in Study 2 enables us the ability to understand whether the mode effects observed in Study 1 for perceptions of intergroup similarity persist when examining perceptions of ideological and affective polarisation. By several accounts, both forms of polarisation have risen within the United States over the past several decades (e.g., Boxell, Gentzkow, & Shapiro, 2022; McCarty, Poole, & Rosenthal, 2016; Pew Research Center, 2017; Druckman & Levy, 2022). Considering the issue chosen (i.e., immigration) is a highly-salient one, where individuals will presumably hold relatively strong and resistant attitudes (and, importantly, will likely have an idea of how polarised Republicans and Democrats tend to be overall), it’s possible that perceptions will not be “movable” via a simple visualisation-shifting intervention. However, in light of the strong effects observed in Study 1, we hypothesise that main

²⁸ University of Cambridge ethics identification number 3131.127

effects of mode of presentation will continue to emerge for perceptions of both ideological and affective polarisation. Additionally, based on the patterns observed between pairwise comparisons of different modes in Study 1, we hypothesise that perceptions of polarisation – both ideological and affective – will be greater for participants exposed to the truncated bar chart than those exposed to either the full-range bar chart or the superimposed normal distributions.

While the superimposed normal distributions *did* produce estimates of intergroup opinion overlap that fell significantly below their accuracy benchmark, Study 1 utilised a particularly “sterile” information environment; namely, the data, beyond being derived from artificial datasets, also was not said to depict a specific issue. In their second study, Hanel et al. (2019) determined that superimposed normal distributions, generally speaking, promoted the most accurate estimates of similarity. Critically, during this study, the researchers had given their (artificial) data *some sort* of label (i.e., “sociability”) as a way to “make the variable more concrete” (p. 549). Consequently, we hypothesised that perhaps the addition of some sort of label (i.e., immigration) to Study 2 would further “concretise” the data and enable the superimposed normal distributions to display similar effects on perception accuracy.

Beyond examining the effects that altering the visual depiction of data will have on perceptions of polarisation, we’re also keen to determine whether certain depictions of the data conform to or undermine one’s *expectations* about the level of (perceived) political agreement. Expectations play a key role in political dynamics as they may prime different types of cross-party attitudes and, consequently, even impact certain behaviours. For example, recent research has found that political partisans may privately be willing to cooperate with one another, but abstain from doing so due to the *expectation* that their opposing-party partner will be unwilling to reciprocate such cooperation (Dimant, 2024). Similarly, expectations may foster self-fulfilling prophecies, wherein the behaviour an individual decides to display is informed by their *expectation* of reciprocation of similar behaviour from a social other. For example, Stinson et al. (2009) determined that one’s *expectations* concerning social acceptance or rejection impacted the level of social warmth they exhibited and, ultimately, the level in which they received in return. However, if expectations can engender self-fulfilling prophecies (e.g., “I expect the other party to strongly disagree with my party on this issue, so I will *behave* towards them in kind”), perhaps *undermining* these expectations may attenuate or even reverse such phenomena. Thus, we opted to measure the degree to which each depiction of data met or violated participant expectations regarding interparty agreement. Due to the fact that individuals already assume higher levels of polarisation than actually exist (e.g., Westfall et al., 2015), we hypothesised that “traditional” visualisation techniques which emphasised group differences (i.e., truncated bar charts) would conform to participant expectations of intergroup agreement more than comparatively novel visualisation techniques (i.e., such as the superimposed normal distributions), which, due to their visual emphasis on agreement, would promote greater levels of expectation violation.

We also sought to determine whether alterations in the way data was depicted could contribute to other pernicious intergroup processes. We hypothesised that certain modes, due to the visual emphasis placed on group *differences*, might exacerbate facets of *essentialist* thinking (i.e., believing certain social and physical categories of the world are defined by underlying “essences” shared by other category members, making their natures both objective and fixed; e.g., Rhodes & Mandalaywala, 2017). Specifically, we believed that bar charts – particularly truncated bar charts – would be likely to promote greater perceptions of intergroup *discreteness*, or a belief that the boundaries between two categories are fixed and non-traversable (i.e., you are either a member of one category or the other) while modes which visually accentuated intergroup agreement (such as superimposed normal distributions and icon array histograms) would be more likely to reduce perceptions of intergroup discreteness, ultimately leading to main effects of mode of presentation on this measure.

Additionally, Study 2 also included an independent variable which altered the *type* of political group to whom the opinions are attributed (i.e., political elites or partisan citizens). The relationship between the two types of political groups and their effect on opinion formation is often complex (e.g. Druckman, Peterson, & Slothuus, 2013; Slothuus & Bisgaard, 2021), and while we do not offer any formal hypotheses for how this variable might interact with mode to alter participant responses, we believe adding it to the design will permit us the opportunity to examine whether the type of political group from which opinions are offered will exert a meaningful impact on response patterns.

Method

Participants

A sample of American participants were recruited via the Prolific Academic platform and paid an approximate rate of £7.50/hour. Following exclusions for non-consent, attention check failure²⁹, or incomplete data, the final sample was 654.³⁰ Table 3 summarises the demographic composition of the sample.

²⁹ The attention check was a question embedded in the dependent variables which began in the same manner as the prior questions (e.g., “200 member of Congress were asked...”) but which explicitly stated – where all other dependent variables asked for an opinion – that the participant *must select* “Slightly Disagree” (amongst a total of six available response options) to prove they were paying attention. 93.56% selected correctly and passed the attention check.

³⁰ A priori power analyses were conducted to determine the sample size required to achieve sufficient statistical power to detect a medium effect size (i.e., Cohen’s *f* of 0.25) with a power level of 0.90 and an alpha threshold of 0.05. The analysis indicated a sample size of approximately 37 participants per group would be required. The effect size was chosen based on the findings from Study 1, which routinely observed large effect sizes, but due to the uncertainty in effect size to be observed as ecological validity is increased, we wished to conduct a more conservative power analysis. Ultimately, given our budget for the project, we recruited more than was required A) to account for potential dropout and B) in the hopes of detecting more subtle effects within the analyses as experimental sophistication increased. Additionally, prior to exclusions, the total sample was 699. For those who provided data, the sample composition prior to exclusions had a mean age of 35.60, was comprised of 321 females, 320 males, and 14 other (with 44 non-responses), and had 186 self-reported Democrats, 234 self-reported Republicans, and 235 self-reported other (with 44 non-responses).

Table 3

Composition of Sample for Study 2

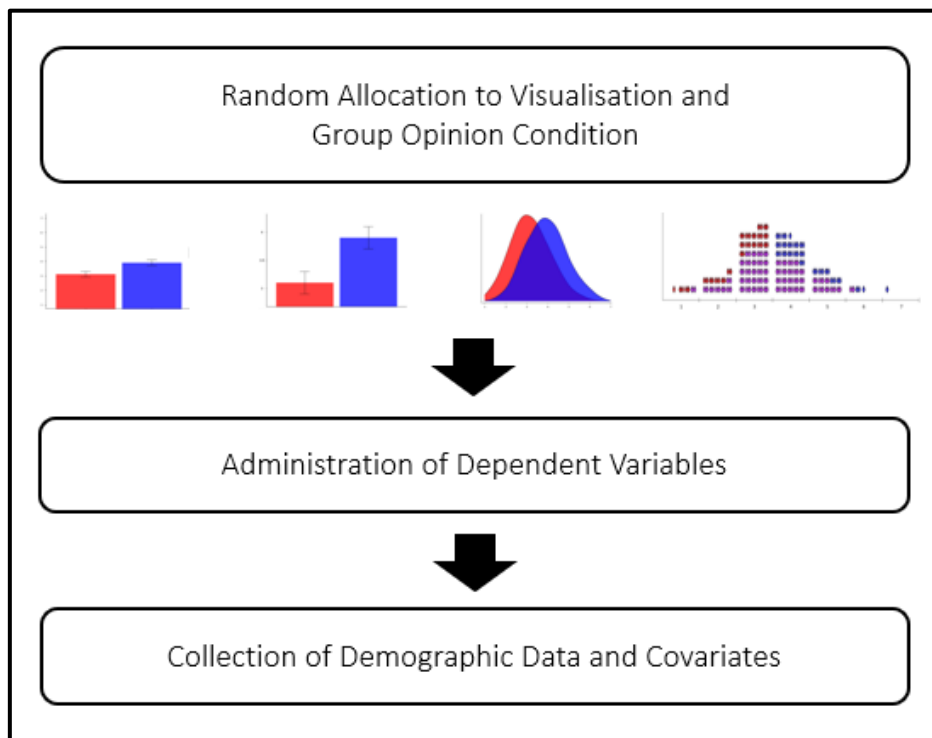
| Identity Characteristic | Sample Composition |
|-------------------------|-----------------------------------------------------------------------------------------------------------|
| Sex | Female: 49.08% ($n = 321$), Male: 48.78% ($n = 319$), Other: 2.14% ($n = 14$) |
| Age | $M = 35.56$, Range: 18.92 |
| Education | 53.21% ($n = 348$) had attained a bachelor's degree or higher |
| Political Affiliation | Democrat: 28.44% ($n = 186$), Republican: 35.78% ($n = 234$), Independent/Other: 35.78% ($n = 234$) |

Among participants who openly identified as Democrat or Republican, 48.39% ($n = 90$) categorised themselves as “strong” Democrats and 45.73% ($n = 107$) categorised themselves as “strong” Republicans. Among those who openly identified as Independent/Other, when forced to choose which of the two major US political parties they considered themselves closest to, 65.81% of this group chose Democrat ($n = 154$) and 34.19% ($n = 80$) chose Republican. Ideologically, the sample skewed slightly liberal ($M = 47.04$ on a 0 to 100 ideological self-placement scale).

Design

The study utilized a 4x2 (mode of presentation by group type), between-subjects design (Figure 6).

Figure 6
Design Outline for Study 2

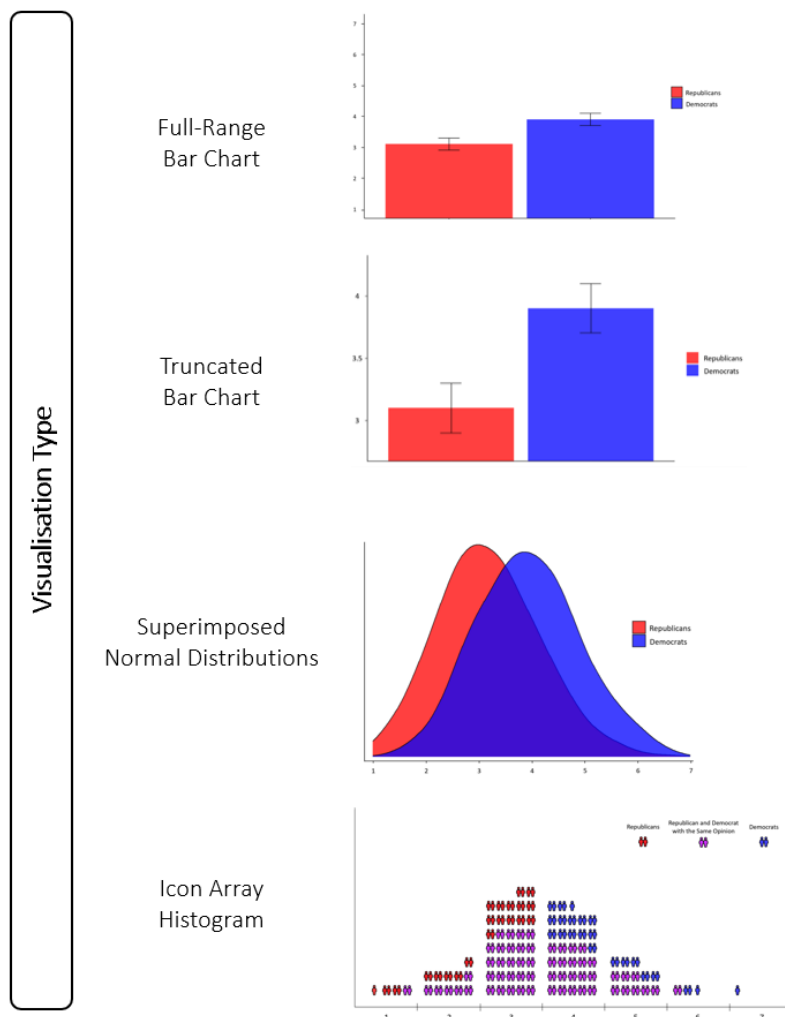


Procedure

The same pair of artificial datasets from Study 1 were used to create four modes of data presentation (see Figure 7), to one of which participants were randomly allocated. These modes included:

1. Full-Range Bar Chart
2. 1.5 SD Bar Chart (henceforth referred to simply as the “truncated bar chart”)
3. Superimposed Normal Distributions
4. Icon Array Histogram

Figure 7
Condition Stimuli for Study 2



Each visualisation presented was said to depict Republican and Democrat opinions on a “newly proposed immigration policy” (though the details of the policy were intentionally unspecified). Within each mode, half the participants were told the data featured was the opinion of partisan elites (i.e., “200 members of

Congress”) whereas half were told it was the opinion of “ordinary” partisans (i.e., “2,000 U.S. citizens”)³¹. Below is the statement which accompanied each visualisation:

“[200 members of Congress OR 2,000 US citizens] (100 Republican and 100 Democrat OR 1,000 Republican and 1,000 Democrat) were asked to rate where they stand - on a scale from 1 (“Do Not Support at All”) to 7 (“Support Completely”) - on a newly proposed immigration policy. The graph above depicts the [average OR full range of] responses from each group.”

Finally, after exposure to the data via one of eight unique visualisation conditions, participants were presented with a set of dependent variables, randomised so as to mitigate possible order effects, which aimed to assess perceptions of ideological and affective polarisation, estimates of intergroup opinion overlap, levels of intergroup agreement *relative to expectations*, and perceptions of group discreteness (see *Measures* for details of item wording and scoring).

Measures

The study featured the same estimate of intergroup opinion overlap item from Study 1, but also contained the following additional measures:

Perceived Issue Polarisation: Participants were asked to indicate their response (via a 101-point scale that ranged from “Not at All Polarising” to “Extremely Polarising”) to the item “Based on the graph, how polarising do you believe this issue is for the two groups?”

Perceived Affective Polarisation: Participants were asked to indicate their response (via a 101-point scale that ranged from “Extremely Cold/Negative” to “Extremely Warm/Positive”) to the item “Based on the graph, how warm (i.e., positive) or cold (i.e., negative) do you think members of each group would feel towards members of the other group on this issue?”

Level of Intergroup Agreement Relative to Expectations: Participants were asked to indicate their response (via a 5-point scale that ranged from “They seem to agree much less than I would have thought” to “They seem to agree much more than I would have thought”) to the item “Based on the graph, how would you classify the level of agreement between the two groups on this issue?”

Perceived Group Discreteness: Participants were asked to indicate their response (via a 6-point scale that ranged from “Strongly Disagree” to “Strongly Agree”) to the item “Based on the graph, how much do you

³¹ The decision was made to use a larger multiple for the citizen polarisation condition as a way to assuage any potential sample size concerns among more well-informed participants. While this causes the design to deviate slightly from being strictly controlled in the sense that all conditions are perfectly identical (with the obvious exception of the other independent variables), we argue that, relative to the total population of members of Congress and the total population of members of the U.S. public, 2,000 is actually a more proportionate figure.

agree with the following statement: when it comes to this issue, Republicans and Democrats just seem to be two completely different types of people?”³²

Hypotheses

Study 2 featured the following pre-registered (OSF URL: <https://osf.io/qy3hx>) hypotheses (see Table 4).

Table 4

*Pre-Registered Hypotheses for Study 2*³³

| Label | Hypothesis |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H1 | Main effects of mode of presentation will be observed on measures of perceived issue polarisation. |
| H1 _A | The truncated bar chart will elicit greater perceptions of issue polarisation than the bar chart with the full-range y-axis. |
| H1 _B | The truncated bar chart will elicit greater perceptions of issue polarisation than the superimposed normal distributions. |
| H2 | Main effects of mode of presentation will be observed on measures of perceived affective polarisation. |
| H2 _A | The truncated bar chart will elicit greater perceptions of affective polarisation than the bar chart with the full-range y-axis. |
| H2 _B | The truncated bar chart will elicit greater perceptions of affective polarisation than the superimposed normal distributions. |
| H3 | Main effects of mode of presentation will be observed for the level of agreement relative to expectations item. |
| H3 _A | Superimposed normal distributions will elicit responses that indicate greater levels of “surprise” about agreement (i.e., “They seem to agree much more than I would have thought”) than the truncated bar chart. |
| H4 | The superimposed normal distributions will elicit estimates of intergroup opinion overlap that will most closely approach the actual level of overlap and estimates. |
| H5 | Main effects of mode of presentation will be observed for the group discreteness measure. |

Results

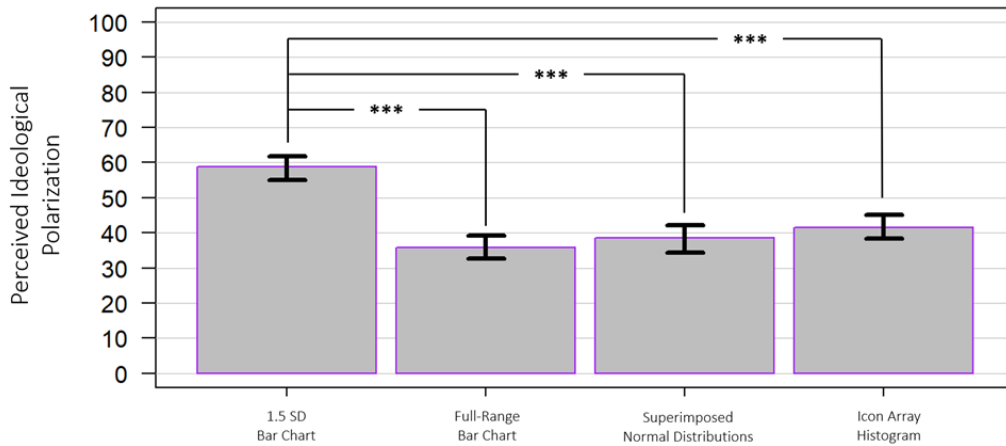
To assess the main effects of mode of presentation on perceived ideological polarisation [H1], a 4x2 (mode by group type) ANOVA was conducted. Results indicated a significant main effect of mode of presentation on perceptions of ideological polarisation [H1] ($F(3, 646) = 34.779, p < .001, \eta_p^2 = 0.139$; see Figure 8). Tukey HSD post-hoc tests confirmed highly-significant (i.e., $p < .001$) pairwise differences between perceptions elicited by the truncated bar chart versus the three other modes, with t-tests confirming the significant differences between the truncated bar chart ($M = 58.76, SD = 23.00$) and the full-range bar chart [H1_A] ($M = 35.91, SD = 21.53; t(325) = 9.280, p < .001, d = 1.03$), the superimposed normal distribution [H1_B] ($M = 38.51, SD = 23.47; t(317) = 7.786, p < .001, d = 0.87$), and icon array histogram ($M = 41.62, SD = 20.93; t(328) = 7.089, p < .001, d = 0.78$).

³² Our discreteness measure was inspired by (and partially adopted from) Bastian and Haslam’s (2006) *Essentialization Index* as well as Dweck’s (2013) *Implicit Theories Index* (and adapted to fit a political context).

³³ Additional hypotheses and their corresponding analyses can be found in Appendix D.

Figure 8

Perceptions of Ideological Polarisation by Mode



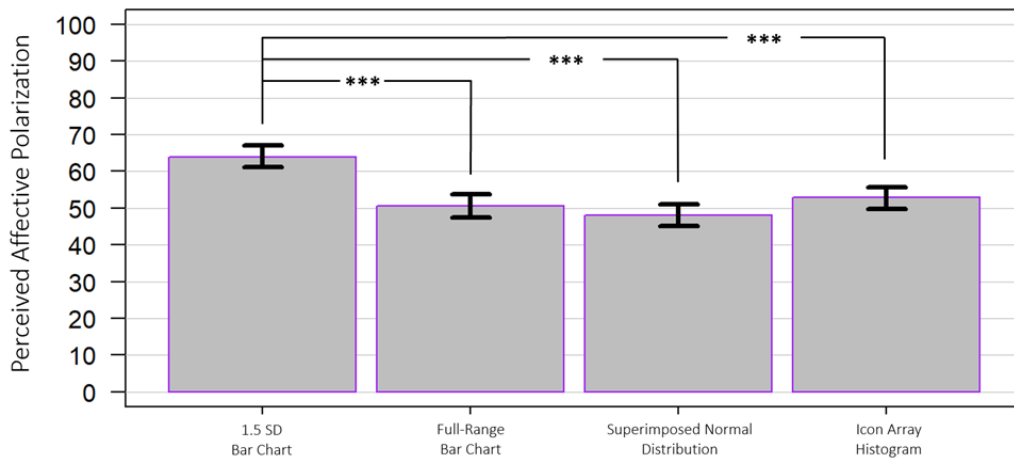
* $p < .05$, ** $p < .01$, *** $p < .001$

The same ANOVA revealed no significant main effect of group type – elite ($M = 42.73$, $SD = 23.83$) versus citizen ($M = 44.56$, $SD = 23.96$) – on perceptions of ideological polarisation ($F(1, 646) = 0.853$, $p = .356$, $\eta_p^2 = 0.001$) nor any significant interactions between mode and perceived ideological polarisation ($F(3, 646) = 0.871$, $p = .456$, $\eta_p^2 = 0.004$).

A similar 4x2 ANOVA assessing the impact of mode and group type on affective polarisation [H2] found similarly significant main effects of mode (albeit with a smaller overall effect size): $F(3, 646) = 20.173$, $p < .001$, $\eta_p^2 = 0.086$ (see Figure 9). Tukey HSD post-hoc tests confirmed highly-significant (i.e., $p < .001$) pairwise differences between perceptions elicited by the truncated bar chart versus the three other modes, with t-tests confirming the significant differences between the truncated bar chart ($M = 63.99$, $SD = 19.33$) and the full-range bar chart [H2_A] ($M = 50.60$, $SD = 21.12$; $t(325) = 5.976$, $p < .001$, $d = 0.66$), the superimposed normal distribution [H2_B] ($M = 47.97$, $SD = 20.20$; $t(317) = 7.237$, $p < .001$, $d = 0.81$), and icon array histogram ($M = 52.86$, $SD = 18.93$; $t(328) = 5.282$, $p < .001$, $d = 0.58$). Once again, no evidence for main effects of group type (elite ($M = 54.93$, $SD = 20.79$) versus citizen ($M = 52.75$, $SD = 20.71$)) ($F(1, 646) = 2.233$, $p = .136$, $\eta_p^2 = 0.003$) or interaction effects ($F(3, 646) = 1.076$, $p = .358$, $\eta_p^2 = 0.005$) were found.

Figure 9

Perceptions of Affective Polarisation by Mode




* $p < .05$, ** $p < .01$, *** $p < .001$

A 4x2 ANOVA on the level of intergroup agreement relative to expectations dependent variable also revealed significant main effects of mode [H3] ($F(3, 646) = 45.425, p < .001, \eta_p^2 = 0.174$; see Figure 10). Notably, “They seem to agree about as much as I would have thought” was the modal response for the truncated bar chart ($M = 3.09, SD = 0.97$) while the full-range bar chart ($M = 4.03, SD = 0.85$), superimposed normal distributions ($M = 4.11, SD = 0.89$), and icon array histogram ($M = 3.86, SD = 0.82$) all had “They seem to agree slightly more than I would have expected” as their modal response. T-tests determined that superimposed normal distributions did elicit greater levels of “surprise” about intergroup agreement relative to truncated bar chart [H3_A] ($t(317) = 9.851, p < .001, d = 1.10$). No main effects of group type (elite ($M = 3.76, SD = 0.99$) versus citizen ($M = 3.79, SD = 0.94$) ($F(1, 646) = 0.232, p = .630, \eta_p^2 = 0.000$) or interaction effects ($F(3, 646) = 0.819, p = .484, \eta_p^2 = 0.004$) were found.

Figure 10

Perceived Levels of Intergroup Agreement Relative to Expectations by Mode

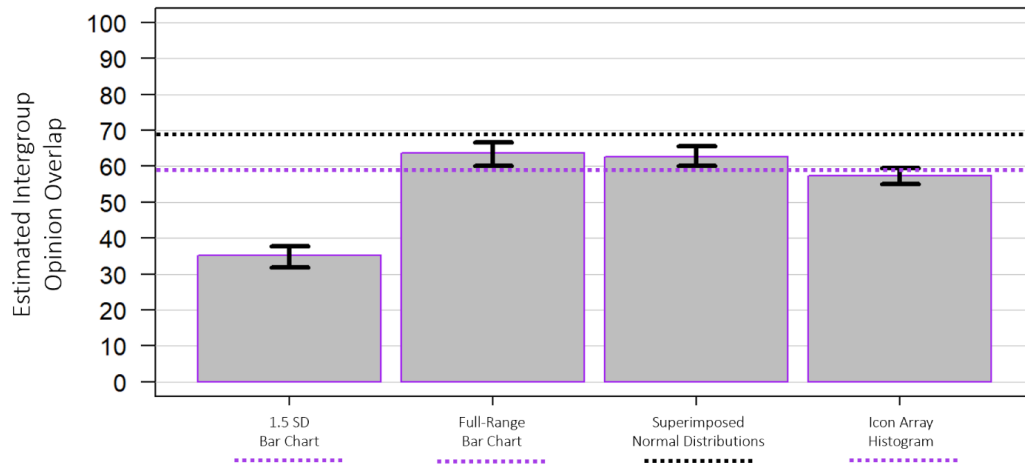


* $p < .05$, ** $p < .01$, *** $p < .001$

Main effects were observed for mode on estimates of intergroup opinion overlap: $F(3, 646) = 89.756$, $p < .001$, $\eta_p^2 = 0.294$ (see Figure 11). When each mode was compared to its' appropriate accuracy benchmark (i.e., a PCR of 69% for the superimposed normal distributions and a PCS of 59% for the remaining three modes), we find that the icon array histogram ($M = 57.25$, $SD = 14.18$) is the only mode that promotes estimates that do not differ significantly from its' benchmark ($t(168) = 1.6002$, $p = .111$, $d = 0.12$). The full-range bar chart ($M = 63.65$, $SD = 21.15$), truncated bar chart ($M = 35.11$, $SD = 18.67$), and superimposed normal distributions [H4] ($M = 62.61$, $SD = 16.75$) all catalyse estimates which differ significantly from their respective accuracy benchmarks ($t(165) = 2.8329$, $p = .005$, $d = 0.22$, $t(160) = 16.237$, $p < .001$, $d = 1.28$, and $t(157) = 4.7911$, $p < .001$, $d = 0.38$, respectively).

Figure 11

Estimates of Intergroup Opinion Overlap by Mode




* $p < .05$, ** $p < .01$, *** $p < .001$

Note. The purple and black horizontal lines refer to the PCS and PCR measures, respectively – the accuracy benchmarks used in the analysis.

Finally, a similar 4x2 ANOVA was conducted for the perceived group discreteness dependent variable (truncated bar chart ($M = 4.20$, $SD = 1.14$), full-range bar chart ($M = 2.84$, $SD = 1.12$), superimposed normal distributions ($M = 2.65$, $SD = 1.10$), icon array histogram ($M = 3.07$, $SD = 1.16$)), which also found significant main effects of mode [H5] ($F(3, 646) = 61.808$, $p < .001$, $\eta_p^2 = 0.223$; see Figure 12) but no main effects of group type (elite ($M = 3.15$, $SD = 1.28$) versus citizen ($M = 3.23$, $SD = 1.27$) ($F(1, 646) = 0.664$, $p = .416$, $\eta_p^2 = 0.001$) nor any significant interactions ($F(3, 646) = 2.511$, $p = .058$, $\eta_p^2 = 0.012$).

Figure 12

Perceptions of Group Discreteness by Mode



* $p < .05$, ** $p < .01$, *** $p < .001$

Note. Significance levels featured are based on Tukey HSD post-hoc tests.

Discussion

Study 2 demonstrated that “mode of presentation” effects, observed primarily in non-political contexts by Hanel et al. (2019), seem to extend to politically-relevant perceptions, as well. More specifically, presenting participants with disparate visualisations based on identical underlying data, attributed to a salient, divisive political issue (i.e., immigration), resulted in significantly different perceptions for both main types of polarisation (i.e., ideological and affective). Such a finding is important as it suggests that individuals can come to significantly different conclusions as to how divided groups are on an issue – an issue for which they likely hold some sort of prior beliefs – simply by altering the visualisation format used to depict partisan opinion data.

When examining the effects of specific modes, a relatively clear pattern emerges between the truncated bar chart and all other modes (i.e., the full-range bar chart, the superimposed normal distributions, and the icon array histogram). For both ideological and affective polarisation, the perceptions elicited by the truncated bar chart were significantly greater than the perceptions elicited by the remaining visualisation conditions, implying that exposure to the former is likely to *exacerbate* polarised perceptions relative to exposure to the remaining modes. However, while mitigating polarised perceptions is likely a desirable outcome (especially in light of the data that suggests a multi-decade pattern of *overestimation* of polarisation; Westfall et al., 2015), it’s important to also couch the performance of each mode in terms of the *accuracy* of perceptions it engenders. Here, while the icon array histogram is the only mode which promotes estimates of intergroup opinion overlap which do not significantly differ from its accuracy benchmark, once again we see that the truncated bar chart promotes estimates which are *significantly* more inaccurate (i.e., vastly *underestimating* the similarity of shared group opinions) than all other modes. Thus, depictions via the truncated bar chart seem to produce a set of especially unfortunate outcomes (relative to the other modes): both worsening perceptions of polarisation while dramatically impairing perception accuracy.

While main effects were observed for both ideological and affective polarisation, the *type* of group to which the opinions were attributed (i.e., elite versus citizen) did not seem to play a meaningful role in perceptions or estimations. Across all analyses, we failed to find evidence for either main effects of group type or any significant interactions between this variable and mode of presentation. While scholars debate the dynamics under which certain groups may exert differential influence over political opinions (e.g., Druckman, Peterson, & Slothuus, 2013; Slothuus & Bisgaard, 2021), the current study finds no data to suggest asymmetrical impacts within this design context.

Beyond perceptions of polarisation, the study also found main effects of mode on level of perceived agreement *relative to expectations*. This item was included as we believed such a metric would shed light on how individuals *expect* to see political data on key divisive issues depicted. Ultimately, it appeared that

intergroup data depicted via the truncated bar chart most closely aligned with how individuals *expected* to see the groups (as evidenced by the fact that it was the only visualisation format which produced a modal response of “They seem to agree about as much as I would have thought” whereas all other visualisations elicited a higher degree of “surprise” at the level of agreement depicted). Such a finding is revealing in that it suggests that visualisations disseminated via truncated bar charts – which, by nature, inherently accentuate group differences – depict the level of intergroup agreement that align with most American expectations. On the contrary, modes such as full-range bar charts and icon array histograms (which seem to promote more *accurate* perceptions of intergroup agreement) generate responses which indicate the level of agreement depicted violates participant expectations (as the groups seem to agree more than they would have thought). This result is troubling in that it implies that the American public may have become so immersed in divisive commentary and rhetoric that now only modes which visually amplify the level of cross-party divergence are seen as “expected.”

Finally, we believe the finding that mode of presentation produces significant main effects for perceptions of group discreteness is of particular note. Our results indicate that a simple alteration of visualisation format can yield significant shifts in how discrete political groups seem to be, with truncated bar charts promoting heightened perceptions of group discreteness relative to other modes. This is troubling as essentialist beliefs, to which discreteness contributes, can intensify the perceptions of category boundaries, which can subsequently decrease the likelihood of (healthy) interaction with members of “other” groups (e.g., Zagefka, Nigbur, Gonzalez, & Tip, 2013; Williams & Eberhardt, 2008). Additionally, such beliefs also appear to be a necessary component for certain forms of dehumanisation (Haslam, 2006) – a particularly disconcerting fact given the hostility and toxicity that already seems to characterise the American political landscape. The optimistic interpretation of the results, however, seems to be that modes like full-range bar charts and icon array histograms, without altering the underlying data or misrepresenting group opinions, can significantly *decrease* perceptions of discreteness relative to truncated bar charts.

While Study 2 was a useful next step in our research, it is not without critique. For instance, the discreteness variable, though borne out of inspiration from validated essentialism measures (i.e., Bastian & Haslam, 2006; Dweck, 2013), is itself likely an imperfect means of capturing attitudes on the concept. But perhaps most notably, while Study 2’s design sought to increase ecological validity relative to Study 1’s design by attributing a specific issue to the data being visualised (i.e., immigration), it’s fair to question whether this addition is specific *enough* to enable meaningful consideration among participants. While immigration is undoubtedly an issue of which virtually all voters are familiar, there are many facets of immigration policy and the immigration debate, some of which individuals may hold conflicting opinions on. Thus, while this design progression is a step towards *greater* specificity, we cannot be sure we will observe similar responses

when, for instance, participants are presented with *actual* specific policies (and the corresponding party opinions on them).

Overall, Study 2 provided the first evidence that “mode of presentation” effects might be capable of extending into the political domain. Alterations in the manner in which data was presented not only resulted in significant shifts in perceptions of polarisation and estimates of intergroup opinion overlap, but also seem to impact perceptions such as group discreteness – a variable which is not inherently political, per se, but which carries potentially serious consequences in regards to overall intergroup relations.

Study 3: Investigating the Intersection of Mode of Presentation and Source Effects³⁴

While Study 2 provided preliminary evidence to suggest that “mode of presentation” effects continue to emerge for politically-relevant perceptions (even when the data is attributed to a salient and divisive political issue for which participants likely hold prior beliefs about the level of intergroup polarisation), Study 3 sought to further contextualise the phenomena in several meaningful ways.

First of all, within Study 3, we’ve added a control condition; specifically, a “no visualisation” control, in which participants assigned to the condition will generate responses to items “based on what they know about Republicans and Democrats on this issue.” We believe the inclusion of this condition is a critical addition to the design if we wish to be able to assess the real-world utility of the intervention. Whereas the prior studies were only able to determine whether (and, if so, to what degree) the different visualisation formats differed *from one another*, adding the no visualisation condition provides us with a meaningful baseline (i.e., a proxy for “here is how individuals might respond to this item *devoid* of any visual depiction of data”), at which point we can juxtapose both how modes impact perceptions relative to visualisation alternatives as well as relative to no visualisation exposure. Broadly, we hypothesise that participants assigned to the no visualisation condition, who will be forced to respond to items based only on their prior knowledge (which, research shows, is likely to lead to an *overestimation* of actual degrees of polarisation; e.g., Westfall et al., 2015), will display response patterns which will not differ to participants exposed to the truncated bar chart, which has consistently catalysed higher perceptions of polarisation through two studies than alternative modes (e.g., the full-range bar chart, the superimposed normal distributions, etc.). Conversely, we believe participants exposed to the full-range bar chart or icon array histograms will come to hold perceptions of polarisation that are significantly less than either those exposed to the truncated bar chart or no visualisation at all.

Secondly, Study 3 sought to determine whether the effects of mode would extend beyond issue-specific perceptions of polarisation to impact *overall* perceptions of polarisation. In other words, could data

³⁴ University of Cambridge ethics identification reference PRE.2022.049

presentation interventions impact not only perceptions of the issue being depicted, but also perceptions of overall levels of polarisation, as well? To test this, we measure participant perceptions of overall ideological polarisation as well as general (i.e., non-issue-specific) affective polarisation both pre- and post-visualisation exposure. While the design is not a “misperception *correction*” paradigm, per se (as it does not seek to deliberately “fix” a skewed perception, but simply seeks to examine how perceptions may be *impacted* by disparate mode exposure, either in a more *or* less accurate manner), there is precedent in the literature that exposure to interventions which alter specific perceptions may cascade to ultimately affect perceptions in adjacent domains, as well. For instance, Ahler and Sood (2018) found that changing perceptions of the *composition* of political outgroups subsequently impacted perceptions concerning outgroup levels of attitudinal extremity as well as one’s *own* feelings of social distance from them. Building off of the rationale posited above, due to the fact that many participants will likely hold exaggerated perceptions of polarisation prior to the intervention, we also venture a hypothesis that participants assigned to the full-range bar chart and the superimposed normal distribution conditions will experience significant reductions in these two polarisation measures due to their mode’s demonstrated ability to promote more *accurate* (and thus less polarised) perceptions, whereas participants assigned to the truncated bar chart condition (as well as those assigned to the no visualisation control) will not exhibit significant pre-post changes in levels of either overall ideological polarisation or general affective polarisation.

Finally, Study 3 features the inclusion of specific *sources* to the data visualisations, enabling us the ability to better understand if and how source attribution may moderate observed effects. A number of studies have concluded that similar messages, when delivered by different political representatives or groups, can alter the efficacy and impact of an identical appeal (Maier, Adam, & Maier, 2017; Baum & Groeling, 2009). Study 3 incorporates a design which gives participants the impression that the data has been disseminated by either one of two partisan news outlets (i.e., Fox News or CNN) or an unspecified source, which permits us the ability to test for possible “source effects” (e.g., Laustsen & Petersen, 2016; Wilson & Sherrell, 1993) and determine whether significant interactions may materialise between mode of presentation and the source to which the data have been attributed. Moreover, by adding a source variable to the design, we’re able to (once again) potentially increase the ecological validity of the experiment (as most political data is coming from *some* identified source). While we do not venture any specific hypotheses related to source effects, we remain interested to see whether main effects of source or interactions between source and mode will be observed.

In addition to the new experimental elements cited above, Study 3 seeks to determine if the overall patterns observed in the prior studies continue to persist after the addition of the no visualisation condition and the source attributions. Broadly, we hypothesise that significant main effects of issue-specific ideological and affective polarisation will still be observed (though we remain open to the fact that, if the no visualisation

condition elicits responses similar to one of the other modes, we may not see such pronounced variance between conditions). Moreover, we also hypothesise that the *relationships* between different modes on these two variables as well as other previously used variables (i.e., estimates of intergroup opinion overlap) will remain consistent (i.e., as observed in our prior studies), with the truncated bar chart eliciting higher perceptions of ideological and affective polarisation than either the full-range bar chart or superimposed normal distributions, lower estimates of intergroup opinion overlap than either of the aforementioned two modes, and less accurate perceptions than either of the alternative modes.

Method

Participants

A sample of American participants were recruited via the Prolific Academic platform and paid an approximate rate of £7.50/hour. Following exclusions for non-consent, attention check failure³⁵, or incomplete data, the final sample was 1,534.³⁶ Table 5 summarises the demographic composition of the sample.

Table 5

Composition of Sample for Study 3

| Identity Characteristic | Sample Composition |
|-------------------------|-----------------------------------------------------------------------------------------------------------|
| Sex | Female: 49.80% ($n = 764$), Male: 48.44% ($n = 743$), Other: 1.76% ($n = 27$) |
| Age | $M = 37.82$, Range: 18-80 |
| Education | 50.39% ($n = 773$) had attained a bachelor's degree or higher |
| Political Affiliation | Democrat: 27.77% ($n = 426$), Republican: 32.14% ($n = 493$), Independent/Other: 40.09% ($n = 615$) |

Among participants who openly identified as Democrat or Republican, 51.41% ($n = 219$) categorised themselves as “strong” Democrats and 44.22% ($n = 218$) categorised themselves as “strong” Republicans. Among those who openly identified as Independent/Other, when forced to choose which of the two major

³⁵ Similar to Study 2, the attention check explicitly asked participants to select a specific response (i.e., “Slightly Disagree”) amongst a total of six possible responses to confirm that they were paying attention. 94.11% of participants responded correctly and passed the attention check.

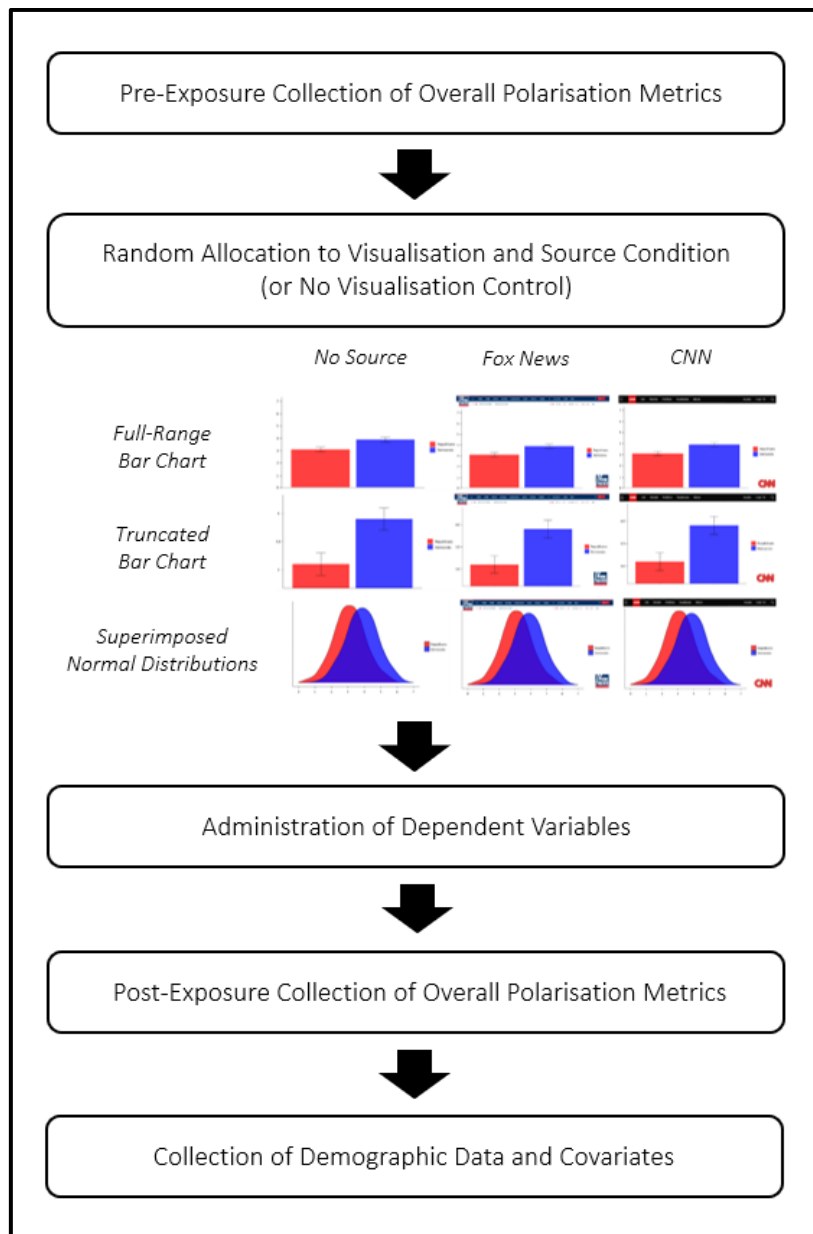
³⁶ A priori power analyses were conducted to determine the sample size required to achieve sufficient statistical power to detect a small-to-medium effect size (i.e., Cohen's f of 0.15) with a power level of 0.90 and an alpha threshold of 0.05. The analysis indicated a sample size of approximately 89 participants per group would be required. The effect size was chosen based on the findings from Study 2, where the effect sizes observed were slightly smaller than those found in Study 1. Were we to assume, as we continued to increase the ecological validity of the design and stimuli, that effect sizes would incrementally decrease, we wished to conduct a more conservative power analysis (which assumed smaller effects). Ultimately, given our budget for the project, we recruited more than was required A) to account for potential dropout and B) in the hopes of detecting more subtle effects within the analyses as experimental sophistication increased. Additionally, prior to exclusions, the total sample was 1,630. For those who provided data, the sample composition prior to exclusions had a mean age of 37.82, was comprised of 764 females, 743 males, and 27 other (with 96 non-responses), and had 426 self-reported Democrats, 493 self-reported Republicans, and 615 self-reported other (with 96 non-responses).

US political parties they considered themselves closest to, 64.88% of this group chose Democrat ($n = 399$) and 35.12% ($n = 216$) chose Republican. Ideologically, the sample skewed slightly liberal ($M = 46.53$ on a 0 to 100 ideological self-placement scale).

Design

The study employed a 3x3 (mode of presentation by source), between-subjects design (with an additional, control condition where participants were provided with no visual depiction of the data; Figure 13).

Figure 13
Design Outline for Study 3



Procedure

Participants were randomly assigned to see the same artificial dataset utilised in Study 1 and Study 2 depicted via one of three modes of presentation³⁷:

1. A full-range bar chart
2. A truncated (i.e., 1.5 SD) bar chart
3. Superimposed normal distributions

Accompanying each data visualisation was a source attribution. Participants were randomly allocated to see one of three source conditions:

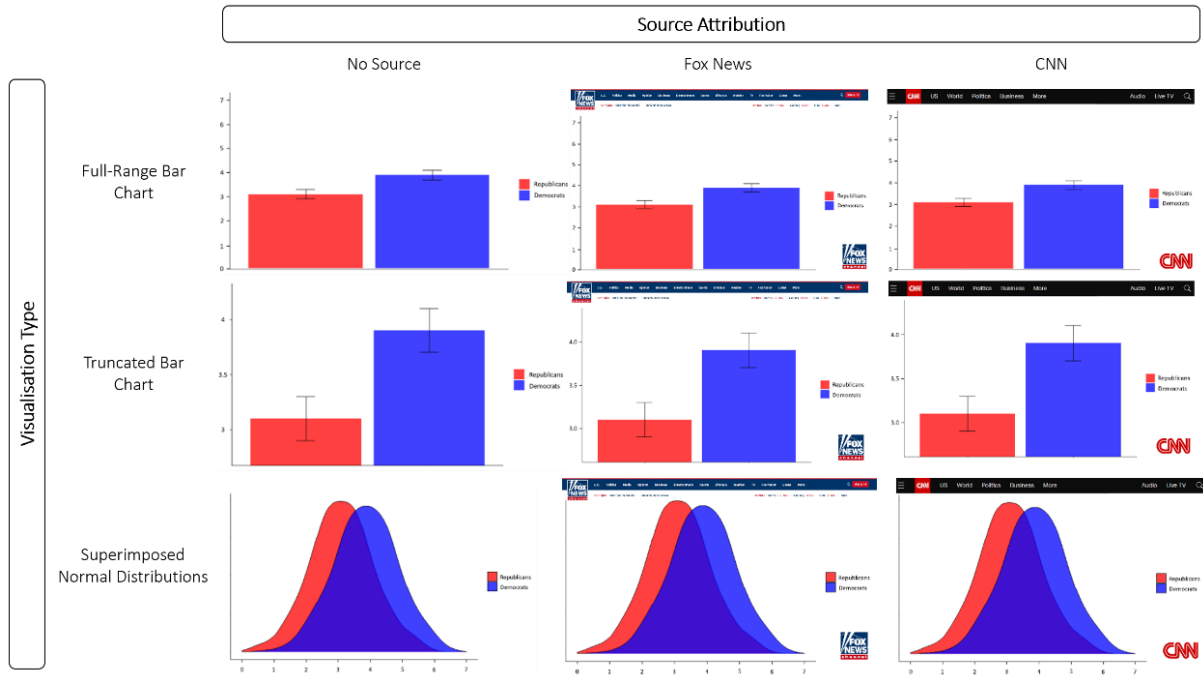
1. No source
2. Fox News
3. CNN

Figure 14 depicts the images of the stimuli used in each of the nine experimental conditions.

³⁷ We believe it's important to note here why icon array histograms were not included in this study. Originally, icon array histograms had not "performed as well" (i.e., they had not reduced polarisation or increased perception accuracy) as much as modes like the full-range bar chart and superimposed normal distributions had. However, during the data collection stage of Wave 3, we realised that our first iteration of the icon array histograms had been constructed in a problematic manner. More specifically, we had designed them in such a way so that we had given the same amount of "visual space" to a *single non-overlapping icon* as we did for a *paired agreement icon* (i.e., the purple ones). Consequently, we had essentially allocated twice the visual space to disagreement as we had to agreement, which likely had been unfairly skewing participant perceptions (which we believed might have contributed to the poor performance). We re-ran the first and second studies with the new, corrected icon array histograms and received the results now presented in the first two studies in this chapter. However, as we had already begun data collection for Study 3 (and since Study 3 was particularly resource-intensive due to the sheer number of conditions and, thus, the number of participants required to achieve sufficient statistical power), we decided to omit the icon array histograms from this study and re-introduce them in Study 4.

Figure 14

Condition Stimuli for Study 3



All participants assigned to an experimental condition were presented with the following statement accompanying their data visualisation:

“1,000 people (500 Republicans and 500 Democrats) were recently asked to rate where they stand – on a scale from 1 (“Do Not Support at All”) to 7 (“Support Completely”) – on a newly proposed immigration policy. The graph above³⁸ was featured on [a major news station/Fox News/CNN], and depicts the [average/full range of] responses from each group.”

Additionally, a portion of the sample was assigned to the “no visualisation” control condition. Here, participants were exposed to no visual depiction of the data and instead were simply asked to respond to questions “based on what [they] know about Republicans and Democrats.”

After exposure to the data via one of nine unique visualisation conditions (or the no visualisation control), participants were presented with a set of dependent variables, randomised so as to mitigate possible order effects, which aimed to assess perceptions of issue-specific ideological polarisation, issue-specific affective polarisation on the issue, and estimates of intergroup opinion overlap. Additionally, a pre-post measure of *overall* ideological and *general* affective polarisation was administered (see *Measures* for details of item wording and scoring).

Measures

³⁸ In the “no visualisation” condition, this sentence was omitted.

The study employed many of the same measures featured in Study 1 and Study 2 (i.e., perceived issue polarisation, perceived affective polarisation, and estimates of intergroup opinion overlap). Additionally, however, while the items used in previous studies captured *issue-specific* assessments (i.e., related to “a newly proposed immigration policy”), the current study also featured a pre- and post-assessment of perceived *overall* issue polarisation and *general* (i.e., non-issue specific) affective polarisation to determine whether any effects of mode of presentation and/or source on issue-specific measures might also generalise more broadly to overall attitudes towards and beliefs about Republicans and Democrats. The items below were presented at the very beginning of the study and then again upon completion of the condition-specific assessment tasks:

Overall Perceived Issue Polarisation: Participants were asked to respond, on a 0 (“Not at All Polarised”) to 100 (“Extremely Polarised”), to the following item: “Overall, how polarised do Republicans and Democrats seem to be on important issues?”

General Affective Polarisation: Participants were asked to respond, on a 0 (“Extremely Cold/Negative”) to 100 (“Extremely Warm/Positive”) scale, to the following item: “How warm (i.e., positive) or cold (i.e., negative) would you say you feel towards [Republicans/Democrats]?”

Hypotheses

Study 3 featured the following pre-registered (OSF URL: <https://osf.io/pgu57>) hypotheses (see Table 6).

Table 6

Pre-Registered Hypotheses for Study 3


| Label | Hypothesis |
|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H1 | Main effects of mode of presentation will be observed on measures of perceived issue polarisation. |
| H1 _A | The truncated bar chart will elicit greater perceptions of issue polarisation than the full-range bar chart. |
| H1 _B | The truncated bar chart will elicit greater perceptions of issue polarisation than the superimposed normal distributions. |
| H1 _C | The no visualisation condition will elicit perceptions of issue polarisation that will be significantly greater than full-range bar chart or the superimposed normal distribution, but which will not differ significantly from the truncated bar chart. |
| H2 | Main effects of mode of presentation will be observed on measures of perceived affective polarisation |
| H2 _A | The truncated bar chart will elicit greater perceptions of affective polarisation than the full-range bar chart. |
| H2 _B | The truncated bar chart will elicit greater perceptions of affective polarisation than the superimposed normal distributions. |
| H3 | The truncated bar chart will promote estimates of intergroup opinion overlap that are significantly lower than those elicited by the full-range bar chart or the superimposed normal distributions. |
| H4 | The superimposed normal distributions and the full-range bar chart will elicit estimates of group opinion overlap that will more closely approach their respective overlap accuracy benchmarks (i.e., PCS of 59% and PCR of 69%, respectively) than the truncated bar chart (i.e., PCS of 59%). |
| H5 _A | The superimposed normal distributions and full-range bar charts will promote a significant pre-post reduction in perceived overall ideological polarisation; the no visualisation condition and the truncated bar chart will show no significant pre-post differences on this measure. |
| H5 _B | The superimposed normal distributions and full-range bar charts will promote a significant pre-post reduction in general affective polarisation; the no visualisation condition and the truncated bar chart will show no significant pre-post differences on this measure. |

Results

To determine if significant differences emerged between the control group and the experimental conditions on perceived issue specific ideological polarisation (Figure 15), pairwise t-tests were conducted between the no visualisation control and each of the three experimental mode of presentation conditions (collapsed to include all intra-mode source conditions). These analyses determined that participants assigned to the no visualisation control exhibited significantly higher perceptions of issue-specific ideological polarisation ($M = 78.62, SD = 21.12$) post-Bonferroni corrections than those assigned to the truncated bar chart condition ($M = 64.01, SD = 21.65; t(610) = 7.335, p < .001, d = 0.68$), the full-range bar chart condition ($M = 37.71, SD = 21.96; t(602) = 20.282, p < .001, d = 1.88$), and the superimposed normal distribution condition ($M = 41.57, SD = 22.94; t(630) = 17.880, p < .001, d = 1.65$) [H1_C]. A 3x3 (mode by source) ANOVA was then conducted to assess both main and interaction effects of mode and source on perceived issue-specific ideological polarisation across each of the experimental conditions, which revealed a significant main effect of mode [H1] ($F(2, 1368) = 185.690, p < .001, \eta_p^2 = 0.214$) but no significant main effect of source (Fox News ($M = 47.92, SD = 25.32$), CNN ($M = 47.79, SD = 25.13$), no source ($M = 47.50, SD = 24.63$); $F(2, 1368) = 0.187, p = .830, \eta_p^2 = 0.000$) nor a significant interaction between the two variables ($F(4, 1368) = 0.601, p = .662, \eta_p^2 = 0.002$). Additionally, pairwise t-tests determined that perceptions of issue polarisation were significantly greater in the truncated bar chart condition than in either the full-range bar chart condition [H1_A] ($t(900) = 18.116, p < .001, d = 1.21$) or the superimposed normal distribution condition [H1_B] ($t(928) = 15.325, p < .001, d = 1.01$).

Figure 15

Perceptions of Issue Polarisation by Mode



* $p < .05$, ** $p < .01$, *** $p < .001$


Note. Only significant differences between the no visualisation condition and each of the three experimental conditions has been plotted.

Pairwise t-tests were conducted between the no visualisation control and each of the three experimental mode of presentation conditions (collapsed to include all intra-mode source conditions) on perceived affective polarisation. These analyses determined that participants assigned to the no visualisation control exhibited significantly higher perceptions of issue-specific affective polarisation ($M = 71.24, SD = 23.90$) than those assigned to the full-range bar chart condition ($M = 55.96, SD = 21.49; t(602) = 7.436, p < .001, d = 0.69$), and the superimposed normal distribution condition ($M = 54.65, SD = 22.66; t(630) = 7.844, p < .001, d = 0.72$). However, post-Bonferroni corrections, the difference between the no visualisation control and participants assigned to the truncated bar chart condition failed to be significant ($M = 66.40, SD = 21.60; t(610) = 2.351, p = .019, d = 0.22$). A 3x3 (mode by source) ANOVA was conducted for perceived affective polarisation across all experimental conditions, which revealed a significant main effect of mode [H2] ($F(2, 1368) = 39.586, p < .001, \eta_p^2 = 0.055$) but also detected no significant main effect of source (Fox News ($M = 58.62, SD = 22.38$), CNN ($M = 58.70, SD = 22.50$), no source ($M = 59.56, SD = 22.78$); ($F(2, 1368) = 0.191, p = .826, \eta_p^2 = 0.000$) nor a significant interaction between the two independent variables ($F(4, 1368) = 1.891, p = .110, \eta_p^2 = 0.005$). Additionally, pairwise t-tests determined that perceptions of affective polarisation were significantly greater in the truncated bar chart condition than in either the full-range bar chart condition [H2_A] ($t(900) = 7.274, p < .001, d = 0.48$) or the superimposed normal distribution condition [H2_B] ($t(928) = 8.091, p < .001, d = 0.53$).

An exploratory 3x3 (mode by source) ANOVA was conducted to assess main and interaction effects of mode of presentation and source on estimates of intergroup opinion overlap. Results indicated a significant main effect of mode ($F(2, 1368) = 245.522, p < .001, \eta_p^2 = 0.264$; see Figure 16) but no significant main effect of source (Fox News ($M = 50.70, SD = 23.17$), CNN ($M = 51.53, SD = 22.67$), no source ($M = 51.04, SD = 22.78$); ($F(2, 1368) = 0.437, p = .646, \eta_p^2 = 0.000$) nor a significant interaction between the two variables ($F(4, 1368) = 0.227, p = .923, \eta_p^2 = 0.003$). As hypothesised [H3], the truncated bar chart promoted estimates of intergroup opinion overlap ($M = 34.40, SD = 17.96$) that were significantly lower than those elicited by the full-range bar chart ($M = 58.71, SD = 21.81; t(862.15) = 18.260, p < .001, d = 1.22$) or the superimposed normal distribution condition ($M = 59.91, SD = 18.97; t(928) = 21.044, p < .001, d = 1.38$). Critically [H4], while the truncated bar chart ($t(454) = 29.227, p < .001, d = 1.37$) and the superimposed normal distributions ($t(474) = 10.443, p < .001, d = 0.48$) significantly underestimated their accuracy benchmarks (i.e., PCS of 59% and PCR of 69%, respectively), the full-range bar chart promoted estimates that did not significantly differ from its' accuracy benchmark (i.e., a PCS of 59%; $t(446) = 0.280, p = .780, d = 0.01$).

Figure 16

Estimates of Intergroup Opinion Overlap by Mode




* $p < .05$, ** $p < .01$, *** $p < .001$

Note. The dotted purple line represents the PCS for the data (the appropriate benchmark against which to assess the accuracy of responses catalysed by the two bar chart modes) and the dotted black line represents the PCR for the data (the appropriate benchmark against which to assess the accuracy of responses catalysed the superimposed normal distribution mode).

Pairwise t-tests were conducted between the no visualisation control and each of the three experimental mode of presentation conditions (collapsed to include all intra-mode source conditions) on pre-post change in perceived overall ideological polarisation (Figure 17). These analyses determined that participants assigned to the no visualisation control exhibited a significantly lower pre-post decrease in perceptions of overall ideological polarisation ($M = -1.15$, $SD = 12.71$) post-Bonferroni corrections than those assigned to the full-range bar chart condition ($M = -8.95$, $SD = 17.62$; $t(377.51) = 5.944$, $p < .001$, $d = 0.47$) and the superimposed normal distribution condition ($M = -9.81$, $SD = 18.96$; $t(398.96) = 6.484$, $p < .001$, $d = 0.49$). However, t-tests determined that the pre-post decrease for participants assigned to the truncated bar chart condition was not significantly different to those in the no visualisation control ($M = -2.77$, $SD = 15.59$; $t(329.52) = 1.302$, $p = .194$, $d = 0.11$). A 3x3 (mode by source) ANOVA was conducted for the post-pre change in perceived overall ideological polarisation across the experimental conditions, which revealed a significant main effect of mode ($F(2, 1368) = 22.100$, $p < .001$, $\eta_p^2 = 0.031$) but no significant main effect of source (Fox News ($M = -7.56$, $SD = 18.22$), CNN ($M = -7.28$, $SD = 17.69$), no source ($M = -6.78$, $SD = 17.31$); $F(2, 1368) = 0.165$, $p = .848$, $\eta_p^2 = 0.000$) nor a significant interaction between the two variables ($F(4, 1368) = 0.512$, $p = .727$, $\eta_p^2 = 0.001$). As hypothesised [H5_A], t-tests determined the full-range bar chart ($t(883.14) = 6.617$, $p < .001$, $d = 0.44$) and the superimposed normal distributions ($t(934.85) = 7.665$, $p < .001$, $d = 0.50$) precipitated a significant pre-post reduction in perceived overall ideological polarisation. While the no visualisation condition did not promote a significant pre-post decline ($t(312) = 0.543$, $p = .588$, $d = 0.06$), the truncated bar chart *did* ($t(908) = 2.094$, $p = .037$, $d = 0.14$).

Figure 17³⁹

Change in Perceptions of Overall Ideological Polarisation by Mode



* $p < .05$, ** $p < .01$, *** $p < .001$

Note. Only levels of statistical significance between experimental conditions and the no visualisation control are depicted.

Finally, exploratory pairwise t-tests were conducted between the no visualisation control and each of the three experimental mode of presentation conditions (collapsed to include all intra-mode source conditions) on pre-post change in general affective polarisation. These analyses determined that participants assigned to the no visualisation control exhibited a significantly lower pre-post decrease in perceptions of overall ideological polarisation ($M = -0.49$, $SD = 11.64$) than those assigned to the superimposed normal distribution condition ($M = -3.48$, $SD = 14.62$; $t(331.4) = 2.612$, $p = .009$, $d = 0.21$). However, t-tests determined that the pre-post decrease for participants assigned to the truncated bar chart condition ($M = -0.78$, $SD = 11.68$; $t(610) = 0.266$, $p = .790$, $d = 0.02$) and the full-range bar chart condition ($M = -2.33$, $SD = 9.79$; $t(238.02) = 1.769$, $p = .078$, $d = 0.18$) were not significantly different to those in the no visualisation control post-Bonferroni corrections. An exploratory 3x3 ANOVA assessed the post-pre change in general affective polarisation across all experimental conditions and found significant main effects of mode ($F(3, 1368) = 5.686$, $p = .003$, $\eta_p^2 = 0.008$) but no significant main effect of source (Fox News ($M = -2.07$, $SD = 13.13$), CNN ($M = -1.92$, $SD = 12.93$), no source ($M = -2.65$, $SD = 10.67$); $F(2, 1368) = 0.461$, $p = .631$, $\eta_p^2 = 0.001$) nor a significant interaction between the two variables ($F(4, 1368) = 0.360$, $p = .837$, $\eta_p^2 = 0.001$). While the no visualisation condition ($t(312) = 0.131$, $p = .896$, $d = 0.01$) and the truncated bar chart ($t(908) = 0.376$, $p = .707$, $d = 0.02$)

³⁹ Please note that we have chosen to truncate the y-axis in this figure. A full-range y-axis for a chart like this would extend from -100 to 100 (as participants perceptions of overall ideological polarisation could *theoretically*, following the intervention, go from 0 to 100 – an increase of 100 points – or 100 to 0 – a decrease of 100 points). However, we say theoretically as, while it's technically *possible* for a participant to undergo a 100 point perception change as a result of the intervention, it is both A) highly improbable for any one participant to report such a drastic swing and B) virtually impossible for *the mean of an entire condition* to realistically approach this number. Consequently, we believe it's not only reasonable to employ a moderate degree of truncation under these circumstances, but we argue that it would be detrimental to overall comprehension to rigidly abstain from truncation. While the former does risk encouraging the audience to *overestimate* the strength of the effect of the intervention, the latter would almost certainly result in a visualisation wherein the highly-significant between-condition differences (of which there are several) were unnecessarily obscured due to a y-axis with a range of 200.

did not promote significant pre-post changes in general affective polarisation (as hypothesised [H5_B]), t-tests determined that the full-range bar chart ($t(892) = 1.146, p = .252, d = 0.08$) and the superimposed normal distributions ($t(948) = 1.785, p = .075, d = 0.12$) *also* did not precipitate any significant pre-post reduction on this measure.

Discussion

Beyond seeking to determine whether the “mode of presentation” effects observed in Study 1 and Study 2 would replicate as we added an additional layer of ecological validity (i.e., source attribution) to the design, Study 3 also sought to determine A) how modes would perform *relative to a control condition* (i.e., the no visualisation baseline) and B) whether mode effects observed for issue-specific perceptions would extend to perceptions of *overall* polarisation (both ideological and affective).

First and foremost, Study 3 provided evidence to suggest that the mode of presentation effects related to perceptions of both (issue-specific) ideological and affective polarisation (as documented in Study 2) continue to emerge even after data has been attributed to particular sources. This is an encouraging finding as it shows that, even as ecological validity is incrementally increased (i.e., adding a specific issue in Study 2, adding real sources in Study 3), mode of presentation effects continue to hold for perceptions for polarisation. While it is premature to declare that such effects would undoubtedly emerge were such an intervention to be deployed in the “real world” (as, for instance, we still have yet to test the effects with non-artificial data), the fact that the effect sizes for both affective ($\eta_p^2 = 0.086$ in Study 2 and $\eta_p^2 = 0.055$ in Study 3) but especially ideological polarisation (i.e., $\eta_p^2 = 0.139$ in Study 2 and $\eta_p^2 = 0.214$ in Study 3) have remained relatively robust is a promising development.

Secondly, however, Study 3 also added to our understanding of the impact of mode of presentation in several notable ways.

First, it showed that the effects of shifting mode do not seem to be dependent on the source via which they are disseminated, as we find little evidence of significant interactions between source and mode (or main effects of source) throughout our analyses. Such a finding is particularly important in understanding the possible real-world utility of the intervention. Had we observed “source effects” (e.g., Laustsen & Petersen, 2016; Maier, Adam, & Maier, 2017), this would have indicated that identical presentations might have been yielding asymmetric impacts depending on *who* had disseminated the data. Consequently, when considering deployment, we would need to understand not simply which *modes* would be most effective at having a desired impact, but instead which modes *depicted by which sources*, adding a layer of complexity. However, the absence of such effects implies that mode effects should emerge similarly across varied sources, making deployment considerations a little more straightforward.

Secondly, Study 3 confirmed that, beyond exerting influence on (among other things) issue-specific measures of perceived affective and ideological polarisation, the effects of mode of presentation can also extend to impact *overall* perceptions of ideological polarisation and – to a lesser degree – *general* affective polarisation, as well. This is perhaps the most consequential finding from a real-world utility perspective as it suggests that an intervention which alters the manner in which partisan data is visualised is capable of not only shifting perceptions of polarisation on that specific issue, but can also produce carry-over effects on how *overall* polarisation is perceived, as well. The implication here is significant in that it theoretically enhances the intervention’s “return on investment” as, for example, exposure to data on immigration depicted via superimposed normal distributions not only mitigates perceptions of how polarised the groups are on this issue ($M = 41.57, SD = 22.94$) relative to both no visualisation ($M = 78.62, SD = 21.12$) or a truncated bar chart ($M = 64.01, SD = 21.65$), but also produces a nearly ten point pre-post decline (i.e., $M = -9.81, SD = 18.96$) in perceptions of *overall* ideological polarisation.

Finally, the inclusion of a “no visualisation” control condition proved revelatory, as it demonstrated that, in many cases, participant perceptions are more exaggerated and less accurate *without* any data visualisations than they would have otherwise been with the aid of even the *least* useful mode (i.e., truncated bar charts). In other words, in certain circumstances, perceptions of issue polarisation and levels of intergroup agreement are so skewed in the absence of depictions of empirical data that *any* visualisation might improve perceptions, leaving individuals better off than if they had been left to simply consider issues based on their pre-existing beliefs. This is a critical insight as it suggests that, while the truncated bar charts have thus far “performed worse” (i.e., catalysing more polarised perceptions and less accurate overlap estimates) than alternative modes, they still might actually *improve* perceptions for those who otherwise might have no exposure to data at all.

While Study 3 represents a useful addition to the ongoing collection of studies, we believe there are important considerations when evaluating the potential issues in the design and methodology. First and foremost, while one of the study’s primary goals was to evaluate the impact of source attribution on mode of presentation effects, we failed to ascertain how *believable* the visual stimuli created to simulate data being disseminated from particular partisan news outlets was perceived to be. While it could be that the stimuli was very believable and still neither main effects of source nor significant interactions emerged, without a manipulation check, it’s possible the null results on these metrics could be due to a failure to convincingly present the stimuli. Secondly, while we believe the no visualisation condition was a particularly revealing addition to the design, it is by no means the only viable control condition against which the experimental conditions could be evaluated. For example, future research might consider a no visualisation control which still communicates the positions of the two parties (i.e., it still provides them with accurate data), but does so without the aid of a visual depiction.

Ultimately, Study 3 confirmed that significant main effects of mode of presentation continue to emerge on politically-relevant perception metrics even under circumstances where the data is attributed to a specific, divisive issue (i.e., immigration) and disseminated by a specified source (e.g., Fox News or CNN). Moreover, it found that the *issue-specific* effects of mode (as observed in Study 1 and Study 2) also seem to extend to impact more *general* perceptions of polarisation. Finally, the study provides evidence to suggest that while certain modes appear to be “better” than others when it comes to mitigating polarisation and promoting accurate perceptions of intergroup agreement, exposure to even “sub-optimal” visualisation formats (e.g., truncated bar charts) might produce more desirable results than permitting individuals to harbour perceptions uninformed by any data whatsoever.

Study 4: Investigating the Immediate and Longitudinal Effects of Mode of Presentation Using Real Partisan Data on a Polarised Issue⁴⁰

For Study 4, our goal was to test whether “mode of presentation” effects persist on both perceptions of polarisation as well as on estimates of intergroup agreement A) when using *actual* data on polarising issues and B) across an extended time period. Additionally, Study 4 aimed to determine whether the effects observed on *perceptions* across the first three studies would also materialise on dependent variables designed to serve as proxies for participant *behaviour*.

Continuing the theme of incrementally increasing the ecological validity of the design with each iteration of the study, in Study 4 we shifted from using a set of artificially-constructed datasets to create the corresponding data visualisations to using *actual* data (taken from nationally-representative surveys) to construct each mode. More specifically, the visualisations depicted partisan responses to items regarding attitudes towards immigration and support for assault rifle bans – each of which is an evergreen source of interparty discord in the United States political context. Such an advance in the design is important as it provides a more reliable simulation of how the intervention might fare in a real-world deployment, where data is unlikely to be perfectly normally-distributed or as “clean” as the artificial datasets had been thus far. Should mode of presentation effects continue to emerge under these circumstances, we can begin to speak more confidently about the possible utility of the technique to significantly alter perceptions of polarisation on real issues. Considering both the strength and consistency of the effects throughout the first three studies, we hypothesise that mode effects will continue to be observed on perceptions of ideological polarisation, with exposure to the full-range bar chart and icon array histogram promoting significant pre-post shifts relative to the truncated bar chart or no visualisation condition. Moreover, as modes such as the full-range bar chart and the icon array histogram demonstrated the ability to catalyse estimates of intergroup opinion overlap which either did not significantly differ from their respective accuracy benchmarks or at least closely

⁴⁰ University of Cambridge ethics reference PRE.2022.069

approached them, we hypothesise that these modes will serve a “corrective” function in pre-post measures of intergroup opinion overlap, changing estimates in the desired direction (i.e., closer to the accuracy benchmark).

As mentioned, Study 4 also sought to examine whether mode of presentation would elicit significant differences in how participants responded to more behaviourally-oriented items. While shifting perceptions is useful, ultimately, an ideal intervention would be one which could also precipitate changes in partisan behaviour, as well. While some might believe that significant perspective shifts concerning beliefs about the extent of political polarisation *should* result in more open, less hostile interparty behaviour, a compendium of literature has documented the “attitude-behaviour gap” (sometimes referred to as the “intention-behaviour gap”; e.g., Nguyen, Nguyen, & Hoang, 2019), wherein attitude change need not necessarily result in corresponding behaviour change (e.g., Park & Lin, 2020; El Haffar, Durif, & Dube, 2020). Consequently, it’s important to empirically investigate whether interventions which show promise in the domain of attitude change will continue to exert effects on related behaviour.

Study 4 incorporates several measures designed to capture relevant facets of interparty behaviour, including 1) an opportunity to donate funds to an organisation which supports political compromise, 2) an opportunity to participate in a civil cross-party discussion on key issues, and 3) an opportunity to respond to an optimistic political statement. We believe each item captures a behavioural response which is likely germane to the health of political relations, with the first two items attempting to capture a participant’s willingness to support political compromise while eschewing partisan hostility and incivility (e.g., Wolak, 2020; Wolf, Strachan, & Shea, 2012) while the final seeks to determine whether exposure to different modes of presentation can moderate the type of pessimism and vitriol which has become characteristic of online forums (e.g., Seely, 2018; Gervais, 2015). Considering the substantial impact mode of presentation has consistently exerted on perceptions through three studies, we hypothesise that similar effects will be observed on the behavioural variables, with modes such as the full-range bar chart and the icon array histogram promoting behaviours which are significantly more “positive” than those engendered by exposure to the truncated bar chart or no visualisation.

Furthermore, Study 4 incorporated a longitudinal design element as a means of determining whether mode of presentation effects, which were routinely observed on participant’s *immediate* post-exposure perceptions, would persist across an extended time period. To do this, participants had their perceptions of overall polarisation, issue-specific polarisation, and estimates of intergroup opinion overlap re-assessed approximately five weeks after the initial collection (i.e., Stage 2) to compare with their prior responses (i.e., during Stage 1). We hypothesised that the effects of exposure to either the full-range bar chart or the icon array histogram on issue-specific perceptions of ideological polarisation would remain robust across the five

week period, as measured by whether participant's responses were significantly lower in the Stage 2 collection than they were in the original Stage 1 collection. Additionally, the second stage of the study featured the introduction of an economic game called "the Trust Game." Economic games (such as the Dictator Game, the Ultimatum Game, Prisoner's Dilemma, etc.; Komorita & Parks, 1995) permit researchers the opportunity to bypass many of the issues inherent to self-report measures and instead observe concrete behavioural responses to situations which simulate elements of real-world interactions (Fehr & Gächter, 2000). In our case, we aim to use the Trust Game – which involves deciding how much money to "trust" to a playing partner (who would be identified as either a Republican or Democrat) – to ascertain how mode of presentation may alter the cooperation patterns of political partisans (see *Measures* for more details on the game).

Prior research has shown that partisanship plays a role in partner trust, with participants sending significantly less money to be entrusted to a playing partner who has been identified as belonging to an opposing political group (Carlin & Love, 2018). However, in addition to partisanship, we wanted to examine how exposure to different modes of data visualisation might affect willingness to trust a political other. While mode effects consistently manifested on one's immediate *perceptions* through the first three studies, we A) did not know whether such effects would persist on behavioural responses and B) were sceptical the effects – which were obtained in a controlled information environment – would not decay after five weeks of immersion in a more realistic, competitive communications environment (see Bolsen & Shapiro, 2017), especially as prior research has determined that the duration, stability, and resilience of framing effects seems to be dependent on whether individuals are also exposed to *competing* frames (Lecheler & de Vreese, 2016). Consequently, participants were allocated to either a "booster" condition (where they would be *re-exposed* to the same data visualisations as Stage 1 immediately before playing the Trust Game) or a non-booster condition (where they would simply play the game, with their only exposure having taken place five weeks prior in Stage 1). Ultimately, we hypothesised that effects were only likely to emerge within the booster condition.

Method

Participants

A sample of American participants were recruited via the Prolific Academic platform and paid an approximate rate of £7.50/hour. Following exclusions for non-consent, attention check failure⁴¹, or

⁴¹ Study 4 featured an attention check which began in a similar fashion to some of the opinion-based dependent variables (i.e., "Some people feel that Republicans and Democrats could find ways to understand each other...") but which, in the third sentence, explicitly stated "If you are paying attention, please select Agree below." Overall, there were seven response options from which to choose, and 93.63% successfully passed the attention check.

incomplete data, the final sample was 529⁴² for Stage 1. Of these, 403 participants took part in Stage 2 five weeks later (i.e., 76.18% retention). Table 7 summarises the demographic composition of the original, Stage 1 sample.

Table 7

Composition of Stage 1 Sample for Study 4

| Identity Characteristic | Sample Composition |
|-------------------------|---------------------------------------------------------------------------------------------------------|
| Sex | Female: 50.09% ($n = 265$), Male: 49.15% ($n = 260$), Other: 00.76% ($n = 4$) |
| Age | $M = 42.54$, Range: 19-85 |
| Education | 51.61% ($n = 273$) had attained a bachelor's degree or higher |
| Political Affiliation | Democrat: 48.58% ($n = 257$), Republican: 48.58% ($n = 257$), Independent/Other: 2.84% ($n = 15$) |

Among participants who openly identified as Democrat or Republican, 62.26% ($n = 160$) categorised themselves as “strong” Democrats and 45.52% ($n = 117$) categorised themselves as “strong” Republicans. Among those who openly identified as Independent/Other, when forced to choose which of the two major US political parties they considered themselves closest to, 46.67% of this group chose Democrat ($n = 7$) and 53.33% ($n = 8$) chose Republican. Ideologically, the sample skewed slightly liberal ($M = 46.11$ on a 0 to 100 ideological self-placement scale).

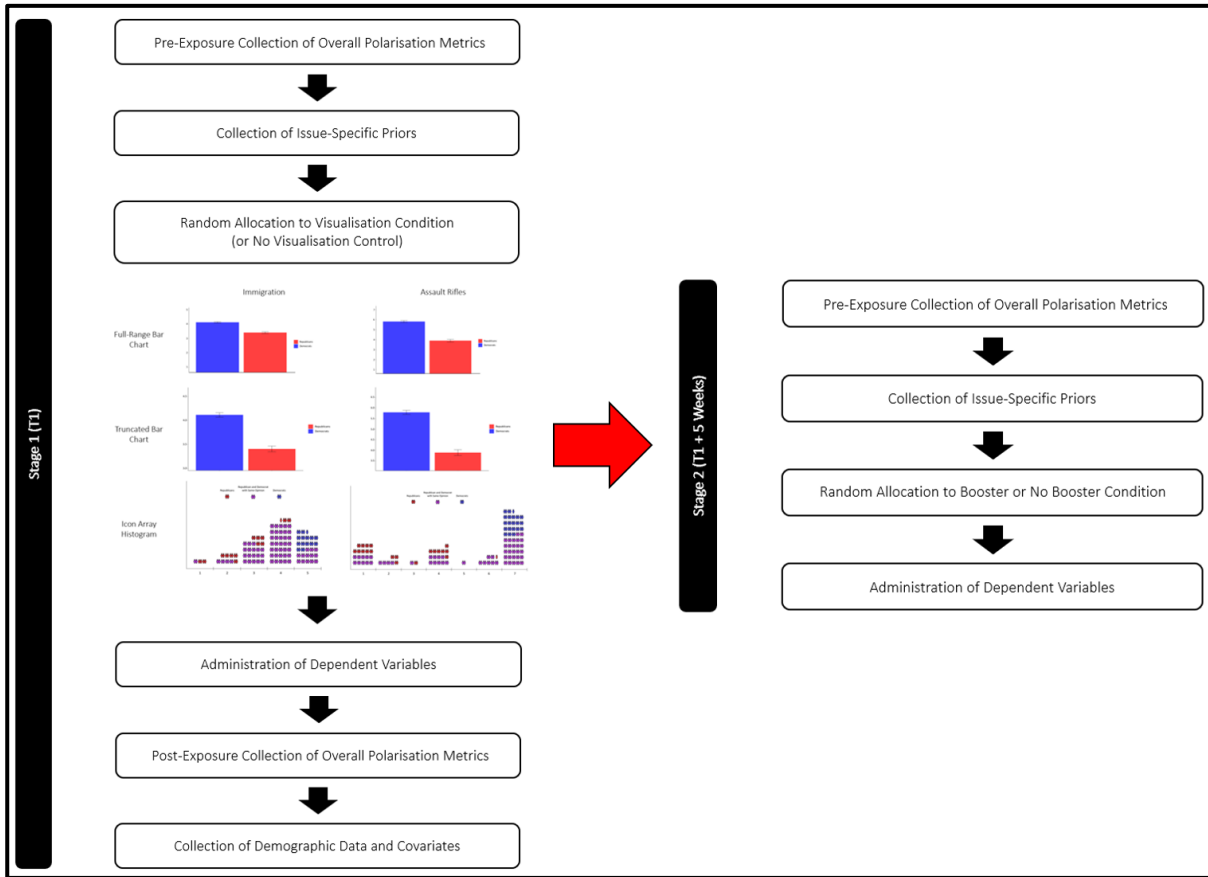
Design

The study employed a between-subjects, longitudinal design (see Figure 18).

⁴² A priori power analyses were conducted to determine the sample size required to achieve sufficient statistical power to detect a medium effect size (i.e., Cohen's f of 0.25) with a power level of 0.90 and an alpha threshold of 0.05. The analysis indicated a sample size of approximately 58 participants per group would be required. The effect size was chosen based on the findings from Study 3, which continued to routinely indicate a large main effect of mode. However, once again we wished to conduct a more conservative power analysis (which assumed smaller effects) given the uncertainty of introducing additional complexity to the design. Moreover, due to the longitudinal nature of the study, we opted to assume a pessimistic dropout rate of about 35% between Stage 1 and Stage 2, and wished for the final (i.e., Stage 2) sample to thus contain a minimum of 232 participants. Additionally, prior to exclusions, the total sample in Stage 1 was 565. For those who provided data, the sample composition prior to exclusions had a mean age of 42.54, was comprised of 265 females, 260 males, and 4 other (with 36 non-responses), and had 257 self-reported Democrats, 257 self-reported Republicans, and 15 self-reported other (with 36 non-responses).

Figure 18

Design Outline for Study 4



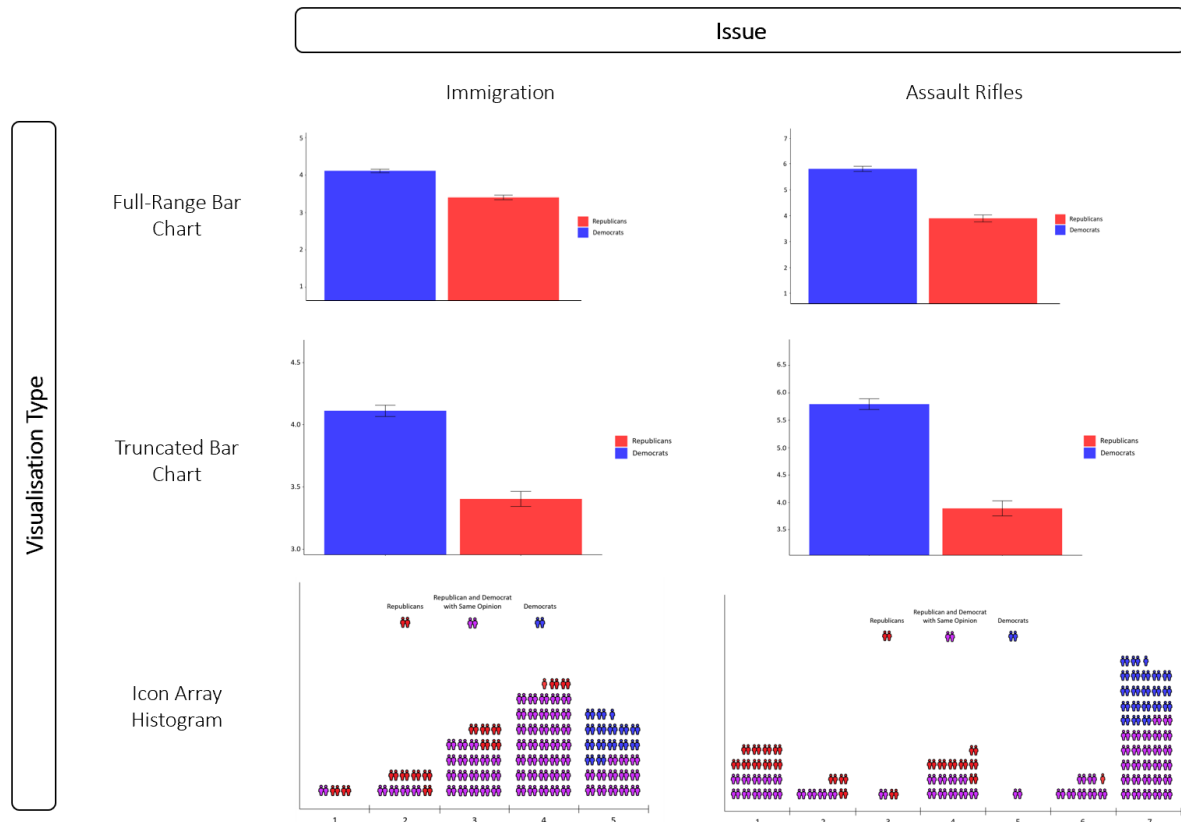
Procedure

The study took place in two parts (henceforth referred to as Stage 1 and Stage 2). Stage 1 began by assessing pre-levels of both general affective and perceived overall ideological polarisation (as described in Study 3). These levels would also be tested following the manipulation at the conclusion of Stage 1 as well as at the beginning of Stage 2 as a means of assessing the change in these levels both from the beginning of Stage 1 to the end of Stage 1 as well as from Stage 1 to Stage 2. Participants would then have their issue-specific prior beliefs and attitudes concerning two polarised issues measured before being exposed to one of three experimental conditions (e.g., full-range bar chart, truncated bar chart, or icon array histogram⁴³; see Figure 19) or a no visualisation condition.

⁴³ Readers may notice that superimposed normal distributions, after being included in the prior three studies, have now been omitted from the design of Study 4. This is due to the fact that, while such a visualisation format is an appropriate way to depict data which is normally-distributed, it is an inappropriate manner with which to depict non-normally-distributed data. As Study 4 uses actual data (which is considerably “messier” than the artificial data used in the past three studies, and which is also decidedly non-normal in its distribution), we opted to remove the superimposed normal distribution mode. Both bar charts and icon array histograms are more capable of faithfully depicting non-normal data.

Figure 19

Condition Stimuli for Study 4



The visualisations used for each mode of presentation were developed using actual data taken from the 2020 version of the American National Election Study. For the immigration visualisations, the data came from Republican and Democrat responses to the item “To what extent do you agree with the following statement: ‘Immigrants are generally good for America’s economy?’” For the assault rifle visualisations, the data came from Republican and Democrat responses to the item “Do you favor or oppose banning the sale of semi-automatic, ‘assault-style’ rifles?”

Participants would be exposed to both issues via the same mode of presentation. Following exposure to each issue, participants would be asked to indicate their perceived levels of issue-specific ideological polarisation between the two groups as well as their estimates of intergroup opinion overlap. Following this, participants would be asked to respond to the three dependent variables meant to serve as proxies for partisan behaviour (see *Measures*).

Stage 2 took place approximately five weeks after participation in Stage 1. It began by re-assessing both general affective polarisation, overall perceived ideological polarisation, and issue-specific measures concerning perceived issue-specific polarisation and estimated intergroup opinion overlap. Participants

were then randomly allocated into either the “booster” or “non-booster” condition, where they were either re-exposed to the mode of presentation visualisations they were originally shown during Stage 1 or shown nothing, respectively. Following this, participants participated in a “trust game” with a partner said to be either a member of their own political party or a member of the opposing political party.

Measures

The study featured the same measures of general affective polarisation and perceived overall ideological polarisation as used in Study 3.

For issue-specific polarisation and estimates of intergroup opinion overlap, whereas prior studies only assessed participants *after* exposure to the modes of presentation, participants in Study 4 also had their prior beliefs about these two measures assessed so as to measure the level of change. For example, prior to exposure to the data via a mode of presentation, participants were asked to indicate how much they *believed* Republicans and Democrats would be polarised on the two issues. Prior beliefs were then subtracted from post-exposure perceptions to measure change.

As noted, three dependent variables used in Stage 1 were designed to serve as a proxy for partisan “behaviour” in response to different mode of presentation conditions. These were:

Donation Opportunity: Participants were provided with an opportunity to donate from an amount of *additional* funds (i.e., 10 cents, specified that it would *not* to be taken out of their original payment) to an organisation whose mission is to investigate areas of cross-party agreement (i.e., More in Common).

Civil Cross-Party Conversation Opportunity: Asked participants how open they would be to participating in a short (i.e., approximately 20 minute) “civil discussion group” where members of both parties would come together to respectfully discuss pressing political issues. Responses ranges from 1 (“Not at All”) to 5 (“Very Interested”).

Response to an Optimistic Political Statement: Participants were asked to indicate – via open text response – their agreement with a statement which essentially said that while the two parties are undoubtedly different, if we looked beyond the extremes of both sides, we’d likely find that we have more in common than we might believe. Responses were coded on a 1 to 5 scale, where 1 was “Extremely Negative/Pessimistic” and 5 was “Extremely Positive/Optimistic.”

Finally, in Stage 2, participants also played the “Trust Game,” wherein they were provided with a certain amount of money (i.e., an *additional* 10 cents) and were forced to make a decision how much of that money to send (i.e., “entrust”) to a partner and how much to keep for themselves. They were informed that any amount sent would be tripled, at which point the partner would decide how to allocate the newly-tripled

pot. Critically, participants were informed their partner was either an own-party member or an opposing-party member prior to the game commencing. Our primary variable of interest was the initial amount participants opted to send to their partner.

Hypotheses

Study 4 featured the following set of pre-registered (OSF URL: <https://osf.io/4fj5u>) hypotheses (Table 8).

Table 8

Pre-Registered Hypotheses for Study 4

| Label | Hypothesis |
|-----------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Effects of Mode on Polarisation and Accuracy Perceptions</i> | |
| H1 | Participants exposed to the full-range bar chart as well as those exposed to the icon array histogram will show a significant reduction in their pre-post measures of perceived issue polarisation for both issues. Participants exposed to the truncated bar chart will show no significant pre-post change on either issue. |
| H2 | Participants exposed to the full-range bar chart as well as those exposed to the icon array histogram will show a significant increase in their pre-post measures of estimated intergroup opinion overlap for both issues. |
| H3 _A | Both the icon array histogram as well as the full-range bar chart will elicit post-exposure estimates of intergroup opinion overlap that do not differ significantly from the actual degree of opinion overlap for the immigration issue (i.e., 71%). All other conditions will elicit estimates that will be significantly different from the accuracy benchmark. |
| H3 _B | Both the icon array histogram as well as the full-range bar chart will elicit post-exposure estimates of intergroup opinion overlap that do not differ significantly from the actual degree of opinion overlap for the assault rifle issue (i.e., 59%). All other conditions will elicit estimates that will be significantly different from the accuracy benchmark. |
| <i>Effects of Mode on Behavioural Dependent Variables</i> | |
| H4 | Main effects of mode of presentation will be observed for responses to the donation opportunity. |
| H5 | Main effects of mode of presentation will be observed for the “willingness to participate in cross-party conversation” dependent variable. |
| H6 | Significant main effects of mode of presentation will be observed on the responses to an optimistic political statement dependent variable. |
| <i>Longitudinal Effects</i> | |
| H7 | Main effects of mode of presentation will be observed for amount sent in the Trust Game (during the booster condition only). |
| H8 | Participants exposed to the full-range bar chart or the icon array histograms will have pre-exposure estimates of perceived issue polarisation that are significantly lower in Stage 2 than they were when originally measured in Stage 1. |


Results

Exploratory one-way ANOVAs on pre-post perceptions of issue polarisation in Stage 1 revealed significant main effects of mode for both the immigration ($F(2, 399) = 14.82, p < .001, \eta_p^2 = 0.069$) and assault rifle ($F(2, 399) = 18.95, p < .001, \eta_p^2 = 0.087$) issues (see Figure 20). As hypothesised [H1], individuals exposed to full-range bar charts experienced a significant pre-post reduction in perceived polarisation for both the immigration ($M = -36.35, SD = 33.39; t(264) = 12.960, p < .001, d = 1.59$) and assault rifle issues ($M = -26.59, SD = 30.94; t(256.59) = 10.381, p < .001, d = 1.27$). Similarly, individuals exposed to icon array histograms also experienced a significant pre-post reduction in perceived polarisation for both the immigration ($M = -29.35, SD = 38.05; t(270) = 10.048, p < .001, d = 1.22$) and assault rifle issues ($M = -19.69, SD = 28.47; t(270) = 7.399, p < .001, d = 0.90$). However, contrary to what was hypothesised [H1], t-tests revealed that exposure to truncated bar charts precipitated a significant pre-post reduction in perceived issue polarisation on both

the immigration ($M = -14.35$, $SD = 28.84$; $t(264) = 5.156$, $p < .001$, $d = 0.63$) and assault rifle ($M = -5.70$, $SD = 24.89$; $t(264) = 1.973$, $p = .049$, $d = 0.24$) issues, as well.

Figure 20

Changes in Beliefs (Post-Visualisation Exposure minus Pre-Visualisation Exposure) about Issue-Specific Polarisation by Mode



* $p < .05$, ** $p < .01$, *** $p < .001$

However, as ANOVAs do not effectively account for baseline (i.e., pre-exposure) differences across groups (and mixed-effects ANOVAs assume balanced data and homogeneity of variances), additional linear mixed-effects modelling was performed to examine the perception changes within each issue. More specifically, each model sought to disentangle the effects of time (pre- vs. post-exposure) and visualisation type (truncated bar chart, full-range bar chart, and icon array histogram) on polarisation scores, while accounting for participant-level variability. The model included time, mode, and their interaction as fixed effects, with random intercepts for participants.


For the immigration issue, there was a significant main effect of time, $b = 14.35$, $SE = 2.85$, $t(798) = 5.04$, $p < .001$, indicating that perceptions of polarisation in the truncated bar chart condition were significantly higher in the pre-exposure condition compared to the post-exposure condition. Additionally, in the post-exposure condition, perceptions of polarisation were significantly lower for the full-range bar chart ($b = -23.71$, $SE = 2.85$, $t(798) = -8.32$, $p < .001$) and the icon array histogram ($b = -17.74$, $SE = 2.83$, $t(798) = -6.26$, $p < .001$), compared to the truncated bar chart. The interaction between time and visualisation type was also significant ($F(2, 798) = 15.56$, $p < .001$). The pre-to-post difference in polarisation was significantly larger for the full-range bar chart compared to the truncated bar chart ($b = 21.99$, $SE = 4.03$, $t(798) = 5.46$, $p < .001$), and similarly larger for the icon array histogram compared to the truncated bar chart ($b = 14.99$, $SE = 4.01$, $t(798) = 3.74$, $p < .001$). For the assault rifle issue, there was also a significant main effect of time, $b = 5.70$, $SE = 2.45$, $t(399) = 2.33$, $p = .020$, indicating that perceptions of polarisation were significantly higher in the pre-exposure condition compared to the post-exposure condition for the truncated bar chart. Additionally, in the post-exposure condition, perceptions of polarisation were significantly lower for the full-range bar chart ($b = -18.21$, $SE = 2.72$, $t(770.32) = -6.70$, $p < .001$) and for the icon array histogram ($b = -9.88$, $SE = 2.70$, $t(770.32) = -3.66$, $p < .001$), compared to the truncated bar chart. The interaction between time and

visualisation type was also significant ($F(2, 399) = 18.95, p < .001$). The pre-to-post difference in polarisation was significantly larger for the full-range bar chart compared to the truncated bar chart ($b = 20.90, SE = 3.46, t(399) = 6.04, p < .001$) and similarly larger for the icon array histogram compared to the truncated bar chart ($b = 13.99, SE = 3.44, t(399) = 4.07, p < .001$).

An exploratory one-way ANOVA on the pre-post (Stage 1) change in estimated intergroup opinion overlap revealed significant main effects of mode for both the immigration ($F(2, 399) = 7.003, p = .001, \eta_p^2 = 0.034$) and assault rifle ($F(2, 399) = 7.848, p < .001, \eta_p^2 = 0.038$) issues (see Figure 21). T-tests determined [H2] that participants exposed to the full-range bar chart ($M = 23.29, SD = 29.06; t(264) = 8.057, p < .001, d = 0.99$), the icon array histogram ($M = 41.98, SD = 22.60; t(270) = 10.641, p < .001, d = 1.29$), and the truncated bar chart ($M = 14.99, SD = 26.88; t(264) = 6.393, p < .001, d = 0.78$) all exhibited significant pre-post increases in estimates of intergroup opinion overlap on the immigration issue. Similarly, participants exposed to the full-range bar chart ($M = 21.45, SD = 25.92; t(251.41) = 8.414, p < .001, d = 1.03$), the icon array histogram ($M = 20.65, SD = 29.35; t(256.85) = 7.326, p < .001, d = 0.89$), and the truncated bar chart ($M = 9.46, SD = 27.53; t(264) = 3.619, p < .001, d = 0.44$) also all exhibited significant pre-post increases in estimates of intergroup opinion overlap on the assault rifle issue.

Figure 21

Change in Estimates of Intergroup Opinion Overlap Following Data Visualisation Exposure by Mode



* $p < .05$, ** $p < .01$, *** $p < .001$

Once again, linear mixed-effects models were also created. For the immigration issue, there was a significant main effect of time, $b = -14.99, SE = 2.49, t(399) = -6.03, p < .001$, indicating that estimates of intergroup opinion overlap were significantly lower in the pre-exposure condition compared to the post-exposure condition for the truncated bar chart. Additionally, in the post-exposure condition, estimates of intergroup opinion overlap were significantly higher for the full-range bar chart ($b = 14.46, SE = 2.64, t(787.87) = 5.48, p < .001$) and for the icon array histogram ($b = 18.49, SE = 2.62, t(787.87) = 7.05, p < .001$) compared to the truncated bar chart. The interaction between time and visualisation type was also significant ($F(2, 399) = 7.00, p = .001$). The pre-to-post difference in estimates of intergroup opinion overlap was significantly smaller for the full-range bar chart compared to the truncated bar chart ($b = -8.29, SE = 3.51, t(399) = -2.36,$

$p = .019$) and similarly smaller for the icon array histogram compared to the truncated bar chart ($b = -12.92$, $SE = 3.49$, $t(399) = -3.70$, $p < .001$). Similarly, for the assault rifle issue, there was a significant main effect of time, $b = -9.46$, $SE = 2.40$, $t(399) = -3.95$, $p < .001$, indicating that estimates of intergroup opinion overlap were significantly lower in the pre-exposure condition compared to the post-exposure condition for the truncated bar chart. Additionally, in the post-exposure condition, estimates of intergroup opinion overlap were significantly higher for the full-range bar chart ($b = 15.44$, $SE = 2.68$, $t(768.05) = 5.77$, $p < .001$) and for the icon array histogram ($b = 15.89$, $SE = 2.66$, $t(768.05) = 5.97$, $p < .001$) compared to the truncated bar chart. The interaction between time and visualisation type was also significant ($F(2, 399) = 7.85$, $p < .001$). The pre-to-post difference in estimates of intergroup opinion overlap was significantly smaller for the full-range bar chart compared to the truncated bar chart ($b = -11.99$, $SE = 3.39$, $t(399) = -3.54$, $p < .001$) and similarly smaller for the icon array histogram compared to the truncated bar chart ($b = -11.20$, $SE = 3.37$, $t(399) = -3.32$, $p = .001$).

Critically, while all three modes across both issues failed to promote estimates that did not differ significantly from their respective accuracy benchmarks [H3_A, H3_B] (i.e., a PCS of 71% for the immigration item and a PCS of 59% for assault rifle item), both the full-range bar chart as well as the icon array histogram catalysed estimates that were significantly *more* accurate than the truncated bar chart on each issue (see Table 9).

Table 9

Results of T-Tests between Post-Exposure Estimates of Intergroup Opinion Overlap for Each Mode and the Corresponding Issue-Specific Accuracy Benchmark

| Mode | Condition-Specific M (SD) and Results of T-Test with Respective Accuracy Benchmark | |
|----------------------|------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| | Immigration (PCS = 71%) | Assault Rifle (PCS = 59%) |
| Truncated Bar Chart | $M = 40.94$ ($SD = 20.37$) $(t(132) = 17.016, p < .001, d = 1.48)$ | $M = 34.29$ ($SD = 21.82$) $(t(132) = 13.065, p < .001, d = 1.13)$ |
| Full-Range Bar Chart | $M = 55.40$ ($SD = 24.65$) $(t(132) = 7.299, p < .001, d = 0.63)$ | $M = 49.72$ ($SD = 18.32$) $(t(132) = 5.837, p < .001, d = 0.51)$ |
| Icon Array Histogram | $M = 59.43$ ($SD = 21.34$) $(t(135) = 6.326, p < .001, d = 0.54)$ | $M = 50.18$ ($SD = 20.45$) $(t(135) = 5.032, p < .001, d = 0.43)$ |

One-way ANOVAs conducted on the responses collected across the three behavioural dependent variables in Stage 1 revealed no significant main effect of mode for the donation opportunity [H4] ($F(3, 525) = 2.020$, $p = .110$, $\eta_p^2 = 0.011$); no visualisation ($M = 3.71$, $SD = 4.34$), truncated bar chart ($M = 3.13$, $SD = 4.04$), full-range bar chart ($M = 2.58$, $SD = 3.82$), icon array histogram ($M = 3.57$, $SD = 4.26$)), no significant main effect of mode for opportunity to participate in cross-party discussions [H5] ($F(3, 525) = 0.342$, $p = .795$, $\eta_p^2 = 0.002$); no visualisation ($M = 2.23$, $SD = 0.75$), truncated bar chart ($M = 2.26$, $SD = 0.79$), full-range bar chart ($M = 2.32$, $SD = 0.76$), icon array histogram ($M = 2.27$, $SD = 0.76$)), and no significant main effect of mode on


responses⁴⁴ to an optimistic political statement [H6] ($F(3, 525) = 1.991, p = .114, \eta_p^2 = 0.011$; no visualisation ($M = 3.64, SD = 1.59$), truncated bar chart ($M = 3.82, SD = 1.47$), full-range bar chart ($M = 3.80, SD = 1.34$), icon array histogram ($M = 4.06, SD = 1.25$)). In Stage 2, a one-way ANOVA for participants assigned to the booster condition in Stage 2 found no significant main effect of mode on amount sent (i.e., “trusted”) in the Trust Game [H7] ($F(2, 154) = 2.226, p = .111, \eta_p^2 = 0.028$; truncated bar chart ($M = 6.47, SD = 3.03$), full-range bar chart ($M = 7.29, SD = 3.17$), icon array histogram ($M = 5.94, SD = 3.69$)).

A one-way ANOVA on the pre-post change in perceptions of *overall* perceived ideological polarisation (within Stage 1) revealed significant main effects of mode (no visualisation ($M = -3.65, SD = 12.83$), truncated bar chart ($M = -4.44, SD = 15.32$), full-range bar chart ($M = -13.36, SD = 21.99$), icon array histogram ($M = -12.76, SD = 22.96$); $F(3, 525) = 10.12, p < .001, \eta_p^2 = 0.055$; see Figure 22). Moreover, a linear mixed-effects model was also conducted to examine the effects of time (pre- vs. post-exposure) and visualisation type on perceptions of overall polarisation while accounting for participant-level variability. The model included time, visualisation type, and their interaction as fixed effects, with random intercepts for participants. We find that there was a significant main effect of time, $b = 3.65, SE = 1.67, t(525) = 2.18, p = .030$, indicating that perceptions of polarisation were significantly higher in the pre-exposure condition compared to the post-exposure condition when analysing across all visualisations. Additionally, in the post-exposure condition, perceptions of polarisation were significantly lower for the full-range bar chart ($b = -8.88, SE = 2.54, t(789.69) = -3.50, p < .001$) and for the icon array histogram ($b = -10.69, SE = 2.52, t(789.69) = -4.24, p < .001$); but not the truncated bar chart, $b = -2.05, SE = 2.53, t(789.69) = -0.81, p = .419$) compared to the no visualisation condition. The interaction between time and visualisation type was also significant ($F(3, 525) = 10.12, p < .001$). The pre-post reduction in perceptions of overall polarisation was significantly larger for the full-range bar chart ($b = 9.72, SE = 2.34, t(525) = 4.15, p < .001$) and the icon array histogram ($b = 9.12, SE = 2.33, t(525) = 3.92, p < .001$); but not the truncated bar chart, $b = 0.79, SE = 2.34, t(525) = 0.34, p = .736$) compared to the no visualisation condition.

⁴⁴ Please note that while the original analysis plan involved allowing the responses to be automatically coded via an R package called *SentimentR*, upon examination of the results we were dissatisfied with the accuracy of the package’s categorisations. Consequently, the decision was made to enlist the help of two independent coders to score the sentiment of each response (on a 1 to 5 scale, with 5 being “extremely supportive/positive”). These separate ratings were then averaged to arrive at a final sentiment score.

Figure 22⁴⁵

Change (Post-Exposure to Data Visualisations minus Pre-Exposure to Data Visualisations) in Perceptions of Overall Ideological Polarisation during Stage 1 by Mode



* $p < .05$, ** $p < .01$, *** $p < .001$

Note. Significance levels featured are based on the results obtained from Tukey HSD post-hoc tests.

In examining the longitudinal effects of the intervention, a one-way ANOVA on change (pre-exposure in Stage 1 to pre-exposure in Stage 2) in perceptions of issue polarisation found no significant main effect of mode for the immigration issue (no visualisation ($M = 1.17$, $SD = 20.91$), truncated bar chart ($M = -2.55$, $SD = 22.00$), full-range bar chart ($M = -3.50$, $SD = 19.99$), icon array histogram ($M = -0.16$, $SD = 21.61$); $F(3, 399) = 0.992$, $p = .397$, $\eta_p^2 = 0.007$) but a significant main effect of mode for the assault rifle issue (no visualisation ($M = 3.16$, $SD = 20.45$), truncated bar chart ($M = -3.49$, $SD = 19.75$), full-range bar chart ($M = -1.70$, $SD = 23.39$), icon array histogram ($M = -5.48$, $SD = 19.29$); $F(3, 399) = 3.043$, $p = .029$, $\eta_p^2 = 0.022$). T-tests determined that participants exposed to the full-range bar chart did not come to have significantly lower pre-exposure perceptions of issue polarisation five weeks later in Stage 2 than they originally reported (pre-exposure) in Stage 1 for either the immigration issue ($t(232) = 0.915$, $p = .361$, $d = 0.12$) or the assault rifle issue ($t(232) = 0.203$, $p = .840$, $d = 0.03$). Similarly, t-tests also showed that participants exposed to the icon array histogram did not come to have significantly lower pre-exposure perceptions of issue polarisation five weeks later in Stage 2 than they originally reported (pre-exposure) in Stage 1 for either the immigration issue ($t(244) = 0.058$, $p = .954$, $d = 0.00$) or the assault rifle issue ($t(244) = 1.922$, $p = .056$, $d = 0.25$) [H8]. Linear mixed-effects models for the immigration issue determined there was no significant main effect of time, $b = -2.12$, $SE = 2.19$, $t(441.81) = -0.97$, $p = .333$, indicating no substantial change in polarisation scores from Stage 1 to Stage 2 across all conditions. In Stage 2, perceived polarisation was not significantly different for the truncated bar chart ($b = -2.61$, $SE = 3.21$, $t(848.75) = -0.81$, $p = .416$), the full-range bar chart ($b = -5.38$, $SE = 3.22$, $t(851.82) = -1.67$, $p = .096$), or the icon array histogram ($b = -3.36$, $SE = 3.17$, $t(842.12) = -1.06$, p

⁴⁵ Please see rationale from Figure 17 in Study 3 in regards to the use of y-axis truncation here.

= .290) compared to the no visualisation condition. The interaction between time and visualisation type was not significant for the truncated bar chart ($b = 4.29, SE = 3.00, t(435.80) = 1.43, p = .153$), the icon array histogram ($b = 2.29, SE = 2.96, t(433.36) = 0.78, p = .439$), or the full-range bar chart ($b = 5.33, SE = 3.01, t(436.97) = 1.77, p = .077$). For the assault rifle issue, there was a significant main effect of time, $b = -4.21, SE = 2.14, t(444.17) = -1.97, p = .050$, indicating that polarisation scores decreased significantly from Stage 1 to Stage 2. In Stage 2, perceptions of polarisation were significantly lower for the truncated bar chart ($b = -9.78, SE = 3.18, t(845.25) = -3.08, p = .002$) and the icon array histogram ($b = -7.59, SE = 3.14, t(838.53) = -2.42, p = .016$) compared to the no visualisation condition, but not significantly lower for the full-range bar chart ($b = -4.77, SE = 3.19, t(848.37) = -1.50, p = .135$). The interaction between time and visualisation type was also significant for the truncated bar chart ($b = 7.80, SE = 2.93, t(438.32) = 2.66, p = .008$) and the icon array histogram ($b = 9.72, SE = 2.89, t(435.95) = 3.36, p < .001$), indicating that these modes led to significantly larger reductions in polarisation from Stage 1 to Stage 2 compared to the no visualization condition. The full-range bar chart did not show a significant pre-post reduction ($b = 5.47, SE = 2.94, t(439.46) = 1.86, p = .064$).

Finally, a one-way ANOVA assessing change (pre-exposure in Stage 1 to pre-exposure in Stage 2) in perceptions of overall perceived ideological polarisation in the approximately five-week period between Stage 1 and Stage 2 found no significant main effects of mode (no visualisation ($M = -1.34, SD = 14.70$), truncated bar chart ($M = -3.69, SD = 16.45$), full-range bar chart ($M = -4.98, SD = 14.59$), icon array histogram ($M = -1.61, SD = 21.39$); $F(3, 399) = 1.018, p = .384, \eta_p^2 = 0.008$). Linear mixed-effects modelling found there was no significant main effect of time, $b = 0.26, SE = 1.79, t(427.78) = 0.15, p = .884$, indicating no meaningful change in polarisation scores from Stage 1 to Stage 2. Additionally, in Stage 2, polarisation scores for the truncated bar chart were, on average, 3.90 points lower than in the no visualisation condition ($b = -3.90, SE = 2.73, t(829.62) = -1.43, p = .153$), though this difference was not statistically significant. Similarly, the full-range bar chart ($b = -3.16, SE = 2.74, t(833.08) = -1.15, p = .249$) and the icon array histogram ($b = -3.13, SE = 2.70, t(822.21) = -1.16, p = .247$) did not significantly differ from the no visualisation condition. The interaction between time and visualisation type was also non-significant for all visualisation types ($F(3, 418.46) = 0.95, p = .416$). For the truncated bar chart, the pre-post change in polarisation scores was 2.64 points larger than the no visualisation condition, but this difference was not statistically significant ($b = 2.64, SE = 2.45, t(422.34) = 1.08, p = .282$). For the full-range bar chart, the pre-post change was 3.99 points larger than the no visualisation condition, but this difference was also not statistically significant ($b = 3.99, SE = 2.46, t(423.39) = 1.62, p = .105$). The icon array histogram pre-post change was 1.56 points larger than the no visualisation condition, but this difference was not significant either ($b = 1.56, SE = 2.42, t(420.12) = 0.65, p = .519$).

Discussion

Overall, Study 4 represented a significant advancement in our mode of presentation investigation for three reasons: 1) it replaced artificial datasets with real-world data on polarising issues, 2) it went beyond examining *perceptions* to determining whether mode of presentation effects could impact relevant partisan *behaviours*, and 3) it added a longitudinal element to the design, which permitted us to test whether mode effects would persist several weeks after initial exposure.

First and foremost, in spite of using “messier” real-world data (as opposed to less “noisy” artificial data), we continued to see significant effects of mode of presentation on a variety of politically-relevant metrics. Main effects of mode of presentation were observed on pre-post changes in perceptions of issue-specific ideological polarisation for both the immigration issue as well as the assault rifle issue. As hypothesised, exposure to data depicted via either the full-range bar chart or the icon array histogram precipitated significant pre-post declines in perceptions of issue-specific ideological polarisation for both issues. However, contrary to our hypotheses, even exposure to the *truncated bar chart* led to a significant pre-post decline in perceptions of polarisation for both issues, suggesting that – while significant between-mode variability exists (as evidenced by the presence of main effects) – using *any* mode of presentation seems to result in significantly reduced perceptions of polarisation on these divisive issues relative to participant’s pre-exposure beliefs. Similar patterns were observed on pre-post changes in estimates of intergroup opinion overlap, with main effects emerging for both issues. However, here we also see that the truncated bar chart catalysed more pronounced *relative* change than the other two experimental conditions – a finding which is wholly unexpected. And while none of the modes of presentation resulted in estimates that did not differ significantly from the accuracy benchmarks, the full-range bar chart and icon array histogram promoted estimates that more closely approached the benchmarks than the truncated bar chart. Finally, as observed in Study 3, mode of presentation also managed to exert significant main effects on the pre-post changes in perceptions of *overall* ideological polarisation, and yet again we see that modes like the full-range bar chart and the icon array histogram, in addition to significantly reducing issue-specific perceptions of polarisation and increasing accuracy of overlap estimates, catalyse significant decreases in beliefs about overall intergroup ideological division, as well. The persistence of these effects under these more realistic data conditions inspires greater confidence in the intervention’s real-world utility.

Secondly, while mode of presentation consistently shifted an array of political *perceptions*, we failed to find any evidence to suggest that it can meaningfully moderate political *behaviour*. More specifically, the mode of presentation to which a participant was exposed failed to exert a significant influence on one’s level of generosity in donating to a bipartisan organisation, one’s willingness to participate in civil cross-party political discussions, or one’s sentiment towards an optimistic political statement. Such a finding undoubtedly has practical significance, as it makes clear the limitations of the intervention. While mode of presentation might prove to be a reliable moderator of political *perceptions*, it does not appear to be a

capable means of altering a moderately broad collection of political *behaviour*. Moreover, it's an important reminder of the criticality of empirical investigation, as one might have intuitively expected that the strong effects mode bore on political perceptions would naturally carry over to political behaviour, and yet evidence shows that intuition to be flawed.

Thirdly, the longitudinal impacts of the intervention appear to be, at best, weak and inconsistent. While we find evidence to suggest that main effects of mode *did* persist over the five week period on perceptions of ideological polarisation for one issue (i.e., assault rifle bans; with icon array histograms and *truncated* bar charts producing significant reductions that held over time), it was not observed for the immigration issue nor was it seen on the measure of overall perceptions of polarisation. This is an important finding as it suggests that such an intervention, while exhibiting a powerful immediate impact on perceptions, may ultimately exert effects that either dissipate quickly or are unable to withstand in a more competitive information and communication environment. This fits into the literature that specifies that certain intervention modalities display more rapid rates of decay in efficacy over time (e.g., Maertens, Anseel, & van der Linden, 2020).

While Study 4 represented an important progression in the exploration of the real-world utility of mode of presentation interventions, there are a few design and methodological choices which warrant further examination. First of all, as the study was intended to go beyond self-reported perceptions and investigate the intervention's potential impacts on partisan *behaviour*, it's fair to wonder whether the responses to the chosen behavioural dependent variables represent a reasonably accurate proxy to real world behaviour. For example, while we believe the open response to an optimistic political statement is a fair (and useful) proxy to understanding how one might respond, for instance, to such a statement on social media, it's less clear whether an opportunity to sign up for a civil cross-party discussion is an opportunity that would present itself in a non-simulated environment. Moreover, we realised that our original way of interpreting certain behavioural responses was probably underdeveloped. For example, we had originally believed that exposure to modes which reduce perceptions of polarisation would encourage greater donations to organisations which promote interparty unity. However, the dynamic could also be the exact *opposite*: it could be that those who are exposed to modes which increase perceptions of polarisation might perceive a greater *need* for the work of such organisations. Future work should not only capture donation amounts, but also the underlying *reasons* for such giving behaviour. Secondly, while we believe the trust game is a useful tool for disentangling how identity might impact decisions to trust, we believe a critical missed opportunity is the failure to account for one's *dispositional* levels of trust (e.g., how much trust one *naturally* possesses for an unidentified other). Ultimately, participants dispositional trust levels could have been a useful covariate for which to control during analyses, as it might explain some variability in the responses observed. Finally, considering the evidence which suggests that (most of) the effects of the intervention decay prior to five

weeks, future research should seek to investigate the precise longevity of the intervention. While mode of presentation effects seem to exert powerful influence on perceptions of polarisation and intergroup agreement in the short-term, it remains unclear just how long these shifts in perception persist.

Ultimately, Study 4 provided us with the most direct test of mode of presentation's utility as an intervention under real-world circumstances. Results were decidedly mixed. On the one hand, altering the mode via which data is presented continued to prove to be a reliable (and moderately powerful means) of shifting perceptions of issue-specific ideological polarisation, estimates of intergroup opinion overlap, and even perceptions of *overall* ideological polarisation – including modes such as the truncated bar chart, which routinely promoted significant alterations in perceptions and estimates relative to not only the no visualisation condition, but occasionally, the other modes, as well. This remained true in spite of the fact that the design incorporated the use of real, “messy” data on specific, divisive political issues (for which individuals would ostensibly hold strong prior beliefs concerning levels of intergroup polarisation). On the other hand, the study also highlighted several noteworthy limitations of the intervention, including its inability to meaningfully affect an array of political behaviours as well as its struggles to perform as a *longitudinal* moderator of perceptions. Taken as a whole, Study 4 provides us with a better understanding of both the strengths and weaknesses of using mode of presentation interventions, finding them to be an effective means of immediately mitigating polarised perceptions on divisive issues, but perhaps unlikely or unable to produce any noticeable changes in partisan behaviour or long-term perspective shifts.

General Discussion

Overall, this series of studies provides evidence that mode of presentation can be influential in changing perceptions of intergroup dynamics, including those related to polarisation and intergroup agreement. Across the four studies, mode of presentation had significant and robust impacts on perceptions of intergroup similarity, estimates of intergroup opinion overlap, perceptions of issue-specific polarisation, perceptions of affective polarisation, perceptions of overall ideological polarisation, degree of general affective polarisation, perceptions of group discreteness, level of intergroup agreement relative to expectations, level of trust concerning other-party fairness, and likelihood of compromise. Effects were most consistently strong on estimates of intergroup opinion overlap and perceptions of group discreteness, with effect sizes (as measured by the partial eta squared) ranging from $\eta_p^2 = 0.114$ (Study 4) and $\eta_p^2 = 0.342$ (Study 1) for the former and $\eta_p^2 = 0.223$ (Study 2) and $\eta_p^2 = 0.225$ (Study 3) for the latter. Less consistent and more marginal effects were observed for the longitudinal effects of mode on perceptions of intergroup opinion overlap and perceived issue-specific polarisation, while the current study provides little evidence to suggest that mode plays a significant role on donation generosity, willingness to engage in civil cross-party discussions, or response sentiment to an optimistic political statement.

Ultimately, across all four studies, the full-range bar chart was the mode which most consistently proved capable of significantly shifting perceptions relative to other modes (e.g., the truncated bar chart) as well as relative to the no visualisation baseline. It routinely mitigated perceptions of polarisation, increased the accuracy of estimated intergroup opinion overlap, reduced beliefs about group discreteness, and bolstered attitudes concerning the likelihood of compromise.

Theoretically, this work provides strong evidence for the *malleability* of partisan perceptions concerning polarisation, even on contentious issues for which individuals would ostensibly hold strong attitudinal priors. Additionally, it establishes that the manner in which data is presented can affect not only perceptions of the *extent* of group disagreement, but can also the perceptions of the groups *themselves*. Evidence to suggest that mode of presentation plays a significant role in perceptions of group discreteness is an important theoretical finding, as it implies that the manner in which individuals choose to depict intergroup data has the ability to moderate a component of essentialised thinking and, by extension, potentially exacerbate cross-party dehumanisation. Moreover, extant literature on data visualisation's role in essentialised perception is, to our knowledge, non-existent, and such a finding could spur greater interest within the academic community.

In regards to practical applications of this research, the current collection of studies should compel any media outlet, government agency, polling organisation, or any other individual interested in reducing polarisation to choose their mode of data visualisation carefully and deliberately. Although perceptions of polarisation are one degree removed from actual polarisation, scholars contend that exaggerated perceptions play a real and crucial role in the self-reinforcing nature of accelerating polarisation (see Lees & Cikara, 2021), meaning correcting inaccurate perceptions should be a priority in work aimed at polarisation mitigation. This series of studies provides a usable framework via which concerned parties can implement a simple, effective, and highly scalable intervention.

While the present collection of studies provides strong evidence that mode of presentation is capable of playing a significant role in moderating political perceptions, it has not exhaustively investigated the capabilities of such an intervention nor has its methodology been without limitations. Firstly, while the effects have been largely robust, many of the outcome variables included the phrase "based on the data" prior to describing the measure of interest, explicitly encouraging participants use the visualisations to inform their responses. While this was a deliberate methodological choice on our part, it is not unreasonable to assume that the effects observed may have been influenced by this intentional refocusing of participant attention. Future research should seek to determine whether similar visualisation effects are obtained without the need to provide such overt cues. Secondly, the studies have all focused on data with large (i.e., Cohen's *d* of approximately 0.80) effect sizes. Once again, this decision was a deliberate one on the part of

those designing the experiment, as the goal was to determine whether mode of presentation could influence political perceptions on issues for which Republicans and Democrats could be described as self-evidently polarised. However, due to the absence of data with different effect sizes, we cannot say whether the pattern of results observed would be replicated for issues, for instance, with substantially less or substantially greater levels of interparty disagreement, or if we would continue to see such a clear differentiation between “good” and “bad” modes (assuming the goal remains making audiences more accurate and less polarised). Thirdly, future research could consider different iterations of the “no visualisation” baseline condition. For instance, in the current studies, we simply asked participants to base their responses on “what they know” about Republicans and Democrats. However, it would not be unreasonable to consider alternative control conditions which, for instance, provide data in a non-visual manner (e.g., describing partisan responses without visualising them). Such an alteration may also provide important insight into the impact on mode over and above different *types* of controls. Fourthly, it will be important for future research to establish how these effects fare in a “competitive” message environment (see Bolsen & Shapiro, 2017). While the current research merely presented the data in a vacuum (i.e., without any “competing” data visualisations), message environments are often crowded and hostile, with other political actors trying to shift perceptions in *more* polarised directions by emphasising group differences. Future research should investigate whether mode effects persist under such conditions. Finally, future researchers may wish to continue investigating how these data presentation interventions may (or may not) impact variables for which our results were weak, inconsistent, or unclear across the present studies: namely, affective polarisation, cross-party behaviour, and longitudinal efficacy.

Overall, the present research suggests that mode of presentation plays a significant role in a myriad of critical intergroup perceptions. Importantly, it demonstrated that partisan perceptions are malleable, as mode of presentation interventions catalysed considerable shifts even on polarised issues for which individuals would presumably hold strong prior attitudes. Ultimately, the collection of studies demonstrates that individuals and institutions interested in reducing polarisation and overall interparty animus may be able to effectively *embed* misperception-correction interventions directly into their communications simply by being more deliberate with their choice of data visualisations.

CHAPTER 2

INVESTIGATING THE REPLICABILITY OF MODE OF PRESENTATION EFFECTS ACROSS GEOPOLITICAL CONTEXTS⁴⁶

Abstract

While the mechanisms underlying polarization are complex, scholars have consistently found a pervasive overestimation of perceptions of polarization to be a contributing factor. We argue that one mitigation strategy that can work at scale to address such misperceptions might be relatively straightforward: better data visualizations of cross-party attitudes on key issues. In a large-scale (N = 6,603), international replication, we find that mode of presentation – or the manner in which data is visually presented – plays a significant role in moderating perceptions of polarization, even for longstanding, divisive issues for which partisans would likely hold strong prior beliefs. Additionally, we find the effects that different modes of presentation have on issue-specific polarization also extend to participant beliefs about overall interparty polarization, with certain modes proving capable of not only promoting less polarized views, but also enabling more accurate estimates of the extent to which political groups agree. Finally, our findings also suggest that the manner in which intergroup data is visualized may also exert influence over the degree to which political groups are “essentialized” – a finding with implications for not only political perception, but also for apolitical social psychological phenomena such as dehumanization.

For years, affective and ideological polarization have been increasing across the United States and certain other western democracies (Garzia, Ferreira da Silva, & Maye, 2023; Druckman & Levy, 2022; McCarty, Poole, & Rosenthal, 2016; Pew Research Center, 2017). Rising levels of polarization in the US threatens effective

⁴⁶ The study contained herein was originally conceptualised as a replication of the work featured in Chapter 1. After being awarded a Cambridge Impact Grant to extend our work, enabling us to broaden both the scope and scale of our first four studies, we partnered with market research firm Ipsos MORI for assistance with this more ambitious undertaking. Though a replication, the following manuscript was submitted as a standalone piece for consideration for publication to *The British Journal of Social Psychology* on February 14th, 2024 under the working title “How the Manner in Which Data is Visualized Affects and Corrects (Mis)Perceptions of Political Polarisation.” It was accepted for publication on June 27th, 2024. The manuscript is included here in its unedited form. Please note that while the study provides a uniquely valuable addition to the four studies described in Chapter 1 (in that it seeks to replicate the general findings in two separate geopolitical contexts – the United States and the United Kingdom – while using large, representative samples), due to the theoretical and methodological similarities of this replication and the four prior studies (as well as the fact that we decided not to submit the original four for possible publication), the Introduction (and portions of the Discussion) are nearly identical to the corresponding sections featured in Chapter 1.

legislative functioning (McCarty, 2016), reduces the willingness of partisans to engage with those outside their political ingroup (Fiorina & Abrams, 2008; Baldassarri & Page, 2021), and has the potential to exacerbate prejudice toward and dehumanization of political opponents (Cassese, 2021; Iyengar, Sood, & Lelkes, 2012). Moreover, rising levels of polarization may reduce support for (Levendusky & Malhotra, 2016) or otherwise inhibit mutually-beneficial compromises between political rivals (Whitt et al., 2021). However, while polarization is undoubtedly increasing in several metrics across the US, a growing number of scholars have begun to address the phenomenon known as *false polarization* (e.g., Fernbach & Van Boven, 2022), wherein individuals overestimate both the magnitude of disagreement within the political sphere as well as the ideological consistency of the actors operating within it.

Research has shown that partisans struggle to identify the actual position of their adversaries on key issues (Chambers, Baron, & Inman, 2006) and often base their assumptions on stereotypic thinking which serves to exaggerate the differences between one's ingroup and outgroup (Keltner & Robinson, 1993; Graham, Nosek, & Haidt, 2012). These propensities can yield profound "perception gaps" (Yudkin, Hawkins, & Dixon, 2019), wherein distorted beliefs about where the other party stands on issues differs dramatically from the actual positions they hold. More broadly, Westfall, Van Boven, Chambers, and Judd (2015) cite evidence showing that Americans have exhibited a pattern of overestimation of polarization which has spanned more than three decades. Consistent with such a finding, recent research has begun to focus on group meta-perceptions, where similar inaccuracies continue to emerge (Lees & Cikara, 2020). Importantly, these results appear to extend well beyond the confines of the United States (Ruggeri et al., 2021).

While political misperceptions have important consequences for interparty relations (Ahler & Sood, 2018; Moore-Berg, Ankori-Karlinsky, Hameiri, & Bruneau, 2020), they may also be capable of engendering *actual* polarization through a process of mutual reinforcement (Lees & Cikara, 2021). Thus, it's important that such misperceptions are corrected whenever possible. Several studies have shown that interventions which expose individuals to the actual positions espoused by one's own as well as opposing political groups can reduce negative outgroup attributions (Lees & Cikara, 2020) and promote political moderation (Ahler, 2014) via the attenuation of beliefs about attitudinal extremity.

However, Druckman and Levy (2022) note that individual interventions, though promising, might still struggle to overcome persistent misperceptions on a large scale due to the fact that many misperceptions are institutionally driven and sustained (see Wilson, Parker, & Feinberg, 2020). More specifically, because exposure to (often exaggerated) polarization narratives is near-ubiquitous in today's modern media landscape, often being promoted by agents with immense scope for influence (e.g., partisan media outlets, political elites, etc.), interventions seeking to effectively counter such influence should have the capacity to be deployed on an equally-large scale, ideally with minimal "training" or requirements of sustained attention

on the part of the target audience. As data visualizations are already an omnipresent feature of print, television, and social media, should certain visualizations demonstrate the ability to reduce political misperceptions, simply shifting current data presentation methodologies to these visualization formats would represent a promising avenue for catalyzing widespread change with minimal effort. Such an intervention would also have an excellent cost-benefit profile and superior scalability to individually-tailored interventions, permitting depolarization agents to easily embed perception-correction mechanisms across multiple domains and mediums in which visualizations already exist.

There are reasons to believe altering visualizations may be a viable mechanism through which one could shift perceptions of groups. Scholars like Hanel, Maio, and Manstead (2019) have argued that traditional means of presenting data about groups – such as bar charts with truncated y-axes – may (purposely or otherwise) accentuate group *differences* while ignoring equally-important similarities. Modes of presentation like truncated bar charts provide little in the way of distributional information, which may inadvertently facilitate oversimplification and categorical thinking (cognitive mechanisms which contribute to false polarization; Fernbach & Van Boven, 2022). Moreover, an oversimplification of group beliefs coupled with exaggerated depictions of group differences may risk potentially facilitating more “essentialized” perceptions of the depicted groups (see Haslam & Whelan, 2008), wherein individuals conceptualize social categories – and, consequently, the delineation between different groups – as both natural and fixed. Beliefs in “natural kinds” have been found to amplify group differences (Rothbart & Taylor, 1992) and the attribution of “essences” to outgroups has been linked to classic conceptions of prejudice (Allport, Clark, & Pettigrew, 1954). However, while the modes that have traditionally been used to visualize data may unintentionally magnify perceived group differences and lead to exaggerated beliefs about political polarization, Hanel et al. (2019) found that simply changing the type of visualization used to depict identical data had a significant impact across an array of intergroup assessments, with modes that more effectively highlight similarity information (e.g., superimposed normal distributions) resulting in both more accurate intergroup perceptions as well as more positive outgroup appraisals.

While such findings might seem modest, we believe they have the potential to be particularly impactful in the context of political polarization. Researchers have consistently shown that minute, sometimes seemingly inconsequential changes to the way in which data is visualized can profoundly alter how that data is perceived and interpreted (e.g., see Nguyen, Jung, & Gupta, 2021; Lo et al., 2022). Considering this, we believe it’s fair to question whether the manner in which we’ve *traditionally* visualized data about political groups may contribute to political misperceptions and false polarization. In addition to expounding upon the disadvantages of traditional modes of presentation and the possible dangers of y-axis truncation in depicting group data, the results reported by Hanel et al. (2019) also provide preliminary evidence for the potential power of using comparatively novel visualization techniques (i.e., those which feature aspects such as full

ranges and distributions of responses, levels of intergroup agreement as opposed to simply disagreement, etc.) to shift intergroup perceptions. We believe this should motivate researchers to consider how expanding their data depiction methodologies beyond the traditional (e.g., bar charts) might similarly alter *political* perceptions. For example, modes such as icon arrays, which feature prominently in the risk communication literature, have shown promise in areas such as increasing the accuracy of risk perception (Galesic, Garcia-Retamero, & Gigerenzer, 2009) and mitigating phenomena such as denominator neglect (Garcia-Retamero, Galesic, & Gigerenzer, 2010),⁴⁷ but have yet to be extensively explored in the context of polarization.

While other researchers have investigated the circumstances and contexts in which data visualizations might prove capable of changing attitudes (e.g., Markant, Rogha, Karduni, Wesslen, & Dou, 2022; Pandey, Manivannan, Nov, Satterthwaite, & Bertini, 2014), few have directly focused on how certain depictions of data may affect beliefs about the extent of political polarization (Alieva, 2023; Santos et al., 2017). There is some precedence for believing that “better” visualizations might constitute useful interventions to reduce political polarization. For example, Rutchick, Smyth, and Konrath (2009) found that maps depicting state-level support for a presidential candidate with a shade of purple along a red and blue continuum (in proportion to the degree of support for the Republican or Democratic candidate) reduced stereotyping as well as perceptions of political division compared to the traditional binary red and blue state-by-state visualization.

The present research will seek to further elucidate the impact that simple shifts in visualization format may have on political perceptions. Broadly, we endeavour to determine A) whether (and, if so, to what degree) y-axis truncation of bar charts that depict intergroup political information shifts perceptions relative to bar charts with a full-range y-axis, B) whether “novel” modes of data visualization (e.g., such as those involving icon arrays) – which depict aspects of intergroup data such as response distributions and intergroup *agreement* – promote different political perceptions than more “traditional” modes of presentation (i.e., bar charts), and C) whether exposure to visualized data leads participants to hold political perceptions which differ significantly from those who are provided with no such visualizations. Should the results indicate that different modes of presentation – despite depicting identical underlying data – yield significant differences across polarization metrics, such findings would have important implications for the process by which individuals and organizations go about making visualization choices, and ultimately may provide a subtle yet highly-scalable addition to the depolarization repertoire.

Method

Participants

⁴⁷ More recently, icon arrays have even shown themselves capable of assisting in pressing, real-world issues such as a reticence to get vaccinated against COVID-19 by reducing the concern over exceedingly rare side effects of vaccines (Fansher et al., 2022).

A sample (N = 6,603) of British (n = 3,303) and American (n = 3,300) participants were recruited via Ipsos MORI's internal recruitment platform. The UK sample had a mean age of 44.64, was 51.04% female, and consisted of 54.50% who had obtained an educational status of NVQ4 or above. Politically, the UK sample contained 1,407 respondents (i.e., 42.60%) who voted "Leave" during the 2016 EU "Brexit" Referendum and 1,048 (i.e., 31.73%) who voted "Remain." The US sample had a mean age of 46.30, was 50.82% female, and consisted of 39.03% who had earned at least a Bachelor's degree. Politically, the US sample contained 1,045 respondents (i.e., 31.67%) who generally consider themselves Democrats and 872 (i.e., 26.42%) who generally consider themselves to be Republican.

Procedure

The study utilized a between-subjects design which featured four conditions to which participants were randomly allocated. Three conditions depicted data via the following distinct modes of presentation (see *Figure 1*):

1. A bar chart with a full-range y-axis (henceforth referred to simply as "full-range bar chart")
2. A bar chart with a 1.5 SD y-axis (henceforth referred to simply as "1.5 SD bar chart")⁴⁸
3. An icon array histogram (essentially a set of histograms constructed using color-coded icons so as to more clearly highlight intergroup response overlap at each level of the survey responses; see Appendix A for more detailed explanation of this novel mode)

The fourth condition – referred to as the "no visualization" condition – did not visually depict the data in any way, but simply asked participants to complete questionnaire items by basing their responses on "what they know" about the two groups in question (i.e., either "Leavers" and "Remainers" for the UK sample or Republicans and Democrats for the US sample). This no visualization condition was designed to serve as a baseline for how individuals might think about the issues devoid of any visual reference point.

The modes of presentation depicted actual data from the 2019 cross-sectional data of the British Election Study and the 2020 American National Election Study. The items were both about immigrants' roles in the economy (see Appendix C for item selection rationale). The item in the UK sample asked "Do you think immigration is good or bad for Britain's economy?" while the item in the US sample asked "How much do you agree or disagree with the following statement: Immigrants are generally good for America's economy?"

⁴⁸ Some researchers have argued that 1.5 standard deviations of the grand mean of the data might represent a degree of truncation that promotes an appropriate sensitivity to effect size (Witt, 2019).




Figure 1: Modes of presentation used within the UK sample (left column) and US sample (right column). The top row depicts the full-range bar charts used for each country, the middle row depicts the 1.5 SD bar charts used for each country, and the bottom row depicts the icon array histograms used for each country.

Following exposure to the randomly-assigned mode of data visualization (or assignment to the no visualization condition), participants were asked to respond to a number of dependent variables (see *Measures* section). Additionally, participants also had their levels of general (i.e., non-issue-specific) affective polarization and overall perceived ideological polarization measured both pre- and post-exposure to the data manipulations.

Measures

When assessing political polarization, there are several metrics researchers may choose to employ. The most common relate to measuring the two primary types of political polarization: ideological (i.e., how far apart do groups perceive themselves to be on policy positions; e.g., see Abramowitz & Saunders, 2008) and affective (i.e., how much people “like” or “dislike” members of their group and the opposing group; e.g., see Iyengar, Leikes, Levendusky, Malhotra, & Westwood, 2019). However, the residual impacts of political

polarization may also be observed in changes in interparty dynamics such as reduced willingness to compromise (e.g., Levendusky & Malhotra, 2016; Whitt et al., 2021). Moreover, for the purposes of the present study, we are also interested in examining how shifting data visualizations may impact variables such as perceived group discreteness (a necessary component of certain types of dehumanization; Haslam, 2006) and one's responses to (and levels of agreement with) optimistic political statements (an especially relevant metric given the hostility that characterizes online political discourse).

This study's primary measures included:

Perceived Ideological Polarization: Participants were asked to indicate their response (via a 100-point scale that ranged from "Not at All Polarizing" and "Extremely Polarizing") to the item "Based on the graph, how polarizing do you believe this issue is for the two groups?"

Perceived Affective Polarization: Participants were asked to indicate their response (via a 100-point scale that ranged from "Extremely Cold/Negative" and "Extremely Warm/Positive") to the item "Based on the graph, how warm (i.e., positive) or cold (i.e., negative) do you think members of each group feel towards members of the other group on this issue?"

Estimates of Intergroup Opinion Overlap: To determine how modes impact the *accuracy* of participant views of intergroup agreement, we asked participants to estimate intergroup opinion overlap by indicating their response (via a 100-point scale) to the item "Based on the graph, what percentage of the two groups share the same opinion on the issue?" The actual levels of overlap (and thus, the benchmarks for these responses) are covered in the *Results* section.

Level of Intergroup Agreement Relative to Expectations: Participants were asked to indicate their response (via a 5-point scale that ranged from "They seem to agree much less than I would have thought" to "They seem to agree much more than I would have thought") to the item "Based on the graph, how would you classify the level of agreement between the two groups on this issue?"

Perceived Group Discreteness: Participants were asked to indicate their response (via a 6-point scale that ranged from "Strongly Disagree" to "Strongly Agree") to the item "Based on the graph, how much do you agree with the following statement: when it comes to this issue, ["Leavers" and "Remainers"/Republicans and Democrats] just seem to be two completely different types of people?"

Perceived Likelihood of Intergroup Compromise: Participants were asked to indicate their response (via a 100-point scale that ranged from "Not at All Likely" and "Extremely Likely") to the item "Based on the graph, how likely do you think it is that ["Leavers" and "Remainers"/Republicans and Democrats] could find ways to compromise on this issue?"

Agreement with Optimistic Political Statement: Participants were presented with the following statement and asked to indicate, via open-text response (which would later be blind-coded on a 1-5 scale by two independent raters), the degree to which they agreed or disagreed with the sentiment expressed:

Imagine you came across the following comment in one of your networks:

"Listen, ["Leavers" and "Remainers"/Republicans and Democrats] are definitely different – I get that. They disagree on a lot and they'll probably always argue over the right way to do things. But I actually think – if we look past some of the more extreme members of each group – we'd probably find that we're not as far apart on some issues than a lot of people think. I don't think we'll ever be fully in agreement on everything, but I do think if we made more of an effort to look for what unites us as opposed to what divides us, we might find areas where we can compromise and cooperate."

Hypotheses

Table 1 below depicts the pre-registered (OSF URL: <https://osf.io/yzq5w>) hypotheses for the study:

Table 1

Study Hypotheses

| Label | Hypothesis |
|-------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H1 | Mode of presentation will have a significant main effect on issue-specific perceptions of ideological polarization but will not have a significant main effect on issue-specific affective polarization. For both measures of polarization, all three modes will promote significantly lower perceptions of both ideological and affective polarization relative to the no visualization condition. |
| H2 | Mode of presentation will have a significant main effect on levels of intergroup agreement relative to expectations. The full-range bar chart and the icon array histogram will promote significantly higher levels of agreement relative to expectations (i.e., participants will be more likely to report that the groups "appear to agree more than they expected") compared to the 1.5 SD bar chart. |
| H3 | Mode of presentation will have a significant main effect on perceived group "discreteness." The full-range bar chart and the icon array histogram will promote significantly lower levels of perceived discreteness compared to the no visualization condition, while no significant differences will emerge between the 1.5 SD bar chart and no visualization condition. |
| H4 | Mode of presentation will have a significant main effect on perceptions of likelihood of intergroup compromise. All three modes will promote significantly higher levels of compromise likelihood than the no visualization condition. |
| H5 | The full-range bar chart and the icon array histogram will promote estimates of intergroup opinion overlap that will be more accurate than the estimates produced within the 1.5 SD bar chart and no visualization conditions. |
| H6 | Mode of presentation will have a significant main effect on pre- and post-exposure measures of general ideological polarization, but will not have a significant main effect on pre- and post-exposure measures of general affective polarization. |

Results

An overview of results can be found in *Table 2* (for the UK sample) and *Table 3* (for the US sample) below:

Table 2*Main Effects of Mode of Data Visualization (United Kingdom Sample)*

| Outcome Measure | No Visualization | | Truncated Bar Chart | | Full-Range Bar Chart | | Icon Array Histogram | | <i>F</i> | η^2 | <i>p</i> |
|--------------------------------------------------------|------------------|-----------|---------------------|-----------|----------------------|-----------|----------------------|-----------|---------------------|----------|----------|
| | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | | | |
| Perceptions of Issue-Specific Ideological Polarization | 60.45 | 24.73 | 60.22 | 22.30 | 55.34 | 20.91 | 56.29 | 22.51 | 11.200 (3, 3299) | .010 | <.001 |
| Perceptions of Issue-Specific Affective Polarization | 55.82 | 24.51 | 55.54 | 23.60 | 51.91 | 23.80 | 54.14 | 22.68 | 4.706 (3, 3299) | .004 | .003 |
| Change in Overall Perceived Ideological Polarization | -0.21 | 19.19 | -0.75 | 19.44 | -4.83 | 20.91 | -4.53 | 21.61 | 11.890 (3, 3299) | .011 | <.001 |
| Change in General Affective Polarization | -6.10 | 28.45 | -8.83 | 30.70 | -6.29 | 31.29 | -8.69 | 28.24 | 1.539 (3, 2451) | .002 | .202 |
| Estimates of Intergroup Opinion Overlap | 43.59 | 23.98 | 44.55 | 21.49 | 51.46 | 19.82 | 54.71 | 19.84 | 52.600 (3, 3299) | .046 | <.001 |
| Intergroup Agreement Relative to Expectations | NA | NA | 2.92 | 1.01 | 3.24 | 1.02 | 3.38 | 1.05 | 42.101 (2, 2475) | .033 | <.001 |
| Perceptions of Intergroup Discreteness | 4.33 | 1.19 | 4.20 | 1.11 | 3.72 | 1.14 | 3.73 | 1.22 | 61.180 (3, 3299) | .053 | <.001 |
| Beliefs about Compromise Likelihood | 43.65 | 23.06 | 43.31 | 21.62 | 51.73 | 21.37 | 51.25 | 21.77 | 36.700 (3, 3299) | .032 | <.001 |
| Response to Optimistic Political Statement | 3.48 | 1.58 | 3.61 | 1.53 | 3.80 | 1.46 | 3.66 | 1.51 | 6.441 (3, 3298) | .006 | <.001 |

Table 3*Main Effects of Mode of Data Visualization (United States Sample)*

| Outcome Measure | No Visualization | | Truncated Bar Chart | | Full-Range Bar Chart | | Icon Array Histogram | | <i>F</i> | η^2 | <i>p</i> |
|--------------------------------------------------------|------------------|-----------|---------------------|-----------|----------------------|-----------|----------------------|-----------|---------------------|----------|----------|
| | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | <i>M</i> | <i>SD</i> | | | |
| Perceptions of Issue-Specific Ideological Polarization | 59.68 | 26.54 | 58.01 | 25.59 | 52.57 | 25.58 | 54.66 | 24.54 | 13.020 (3, 3296) | .012 | <.001 |
| Perceptions of Issue-Specific Affective Polarization | 56.75 | 25.86 | 57.24 | 26.10 | 55.41 | 25.07 | 53.66 | 24.43 | 3.276 (3, 3296) | .003 | .020 |
| Change in Overall Perceived Ideological Polarization | -1.07 | 21.04 | -2.12 | 21.71 | -5.05 | 21.64 | -2.58 | 23.48 | 4.862 (3, 3296) | .004 | .002 |
| Change in General Affective Polarization | -5.68 | 26.16 | -1.78 | 24.41 | -4.89 | 22.28 | -4.13 | 21.63 | 2.366 (3, 1856) | .004 | .069 |
| Estimates of Intergroup Opinion Overlap | 42.59 | 25.40 | 43.51 | 22.62 | 54.27 | 23.58 | 54.46 | 23.00 | 63.070 (3, 3296) | .054 | <.001 |
| Intergroup Agreement Relative to Expectations | NA | NA | 3.01 | 1.13 | 3.42 | 1.24 | 3.47 | 1.24 | 36.460 (3, 2472) | .029 | <.001 |
| Perceptions of Intergroup Discreteness | 4.58 | 1.27 | 4.46 | 1.21 | 3.71 | 1.32 | 3.88 | 1.36 | 90.800 (3, 3296) | .076 | <.001 |
| Beliefs about Compromise Likelihood | 41.35 | 24.70 | 41.39 | 23.81 | 51.07 | 24.95 | 50.04 | 25.54 | 38.120 (3, 3296) | .034 | <.001 |
| Response to Optimistic Political Statement | 3.47 | 1.63 | 3.60 | 1.57 | 3.78 | 1.50 | 3.71 | 1.58 | 6.044 (3, 3296) | .005 | <.001 |

Results for the one-way ANOVA for perceived ideological polarization indicated a significant main effect of mode of presentation on perceptions of ideological polarization for both the UK sample ($F(3, 3299) = 11.2$, $p < .001$, $\eta_p^2 = 0.010$) as well as for the US sample ($F(3, 3296) = 13.02$, $p < .001$, $\eta_p^2 = 0.012$). For the UK sample, Tukey HSD post-hoc tests determined that exposure to the full-range bar chart and icon array histogram promoted perceptions of ideological polarization that differed significantly from the no visualization condition ($p < .001$ and $p = .001$, respectively), but exposure to the 1.5 SD bar chart did not ($p = .997$). For the US sample, Tukey HSD post-hoc tests determined that exposure to the full-range bar chart and icon array histogram promoted perceptions of ideological polarization that differed significantly from the no visualization condition ($p < .001$ and $p < .001$, respectively), but exposure to the 1.5 SD bar chart did not ($p = .543$). Results of the one-way ANOVA for affective polarization also indicated a significant main effect of mode of presentation for both the UK sample ($F(3, 3299) = 4.706$, $p = .003$, $\eta_p^2 = 0.004$) as well as for the US sample ($F(3, 3296) = 3.276$, $p = .020$, $\eta_p^2 = 0.003$), albeit with a smaller effect size than that of the effects of mode on perceived ideological polarization (as measured by partial eta squared; see *Figure 2*). For the UK sample, Tukey HSD post-hoc tests determined that exposure to the full-range bar chart promoted perceptions of affective polarization that differed significantly from the no visualization condition ($p = .004$), but exposure to the icon array histogram and the 1.5 SD bar chart did not ($p = .472$ and $p = .995$, respectively). For the US sample, Tukey HSD post-hoc tests determined that neither exposure to the full-range bar chart, the icon array histogram, nor the 1.5 SD bar chart promoted perceptions of affective polarization that differed significantly from the no visualization condition ($p = .705$, $p = .065$, and $p = .980$, respectively).




Figure 2: Perceptions of ideological polarization (top row) and affective polarization (bottom row) by mode for the UK sample (left column) and the US sample (right column).

To determine whether the issue-specific effects of mode extended to more general measures of polarization, differences between the pre- and post-measures of overall (i.e., non-issue-specific) perceived ideological polarization and affective polarization were calculated to determine whether changes varied as a function of condition.

A one-way ANOVA revealed significant main effects of mode of presentation for changes (i.e., post-pre) in overall perceived issue polarization for both the UK sample ($F(3, 3299) = 11.89, p < .001, \eta_p^2 = 0.011$) and the US sample ($F(3, 3296) = 4.862, p = .002, \eta_p^2 = 0.004$; see *Figure 3*).




Figure 3: Change in the pre- and post-measurements of perceived overall ideological polarization by mode for the UK sample (left) and the US sample (right).

A one-way ANOVA on changes in self-reported general affective polarization pre- and post-mode exposure for self-identified Leavers and Remainers in the UK (a subsample of 2,455) revealed no significant change in overall affective polarization ($F(3, 2451) = 1.539, p = .202, \eta_p^2 = 0.002$). In the US, pre- and post-measures of general affective polarization for Democrats and Republicans (a sub sample of 1,860) revealed no significant change in overall affective polarization based on a one-way ANOVA of the difference in pre-post scores ($F(3, 1856) = 2.366, p = .069, \eta_p^2 = 0.004$; see *Figure 4*).




Figure 4: Change in the pre- and post-measurements of general affective polarization by mode for the UK sample (left) and the US sample (right).

An exploratory one-way ANOVA on estimates of intergroup opinion overlap revealed a significant main effect of mode of presentation for both the UK sample ($F(3, 3299) = 52.6, p < .001, \eta_p^2 = 0.046$) and the US sample ($F(3, 3296) = 63.07, p < .001, \eta_p^2 = 0.054$; see *Figure 5*). However, beyond main effects (which can only detect whether the modes differ significantly *from one another*), the modes were also assessed relative to an accuracy benchmark known as *percentage of common scores* (PCS; Hanel et al., 2019). For the UK data, the PCS was approximately 63%; for the US data, the PCS was approximately 71%. Across both samples, all four conditions yielded estimates that fell significantly below their respective accuracy benchmarks. However, the full-range bar chart and the icon array histogram (but not the 1.5 SD bar chart) elicited estimates that differed significantly from the no visualization condition in the desired direction (i.e., more closely approaching the accuracy benchmark) in both the UK ($t(1591.6) = 7.271, p < .001, d = 0.36$ and $t(1592.6) = 10.27, p < .001, d = 0.51$, respectively) and US ($t(1639.7) = 9.686, p < .001, d = 0.48$ and $t(1632) = 9.951, p < .001, d = 0.49$, respectively) samples.




Figure 5: Estimates of intergroup opinion overlap by mode for the UK sample (left) and the US sample (right). The horizontal purple line indicates the actual level of overlap for the data used in each sample.

A one-way ANOVA was conducted to assess the main effect of mode of presentation on level of intergroup agreement relative to expectations, which revealed a significant main effect for both the UK sample ($F(2, 2475) = 42.01, p < .001, \eta_p^2 = 0.033$) and US sample ($F(2, 2472) = 36.46, p < .001, \eta_p^2 = 0.029$; see *Figure 6*). For the UK sample, Tukey HSD post-hoc tests determined that exposure to the full-range bar chart and icon array histogram promoted greater perceptions of intergroup agreement relative to expectations than the 1.5 SD bar chart condition ($p < .001$ and $p < .001$, respectively). For the US sample, Tukey HSD post-hoc tests also determined that exposure to the full-range bar chart and icon array histogram promoted greater perceptions of intergroup agreement relative to expectations than the 1.5 SD bar chart condition ($p < .001$ and $p < .001$, respectively).




Figure 6: Perceptions concerning levels of intergroup agreement relative to participant expectations by mode for the UK sample (left) and the US sample (right). Higher responses correspond to perceived levels of agreement greater than what participants expected.

A one-way ANOVA was conducted to assess the main effect of mode of presentation on perceived group discreteness, which revealed a significant main effect for both the UK sample ($F(3, 3299) = 61.18, p < .001, \eta_p^2 = 0.053$) and US sample ($F(3, 3296) = 90.8, p < .001, \eta_p^2 = 0.076$; see Figure 7). For the UK sample, Tukey HSD post-hoc tests determined that exposure to the full-range bar chart and icon array histogram promoted significantly lower perceptions of intergroup discreteness than the no visualization condition ($p < .001$ and $p < .001$, respectively) but exposure to the 1.5 SD bar chart did not ($p = .134$). For the US sample, Tukey HSD post-hoc tests determined that exposure to the full-range bar chart and icon array histogram promoted significantly lower perceptions of intergroup discreteness than the no visualization condition ($p < .001$ and $p < .001$, respectively) but exposure to the 1.5 SD bar chart did not ($p = .236$).




Figure 7: Perceptions of group discreteness by mode for the UK sample (left) and the US sample (right).

A one-way ANOVA was conducted to assess the main effect of mode of presentation on the perceived likelihood of intergroup compromise, which revealed a significant main effect for both the UK sample ($F(3, 3299) = 36.7, p < .001, \eta_p^2 = 0.032$) and US sample ($F(3, 3296) = 38.12, p < .001, \eta_p^2 = 0.034$; see Figure 8). For the UK sample, Tukey HSD post-hoc tests determined that exposure to the full-range bar chart and icon array histogram promoted significantly higher perceptions of compromise likelihood than the no visualization condition ($p < .001$ and $p < .001$, respectively) but exposure to the 1.5 SD bar chart did not (p

= .990). For the US sample, Tukey HSD post-hoc tests determined that exposure to the full-range bar chart and icon array histogram promoted significantly higher perceptions of compromise likelihood than the no visualization condition ($p < .001$ and $p < .001$, respectively) but exposure to the 1.5 SD bar chart did not ($p > .999$).




Figure 8: Perceptions of the likelihood of compromise by mode for the UK sample (left) and the US sample (right).

Finally, an exploratory one-way ANOVA was conducted to assess whether main effects of mode would be observed on participant response valence to an optimistic political statement. To run the analysis, open-ended participant responses were manually-coded (on a 1 to 5 scale) by two third-party raters blind to the conditions of the sample. This analysis revealed a significant main effect for the UK sample ($F(3, 3298) = 6.441, p < .001, \eta_p^2 = 0.006$) as well as for the US sample ($F(3, 3296) = 6.044, p < .001, \eta_p^2 = 0.005$; see Figure 9).




Figure 9: Average sentiment score following exposure to an optimistic political statement by mode for the UK sample (left) and the US sample (right).

Discussion

The results demonstrate that perceptions of polarization are not fixed and can be influenced by the mode in which data about group differences are graphically represented. In both the UK and the US, mode of presentation had a significant impact on perceptions of how ideologically polarizing the issue seemed to be, how affectively polarized the two groups were likely to be on the issue, pre- and post-measures of perceived

overall ideological polarization, perceptions of intergroup opinion overlap, perceptions of intergroup agreement relative to expectations, perceptions of group discreteness, perceptions of likelihood of compromise, and expressions of positive sentiment in response to political optimism. The most robust results concerned mode's impact on estimates of intergroup opinion overlap, perceived intergroup agreement relative to expectations, perceived group discreteness, and perceptions of the likelihood of compromise.

Importantly, the findings indicate that certain modes consistently yield more desirable results than others. More specifically, the full-range bar chart and the icon array histogram routinely reduced polarized perceptions, increased accuracy of intergroup opinion overlap estimates, and decreased perceptions of group discreteness relative to both the 1.5 SD bar chart and the no visualization condition. Such an insight is critical as it implies that deliberate choices related to data visualization can produce meaningful shifts in partisan perceptions. For example, in the US, the choice to not depict data and to simply allow individuals to make assumptions based on their prior knowledge of the issue (i.e., whether immigrants are generally good for America's economy) as opposed to utilizing a full-range bar chart to depict the data would result, on average, in a 12.67%⁴⁹ difference in perception of issue-specific ideological polarization, 24.13% difference in estimates of intergroup opinion overlap, a 20.98% difference in perceptions of group discreteness, and a 21.05% difference in estimates of compromise likelihood – all of which represent highly significant statistical differences. Relatedly, it's not just the choice between including a visualization versus not including one that proves consequential; a "poor" choice of mode of data presentation (versus a more sensible one) can also have profound effects on partisan perceptions and beliefs. For instance, if an agency in the UK were to opt for a 1.5 SD bar chart as opposed to an icon array histogram to depict the polarized issue featured within the study, such a choice would result, on average, in a 6.74% difference in perceptions of issue-specific ideological polarization, a 20.48% difference in estimates of intergroup opinion overlap, a 11.96% difference in perceptions of group discreteness, and a 16.78% difference in estimates of compromise likelihood.

Not only does the current study suggest that perceptions of levels of polarization are malleable (even for issues where degrees of polarization may be historically entrenched), it also shows that perceptions of *political groups themselves* are malleable. Our findings suggest that mode of presentation can also moderate the degree to which we see political opponents as discrete and thus fundamentally different from one another. Such a finding should not be overlooked as the perceived discreteness of groups has been implicated in prior literature as a contributor to essentialist beliefs (Haslam, Bastian, & Bissett, 2004), which appear to be a necessary component for certain forms of dehumanization (Haslam, 2006). Consequently,

⁴⁹ Differences were calculated by taking the absolute value of the change in value (e.g., between the full-range bar chart measure and the no visualization measure), dividing it by the average of the two figures, and multiplying this figure by 100.

the ability of certain modes to mitigate perceptions of discreteness may prove to be a subtle yet valuable lever by which agents may be able to prevent the exacerbation of interparty hostility and alter the trajectory of intergroup relations.

Moreover, the study also offered a unique, serendipitous insight. Exploratory analyses revealed that perceptions of ideological polarization were greatest for participants with the most education (see Appendix D). While such a finding may offer a modest contribution to the literature on the relationship between education and polarization, we believe the more relevant takeaway for the current study is that, if overestimations of polarization are being disproportionately driven by the highly educated, and should this group be more likely to consume certain types of news (Schulz, Levy, & Nielsen, 2019), perhaps they would be more likely to encounter interventions such as the one being proposed here (i.e., embedding “better” visualizations of inter-party data into communications). It would be important to test whether interventions which embed “better” visualizations across these news and polling platform mediums could yield stronger effects when sustained over time.

The current study provides a framework by which any media outlet, polling organization, government agency, or any other political actor interested in reducing polarized perceptions can implement a simple yet effective intervention. Moreover, the potential scalability of this intervention may make it a more defensible candidate for widespread application than other options found in the academic literature. Though perceptions of polarization are one step removed from actual polarization, scholars in the field have argued that the two processes might be self-reinforcing (Lees & Cikara, 2021), lending greater urgency to the need to curtail exaggerated beliefs about political others. However, we strongly suspect the implications of this research may also extend beyond the domain of political polarization. The issue of persistent, seemingly intractable disagreement is not exclusive to politics. Factions may – and often do – form both within and between organizations and institutions based on disputes stemming from divergent philosophies or incompatible beliefs. Should such disputes share similar dynamics to political polarization, *perceptions* of disagreement may outpace the *actual levels* of disagreement, but may also, critically, create a self-reinforcing cycle wherein mere perceptions may ultimately drive the true progression of the divide. Ultimately, over time, differences of opinion risk becoming “essentialized,” at which point the motivation for reconciliation suffers as groups come to view one another as *fundamentally* different. The current research finds that certain modes (such as truncated bar charts) may do little to disrupt the momentum of this cycle, and often simply provide justification for the antagonism via depictions which support the irreconcilable disagreement narrative while bolstering the perceived “otherness” of the opposition. However, modes such as full-range bar charts and icon array histograms, via their non-exaggerated depiction of group differences and more balanced coverage of group agreement, respectively, may be capable of violating expectations regarding the extent of group differences. This violation may effectively disrupt the

cyclical escalation of intergroup animus, permitting groups the opportunity to reassess their actual levels of disagreement, acknowledge their (often substantial but under-emphasized) areas of agreement, and potentially shift the trajectory of the relationship in a more positive direction.

While the current study highlights promising evidence for the role of mode of presentation in shifting perceptions of polarization, it is not without its limitations. Firstly, while mode was shown to be capable of shifting myriad intergroup perceptions relevant to polarization (including multiple measures of polarization itself), the current study provides only weak evidence that it might be capable of impacting political sentiment. Such a finding is consistent with our prior research, which found little evidence to suggest that mode exerts a significant impact on partisan behavior, and inconsistent evidence for the longevity of the effects of the single-exposure intervention on perceptions over a multi-week time period. While we believe that shifting perceptions is undoubtedly an important step, more work is required, both in academia as well as in partnership with real-world stakeholders, to determine whether (and, if so, to what degree) the effects of mode of presentation might extend beyond perceptions to impact partisan behavior, and whether shifts in perception may become more permanent when repeated-exposure paradigms are employed. Secondly, many of the outcome variables used in this investigation contained the phrase “based on the data,” explicitly encouraging participants use the visualizations to inform their responses. While this was a deliberate methodological choice on our part, it is not unreasonable to assume that the effects observed may have been influenced by this deliberate refocusing of participant attention. Future research should seek to determine whether similar visualization effects are obtained without the need to provide such overt cues. Thirdly, while some mode effects were pronounced, the interventions were not longitudinal in nature, and thus future research should investigate their long-term effects and rates of decay. Finally, it will be important for future research to establish how these effects fare in a “competitive” message environment (see Bolsen & Shapiro, 2017) where other political actors might be trying to shift perceptions in *more* polarized directions by emphasizing group differences.

Conclusion

Overall, the present study confirmed that changing mode of presentation – a simple and highly-scalable intervention – is capable of exerting significant influence on an array of critical interparty perceptions and beliefs, even on a divisive and highly-polarized issue for which individuals ostensibly hold strong attitudinal priors.

Between partisan “echo chambers” (Terren & Borge-Bravo, 2021), selective, ideologically-aligned media exposure (Stroud, 2010), and the rapid emergence of new, polarized, and otherwise hostile information ecosystems online, it’s reasonable to assume that individuals might have disproportionate exposure to biased data. Truncated bar charts, by virtue of their ability to visually accentuate of group *differences*, are a

commonly-used method of data presentation (Hanel et al., 2019) and may also constitute a depiction of political data more commonly used in such ecosystems, and thus more congruent with participant expectations. Conversely, the full-range bar chart's *reduced* exaggeration of group differences and the icon array histogram's inclusion of distributional and similarity information may induce expectation violations and surprise, which can play an important role in belief updating (e.g., Nassar, Wilson, Heasley, & Gold, 2010). Ultimately, we believe that many of the results contained herein may be products of the interaction between participant *expectations* about what intergroup agreement *should* look like (informed by inputs which tend to highlight disagreement in a highly-polarized environment) and what intergroup data depicted in less biased, more balanced ways actually *does* look like.

While further work is needed to determine the extent and boundary conditions of the effects of mode (particularly as they relate to "real world" behavior), the evidence presented clearly demonstrates that those who disseminate data about political groups should choose their visualizations carefully and deliberately, as the manner in which the data is presented plays a significant role in a multitude of intergroup-relevant perceptions and attitudes.

CHAPTER 3

INVESTIGATING WHETHER DATA SUMMARISATION CHOICES SHIFT VIEWS OF POLARISATION AND INTERPARTY DYNAMICS OVER TIME⁵⁰

Abstract

While research investigating the impact of data visualisations is critical, it's important to recognise that not all disseminations of political information are accompanied by visual depictions. In fact, the vast majority of political content is consumed in the form of either written or spoken communication (e.g., articles, soundbites, online and in-person discussions, etc.), all of which may fail to provide any visual depiction of the data being discussed. Consequently, the current research sought to determine whether the manner in which individuals choose to summarise political data plays a role in moderating perceptions of polarisation and shifting intergroup dynamics. Across an eight-week period, a sample of US participants (n = 594) was repeatedly exposed to data describing where political partisans stand on key divisive issues, with each condition seeing the same underlying data summarised in a different fashion. Results suggest that while different summarisation techniques may not result in divergent perceptions of polarisation or attitudes toward one's political outgroup, repeated exposure to any form of data summarisation – even on highly-polarised issues – may promote more positive intergroup perceptions relative to participants who receive no exposure at all.

Americans get their political news through a number of mediums. According to a 2020 report by Pew Research Center, approximately 45% of Americans get their political news from some form of television (i.e., cable, local, or network TV), 25% from a news website or app, 18% from social media, 9% from radio programs, and 3% from print (i.e., newspapers or magazines; Pew Research Center, 2020). While most of these information channels present the *opportunity* for sources to include visualisations of data, such modes of presentation are – in the event they *are* included – generally complementary pieces to the story being discussed or reported on. Consequently, although data visualisations undoubtedly represent one pathway through which political information may be transmitted and through which political perception can be moderated, we contend that data *summarisation* – or the verbal or written manner in which sources choose

⁵⁰ University of Cambridge ethics identification number 2581.99

to summarise sets of political data ultimately featured in online articles, TV programs, or radio shows – may be the more widely-used (and thus, arguably more consequential) of the data presentation formats.

Due to time, attention, and access constraints (which may prevent a comprehensive examination of complex datasets), the majority of the general public relies on the media to summarise data in a way that is both succinct and which provides the audience with a general “gist” of the total available information. Consequently, sources will often utilise common “summary statistics,” or measures which aim to communicate overall patterns in data in as concise a manner as possible. While by no means exhaustive, the most commonly used summary statistics typically take two forms: measures of central tendency (e.g., mean, median, and mode) and measures of data dispersion or “spread” (e.g., range and standard deviation). Although each type of summary statistic succeeds in providing valuable information about the data, it is often imperative to have access to both types of measurements in order to form a holistic (and, ultimately, more accurate) impression of the dataset. An illustrative example is Anscombe’s quartet, wherein four datasets feature identical means for both the x and y coordinates, but whose ranges of responses and heterogeneity of dispersion patterns paint a picture of dramatically different compositions of data – differences which otherwise would not have been revealed by access to the means alone.

Unfortunately, measures of central tendency – especially means – seem to consistently take precedence over measures of dispersion when summarising data across formats. For example, Hanel, Maio, and Manstead (2019) note that bar charts depicting group means continue to be the dominant method of data presentation in psychology. In an analysis of prominent psychological journals, Lane and Sándor (2009) found that only 10% of graphs depicted distributional information that went beyond central tendency. Importantly, even when visualisations like bar charts *do* present distributional information, they fail to capture the complete range of responses and instead utilise sample-size-dependent measures such as standard errors and confidence intervals. This prioritisation of group mean data and the corresponding failure to sufficiently incorporate characteristics of dispersion can also be seen in the way researchers have traditionally chosen to measure polarisation. Building on work by scholars like Esteban and Ray (1994), who sought to critically examine the correct manner in which to quantify polarisation, Levendusky and Pope (2011) have critiqued how the vast majority of prior research within the field has tended to exclusively measure polarisation via difference-of-means tests; that is, if the average responses of Group A and the average responses of Group B are measured to be significantly different, then the two groups are effectively considered to be “polarised.” However, a unilateral reliance on group means ignores aspects of the group compositions which should arguably factor into a more holistic and nuanced analysis of polarisation. For example, Mehlhaff (2023) argues that to accurately assess polarisation, one must not rely exclusively on the difference *between* groups (i.e., as assessed by calculating the difference between group means) but should instead also consider factors like the degree of response *variability within* groups (i.e., levels of intragroup dispersion). What’s

more, by taking into account the full distributional information of both groups, researchers can go beyond a sole focus on group differences to also begin to effectively consider metrics aimed at estimating intergroup *similarity* (e.g., “percentage of common scores,” or PCS; see Hanel et al., 2019).

Arguments about the proper way to measure polarisation are important as the authors engaged in this debate are essentially arguing that *the parts of the data we choose to focus on* (e.g., group means, inter- and intragroup dispersion, intergroup overlap, etc.) will play the defining role in the *conclusions we ultimately draw* (i.e., whether or not groups should be categorised as polarised). The premise of the current study rests on a similar principle: the *perceptions partisans come to hold* may be influenced by the *parts of the data to which they are exposed* (e.g., means versus level of agreement).

While, to the best of our knowledge, little empirical work has examined how political perceptions may be impacted by disparate data summarisation techniques *per se*, a plethora of research has been dedicated to better understanding how subtle shifts in the way messages are presented can yield meaningful shifts in political attitudes. Specifically, a robust body of research on *message framing* has determined that minor, sometimes seemingly insignificant alterations in the way a communication is structured is capable of bringing about significant changes in target attitudes, decisions, and preferences (e.g., Chong & Druckman, 2007).

Critically, while message framing has been shown to consistently impact more general, apolitical decision-making (e.g., Tversky & Kahneman, 1974), a large cohort of scholars have also documented its influence on political attitudes. For instance, Berinsky and Kinder (2006) argue that individuals utilise frames as means by which to organise complex, often dynamic political issues into a coherent narrative. Using the 1998-1999 war in Kosovo as an example, they exposed participants to a series of articles that featured either a pro-interventionist framing (using a story structure) of the unfolding crisis (i.e., the US *should* intervene) or an anti-interventionist framing. Their findings provided strong support to suggest that disparate media frames not only impacted the way participants *remembered* particular information about the crisis (i.e., exerting an effect on cognitive faculties such as memory) but also their opinions about the correct courses of current and future political action. Other research has shown that *repeated exposure* to particular frames produce even stronger impacts on attitude change than single exposure paradigms, with effects being both more powerful while also promoting more persistent attitude change (Lecheler, Keer, Schuck, & Hänggli, 2015). However, not all research concludes that repeated exposure necessarily strengthens framing effects. For instance, Lecheler and de Vreese (2013) conducted studies which showed that, while the power of repetitive framing techniques could be improved by reducing the delay between sequential exposures, the overall pattern of results did not provide sufficient evidence to suggest that repeated exposure alone could systematically impact opinion formation. Moreover, others have found that the duration, stability, and

resilience of framing effects seems to be dependent on whether individuals are also exposed to *competing* frames (Lecheler & de Vreese, 2016).

In spite of the mixed evidence for the efficacy of repeated exposure paradigms, we believe such a design more faithfully simulates the information environments in which most participants exist, where they may be repeatedly exposed to a particular narrative, perspective, or ideologically-consistent angle depending on how frequently they return to a select few preferred sources of information (e.g., TV channel, website, etc.). In addition to the fact that repeated exposure designs may offer enhanced ecological validity versus single-exposure interventions in the domain of news framing, we also conjecture that new or unfamiliar frames may require multiple exposures in order for participants to have an opportunity to properly acclimatise themselves. For instance, a single exposure to a new or unfamiliar way of looking at data might engender confusion, disbelief, or reactive rejection (especially if the new perspective challenges a long-held old one). While multiple exposures may not necessarily promote a more positive reappraisal of the new perspective, phenomena such as the “mere exposure effect” (e.g., Fang, Singh, & Ahluwalia, 2007) suggest that such repeated exposures, beyond providing more *opportunities* to reconsider the new perspective, may create a greater sense of familiarity and consequently comfort with it.

While framing can be a powerful tool to bring about attitude and opinion shifts, the mechanisms by which these processes operate remain debated. Nelson, Clawson, and Oxley (1997) argue that providing audiences with different frames may “affect the balance of considerations that individuals weigh when contemplating political issues” (p. 235). Slothuus (2008) posits a framework for a dual-process model, wherein framing effects operate via two distinct routes (i.e., either changing the *importance* of certain issue considerations or the *content* of those considerations), and where the relative influence of each route is dependent on variables such as level of political awareness and the strength of one’s values.

Scholars broadly separate framing techniques into two distinct categories: equivalence framing and emphasis framing. While equivalence framing involves “semantically distinct but logically equivalent portrayals” of information (Druckman, 2011, p. 282), emphasis framing refers to how a communicator may choose to emphasise different aspects of the same issue, such as a foetus’ right to life versus a woman’s right to choose in the abortion debate (McCaffrey & Keys, 2000). Germane to the current study, we contend that the type of summary statistics that sources choose to disseminate represent a form of emphasis framing. Specifically, when a source chooses to disseminate group means, we hypothesise that they may (intentionally or otherwise) emphasise group *differences* more so than similarities; conversely, when they choose to disseminate information about group opinion overlap and non-overlap (e.g., “percentage of common scores,” or PCS; Hanel et al., 2019), they may instead promote a comparative shift away from exclusively emphasising

group differences to creating a more balanced emphasis on levels of both disagreement *and* agreement between groups.

Ultimately, the current study seeks to quantify the longitudinal impact that repeated exposure to different summarisations of partisan data has on political perceptions and behaviour. Broadly, though speculative, we hypothesise that certain summary statistics, such as group means – which provide little to no information concerning group response spread or dispersion – may promote exaggerated perceptions of polarisation and other opposition group-level characteristics such as outgroup homogeneity (e.g., Boldry, Gaertner, & Quinn, 2007) and outgroup extremism (e.g., Ahler, 2014) relative to statistics like PCS. Moreover, we hypothesise that individuals exposed solely to group mean data, being devoid of information pertaining to the range or dispersion of intragroup opinion and, critically, the degree of intergroup opinion overlap, may ultimately come to hold heightened perceptions of polarisation and beliefs about “ideological distance” than individuals repeatedly exposed to PCS data.

Method

Participants

A sample of American participants were recruited via the Prolific Academic platform and paid an approximate rate of £7.50/hour. Following exclusions for non-consent, attention check failure⁵¹, or incomplete data (either due to failure to complete all items each week or failure to complete all eight consecutive weeks of participation), the final sample was 594⁵² participants. Table 1 summarises the demographic composition of the sample.

⁵¹ Each week of the eight week study, an attention check was embedded in the survey. As each week depicted data from the results of topical opinion polls, the majority of weeks featured an attention check which asked participants – following exposure to the results of that week’s poll – to confirm what the poll was about (e.g., immigration, abortion, education, etc.). Each week’s attention check question had four options from which participants could choose (with only one being correct). Participants were excluded from future waves for any failed attention check, and only participants who successfully completed attention checks across all eight weeks were included in the final sample. As an example, during Wave 1, 93.29% passed the attention check.

⁵² A priori power analyses were conducted to determine the sample size required to achieve sufficient statistical power to detect a small-to-medium effect size (i.e., Cohen’s *f* of 0.20) with a power level of 0.90 and an alpha threshold of 0.05. The analysis indicated a sample size of approximately 63 participants per group would be required. The effect size was chosen based on the findings from Chapter 1 and Chapter 2, where the effect sizes observed were routinely large, but where the stimuli (i.e., data visualisations) might yield more pronounced effects than data summarisations (which we hypothesised to exert a more subtle influence on perceptions). Consequently, we wished to conduct a more conservative power analysis (which assumed slightly smaller effects). Ultimately, we estimated a 5% attrition rate per week, and thus sought to recruit approximately 700 participants to achieve the desired number (post-dropout) in the final sample. However, we underestimated the level of attrition, and thus needed to conduct two additional rounds of data collection. Prior to exclusions, the total sample size at Week 1 (cumulatively, for all three waves of data collection) was 1,285. For those who provided data, the sample composition prior to exclusions had a mean age of 41.06, was comprised of 662 females, 558 males, and 17 other (with 48 non-responses), and had 697 self-reported Democrats, 540 self-reported Republicans, and 46 self-reported other (with 2 non-responses).

Table 1*Composition of Study Sample*

| Identity Characteristic | Sample Composition |
|-------------------------|--------------------------------------------------------------------------------------------------------|
| Sex | Female: 50.34% ($n = 299$), Male: 48.15% ($n = 286$), Other: 1.52% ($n = 9$) |
| Age | $M = 43.53$, Range: 19-84 |
| Education | 60.77% ($n = 361$) had attained a bachelor's degree or higher |
| Political Affiliation | Democrat: 53.03% ($n = 315$), Republican: 46.97% ($n = 279$), Independent/Other: 0.00% ($n = 0$) |

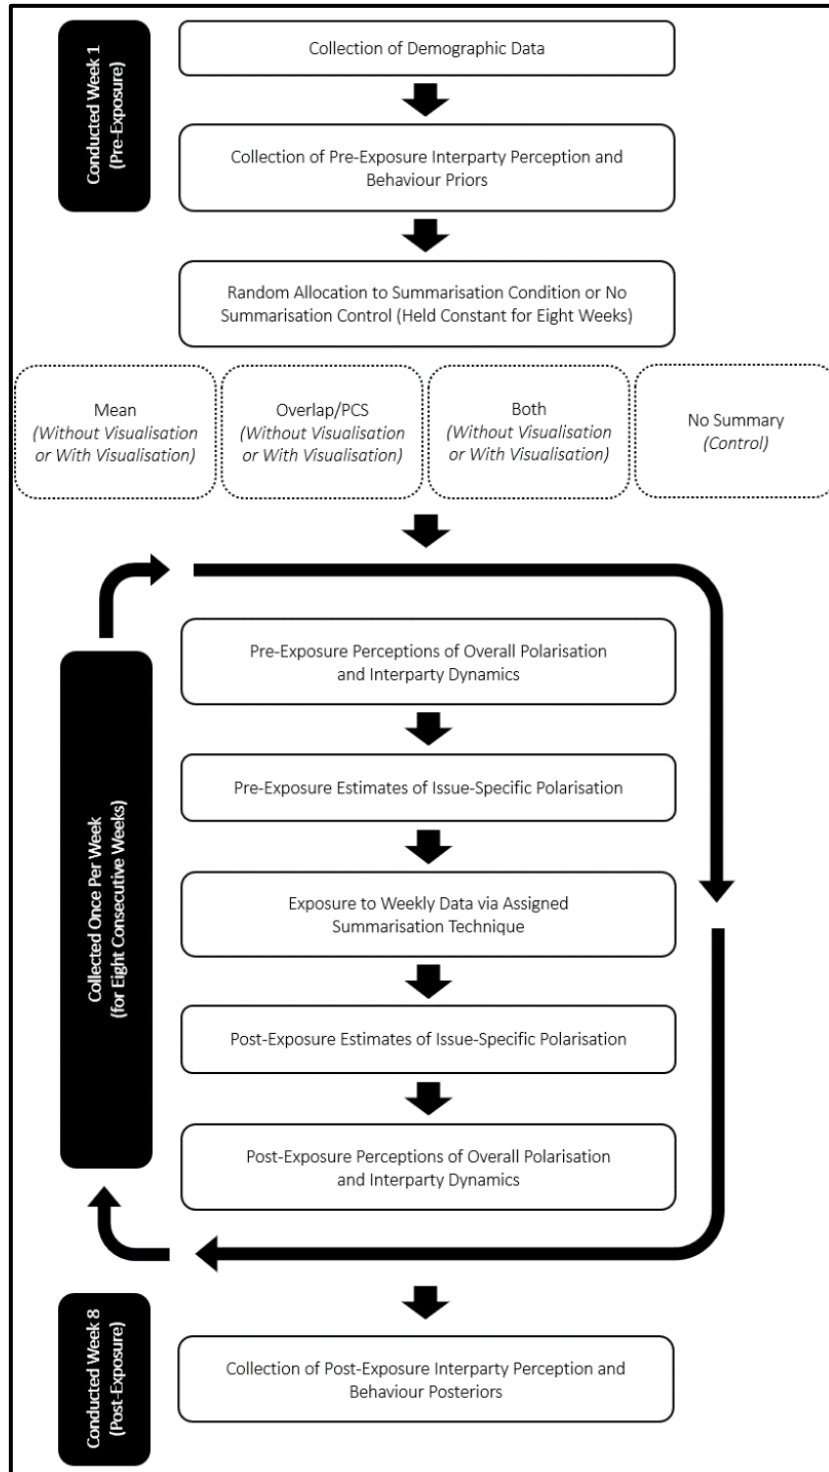
Among participants who openly identified as Democrat or Republican, 53.65% ($n = 169$) categorised themselves as “strong” Democrats and 43.37% ($n = 121$) categorised themselves as “strong” Republicans. Due to the study’s stated intent (i.e., to examine the longitudinal effects of disparate forms of summary statistics on political *partisans*), participants who indicated their political identity as anything other than Republican or Democrat were not permitted to take part in the study.

Design

The study employed a 3x2 (summary statistic used by visualisation inclusion/omission), between-subjects, longitudinal design (with an additional, minimal-exposure control condition; see Figure 1).

Figure 1

Design Outline for Longitudinal Summarisation Technique Study



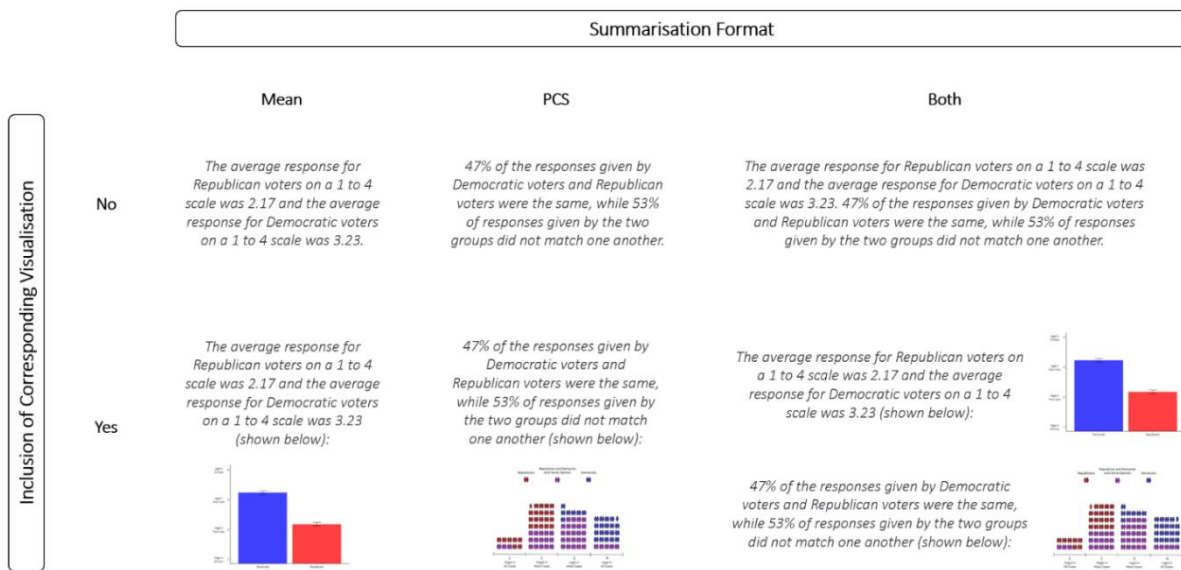
Procedure

Once a week for a period of eight consecutive weeks, participants were exposed to actual data concerning partisan attitudes (i.e., Republican and Democratic voters) on divisive political issues (sourced from

YouGov.com). Based on the condition to which the participant had been randomly assigned, each week’s data was summarised either via dissemination of group means, percentage of common scores (i.e., PCS; see Hanel et. al., 2019), or a combination of both (randomly counterbalanced to mitigate order effects). Moreover, the summary statistics were either shown in isolation or accompanied by a visualisation which corresponded with the data (i.e., full-range bar charts for group means, icon array histograms for PCS; see Figure 2). For the control condition, participants were shown no summaries of data, but were simply asked to provide their opinions on the issues “based on what they know about Republican and Democrat voters.”

Figure 2

Example of Condition Stimuli



Note: The condition-specific stimuli above were taken directly from Week 1 of the study and concern partisan perceptions on the legality of abortion. Subsequent weeks feature data from different divisive issues, but the formatting of the data remains identical to the example above.

Each week, participants were asked a battery of questions both before and after exposure to that week’s data summary (see *Measures*). There was a set of *issue-agnostic* questions (e.g., related to *overall* perceptions of affective and ideological polarisation, regardless of issue) as well as a set of *issue-specific* questions (i.e., directly related to intergroup perceptions concerning that week’s issue). Additionally, at Week 1 and Week 8, a unique set of pre/post items were administered to assess the possibility of emergent participant attitude shifts over the course of the longitudinal study.

Measures


Each week, participants were administered sets of issue-agnostic and issue-specific items both before and after data exposure. The issue-agnostic items included:

General Affective Polarisation: Participants were asked to indicate, on a scale from 0 (“Extremely Cold/Negative”) to 100 (“Extremely Warm/Positive”), how they felt towards voters⁵³ from each party.

Perceived Overall Ideological Polarisation: Participants were asked to indicate, on a scale from 0 (“Not at All Polarised”) to 100 (“Extremely Polarised”), how polarised Republican and Democratic voters seem to be overall on key issues.

Perceived Ideological Distance of Average Partisan Voters: Participants were asked to indicate, using a six-point visual scale (see Figure 3), which image they believed best represented how far apart the average Republican voter and average Democratic voter are on key issues.


Figure 3
Visual Scale Used for Perceived Ideological Distance of Average Partisan Voters Item



Perceived Level of Agreement of Average Partisan Voters: Participants were asked to indicate, using a six-point visual scale (see Figure 4), which image they believed best represented the typical level of agreement/disagreement between Republican and Democratic voters on key issues.

⁵³ Due to findings which suggest that Republican and Democrat legislators may hold more extreme policy positions than their corresponding constituent bases (e.g., Bafumi & Herron, 2010), and that participants may spontaneously think of party *elites* – against whom they generally harbour more animus – when asked to provide estimates of interparty polarisation (Druckman & Levendusky, 2019), we opted to specify that the targets of interest during participant responses should be *partisan voters*.

Figure 4
*Visual Scale Used for Perceived Ideological Overlap
of Average Partisan Voters Item*



Note: Participants were told that the purple area represented interparty agreement.

The issue-specific items included:

Perceived Issue-Specific Ideological Polarisation: Participants were asked to indicate, on a scale from 0 (“Not at All Polarised”) to 100 (“Extremely Polarised”), how polarised they believed Republican and Democratic voters seemed to be on that week’s issue.⁵⁴

Additionally, at Week 1 and Week 8, participants were asked to complete the following items (used as pre-post measures):

Dictator Game: Participants were told they would be playing the “Dictator Game” and would be “randomly assigned” both a role (i.e., the “decider” or the “receiver”) and a partner. In reality, in both iterations of the game (i.e., at Week 1 and Week 8), participants were assigned to the role of the “decider” and partnered with an individual whom they were told was a member of the opposing political party.

Perceived Outgroup Homogeneity: Participants were asked to indicate, on a scale from 0 (“Not at All Similar”) to 100 (“Extremely Similar”), how similar they believed outgroup party members were to one another in their opinions on key issues.

⁵⁴ Participants were also asked to estimate the average responses of each party to each week’s issue as well as the PCS, however these estimates were not used in any subsequent analyses.

Perceived Outgroup Extremism: Participants were asked to indicate their level of agreement, on a scale from 0 (“Strongly Disagree”) to 100 (“Strongly Agree”), to the statement “It’s not just a small percentage: the vast majority of [outgroup] voters are extremists.”

Perceived Difference from Average Outgroup Member: Participants were asked to indicate, on a scale from 0 (“Not at All Different”) to 100 (“Extremely Different”), how different they believed they were from the average outgroup voter.

Anti-Democratic Attitudes: A collection of four items, adopted from Voelkel et al. (2022), which served to measure the degree to which participants endorsed anti-democratic positions. These were:

1. [Ingroup] should reduce the number of polling stations in areas that support [outgroup].
2. [Ingroup] governors should ignore unfavourable court rulings by [outgroup]-appointed judges.
3. [Ingroup] governors should prosecute journalists who accuse [ingroup] politicians of misconduct without revealing sources.
4. [Ingroup] should not accept the results of elections if they lose.

Willingness to Engage in Civil Cross-Party Conversation: Participants were asked to indicate how open they would be to participating in a short (i.e., approximately 20-30 minute) discussion where members of both parties would come together to respectfully discuss pressing political issues; an opportunity to “engage with members of the other party on critical issues in a constructive, civil, and non-combative way; to listen to one another’s perspectives and try to find common ground.”

Hypotheses

Table 2 depicts the pre-registered (OSF URL: <https://osf.io/98bx2>) hypotheses for the study.

Table 2*Pre-Registered Hypotheses for the Study*

| Label | Hypothesis |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H1 | The pre-post change (i.e., between Week 1 and Week 8) in measures of perceived overall ideological polarisation will be significantly greater for participants in the PCS-only condition as opposed to participants in the mean-only condition. More specifically, those in the PCS-only condition will reduce their perceptions of overall ideological polarisation significantly more than those in the mean-only condition. |
| H2 | The change between the first (i.e., Week 1) and last (i.e., Week 8) measures of the perceived level of agreement measure will be significantly greater for participants in the PCS-only condition than participants in the mean-only condition. More specifically, those in the PCS-only condition will display a significantly greater shift towards perceiving more intergroup overlap than those in the mean-only condition. |
| H3 | Participants in the PCS-only condition will have significantly greater pre-post differences in the perceived outgroup homogeneity measure (i.e., they will display a significant reduction in perceived outgroup homogeneity post-intervention) than participants in the mean-only condition. |
| H4 | Participants in the PCS with visualisation condition will have significantly greater pre-post differences in the perceived outgroup homogeneity measure (i.e., they will display a significant reduction in perceived outgroup homogeneity post-intervention) than participants in the group mean with visualisation condition. |
| H5 | Participants in the PCS with visualisation condition will see a significantly greater pre-post reduction in the perceived ideological “distance” between average group members measure than participants in the group mean with visualisation condition. |
| H6 | Participants in the PCS with visualisation condition will see a significantly greater pre-post reduction in the perceived ideological “distance” between average group members measure than participants in the control condition. |
| H7 | The pre-post change in the perceived group extremism measure will be significantly greater (i.e., a greater reduction) for participants in the PCS-only condition compared to those in the control condition. |
| H8 _A | Time will be a significant predictor of change in the weekly pre-data exposure measure of perceived overall ideological polarisation for participants in the PCS-only condition but will not be a significant predictor of change in the weekly pre-data-exposure measure of perceived overall ideological polarisation for participants in the mean-only condition. |
| H8 _B | Time will be a significant predictor of change in the weekly pre-data exposure measure of affective polarisation for participants in the PCS-only condition but will not be a significant predictor of change in the weekly pre-data-exposure measure of affective polarisation for participants in the mean-only condition. |
| H8 _C | Time will be a significant predictor of change in the weekly pre-data exposure measure of perceived issue-specific polarisation for participants in the PCS-only condition but will not be a significant predictor of change for participants in the mean-only condition. |
| H9 | Main effects of summary statistic type will be observed on pre-post changes of anti-democratic attitudes. |
| H10 | The summary statistic type will be a significant predictor of willingness to engage in civil cross-party discussions on compromise. |
| H11 | Change in the perceived personal difference to average opposing-party member will be a significant predictor of change in cross-party generosity in the Dictator Game (even after controlling for demographic variables). |

Results


Contrary to our hypotheses, significant differences failed to emerge between experimental conditions on a number of outcome measures. For example, t-tests determined that there was no significant differences between participants in the mean-only condition and participants in the PCS-only condition on pre-post changes in perceived overall ideological polarisation [H1] ($t(162) = 0.187, p = .852, d = 0.03$), pre-post changes in perceived intergroup agreement [H2] ($t(137.06) = 0.0871, p = .931, d = 0.01$), or pre-post changes in perceived outgroup homogeneity [H3] ($t(161.61) = 0.965, p = .336, d = 0.15$). Similarly, t-tests determined that there were no significant differences between participants in the mean with visualisation condition and participants in the PCS with visualisation condition on pre-post changes in perceived outgroup homogeneity [H4] ($t(158.85) = -1.7083, p = .090, d = 0.27$) or pre-post changes in the perceived ideological distance of

average partisan voters [H5] ($t(160.37) = 1.1372, p = .257, d = 0.18$). However, significant differences were detected between participants in the PCS with visualisation condition and participants in the control condition on pre-post changes in the perceived ideological distance of average partisan voters) [H6] ($t(155.77) = 5.1416, p < .001, d = 0.81$), but t-tests determined that no significant differences emerged between participants in the PCS-only condition and participants in the control condition on pre-post changes of perceived outgroup extremism [H7] ($t(155.93) = 1.0199, p = .309, d = 0.16$).

In examining the impact of time on the effects of summary statistic, exploratory single-predictor linear regression models determined that time was a significant predictor of change in the weekly pre-data exposure measure of perceived overall ideological polarisation for participants in the PCS-only condition ($F(1, 590) = 46.84, \beta = -2.369, p < .001, \text{Adj. } R^2 = .072$) as well as for participants in the mean-only condition ($F(1, 718) = 12.54, \beta = -1.144, p < .001, \text{Adj. } R^2 = .016$; see Figure 5). Time remained a significant predictor of perceived overall ideological polarisation for both participants in the PCS-only condition ($t = -6.960, \beta = -2.404, p < .001$) and participants in the mean-only condition ($t = -3.539, \beta = -1.122, p < .001$) even after demographic explanatory variables (i.e., age, sex, education, and party) were added into the regression models [H8_A].

Figure 5

Regression Lines for Pre-Exposure Measures of Overall Perceived Ideological Polarisation by Week for Participants in the Mean-Only and PCS-Only Conditions



Exploratory single-predictor linear regression models determined that time was not a significant predictor of change in the weekly pre-data exposure measures of general affective polarisation for participants in the PCS-only condition ($F(1, 590) = 0.355, \beta = -0.377, p = .552, \text{Adj. } R^2 = .000$) or for participants in the mean-only condition ($F(1, 718) = 0.171, \beta = -0.207, p = .680, \text{Adj. } R^2 = .000$). Time remained a non-significant

predictor of general affective polarisation for both participants in the PCS-only condition ($t = -0.629$, $\beta = -0.385$, $p = .530$) and participants in the mean-only condition ($t = -0.343$, $\beta = -0.167$, $p = .731$) after demographic explanatory variables (i.e., age, sex, education, and party) were added into the regression models [H8_B].

Additionally, exploratory single-predictor linear regression models determined that time was a significant predictor of change in the weekly pre-data exposure measures of perceived issue-specific polarisation for participants in the PCS-only condition ($F(1, 590) = 30.95$, $\beta = -2.210$, $p < .001$, Adj. $R^2 = .048$) and participants in the mean-only condition ($F(1, 718) = 5.22$, $\beta = -0.877$, $p = .023$, Adj. $R^2 = .006$). Time remained a significant predictor of change in the weekly pre-post changes of perceived issue-specific polarisation for participants in the PCS-only condition ($t = -5.572$, $\beta = -2.235$, $p < .001$) and participants in the mean-only condition ($t = -2.216$, $\beta = -0.857$, $p < .027$) even when demographic explanatory variables (i.e., age, sex, education, and party) were included within the regression model [H8_C].

Pairwise t-tests conducted on pre (i.e., Week 1) and post (i.e., Week 8) changes in perceptions of overall ideological polarisation for participants in the control condition ($M_{CHANGE} = 4.01$, $SD = 20.57$) versus those in the six experimental conditions determined that exposure to each of the experimental conditions yielded a significantly greater decrease (post-Bonferroni corrections to alpha thresholds) in perceptions of polarisation relative to a no summarisation baseline (see Table 3).⁵⁵

Table 3

Results of T-Tests between the Control Condition and Each Experimental Condition on Pre-Post Change in Perceived Overall Ideological Polarisation

| Summary Condition | M_{CHANGE} (Week 8 – Week 1) | SD | t -Statistic (df) | p -Value |
|---------------------------|-----------------------------------|-------|------------------------|------------|
| Control | 4.01 | 20.57 | | |
| Mean-Only | -13.20 | 23.48 | 5.129 (172) | <.001*** |
| Mean with Visualisation | -12.05 | 30.02 | 4.028 (144.9) | <.001*** |
| PCS-Only | -13.86 | 21.69 | 5.314 (156) | <.001*** |
| PCS with Visualisation | -16.15 | 22.78 | 5.954 (162) | <.001*** |
| Both (No Visualisation) | -16.70 | 24.79 | 6.012 (175) | <.001*** |
| Both (with Visualisation) | -17.94 | 24.86 | 6.323 (172) | <.001*** |

* $p < .008$, ** $p < .002$, *** $p < .0002$

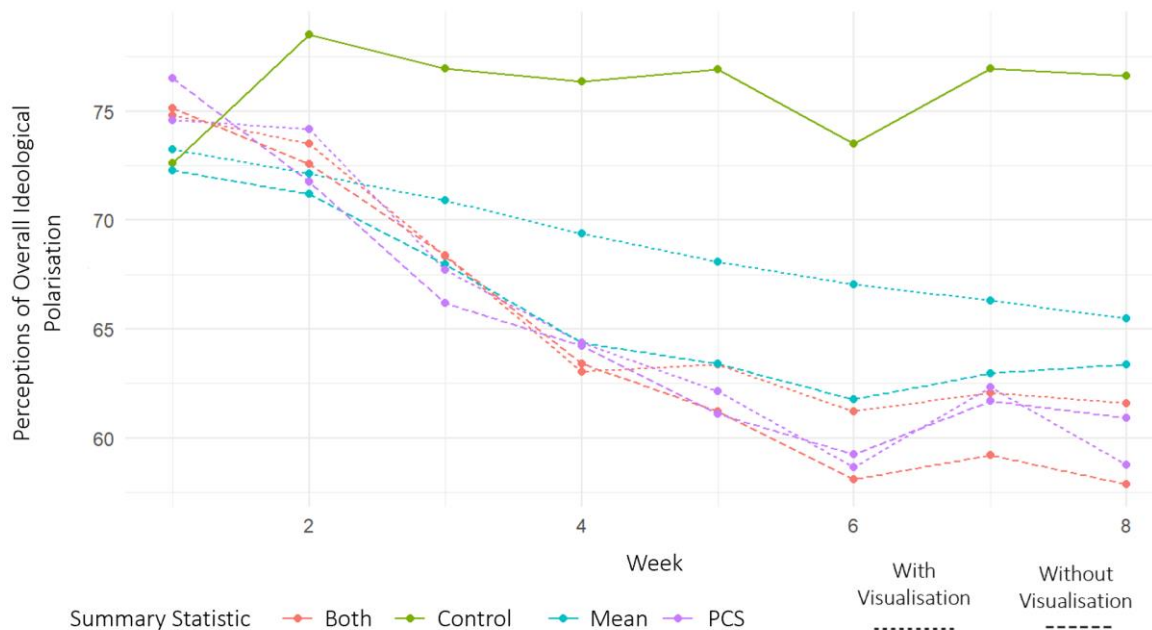
Note. Bonferroni corrections have been applied to significance thresholds.

⁵⁵ See Appendix K for results of t-tests between each pair of experimental conditions on this measure.

However, an exploratory ANOVA on the pre- and post-changes in perceptions of overall ideological polarisation across the six experimental summarisation conditions revealed no significant main effect of summary statistic used ($F(2, 504) = 1.573, p = .208, \eta_p^2 = 0.006$), no significant main effect of visualisation inclusion ($F(1, 504) = 0.115, p = .734, \eta_p^2 = 0.000$), nor a significant interaction between the two variables ($F(2, 504) = 0.210, p = .811, \eta_p^2 = 0.001$). Figure 6 depicts the weekly of trend of measures of perceived overall ideological polarisation (pre-data exposure) by condition.

Figure 6⁵⁶

Weekly Pre-Data Exposure Measures of Perceived Overall Ideological Polarisation by Condition



A linear mixed-effects model was conducted to examine the effects of time, summarisation format (excluding the control condition), and visualisation inclusion on weekly pre-exposure perceptions of overall ideological polarisation. The model included fixed effects for time, summarisation format, and visualisation inclusion, as well as their interactions, with random intercepts for participants. The “both” condition, wherein participants were exposed to both mean and overlap data, served as the reference group. The model determined was a significant main effect of time, $b = -2.03, SE = 0.22, t(3564) = -9.35, p < .001$,

⁵⁶ Please note that for Figure 6, Figure 7, and Figure 9, we’ve opted to truncate the y-axis as a means of permitting readers the opportunity to more meaningfully examine the changes between inter-summary type and intra-summary type. We acknowledge that, whenever possible, it’s likely prudent to adhere to a more conservative approach to truncation (i.e., avoiding it unless a compelling reason to do so exists). This is because, as observed in Chapter 1 and Chapter 2, truncation can visually accentuate differences, leading viewers to *overestimate* differences, ultimately skewing the accuracy of perceptions. However, we also contend that *not* truncating in certain circumstances carries its own risks (see *General Discussion*). For instance, although Figure 6 above, due to its degree of y-axis truncation, risks overemphasising the differences between the control condition and the experimental condition, an alternative graphic which utilised a full-range y-axis (i.e., 0-100) resulted in a visualisation where these differences, despite being highly-significant, were virtually imperceptible. Consequently, in certain circumstances, we contend that *not* truncating the y-axis may, conversely, result in an *underestimation* of condition differences. Ultimately, while we are staunchly against most instances of y-axis truncation when depicting political rivals from a polarised political ecosystem (see *General Discussion*), we recognise that the choice to truncate in alternate circumstances may be defensible as long as a communicator understands (and, ideally, discusses) the trade-offs and accompanying risks.

indicating that polarisation scores decreased by approximately 2.03 points per week across all conditions. The main effects of summarisation format and visualisation inclusion were not significant. Specifically, the mean-only format did not differ significantly from the both condition, $b = -0.92$, $SE = 2.70$, $t(917.10) = -0.34$, $p = .734$, nor did the PCS format, $b = 0.87$, $SE = 2.85$, $t(917.10) = 0.31$, $p = .759$. Including a visualisation did not significantly affect polarisation scores compared to no visualisation, $b = 1.16$, $SE = 2.70$, $t(917.10) = 0.43$, $p = .669$. There was a significant interaction between time and the mean-only summarisation format, $b = 0.89$, $SE = 0.31$, $t(3564) = 2.86$, $p = .004$, suggesting that participants in the mean-only condition experienced a less pronounced decline in polarisation over time compared to those in the both condition. However, the interaction between time and the PCS condition was not significant, $b = -0.34$, $SE = 0.33$, $t(3564) = -1.04$, $p = .299$. A marginally significant interaction was observed between time and visualisation inclusion, $b = -0.59$, $SE = 0.31$, $t(3564) = -1.92$, $p = .055$, indicating that including a visualisation led to a slightly steeper weekly decline in polarisation scores compared to no visualisation.

Table 4 provides a breakdown of weekly pre- and post-perceptions of overall ideological polarisation (along with a measure of intra-weekly change) by summary statistic.

Table 4

Weekly Measures of Pre- and Post-Perceptions of Overall Ideological Polarisation by Summary Statistic

| Summary Statistic | Measure | Week | | | | | | | |
|--------------------|---------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|---------------------------|
| | | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| | | <i>M</i> (<i>SD</i>) | <i>M</i> (<i>SD</i>) | <i>M</i> (<i>SD</i>) | <i>M</i> (<i>SD</i>) | <i>M</i> (<i>SD</i>) | <i>M</i> (<i>SD</i>) | <i>M</i> (<i>SD</i>) | <i>M</i> (<i>SD</i>) |
| <i>Control</i> | Pre | 72.60 (21.54) | 78.52 (19.07) | 76.95 (17.25) | 76.35 (20.13) | 76.93 (16.72) | 73.49 (21.75) | 76.95 (17.59) | 76.61 (17.43) |
| | Post | - | - | - | - | - | - | - | - |
| | Change | - | - | - | - | - | - | - | - |
| <i>Mean</i> | Pre | 72.76 (21.84) | 71.67 (20.80) | 69.51 (19.41) | 66.98 (19.93) | 65.83 (18.13) | 64.53 (20.68) | 64.69 (19.63) | 64.48 (19.39) |
| | Post | 70.08 (20.13) | 66.62 (20.54) | 64.14 (18.08) | 61.29 (21.77) | 60.45 (20.07) | 63.73 (19.52) | 58.86 (20.14) | 60.12 (19.92) |
| | Change | -2.68 | -5.05 | -5.37 | -5.69 | -5.38 | -0.80 | -5.83 | -4.36 |
| <i>Overlap/PCS</i> | Pre | 75.58 (19.45) | 72.90 (19.39) | 66.84 (19.61) | 64.29 (18.49) | 61.62 (19.38) | 58.97 (22.23) | 61.99 (19.91) | 59.90 (19.54) |
| | Post | 71.14 (17.63) | 64.05 (19.82) | 61.41 (19.21) | 58.01 (23.24) | 55.31 (24.06) | 61.54 (18.84) | 55.04 (20.01) | 60.53 (17.80) |
| | Change | -4.44 | -8.85 | -5.43 | -6.28 | -6.31 | 2.57 | -6.95 | 0.63 |
| <i>Both</i> | Pre | 74.96 (19.30) | 73.03 (19.99) | 68.36 (20.70) | 63.23 (20.26) | 62.32 (20.82) | 59.69 (21.58) | 60.68 (20.77) | 59.77 (20.14) |
| | Post | 69.35 (19.03) | 62.21 (21.95) | 60.41 (18.42) | 56.56 (22.14) | 52.84 (22.81) | 59.35 (20.21) | 52.66 (21.34) | 57.64 (19.42) |
| | Change | -5.61 | -10.82 | -7.95 | -6.67 | -9.48 | -0.34 | -8.02 | -2.13 |

Pairwise t-tests conducted on pre (i.e., Week 1) and post (i.e., Week 8) changes in perceptions of intergroup agreement between average partisan voters among participants assigned to the control condition ($M_{CHANGE} = -0.15$, $SD = 2.07$) versus those in the six experimental conditions determined that exposure to each of the experimental conditions yielded a significantly greater increase (post-Bonferroni corrections to alpha thresholds) in perceptions of intergroup agreement after eight weeks relative to the no summarisation baseline (see Table 5).⁵⁷

Table 5

Results of T-Tests between the Control Condition and Each Experimental Condition on Pre-Post Change in Perceived Intergroup Agreement

| Summary Condition | M_{CHANGE} (Week 8 – Week 1) | SD | t-Statistic (df) | p-Value |
|---------------------------|-----------------------------------|------|---------------------|----------|
| Control | -0.15 | 1.07 | | |
| Mean-Only | 0.56 | 0.95 | -4.639 (172) | <.001*** |
| Mean with Visualisation | 0.43 | 1.17 | -3.392 (165) | <.001** |
| PCS-Only | 0.54 | 1.21 | -3.839 (156) | <.001*** |
| PCS with Visualisation | 0.84 | 1.14 | -5.747 (162) | <.001*** |
| Both (No Visualisation) | 0.43 | 1.26 | -3.306 (175) | <.001** |
| Both (with Visualisation) | 0.81 | 1.15 | -5.724 (172) | <.001*** |

* $p < .008$, ** $p < .002$, *** $p < .0002$


Note. Bonferroni corrections have been applied to significance thresholds.

However, once again, and exploratory ANOVA on the pre- and post-changes in perceptions of intergroup agreement between average partisan voters across the six experimental conditions revealed no significant main effect of summary statistic used ($F(2, 504) = 1.237$, $p = .291$, $\eta_p^2 = 0.005$), no significant main effect of visualisation inclusion ($F(1, 504) = 3.302$, $p = .070$, $\eta_p^2 = 0.007$), and no significant interaction between the two variables ($F(2, 504) = 2.384$, $p = .093$, $\eta_p^2 = 0.009$). Figure 7 depicts the weekly of trend of measures of perceived intergroup agreement between average partisan voters (pre-data exposure) by condition.

⁵⁷ See Appendix K for results of t-tests between each pair of experimental conditions on this measure.

Figure 7

Weekly Pre-Data Exposure Measures of Perceived Intergroup Agreement by Condition



Pairwise t-tests conducted on the pre (i.e., Week 1) and post (i.e., Week 8) changes in perceived ideological distance of average partisan voters revealed a significant difference (post-Bonferroni corrections to alpha thresholds) between participants assigned to the control condition ($M_{CHANGE} = 0.42$, $SD = 1.16$) and participants in each of the six experimental conditions. Specifically, these analyses determined that each of the experimental conditions promoted a decrease in perceived ideological distance of average partisan voters across the eight weeks that was significantly greater than the change catalysed by assignment to the control condition (see Table 6).⁵⁸

⁵⁸ See Appendix K for results of t-tests between each pair of experimental conditions on this measure.

Table 6

Results of T-Tests between the Control Condition and Each Experimental Condition on Pre-Post Change in Perceived Ideological Distance between Average Partisan Voters

| Summary Condition | M_{CHANGE} (Week 8 – Week 1) | SD | t-Statistic (df) | p-Value |
|---------------------------|-----------------------------------|------|---------------------|----------|
| Control | 0.42 | 1.16 | | |
| Mean-Only | -0.13 | 0.91 | 3.450 (157.44) | <.001** |
| Mean with Visualisation | -0.36 | 1.32 | 4.040 (165) | <.001*** |
| PCS-Only | -0.41 | 1.10 | 4.551 (156) | <.001*** |
| PCS with Visualisation | -0.60 | 1.36 | 5.161 (162) | <.001*** |
| Both (No Visualisation) | -0.42 | 1.24 | 4.619 (175) | <.001*** |
| Both (with Visualisation) | -0.51 | 1.37 | 4.803 (172) | <.001*** |

* $p < .008$, ** $p < .002$, *** $p < .0002$


Note. Bonferroni corrections have been applied to significance thresholds.

Yet again, an exploratory ANOVA on the pre- and post-changes in the perceived ideological distance of average partisan voters across the six experimental conditions revealed no significant main effect of summary statistic used ($F(2, 504) = 2.255, p = .106, \eta_p^2 = 0.009$), no significant main effect of visualisation inclusion ($F(1, 504) = 2.420, p = .120, \eta_p^2 = 0.005$), and no significant interaction between the two variables ($F(2, 504) = 0.149, p = .861, \eta_p^2 = 0.001$). A mixed-effects model determined there was a significant main effect of time, $b = 0.0812, SE = 0.0098, t(3564) = 8.27, p < .001$, indicating that perceived agreement increased by approximately 0.08 points per week across all conditions. The main effects of summarisation format and visualisation inclusion were not significant. Specifically, the mean-only format did not differ significantly from the both condition, $b = -0.0548, SE = 0.1376, t(798.20) = -0.40, p = .690$, nor did the PCS format, $b = 0.0624, SE = 0.1450, t(798.20) = 0.43, p = .667$. Including a visualisation did not significantly affect perceived agreement compared to no visualisation, $b = -0.0354, SE = 0.1376, t(798.20) = -0.26, p = .797$. However, there was a significant interaction between time and visualisation inclusion, $b = 0.0444, SE = 0.0140, t(3564) = 3.17, p = .0015$, suggesting that participants in the visualisation condition experienced a greater increase in perceived agreement over time compared to those without visualisations. The interaction between time and the mean-only summarisation format was not significant, $b = -0.0075, SE = 0.0140, t(3564) = -0.54, p = .593$, nor was the interaction between time and the PCS format, $b = 0.0102, SE = 0.0147, t(3564) = 0.69, p = .488$.

Figure 8 depicts the regression lines for the pre-post differences in this outcome variable by summary statistic.

Figure 8

Regression Lines for Perceived Ideological Distance between Average Partisan Voters by Summary Statistic




A two-way ANOVA across all experimental conditions found no significant main effect of summary statistic used ($F(2, 504) = 1.883, p = .153, \eta_p^2 = 0.007$), visualisation inclusion ($F(1, 504) = 0.143, p = .705, \eta_p^2 = 0.000$), or interaction ($F(2, 504) = 0.218, p = .804, \eta_p^2 = 0.001$) on pre-post changes in anti-democratic attitudes [H9]. An exploratory two-way ANOVA across all experimental conditions also found no significant main effect of summary statistic used ($F(2, 504) = 0.870, p = .420, \eta_p^2 = 0.003$), visualisation inclusion ($F(1, 504) = 2.153, p = .143, \eta_p^2 = 0.004$), or interaction ($F(2, 504) = 0.897, p = .408, \eta_p^2 = 0.004$) on pre-post changes perceived outgroup extremism, nor were significant main effects of summary statistic used ($F(2, 504) = 0.841, p = .432, \eta_p^2 = 0.003$), visualisation inclusion ($F(1, 504) = 1.725, p = .190, \eta_p^2 = 0.003$), or interaction ($F(2, 504) = 2.134, p = .119, \eta_p^2 = 0.008$) observed for pre-post measures of perceived outgroup homogeneity.

Finally, in examining the impact of condition on our behavioural and pseudo-behavioural dependent variables, an ANOVA examining willingness to engage in civil cross-party conversations found no significant interaction effects between one's response at Week 1 and Week 8 and the summarisation condition to which they had been exposed [H10]⁵⁹ ($F(3, 1180) = 0.393, p = .758, \eta_p^2 = 0.001$). An exploratory ANOVA on pre-post changes in perceived difference between oneself and the average out-party member observed a

⁵⁹ This analysis was originally supposed to be conducted via a logistic regression, but post data collection it was decided that a two-way ANOVA examining interaction effects would be more straightforward approach.

significant main effect of summary statistic used ($F(3, 587) = 2.803, p = .039, \eta_p^2 = 0.014$; see Figure 9). However, linear regressions determined that change in perceived personal difference to average opposing-party member was not a significant predictor of change in cross-party generosity in the Dictator Game [H11] ($F(1, 592) = 0.000, \beta = -0.000, p = .984, \text{Adj. } R^2 = -0.002$).

Figure 9
Week 1 to Week 8 Change in Perceived Difference between Oneself and Average Out-party Member by Summary Statistic



Discussion

The current investigation provides mixed evidence concerning the impact of repeated exposure to different summarisation techniques (as well as summarisation and visualisation combinations) on political perceptions and behaviour.

Firstly, it's important to acknowledge that a majority of null hypotheses failed to be rejected, suggesting that different summary statistics do not appear to yield disparate impacts on many political perception and behavioural variables. Across an eight-week period, main effects of summary statistic were not observed for items such as perceived outgroup homogeneity, perceived outgroup extremism, or anti-democratic attitudes. Additionally, different summary statistics failed to promote significantly different changes (relative to one another as well as relative to control conditions) on proxies of partisan behaviour such as generosity in a cross-party Dictator game or willingness to engage in civil political conversations on compromise with opposing groups. Such findings suggest that, at least within this particular timeframe and at this rate of exposure, choosing between different summary statistics and summary statistic and visualisation

combinations may be inconsequential in terms of promoting a material improvement in certain intergroup perceptions or behavioural dynamics. Moreover, none of the analyses conducted found results to support the contention that the inclusion of visual depictions in conjunction with summary statistics yields differential impacts on political items versus simply disseminating the relevant summary data. Even for certain metrics where main effects of summary statistics *were* observed (e.g., overall perceptions of ideological polarisation; discussed momentarily), main effects of visualisation inclusion failed to materialise.

The second finding of interest concerned the impact of experimental conditions relative to the control condition. While main effects of summary statistic were rarely observed when analyses were limited to the six experimental conditions (implying no statistically-significant differences across the different summary types), when the control condition *was* included, main effects often emerged. This seems to be due to the control condition consistently exhibiting significant differences relative to the six experimental conditions, with exposure to the latter set almost always promoting change in the *desired* direction (i.e., reducing perceptions of overall ideological polarisation, perceived intergroup agreement, and perceived ideological distance) relative to the control. Consequently, when it comes to reducing perceptions of political polarisation and intergroup division, it seems that providing audiences with *any* of the six summary statistics and visualisation combinations featured in the study is superior to permitting partisans to conjecture as to the state of interparty relations based on nothing but their pre-existing private knowledge.

We believe the implications of this finding are quite important. Devoid of such results, it would be fair to wonder whether consistently providing individuals with data detailing where partisans stand on highly-divisive issues might *exacerbate* perceptions of polarisation. For instance, it would not be unreasonable to hypothesise that weekly exposure to data summarising areas of such profound ideological conflict (e.g., abortion, immigration, gun control, etc.; many of which were featured in the study) might fuel an *acceleration* of perceived division. However, the research herein seems to demonstrate that, for several perception variables measured, essentially *any* combination of data summarisation type and visualisation yields preferable results (i.e., results which indicate reduced perceptions of intergroup difference) to allowing individuals to form their intergroup perceptions and opinions in the absence of such data.

Thirdly, a large portion of hypotheses sought to determine whether significant differences would emerge between the two individual summary statistics of interest (e.g., mean-only and PCS-only summarisations of data). These summarisation techniques were meant to roughly represent the two broad types of summary statistics (i.e., measures of central tendency and measures of data dispersion, respectively), and a number of hypotheses postulated that repeated exposure to mean-only data would lead to significant differences in perceptions such as overall perceived ideological polarisation, perceived intergroup agreement, and perceived outgroup homogeneity relative to participant exposed to the PCS-only data. However, t-tests

repeatedly indicated no significant differences between the two conditions, implying that, for these metrics utilising the present methodology, choosing between mean-only exposure and PCS-only exposure doesn't seem to promote heterogeneous perceptions.

While providing a modest contribution to our understanding of how repeated exposure to particular types of summary statistics may impact political perception, the current study (and any future research which seeks to build off of it) could be improved in a number of ways. First and foremost, we did not investigate the impact that *competing frames* may have on the overall patterns of results. As mentioned, while many studies may find evidence for framing effects, fewer consider whether such effects would persist if studied in a context which permitted the presence of competing frames (e.g., Lecheler & de Vreese, 2016). While biased information ecosystems, including partisan “echo chambers” and “filter bubbles” (e.g., Flaxman, Goel, & Rao, 2016) may at times limit exposure to competing viewpoints, the wider political sphere is populated by a seemingly infinite number of perspectives and frames, meaning that individuals are likely to come into contact with frames which directly oppose one another. Consequently, research should aim to understand how effects – if observed in isolation – may persist or dissipate in the face of more realistic political communication dynamics. Secondly, the current study failed to measure if (and, if so, how long) any of the observed framing effects persisted. Findings on the duration of framing effects – both of the “single shot” and repeated exposure variety – have been inconsistent (see Lecheler & de Vreese, 2013), with minimal consensus in the literature as to how resilient one can expect a frame to be and what drives rates of decay for the observed effects. Consequently, the integration of a “follow-up” period several days and/or weeks after final exposure into the experimental design would represent a useful addition to future iterations of the study. Finally, the study failed to account for the pre-existing media consumption habits of participants, leaving us unable to control for heterogeneity in such habits across conditions or investigate how these habits may drive overall patterns of results. Of particular note, it would be interesting for future research to determine A) how media consumption habits – including type of media (e.g., radio, television, Internet, etc.), lean of programming (i.e., liberal versus conservative), and frequency of exposure (e.g., daily, weekly, etc.) – may interact with specific exposure conditions and B) how heterogeneity in such habits may shift the “baseline” levels provided by participants in the control group.

Overall, the current study advanced our understanding of how the use of different summarisation techniques – conceptualised as a form of emphasis framing – impacts political perceptions and behaviour. While a large portion of our results suggested that different summarisation formats do not differ significantly from one another in the influence they exert on political perceptions and behaviour, we also observed evidence to suggest that – across perception items like overall ideological polarisation, intergroup agreement, and ideological distance – exposure to *some form* of data summarisation is often superior to no exposure at all. Such a finding is critical as it implies that, although one might intuitively assume that

repeated exposure to the position of opposing political groups on divisive issues might serve to further *exacerbate* polarisation, it appears that such exposure – at least when conducted in the form of group mean data, PCS data, or a combination of the two – may actually *reduce* certain forms of polarised perception.

SHIFTING FOCUS

FROM IDENTIFYING CHALLENGES TO INVESTIGATING POTENTIAL SOLUTIONS

Thus far, the research contained herein has detailed how one's choice to either visualise or summarise data in particular manners may shift perceptions of polarisation. Concerningly, it has determined that certain modes of presentation can significantly skew perceptions, both relative to alternative modes (despite depicting identical underlying data) as well as relative to objective benchmarks of accuracy. We might refer to this as utilising “questionable presentation practices” (similar to “questionable research practices”; Xie, Wang, & Kong, 2021). A natural next question becomes: if certain means of depicting and disseminating data might be inferior to others, how might we defend against the use of questionable presentation practices?

Broadly, there are two targets researchers might consider when designing an intervention. For several decades, most researchers within psychology and behavioural science have seemed to prioritise intervention strategies in which the impetus to change was placed squarely on the *individual*. For example, the work of Thaler and Sunstein (2021) largely attempts to address the myriad types of “inferior” decision-making propensities exhibited by humans (relative to their uber-rational, utility-maximising counterparts which comprise the bedrock of classical economic theory, sometimes referred to as “econs”; see Thaler, 2015). While such an approach is not without merit (e.g., interventions such as “nudges” have proven capable of consistently moderating both choice and action in domains like health and consumer behaviour; e.g., Mertens, Herberz, Hahnel, & Brosch, 2022), recent contributions to the literature have questioned both the efficacy as well as the downstream consequences of these individual-focused approaches. For example, Chater and Loewenstein (2023) contend that while interventions designed with the intention of altering individual behaviour may be useful, they are often an insufficient means by which to bring about meaningful, widespread change on some of our most pressing socioeconomic and geopolitical issues (e.g., climate change). Moreover, they argue that by consistently targeting *individuals*, practitioners implicitly endorse the narrative that the root of societal problems lies in *individual* faults or weaknesses – a narrative that is perhaps uncoincidentally also forwarded by corporations whose interests are not served by the prospect of large-scale social reforms.

Consequently, scholars have begun to advocate for changemakers to make a more concerted effort to complement “I-frame” interventions (i.e., those targeting the individual) with similar proposals for “S-frame” interventions (i.e., those targeting the systems in which individuals operate), with some arguing for strategies which combine the two approaches (e.g., Newell et al., 2023). In a similar vein, we have attempted to explore two potential means of defending against questionable presentation practices: one built upon S-

frame logic, the other employing an I-frame approach. Chapter 4 explores an S-frame intervention, which aims to understand if one's perception of media sources can be influenced by the manner in which these sources choose to present data. If this turns out to be true, and perceptions of source characteristics (such as trust and credibility) can be moderated by the data presentation choices media companies make, it could motivate institutions within "the system" to be more discerning (and, ideally, more responsible) with how they visualise and/or summarise data. Chapter 5 explores the I-frame intervention, which borrows heavily from the literature on inoculation theory, where we aim to determine the viability of an intervention which "protects" the individual against a particularly widespread questionable presentation practice (i.e., y-axis truncation) – an especially important line of defence should the former, institutionally-focused intervention fail to motivate wider systemic change.

CHAPTER 4

EXPLORING A SYSTEMICALLY-FOCUSED INTERVENTION: ELUCIDATING THE IMPACT OF DATA PRESENTATION CHOICES ON MEDIA REPUTATION AND TRUST⁶⁰

Abstract

When reporting political results (e.g., from a recent poll, survey, etc.), outlets have a number of ways in which they can choose to either visualise or summarise the data. While prior research has investigated how disparate visualisation formats and phenomena analogous to employing different summarisation techniques (such as framing effects) can asymmetrically impact audience perception of and responses to identical data, there is little evidence to suggest how using certain data presentation techniques (as opposed to available alternatives) may impact the perceptions of sources themselves. In a pair of methodologically-identical studies (N = 1,180), we provide evidence to suggest that while different data presentation approaches may not exert differential impacts on engagement with source content, they can exert significant effects on audience perception of source trust and credibility, level of appreciation for source content, and intent to return to the source in the future.

Several decades of research have attempted to disentangle the precise relationship between engagement with media (particularly partisan media sources) and political polarisation (e.g., Campbell, 1980; Zaller, 1992). The extant literature is notably devoid of consensus, however, both concerning *whether* partisan news media contributes to polarisation within the American public, and – if one is to accept that it can – the mechanisms by which the process of polarisation occurs.

In regard to the former, while there exists a large cohort of researchers who insist that exposure to partisan media can lead to polarisation (discussed momentarily), there are some within the field who question this assertion. For example, recently, Wojcieszak et al. (2023) analysed a year's worth of web browsing data, encompassing nearly 40 million cumulative website visits, and found virtually no evidence to suggest that exposure to news websites – partisan or centrist in leaning, congenial or cross-cutting in nature – impacted levels of polarisation. Similarly, Guess, Barberá, Munzert, and Yang (2021) found that while manipulating the

⁶⁰ University of Cambridge ethics identification number 8381.228

propensity of exposure to either left- or right-leaning online news sources (via the alteration of browser settings and social media following patterns) can cause short-term spikes in *knowledge* of political events, data did not indicate any resultant shifts in political opinions or partisan affect.

However, the majority of researchers within the field seem to believe there exists *some* relationship between partisan media and polarisation, but disagree about the underlying dynamics of this relationship (see de Benedictis-Kessner, Baum, & Berinsky, 2019). More specifically, competing cohorts within the community debate about whether political polarisation is *caused* by partisan news exposure or whether the divergent agendas of partisan news sources are merely a *reflection* of the pre-existing composition of political biases and preferences already harboured by the general public. Proponents of the former explanation suggest that partisan news sources can actively persuade audiences to become more extreme in their views (e.g., Levendusky, 2013). For example, Druckman, Levendusky, and McLain (2018) found that exposure to both pro- and counter-attitudinal partisan media can lead to increases in political extremity. Conversely, defenders of the alternative perspective argue that it is not partisan media that is actively driving polarisation but rather that an already polarised public creates a *natural demand* for partisan media. Proponents of this view suggest that individuals engage in self-sorting into ideologically-aligned camps based on their pre-existing political dispositions (e.g., Arceneaux, Johnson, & Murphy, 2012). In fact, some researchers have suggested that expanding landscape of partisan news options may have produced “symptoms” of greater polarisation while not actually increasing the levels of polarisation as they simply caused partisan viewers to grow more politically involved while leading to a comparative reduction in the share of more moderate yet less politically-interested individuals participating in political dialogue (Prior, 2013). While this type of selective exposure may not exacerbate polarisation per se, it may still precipitate deleterious effects on democratic functioning, as research has demonstrated that consistent exposure to just one side of an argument (e.g., via partisan “echo chambers” and “filter bubbles”; e.g., Flaxman, Goel, & Rao, 2016) can make individuals less likely to entertain compromises from opposing parties (Stroud, 2010). Critical to the present research, Levendusky and Malhotra (2016a) find that media coverage which focuses on polarisation narratives (e.g., depicting the electorate as highly-divided) can significantly increase *perceptions* of political polarisation. As polarisation is currently a sociopolitical reality within the United States (and elsewhere), completely eliminating polarisation-focused news coverage is unlikely to be a realistic option. However, media sources interested in reducing their contribution to exaggerated perceptions of polarisation may instead begin to consider *how* they can discuss and depict the state of polarisation whilst minimising the undesirable downstream consequences of such coverage. Here, data presentation has the potential to play a pivotal role, especially as research is beginning to show that the inclusion of accurate data visualisations in news stories may be a promising way to counteract misperceptions (e.g., Mena, 2023).

While our prior research has examined how data presentation affects intergroup political perceptions (e.g., polarisation, likelihood of compromise, etc.), we now shift our focus to understanding whether it might affect perceptions of *the sources that choose to use them*. More specifically, we aim to determine whether a source's choice to employ different visualisation and summarisation methods will exert meaningful impacts on metrics linked to a source's reputation and their audience's engagement with source content – areas which, should data presentation prove influential, might cause sources to reevaluate their data presentation practices.

Thus, the relevant question becomes: which factors, if heterogeneity were to be observed as a product of data presentation choice, might motivate sources to reconsider their data presentation strategies? We believe that four key levers could include A) changes in an audience's tendency to engage with content, B) changes in perceptions of source trustworthiness and credibility, C) changes in audience appreciation of and/or satisfaction with the source's content, and D) changes in audience intentions regarding future engagement with the source.

Beginning with content engagement, the economics of social media place a premium on “engagement rate,” or the amount of interaction (e.g., via reactions, comments, shares, etc.) one's content can generate. Considering its importance, practitioners in disciplines from consumer psychology to communications research have begun developing different ways to categorise and calculate types of engagement (e.g., “CEBs,” or consumer engagement behaviours; Van Doorn et al., 2010). Engagement rate plays such a major role in the online marketplace (e.g., dictating how much advertising companies can demand for their services or how much social media “influencers” can charge for product promotions) that tactics such as “clickbait” (wherein previews of content are provided in a manner that is sensationalised, mysterious, or otherwise misleading in order to encourage user interest and, consequently, “clicks”) have become mainstays of media toolkits (even among world governments; see Lu & Pan, 2021) due to their ability to artificially inflate engagement metrics. More recently, other forms of “engagement bait” have also emerged, including “rage bait,” wherein actors will deliberately post content intended to agitate and elicit negative responses, thus fostering higher engagement rates. As Elie Wiesel once wrote, “the opposite of love is not hate, it's indifference.” On the internet, when it comes to remaining economically viable, a similar logic applies: negative reactions are not the opposite of positive reactions, no reactions are. Consequently, should media sources find that different means of data presentation are adversely impacting their rates of engagement, such a realisation is likely to spur corrective action.

Secondly, changes in perceived source trust and credibility precipitated by data presentation choices may motivate a reconsideration of presentation practices. Granted, not all sources would be equally motivated by altered perceptions along these domains. For example, certain partisan news sources, which sustain their

viewership or readership not due to their objectivity but rather due to the ideological alignment between their reporting and their audience's worldview, may be less inclined to alter their behaviour based on the prospect that they might be viewed as biased, polarising, or untrustworthy. In fact, for many of these outlets, it is arguably their biased reporting which *facilitates* audience acquisition and their divisive narratives which *sustain* audience attention and bolster retention (e.g., Mullainathan & Shleifer, 2005). And though perceived trustworthiness may be an important factor to many of them, it is a decidedly specific *type* of trust that they seek to promote: the trust of those on “their side” of the political aisle. Distrust from those on the “other side,” rather than being a detriment to the viability of their business, is irrelevant at best and a badge of honour at worst.

But trust and credibility have time and again been implicated as critical factors in the reputation and influence of news sources (e.g., Baum & Gussin, 2008; Stroud & Lee, 2013). Moreover, not all sources operate under the (arguably perverse) incentive systems elaborated above. Some news outlets strive to be a bias-free source of information; aiming to maintain both objectivity and political neutrality. For these outlets, metrics like public trust carry immense importance. For example, an article from Oxford's Reuters Institute for the Study of Journalism spoke to the importance of maintaining public trust among media outlets. The authors suggest that sources like the BBC derive much of their value from the fact that audiences perceive them to be a trustworthy source of information (Nielsen, Schulz, & Fletcher, 2020), noting:

“...the UK's independent communications regulator Ofcom has found that large majorities of audiences in the UK value public service media providers like the BBC very highly for providing trustworthy news programmes that help people understand what is going on in the world.”

For sources like the BBC, avoiding reporting practices which might compromise their perceived trustworthiness or credibility is a top-priority. As recent research has suggested that audiences seem more inclined to engage in critical manners with data visualisations perceived as misleading (i.e., providing corrections and critiques; Lisnic, Lex, & Kogan, 2024), such practices are likely the type that sources concerned with reputation preservation will want to avoid.

Thirdly, changes in audience appreciation of or satisfaction with the source and its content could motivate change. Within a media landscape saturated with options, sources are subject to similar competitive pressures as found in other consumer markets. Consequently, should audiences find themselves dissatisfied with the “product being put out” by certain outlets, market forces might compel them to consider alternative “similar products” (i.e., other news outlets; e.g., Vettehen, Zhou, Kleemans, d'Haenens, & Lin, 2012) – a fact that would likely motivate sources to highly-value their ability to keep their audiences satisfied with their content and coverage.

Finally, should audiences come to significantly shift their intentions to return to the source in the future as a result of the manner in which the source has chosen to present data, such a result would naturally motivate sources to critically evaluate their reporting guidelines. While we believe that most readers will find audience retention to be self-evidently important objective for news outlets, in expounding on the rapidly-evolving media ecosystem, researchers have confirmed that “While the main function of the media used to be informational, now the task of attracting and retaining readers in the rapidly changing digital environment is of primary importance” (Hradziushka, Argylov, & Babuk, 2024, p. 63).

To examine these relationships, we devised a pair of methodologically-identical studies: one exploring the impact of different visualisation techniques on the aforementioned areas of interest, the other exploring the impact of different summarisation styles. Each study sought to determine A) whether different modes of data presentation might asymmetrically impact audience engagement with media content or perceptions of the sources who choose to disseminate that content in a particular manner, and B) if (and if so, how) exposure to alternative modes of data presentation may affect these perceptions.

Broadly, we have two overarching sets of hypotheses, one for the visualisation study and one for the summarisation study. For the visualisation study, due to the results obtained throughout Chapter 1 and Chapter 2 (i.e., regarding how different visualisations affect perceptions *relative to one another*), we have confidence that, at least in the explicit alternative choice revelation condition, sources that opt to use full-range bar charts and icon array histograms will exhibit more “positive” changes to source perceptions and future anticipated behaviour (e.g., greater perceptions of source trust, higher levels of audience appreciation, etc.) compared to participants exposed to sources choosing to use truncated bar charts. In the summarisation study, however, we are less confident about how each condition will fare *relative to one another*. Chapter 3 indicated that while many summarisation formats can alter perceptions relative to a no summarisation control, they may not be capable of producing much significant variance in responses relative to one another. However, Chapter 3 never *explained* the pros and cons of each summarisation technique. We believe that certain summarisation techniques – such as percentage of common scores (PCS; e.g., Hanel et al., 2019) – due to their comparative novelty (and thus unfamiliarity) might not be fully appreciated by participants exposed to them with no context. However, should they have their relative advantages and disadvantages explained and be juxtaposed against alternatives, we believe there is potential for substantial pre-post shifts in perception. Consequently, we hypothesise that, in the explicit alternative choice revelation condition, sources choosing to use the PCS summarisation format will experience a significant *positive* pre-post shift in a number of perception and anticipated behaviour metric (e.g., source trust, desire to see more data summarised in this manner, etc.) while sources choosing to use a generic summarisation format will experience significant *negative* pre-post shifts along these measures.

Method

Participants

In both the visualisation and summarisation studies, samples of American citizens over the age of 18 were recruited via the Prolific platform and paid at an approximate rate of 7.50 GBP per hour for their participation. Following exclusions for non-consent, attention check failure⁶¹, or incomplete data, the final sample for the visualisation study was 495 participants and the final sample for the summarisation study was 685⁶² participants. Table 1 summarises the demographic composition of the sample.

Table 1

Sample Composition for the Pair of Perceptions of Media Studies

| Identity Characteristic | Sample Composition | |
|-------------------------|-----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------|
| | <i>Visualisation Study</i> | <i>Summarisation Study</i> |
| Sex | Female: 52.53% ($n = 260$), Male: 44.65% ($n = 221$), Other: 2.83% ($n = 14$) | Female: 58.10% ($n = 398$), Male: 40.15% ($n = 275$), Other: 1.75% ($n = 12$) |
| Age | $M = 43.68$, Range: 18-78 | $M = 42.90$, Range: 18-82 |
| Education | 55.76% ($n = 276$) had attained a bachelor's degree or higher | 56.50% ($n = 387$) had attained a bachelor's degree or higher |
| Political Affiliation | Democrat: 26.06% ($n = 129$), Republican: 24.24% ($n = 120$), Independent/Other: 49.70% ($n = 246$) | Democrat: 34.89% ($n = 239$), Republican: 23.36% ($n = 160$), Independent/Other: 41.75% ($n = 286$) |

⁶¹ Both the visualisation and summarisation study featured two separate attention checks, each of which would result in expulsion from the study if not passed. The first, shown immediately after a screen which described a recent opinion poll, asked participants to describe what the poll was about (e.g., opinions on gun laws, opinions on abortion, etc.). There were four response options. The second was disguised as a typical question (i.e., it began with "Republican voters and Democratic voters sometimes have trouble agreeing...") but the third sentence (of only four total sentences) read "If you're paying attention, select Slightly Disagree below." There were seven total response options. In total, 68.28% of participants in the visualisation study passed both attention checks, and 73.66% of participants in the summarization study passed both attention checks.

⁶² A priori power analyses were conducted to determine the sample size required to achieve sufficient statistical power to detect a medium effect size (i.e., Cohen's f of 0.25) with a power level of 0.90 and an alpha threshold of 0.05 in the visualisation study and a small effect size (i.e., Cohen's f of 0.15) with a power level of 0.90 and an alpha threshold of 0.05 in the summarisation study. The analysis indicated a sample size of approximately 45 participants per group would be required for the former, and approximately 102 per group for the latter. The effect size for the visualisation study was chosen based on the findings from Chapter 1 and Chapter 2, where the effect sizes observed were routinely large. We assumed the probability of a large effect size, but opted to conduct a more conservative power calculation just to be safe. The effect size for the summarization study was chosen primarily due to the effects observed in Chapter 3. While effect sizes between experimental conditions were often quite small in that design, we anticipated them being slightly large in the current study due to the fact that participants would have the opportunity to receive *simultaneous exposure* to each format (with half receiving explicit explanations of their differences), likely producing more pronounced between-condition effects. Due to budget constraints, however, we were unable to achieve the full ideal sample size for the summarization study (i.e., 808). Prior to exclusions, the total sample size for the visualisation study was 725. For those who provided data, the sample composition prior to exclusions had a mean age of 42.28, was comprised of 364 females, 341 males, and 19 other (with 1 non-response), and had 186 self-reported Democrats, 189 self-reported Republicans, and 349 self-reported other (with 1 non-response). Prior to exclusions, the total sample size for the summarisation study was 930. For those who provided data, the sample composition prior to exclusions had a mean age of 42.17, was comprised of 518 females, 387 males, and 19 other (with 6 non-responses), and had 351 self-reported Democrats, 212 self-reported Republicans, and 361 self-reported other (with 6 non-responses).

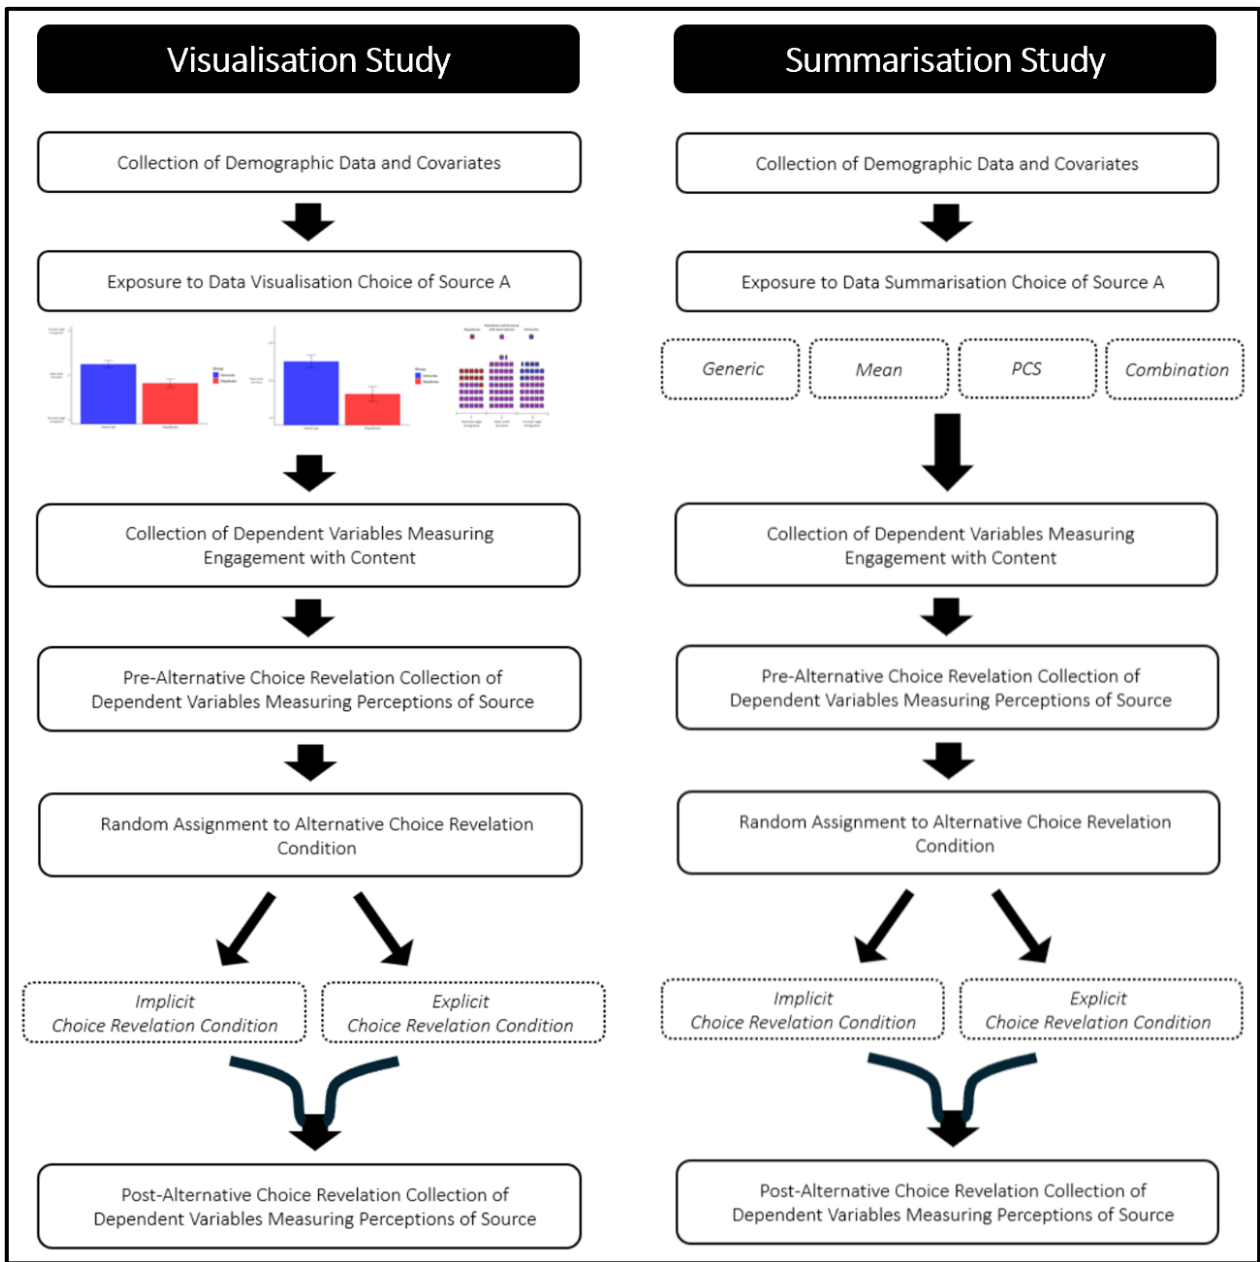
In the visualisation study, among participants who openly identified as Democrat or Republican, 55.04% ($n = 71$) categorised themselves as “strong” Democrats and 45.83% ($n = 55$) categorised themselves as “strong” Republicans. Among those who openly identified as Independent/Other, when forced to choose which of the two major US political parties they considered themselves closest to, 70.33% of this group chose Democrat ($n = 173$) and 29.67% chose Republican ($n = 73$). In the summarisation study, among participants who openly identified as Democrat or Republican, 55.65% ($n = 133$) categorised themselves as “strong” Democrats and 46.88% ($n = 75$) categorised themselves as “strong” Republicans. Among those who openly identified as Independent/Other, when forced to choose which of the two major US political parties they considered themselves closest to, 65.38% of this group chose Democrat ($n = 187$) and 34.62% chose Republican ($n = 99$).

Design

The visualisation study featured a 3x2 (source choice of data visualisation format by alternative choice revelation condition), between-subjects design while the summarisation study featured a 4x2 (source choice of data summarisation format by alternative choice revelation condition). Figure 1 depicts an outline of both study designs:

Figure 1

Design Outlines for the Visualisation Study (left) and the Summarisation Study (right)



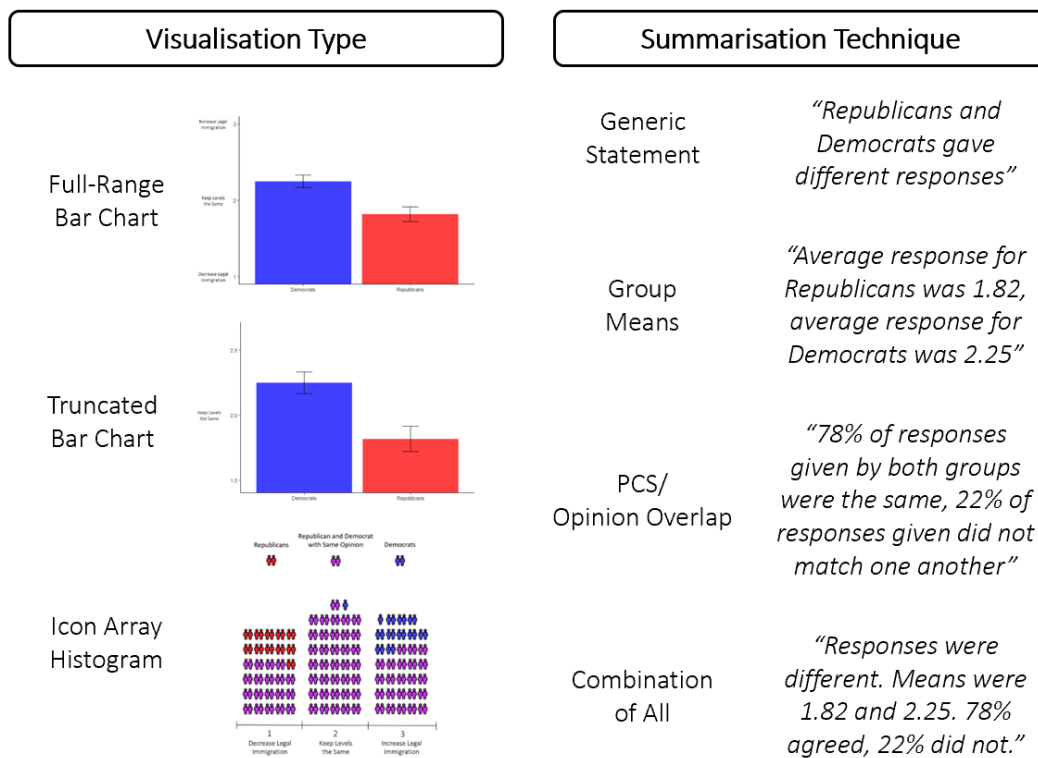
Procedure

The two studies were carried out independently of one another (using separate, non-duplicative participant samples). In the data visualisation line of research, participants were first randomly allocated to see actual data (concerning the views of Republicans and Democrats on the appropriate level of legal immigration, taken from a survey conducted by YouGov between October 16th and 19th, 2023) depicted in one of three visualisation conditions: as a full-range bar chart, as a truncated bar chart, or as an icon array histogram (see Figure 2). In the data summarisation line of research, participants were first randomly allocated to be exposed to data summaries (of the same data set referenced above) which were either 1) generic (e.g., “The

two groups disagreed...”), 2) descriptions of the average responses of the two groups, 3) descriptions of the degree of response overlap/non-overlap (i.e., “percentage of common scores,” or PCS; see Hanel et al., 2019) between the two groups, or 4) a collection of all three summarisation types (see Figure 2).

Figure 2

Condition Stimuli for Perceptions of Media Studies Based on Visualisation Choice (Left) and Summarisation Choice (Right)



Note. The summarisations above are merely abbreviated versions of the condition-specific wording used in the study, meant only to reflect the general spirit of the actual summaries to which participants were exposed.

Participants were informed that a major news outlet (referred to only as “Source A”) recently published an article describing how Republican and Democrat voters responded to a recent survey about legal immigration. They were provided with the original wording of the survey question, and then told “this is how Source A depicted (for the visualisation study)/summarised (for the summarisation study) the results in their article.”

Following initial exposure to the data, participants were asked to complete two sets of dependent variables. The first was designed to assess levels of engagement with the media as-presented, while the second was designed to assess perceptions of, opinions about, and intent to engage with the source depicting or describing the data (see *Measures*). The latter set of dependent variables was employed as a pre-post measure, and thus was also administered at the end of the study (following exposure to presentation alternatives; see below).

After completing the first sets of dependent variables, participants were then randomly allocated into one of two “alternative choice revelation conditions” (where they were made aware of several alternative manners in which the data had been presented by “other sources”). In the implicit alternative choice revelation condition, participants were simply shown the two (for the visualisation study) or three (for the summarisation study) other ways the data had (ostensibly) been reported by different sources (e.g., “*Source B depicted the results of the survey like this [truncated bar chart], and Source C depicted the results of the survey like this [icon array histogram]*”). In the explicit alternative choice revelation condition, as opposed to simply showing participants the other ways in which the data had been presented by other sources, participants were explicitly told that “*Source A chose (italics added for emphasis; not present in final design) to depict/summarise the results of the survey like this.*” They were then shown the alternative ways other sources had depicted the data *and* briefly explained the advantages and disadvantages of each data presentation format. Thus, whereas the implicit condition serves as a proxy for how audiences may respond after simply becoming *aware* of alternative data presentation practices, the explicit condition serves as a proxy for how audiences may respond if they both learned about alternative practices *and* were taught about the relative merits and drawbacks of each approach. For example, participants in the explicit condition of the summarisation study who had originally seen Source A summarise the survey using a generic statement would read the following:

Source A chose not to give any specific data about the responses given by the two groups, simply writing: "Republican voters and Democratic voters gave substantially different responses to the survey." This summarisation technique is succinct but doesn't provide any actual information about how much the two groups agreed or disagreed. But Source A was not the only news media outlet to publish summaries of the immigration survey. Below is how three other major news outlets - Source B, Source C, and Source D - summarised the results in their own articles. Source B chose to focus on the average responses given by Republican voters and Democratic voters, writing: "The average response for Republican voters on a 1 to 3 scale was 1.82 and the average response for Democratic voters on a 1 to 3 scale was 2.25." This summarisation technique is widely-used but some argue it may overly emphasise group differences while ignoring group similarities. Source C chose to focus on the degree of response overlap/similarity between the two groups, writing: "78% of the responses given by Democratic voters and Republican voters were the same, while 22% of responses given by the two groups did not match one another." This summarisation technique is less familiar to some readers but does a good job of describing both the level of disagreement and agreement between the two groups. And Source D chose to include multiple types of summarisations in their article, writing: "Republican voters and Democratic

voters gave substantially different responses to the survey. The average response for Republican voters on a 1 to 3 scale was 1.82 and the average response for Democratic voters on a 1 to 3 scale was 2.25. Additionally, 78% of the responses given by Democratic voters and Republican voters were the same, while 22% of responses given by the two groups did not match one another." This summarisation technique is more time-consuming, but it provides the reader with the greatest amount of detail on the issue.

Following revelation of alternative choices, participants were then asked to once again complete the set of dependent variables focused on how presenting data in this manner had impacted their perceptions of the original media source (i.e., "Source A").

Measures

For assessing levels of engagement with the media, the following items were used:

Opportunity to React: Participants were provided with an opportunity to "like," "dislike," or abstain from reacting to the content.

Opportunity to Comment: Participants were provided with an opportunity to indicate whether or not they would comment on the content (with three response options: "yes," "maybe," and "no"). For those who indicated affirmatively, an actual comment box was provided.

Opportunity to Share: Participants were asked to indicate – using a 0 ("Not at All Interested") to 100 ("Extremely Interested") scale – how interested they would be in sharing the content with their network(s).

Opportunity to View Entire Article (Serving as a Proxy to Measure Click-Through Engagement): Participants were provided with an option to be sent the full article from which the data visualisation/summarisation was (ostensibly) derived. They could either choose to be sent the article, decline being sent the article, or declare that they were unsure. This measure was designed to simulate the "click-through" rate often tracked by online sources.

For assessing perceptions of, opinions about, and intent to engage with the source, the following items were used (in a pre-post manner, collected after initial, isolated exposure to the Source A's presentation of the data and then again following exposure to alternative data presentation formats used by other sources):

Level of Trust in Source: Participants were asked to indicate, on a 0 ("Not at All") to 100 ("Completely") scale, how much they trusted Source A based on the way it visualised/summarised the data.

Perception of Source Credibility: Participants were asked to indicate, on a 0 ("Not at All") to 100 ("Completely") scale, how credible they believed Source A to be based on the way it visualised/summarised the data.

Perception of Source Intent: Based on the way in which the data was visualised/summarised, participants were asked to indicate the perceived intent of Source A on a scale from 0 (“To Divide Completely”) to 100 (“To Unite Completely”), where the midpoint (i.e., 50) was labelled as “To Strictly Inform/Remain Neutral.”

Level of Appreciation for Presentation Choice: Participants were asked to indicate, on a scale from 0 (“Completely Disagree”) to 100 (“Completely Agree”), how much they agreed or disagreed with the following statement: “I appreciate the way Source A has [visualised/summarised] the data for its readers.”

Desire to See More Instances of this Presentation Style: Participants were asked to indicate, on a scale from 0 (“Completely Disagree”) to 100 (“Completely Agree”), how much they agreed or disagreed with the following statement: “In the future, I would like to see more data [visualised/summarised] like this data was.”

Likelihood of Returning to Source: Participants were asked to indicate, on a scale from 0 (“Not at All Likely”) to 100 (“Extremely Likely”), how likely or unlikely they were to return to Source A based on the way it had chosen to visualise/summarise the data.

Hypotheses

Table 2 depicts the pre-registered (OSF URL: <https://osf.io/93pj2>) hypotheses for the study:

Table 2

Pre-Registered Hypotheses for Pair of Perception of Media Studies

| Label | Hypothesis |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Visualisation Study</i> | |
| [VIS] H1 | For participants in the explicit choice condition, sources using the icon array histogram and the full-range bar chart will exhibit significantly greater pre-post increases in perceptions of source trust than sources using the truncated bar chart. |
| [VIS] H2 | For participants in the explicit choice condition, sources using the icon array histogram and the full-range bar chart will exhibit significantly greater pre-post increases in perceptions of source credibility trust than sources using the truncated bar chart. |
| [VIS] H3 | For participants in the explicit choice condition, sources using the icon array histogram and the full-range bar chart will exhibit significantly greater pre-post increases in perceptions of “unifying” intent than sources using the truncated bar chart (which will be perceived as having relatively greater “divisive” intent). |
| [VIS] H4 | For participants in the explicit choice condition, sources using the icon array histogram and the full-range bar chart will exhibit significantly greater pre-post increases in desire to see more similar depictions than sources using the truncated bar chart. |
| [VIS] H5 | For participants in the explicit choice condition, sources using the icon array histogram and the full-range bar chart will exhibit significantly greater pre-post increases in perceptions of audience appreciation than sources using the truncated bar chart. |
| <i>Summarisation Study</i> | |
| [SUM] H1 _A | Participants in the PCS condition who are exposed to the explicit choice revelation condition will exhibit a significant pre-post increase in perceptions of source trust whereas participants in the PCS condition who are exposed to the implicit choice revelation condition will not exhibit significant changes in their perceptions of source trust. |
| [SUM] H1 _B | Participants in the generic description condition who are exposed to the explicit choice revelation condition will exhibit a significant pre-post decrease in perceptions of source trust whereas participants in the generic description condition who are exposed to the implicit choice revelation condition will not exhibit significant changes in their perceptions of source trust. |
| [SUM] H2 _A | Participants in the PCS condition who are exposed to the explicit choice revelation condition will exhibit a significant pre-post increase in level of source appreciation whereas participants in the PCS condition who are exposed to the implicit choice revelation condition will not exhibit significant changes in their levels of source appreciation. |
| [SUM] H2 _B | Participants in the generic description condition who are exposed to the explicit choice revelation condition will exhibit a significant pre-post decrease in level of source appreciation whereas participants in the generic description condition who are exposed to the implicit choice revelation condition will not exhibit significant changes in their levels of source appreciation. |
| [SUM] H3 _A | Participants in the PCS condition who are exposed to the explicit choice revelation condition will exhibit a significant pre-post shift in their perceptions of source intent (towards “unite”) whereas participants in the PCS condition who are exposed to the implicit choice revelation condition will not exhibit significant changes in their perceptions of source intent. |
| [SUM] H3 _B | Participants in the generic description condition who are exposed to the explicit choice revelation condition will exhibit a significant pre-post shift in their perceptions of source intent (towards “divide”) whereas participants in the generic description condition who are exposed to the implicit choice revelation condition will not exhibit significant changes in their perceptions of source intent. |
| [SUM] H4 _A | Participants in the PCS condition who are exposed to the explicit choice revelation condition will exhibit a significant pre-post increase in their desire to see more data summarised in this manner whereas participants in the PCS condition who are exposed to the implicit choice revelation condition will not exhibit significant changes in their levels of desire. |
| [SUM] H4 _B | Participants in the generic description condition who are exposed to the explicit choice revelation condition will exhibit a significant pre-post decrease in their desire to see more data summarised in this manner whereas participants in the generic description condition who are exposed to the implicit choice revelation condition will not exhibit significant changes in their levels of desire. |

Results

Visualisation Study

[H1] A linear mixed-effects model were created for the sub-sample of participants assigned to the explicit alternative revelation condition. Such an analytical approach was chosen for its ability to handle repeated measures (i.e., pre- and post-measures of a variable), account for differences in individual-level variability, and interpret interaction effects (e.g., visualisation and time) without relying on assumptions of sphericity (as repeated ANOVAs do). [H1] The model revealed that, pre-alternative choice revelation, participants who saw the full-range bar chart had a non-significant 1.07-point higher trust score, $b = 1.07$, $SE = 3.58$, $t(355) = 0.30$, $p = .765$, than those who saw the truncated bar chart while those who saw the icon array histogram had a non-significant 2.92-point lower trust score, $b = -2.92$, $SE = 3.57$, $t(355) = -0.82$, $p = .414$, relative to those who saw truncated bar chart. However, when examining pre-post scores, significant interaction effects between time and visualisation type were found. Specifically, participants in the full-range bar chart condition showed a 19.56-point increase in trust from pre- to post-test, $b = 19.56$, $SE = 3.15$, $t(244) = 6.21$, $p < .001$, compared to the truncated bar chart condition. Participants in the icon array histogram condition exhibited a 31.11-point increase in trust from pre- to post-test, $b = 31.11$, $SE = 3.14$, $t(244) = 9.90$, $p < .001$, compared to the truncated bar chart condition. An exploratory mixed-effects model for the implicit alternative revelation sub-sample revealed that, pre-alternative choice revelation, participants who saw the full-range bar chart had a non-significant 1.90-point lower trust score, $b = -1.90$, $SE = 3.67$, $t(304) = -0.52$, $p = .604$, compared to those who saw the truncated bar chart and those who saw the icon array histogram had a non-significant 4.88-point lower trust score, $b = -4.88$, $SE = 3.68$, $t(304) = -1.33$, $p = .185$, compared to those who saw the truncated bar chart. Moreover, when examining pre- and post-data, no significant interaction effect was found between time and visualisation type. Participants in the full-range bar chart condition experienced a 4.01-point increase in trust from pre- to post-test, $b = 4.01$, $SE = 2.41$, $t(245) = 1.66$, $p = .098$, while participants in the icon array histogram condition exhibited a 4.76-point increase in trust from pre- to post-test, $b = 4.76$, $SE = 2.42$, $t(245) = 1.97$, $p = .051$, compared to the truncated bar chart condition. See Table 3 for this item's descriptive statistics.

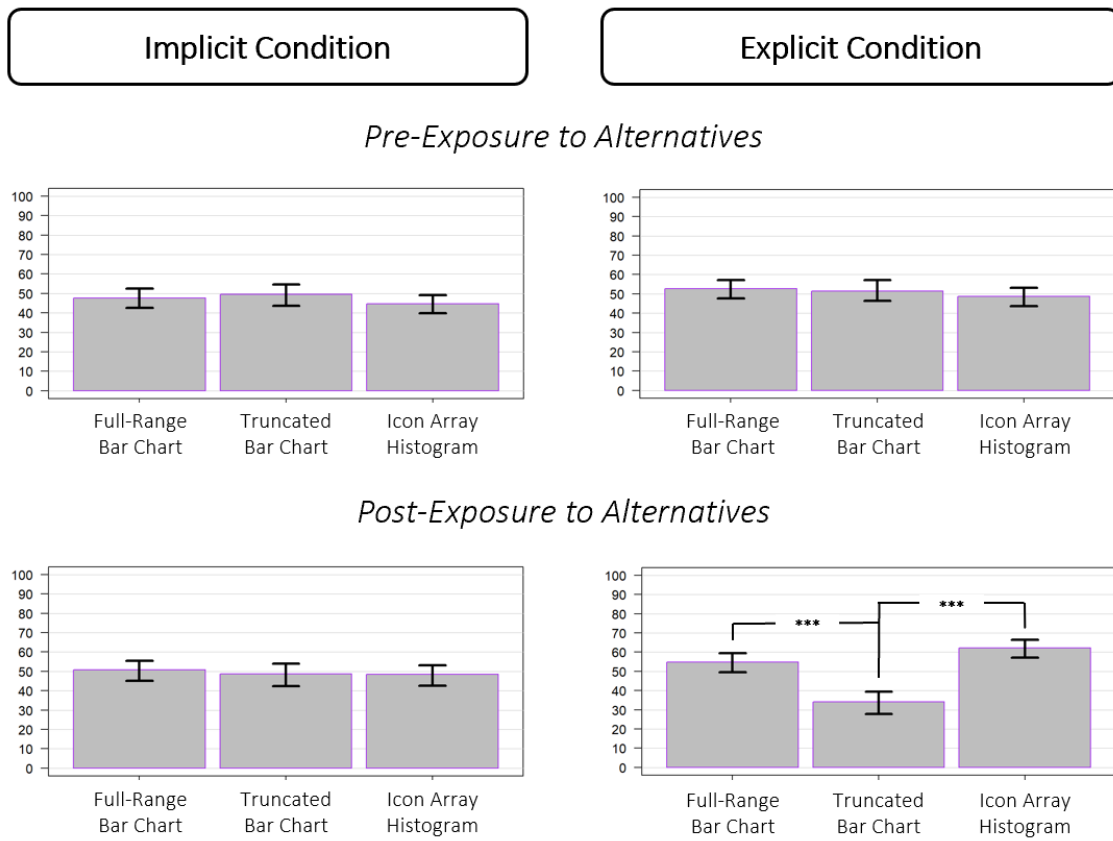
Table 3*Descriptive Statistics for Perceptions of Source Trust*

| Visualisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|---------------------------------|--------------------|--------------------------------|---------------------------------|--------------------|
| | <i>Implicit</i> | | | <i>Explicit</i> | | |
| | <i>Pre-Exposure M (SD)</i> | <i>Post-Exposure M (SD)</i> | <i>Mean Change</i> | <i>Pre-Exposure M (SD)</i> | <i>Post-Exposure M (SD)</i> | <i>Mean Change</i> |
| Truncated Bar Chart | 49.60 (24.64) | 48.67 (24.71) | -0.93 | 51.54 (24.97) | 34.09 (24.45) | -17.45 |
| Full-Range Bar Chart | 47.70 (22.88) | 50.78 (24.42) | +3.08 | 52.61 (21.53) | 54.72 (22.80) | +2.11 |
| Icon Array Histogram | 44.72 (21.02) | 48.55 (23.77) | +3.83 | 48.61 (22.84) | 62.28 (20.75) | +13.67 |

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of visualisation type on perceptions of source trust pre-exposure to alternatives (truncated bar chart ($M = 49.60, SD = 24.64$), full-range bar chart ($M = 47.70, SD = 22.88$), icon array histogram ($M = 44.72, SD = 21.02$); $F(2, 245) = 0.952, p = .387, \eta_p^2 = 0.008$) or post-exposure (truncated bar chart ($M = 48.67, SD = 24.71$), full-range bar chart ($M = 50.78, SD = 24.42$), icon array histogram ($M = 48.55, SD = 23.77$); $F(2, 245) = 0.221, p = .802, \eta_p^2 = 0.002$; Figure 3). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of visualisation type on perceptions of source trust pre-exposure to alternatives (truncated bar chart ($M = 51.54, SD = 24.97$), full-range bar chart ($M = 52.61, SD = 21.53$), icon array histogram ($M = 48.61, SD = 22.84$); $F(2, 244) = 0.659, p = .518, \eta_p^2 = 0.005$) but a significant main effect post-exposure ($F(2, 244) = 33.990, p < .001, \eta_p^2 = 0.218$; Figure 3).

Figure 3

Pre- and Post-Exposure Measures of Perceptions of Source Trust for Each Condition



Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

[H2] A linear mixed-effects model created for participants assigned to the explicit revelation condition revealed that, pre-alternative choice exposure, participants who saw the full-range bar chart had a non-significant 0.68-point higher credibility score, $b = 0.68$, $SE = 3.60$, $t(361) = 0.19$, $p = .850$, than those who saw the truncated bar chart and those who saw the icon array histogram had a non-significant 2.86-point lower credibility score, $b = -2.86$, $SE = 3.59$, $t(361) = -0.80$, $p = .425$, compared to those who saw the truncated bar chart. However, when examining pre- and post-data, significant interaction effects between time and visualisation type were found. Participants in the full-range bar chart condition showed a significant 15.37-point increase in perceived source credibility from pre- to post-test, $b = 15.37$, $SE = 3.24$, $t(244) = 4.74$, $p < .001$, compared to the truncated bar chart condition. Similarly, participants in the icon array histogram condition experienced a significant 26.62-point increase in perceived source credibility from pre- to post-test, $b = 26.62$, $SE = 3.23$, $t(244) = 8.23$, $p < .001$, compared to the truncated bar chart condition. An exploratory linear mixed-effects model for participants in the implicit revelation condition revealed that, pre-alternative choice exposure, participants who saw the full-range bar chart had a non-significant 3.48-point lower credibility score, $b = -3.48$, $SE = 3.46$, $t(324) = -1.01$, $p = .315$, than those who saw the truncated

bar chart and those who saw the icon array histogram had a non-significant 4.93-point lower credibility score, $b = -4.93$, $SE = 3.47$, $t(324) = -1.42$, $p = .156$, compared to those that saw the truncated bar chart. However, when examining pre- and post-data, significant interaction effects between time and visualisation type were found. Participants in the full-range bar chart condition showed a significant 7.28-point increase in credibility from pre- to post-test, $b = 7.28$, $SE = 2.61$, $t(245) = 2.78$, $p = .006$, compared to the truncated bar chart condition. Participants in the icon array histogram condition experienced a significant 6.52-point increase in credibility from pre- to post-test, $b = 6.52$, $SE = 2.62$, $t(245) = 2.49$, $p = .014$, compared to the truncated bar chart condition. See Table 4 for this item’s descriptive statistics.

Table 4

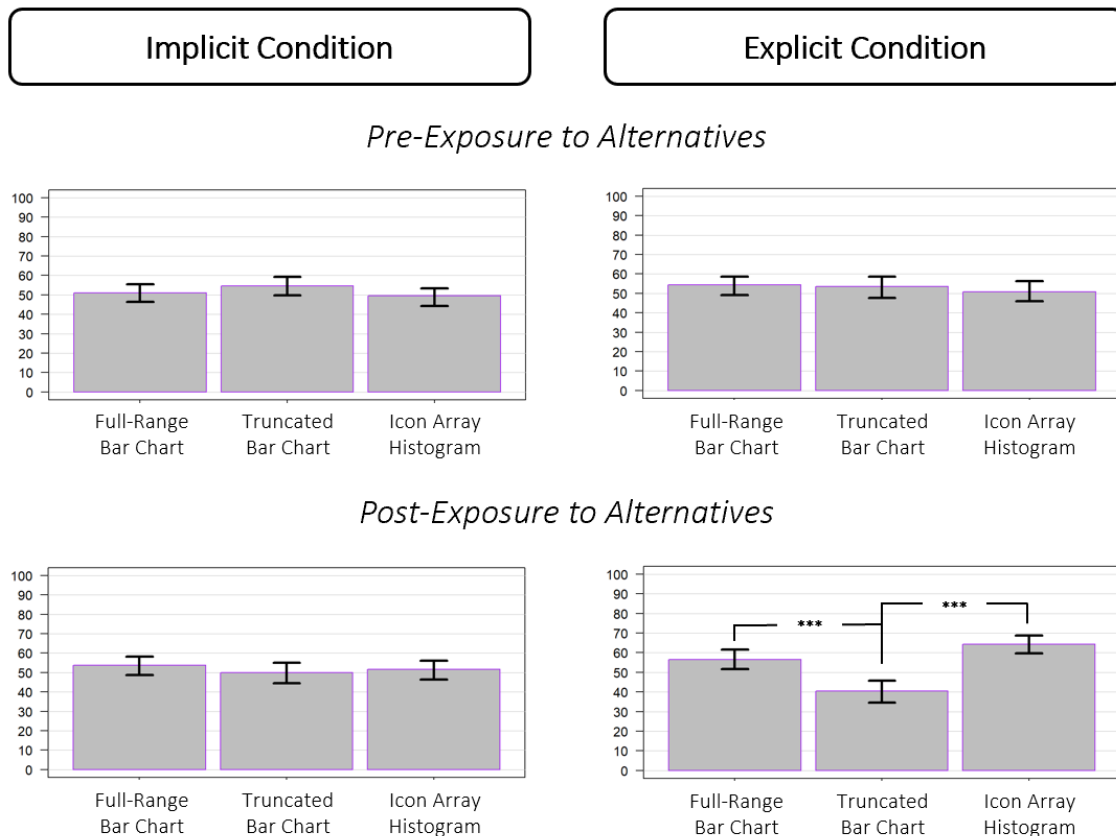
Descriptive Statistics for Perceptions of Source Credibility

| Visualisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|-------------------------|-------------|------------------------|-------------------------|-------------|
| | Implicit | | | Explicit | | |
| | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change |
| Truncated Bar Chart | 54.52 (21.96) | 50.04 (24.00) | -4.48 | 53.63 (24.33) | 40.51 (25.90) | -13.12 |
| Full-Range Bar Chart | 51.04 (21.53) | 53.83 (22.76) | +2.79 | 54.32 (21.61) | 56.56 (22.87) | +2.24 |
| Icon Array Histogram | 49.59 (21.16) | 51.62 (22.28) | +2.03 | 50.77 (22.14) | 64.27 (21.10) | +13.50 |

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of visualisation type on perceptions of source credibility pre-exposure to alternatives (truncated bar chart ($M = 54.52$, $SD = 21.96$), full-range bar chart ($M = 51.04$, $SD = 21.53$), icon array histogram ($M = 49.59$, $SD = 21.16$); $F(2, 245) = 1.143$, $p = .321$, $\eta_p^2 = 0.009$) or post-exposure (truncated bar chart ($M = 50.04$, $SD = 24.00$), full-range bar chart ($M = 53.83$, $SD = 22.76$), icon array histogram ($M = 51.62$, $SD = 22.28$); $F(2, 245) = 0.569$, $p = .567$, $\eta_p^2 = 0.005$; Figure 4). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of visualisation type on perceptions of source credibility pre-exposure to alternatives (truncated bar chart ($M = 53.63$, $SD = 24.33$), full-range bar chart ($M = 54.32$, $SD = 21.61$), icon array histogram ($M = 50.77$, $SD = 22.14$); $F(2, 244) = 0.567$, $p = .568$, $\eta_p^2 = 0.005$) but a significant main effect post-exposure ($F(2, 244) = 22.16$, $p < .001$, $\eta_p^2 = 0.154$; Figure 4).

Figure 4

Pre- and Post-Exposure Measures of Perceptions of Source Credibility for Each Condition



Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

[H3] A linear mixed-effects model created for participants assigned to the explicit revelation condition revealed that, pre-alternative choice exposure, participants who saw the full-range bar chart had a significant 8.43-point higher intent score compared to those who saw the truncated bar chart, $b = 8.43$, $SE = 2.70$, $t(371) = 3.12$, $p = .002$, and participants who saw the icon array histogram had a significant 11.10-point higher intent score compared to the truncated bar chart, $b = 11.10$, $SE = 2.69$, $t(371) = 4.12$, $p < .001$, indicating that participants saw sources that used the full-range bar chart and icon array histogram as having a significantly greater intent to “unite” than sources that used the truncated bar chart. Moreover, when examining pre- and post-data, there were also significant interaction effects between time and visualisation type. Participants in the full-range bar chart condition experienced a significant 9.20-point increase from pre- to post-test compared to the truncated bar chart condition, $b = 9.20$, $SE = 2.53$, $t(244) = 3.64$, $p < .001$. Participants in the icon array histogram condition showed a significant 12.72-point increase from pre- to post-test, $b = 12.72$, $SE = 2.52$, $t(244) = 5.05$, $p < .001$, compared to the truncated bar chart condition. For participants in the implicit revelation condition, an exploratory linear mixed-effects model revealed that, pre-alternative choice exposure, participants who saw the icon array histogram had a significant 5.87-point

higher intent score compared to those who saw the truncated bar chart, $b = 5.87$, $SE = 2.81$, $t(356) = 2.09$, $p = .037$, but there were no significant differences between the full-range bar chart and the truncated bar chart at baseline, $b = -0.19$, $SE = 2.80$, $t(356) = -0.07$, $p = .945$. However, when examining pre- and post-data, there were no significant interaction effects between time and visualisation type. Participants in the full-range bar chart condition experienced a non-significant 4.58-point increase from pre- to post-test compared to the truncated bar chart condition, $b = 4.58$, $SE = 2.45$, $t(245) = 1.87$, $p = .063$, and participants in the icon array histogram condition showed a non-significant 4.11-point increase from pre- to post-test, $b = 4.11$, $SE = 2.46$, $t(245) = 1.67$, $p = .096$. See Table 5 for this item's descriptive statistics.

Table 5

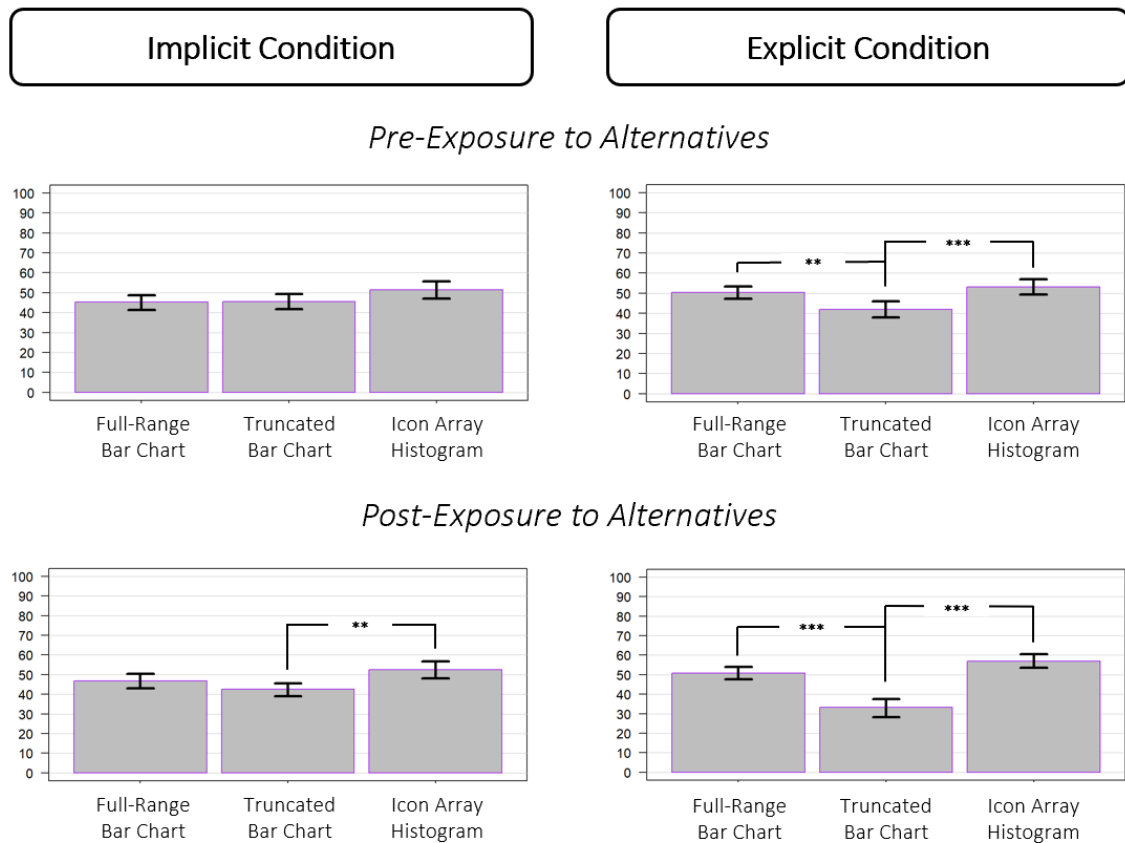
Descriptive Statistics for Perceptions of Source Intent

| Visualisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|-------------------------|-------------|------------------------|-------------------------|-------------|
| | Implicit | | | Explicit | | |
| | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change |
| Truncated Bar Chart | 45.47 (17.59) | 42.48 (16.08) | -2.99 | 41.98 (18.64) | 33.21 (21.92) | -8.77 |
| Full-Range Bar Chart | 45.28 (17.67) | 46.87 (16.97) | +1.59 | 50.40 (14.31) | 50.83 (14.11) | +0.43 |
| Icon Array Histogram | 51.34 (19.63) | 52.46 (19.91) | +1.12 | 53.07 (17.00) | 57.02 (16.48) | +3.95 |

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of visualisation type on perceptions of source intent pre-exposure to alternatives (truncated bar chart ($M = 45.47$, $SD = 17.59$), full-range bar chart ($M = 45.28$, $SD = 17.67$), icon array histogram ($M = 51.34$, $SD = 19.63$); $F(2, 245) = 2.916$, $p = .056$, $\eta_p^2 = 0.023$) but a significant main effect post-exposure (truncated bar chart ($M = 42.48$, $SD = 16.08$), full-range bar chart ($M = 46.87$, $SD = 16.97$), icon array histogram ($M = 52.46$, $SD = 19.91$); $F(2, 245) = 6.571$, $p = .002$, $\eta_p^2 = 0.051$; see Figure 5). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was a significant main effect of visualisation type on perceptions of source intent pre-exposure to alternatives (truncated bar chart ($M = 41.98$, $SD = 18.64$), full-range bar chart ($M = 50.40$, $SD = 14.31$), icon array histogram ($M = 53.07$, $SD = 17.00$); $F(2, 244) = 9.846$, $p < .001$, $\eta_p^2 = 0.075$) as well as a significant main effect post-exposure ($F(2, 244) = 39.670$, $p < .001$, $\eta_p^2 = 0.245$).

Figure 5

Pre- and Post-Exposure Measures of Perceptions of Source Intent for Each Condition



Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

[H4] A linear mixed-effects model created for participants assigned to the explicit revelation condition revealed that, pre-alternative choice exposure, participants who saw the full-range bar chart had a non-significant 3.42-point higher score, $b = 3.42$, $SE = 3.86$, $t(384) = 0.89$, $p = .377$, and those who saw the icon array histogram had a non-significant 3.13-point higher score, $b = 3.13$, $SE = 3.85$, $t(384) = 0.81$, $p = .417$, on the “desire to see more data depicted like this” variable compared to the truncated bar chart. However, when examining pre- and post-data, significant interaction effects were found between time and visualisation type. Participants in the full-range bar chart condition showed a significant 19.94-point increase in their desire to see more visualisations depicted like this from pre- to post-test compared to the truncated bar chart condition, $b = 19.94$, $SE = 3.77$, $t(244) = 5.29$, $p < .001$, and participants in the icon array histogram condition also experienced a significant 32.18-point increase in their desire to see more data visualized in this manner from pre- to post-test, $b = 32.18$, $SE = 3.76$, $t(244) = 8.56$, $p < .001$, compared to the truncated bar chart condition. An exploratory model created for participants in the implicit revelation condition revealed that, pre-alternative choice exposure, participants who saw the full-range bar chart had a non-significant 5.04-point higher score, $b = 5.04$, $SE = 4.02$, $t(354) = 1.25$, $p = .212$, and those who saw the icon

array histogram had a non-significant 0.85-point lower score, $b = -0.85$, $SE = 4.04$, $t(354) = -0.21$, $p = .833$, on desire to see more data depicted in this manner than those who saw the truncated bar chart. However, when examining pre- and post-data, there was a significant interaction between time and visualisation type for the icon array histogram condition. Participants in this condition showed a significant 9.82-point increase in their desire to see more visualisations depicted in this manner from pre- to post-test, $b = 9.82$, $SE = 3.52$, $t(245) = 2.79$, $p = .006$, compared to the truncated bar chart condition. However, there was no significant interaction between time and full-range bar chart, $b = 2.00$, $SE = 3.51$, $t(245) = 0.57$, $p = .569$, compared to the truncated bar chart condition. See Table 6 for this item's descriptive statistics.

Table 6

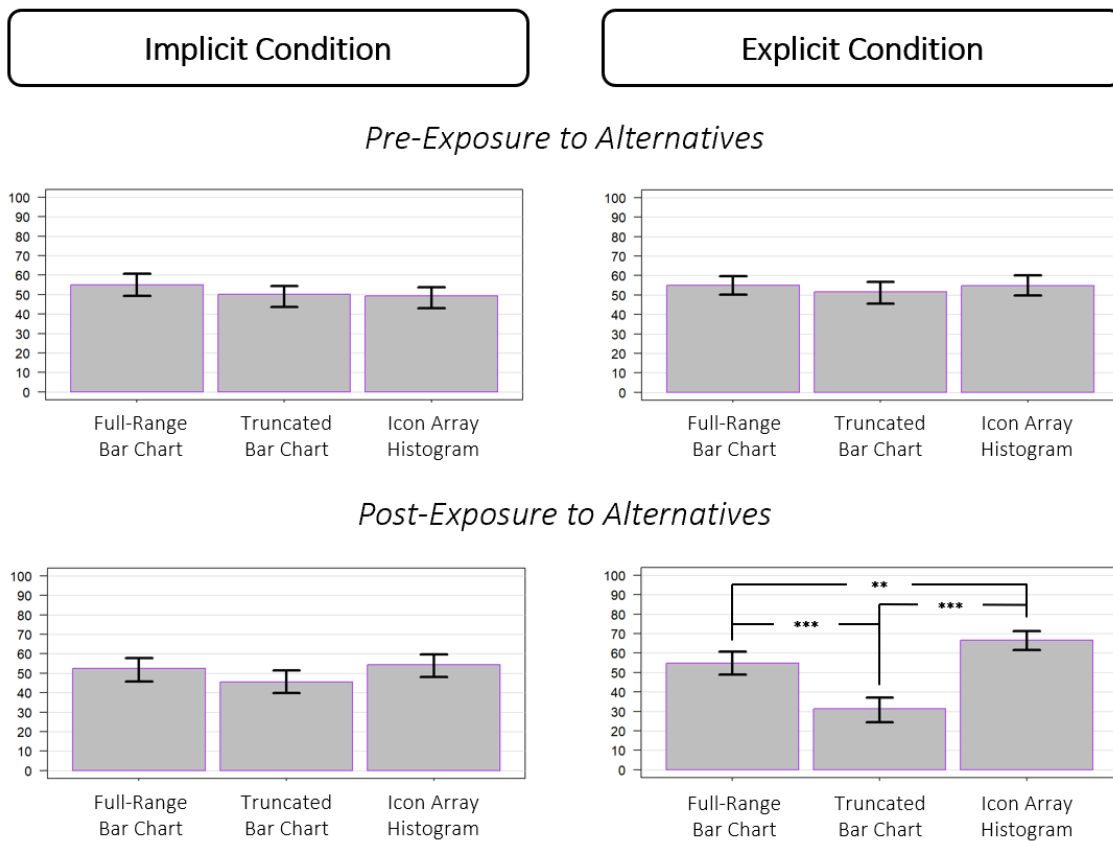
Descriptive Statistics for Desire to See More Data Visualised in this Manner

| Visualisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|-------------------------|-------------|------------------------|-------------------------|-------------|
| | Implicit | | | Explicit | | |
| | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change |
| Truncated Bar Chart | 50.08 (23.19) | 45.48 (26.32) | -4.60 | 51.65 (25.87) | 31.37 (27.89) | -20.28 |
| Full-Range Bar Chart | 55.12 (25.76) | 52.52 (27.82) | -2.60 | 55.06 (21.56) | 54.72 (25.77) | -0.34 |
| Icon Array Histogram | 49.23 (25.01) | 54.45 (27.19) | +5.22 | 54.77 (23.36) | 66.67 (23.22) | +11.90 |

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of visualisation type on desire to see more data depicted in such a manner pre-exposure to alternatives (truncated bar chart ($M = 50.08$, $SD = 23.19$), full-range bar chart ($M = 55.12$, $SD = 25.76$), icon array histogram ($M = 49.23$, $SD = 25.01$); $F(2, 245) = 1.377$, $p = .254$, $\eta_p^2 = 0.011$) or post-exposure (truncated bar chart ($M = 45.48$, $SD = 26.32$), full-range bar chart ($M = 52.52$, $SD = 27.82$), icon array histogram ($M = 54.45$, $SD = 27.19$); $F(2, 245) = 2.506$, $p = .084$, $\eta_p^2 = 0.020$; see Figure 6). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of visualisation type on desire to see more data depicted in such a manner pre-exposure to alternatives (truncated bar chart ($M = 51.65$, $SD = 25.87$), full-range bar chart ($M = 55.06$, $SD = 21.56$), icon array histogram ($M = 54.77$, $SD = 23.36$); $F(2, 244) = 0.526$, $p = .592$, $\eta_p^2 = 0.004$) but a significant main effect post-exposure ($F(2, 244) = 40.260$, $p < .001$, $\eta_p^2 = 0.248$).

Figure 6

Pre- and Post-Exposure Measures of Desire to See More Data Visualised in this Manner for Each Condition



Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

[H5] A linear mixed-effects model created for participants assigned to the explicit revelation condition revealed that, pre-alternative choice exposure, participants who saw the full-range bar chart had a 7.10-point higher appreciation of the source compared to those who saw the truncated bar chart, $b = 7.10$, $SE = 3.90$, $t(387) = 1.82$, $p = .069$. There were no significant baseline differences in levels of appreciation for participants who saw the icon array histogram, $b = -1.41$, $SE = 3.89$, $t(387) = -0.36$, $p = .717$, compared to the truncated bar chart. However, examining pre- and post-data, there were significant interaction effects between time and visualisation type for both conditions compared to the truncated bar chart. Participants in the full-range bar chart condition showed a significant 17.98-point increase in appreciation from pre- to post-test compared to those in the truncated bar chart condition, $b = 17.98$, $SE = 3.85$, $t(244) = 4.67$, $p < .001$, and participants in the icon array histogram condition experienced a significant 35.38-point increase in appreciation from pre- to post-test, $b = 35.38$, $SE = 3.84$, $t(244) = 9.21$, $p < .001$, compared to the truncated bar chart condition. An exploratory model created for participants in the implicit revelation condition revealed that, pre-alternative choice exposure, participants who saw the full-range bar chart had a non-significant 5.36-point higher appreciation score, $b = 5.36$, $SE = 3.92$, $t(335) = 1.37$, $p = .172$, and those who

saw the icon array histogram had a non-significant 1.51-point higher score, $b = 1.51$, $SE = 3.93$, $t(335) = 0.39$, $p = .700$, compared to those who saw the truncated bar chart. However, when examining pre- and post-data, significant interaction effects were found for the icon array histogram. Participants in this condition showed a significant 7.97-point increase in appreciation from pre- to post-test compared to the truncated bar chart condition, $b = 7.97$, $SE = 3.14$, $t(245) = 2.54$, $p = .012$, but no significant interaction effect was found for the full-range bar chart, $b = 4.47$, $SE = 3.13$, $t(245) = 1.43$, $p = .154$. See Table 7 for this item’s descriptive statistics.

Table 7

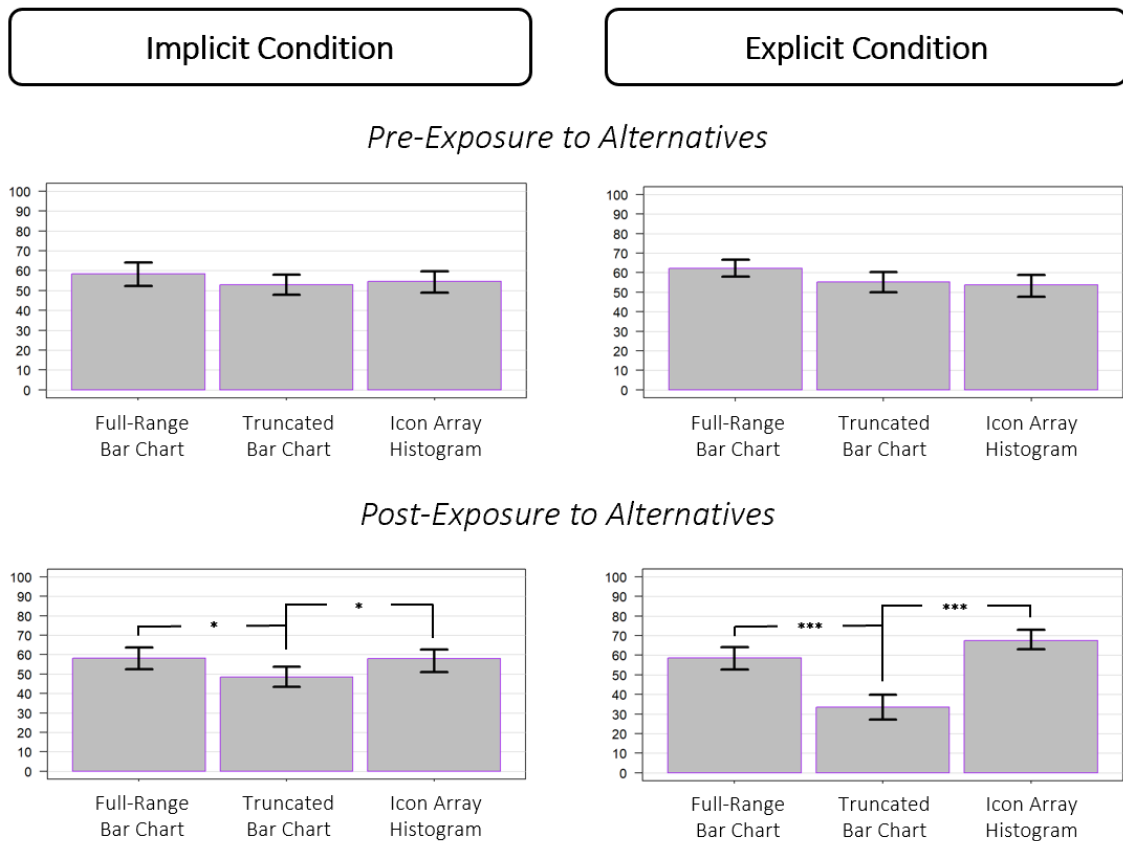
Descriptive Statistics for Level of Audience Appreciation

| Visualisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|--------------------------------|-------------|-------------------------------|--------------------------------|-------------|
| | Implicit | | | Explicit | | |
| | Pre-Exposure <i>M (SD)</i> | Post-Exposure <i>M (SD)</i> | Mean Change | Pre-Exposure <i>M (SD)</i> | Post-Exposure <i>M (SD)</i> | Mean Change |
| Truncated Bar Chart | 53.00 (23.43) | 48.40 (25.10) | -4.60 | 55.20 (25.23) | 33.55 (27.83) | -21.65 |
| Full-Range Bar Chart | 58.36 (26.38) | 58.23 (26.68) | -0.13 | 62.29 (20.09) | 58.62 (25.45) | -3.67 |
| Icon Array Histogram | 54.51 (23.91) | 57.88 (25.59) | +3.37 | 53.78 (25.81) | 67.52 (24.71) | +13.74 |

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of visualisation type on levels of audience appreciation pre-exposure to alternatives (truncated bar chart ($M = 53.00$, $SD = 23.43$), full-range bar chart ($M = 58.36$, $SD = 26.38$), icon array histogram ($M = 54.51$, $SD = 23.97$); $F(2, 245) = 1.045$, $p = .353$, $\eta_p^2 = 0.008$) but a significant main effect post-exposure (truncated bar chart ($M = 48.40$, $SD = 25.10$), full-range bar chart ($M = 58.23$, $SD = 26.68$), icon array histogram ($M = 57.88$, $SD = 25.59$); $F(2, 245) = 3.873$, $p = .022$, $\eta_p^2 = 0.031$; see Figure 7). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of visualisation type on levels of audience appreciation pre-exposure to alternatives (truncated bar chart ($M = 55.20$, $SD = 25.23$), full-range bar chart ($M = 62.29$, $SD = 20.09$), icon array histogram ($M = 53.78$, $SD = 25.81$); $F(2, 244) = 3.006$, $p = .051$, $\eta_p^2 = 0.024$) but there was a significant main effect post-exposure ($F(2, 244) = 37.720$, $p < .001$, $\eta_p^2 = 0.236$).

Figure 7

Pre- and Post-Exposure Measures of Level of Appreciation of Visualisation Format for Each Condition



Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

Finally, exploratory analyses across the four outcome variables related to engagement with the content (e.g., “liking,” sharing, etc.) failed to find any evidence to suggest that differences in visualisation choices yield any significant asymmetries in engagement tendencies. Specifically, a one-way ANOVA on likelihood of sharing the article found no significant main effect of visualisation format ($F(2, 492) = 0.852, p = .427, \eta_p^2 = 0.003$), and chi-square tests of independence (chosen due to the variables being categorical and response variables being discrete) determined no significant association between visualisation format and liking ($\chi^2 = 1.473, df = 4, p = 0.831$), commenting ($\chi^2 = 1.522, df = 2, p = 0.467$), or clickthrough behaviour ($\chi^2 = 8.555, df = 4, p = 0.073$).

Summarisation Study

To test whether participants exposed to the PCS summarisation format exhibited a significant pre-post increase in source trust in the explicit alternative choice revelation condition and exhibited no significant pre-post increase in source trust in the implicit alternative choice revelation condition, a linear mixed-effects model was fitted using R. The model included fixed effects for summarisation format, revelation condition,

and time (i.e., pre and post), along with their interactions. Participant ID was modelled as a random intercept to account for repeated measures. Pairwise comparisons of estimated marginal means (EMMs) were conducted to assess pre-post changes within each PCS condition (i.e., explicit and implicit). Perceptions of source trust significantly increased from pre ($M = 43.8$, $SE = 2.69$, $CI_{95} = [38.5, 49.0]$) to post ($M = 50.6$, $SE = 2.69$, $CI_{95} = [45.3, 55.8]$) in the explicit revelation condition, with an estimated mean difference of 6.81, $SE = 2.35$, $CI_{95} = [11.45, 2.17]$, $t(336) = 2.895$, $p = 0.004$. Perceptions of source trust did not significantly change from pre ($M = 42.1$, $SE = 2.64$, $CI_{95} = [36.9, 47.3]$) to post ($M = 41.1$, $SE = 2.64$, $CI_{95} = [35.9, 46.3]$) in the implicit revelation condition, with an estimated mean difference of -1.02 , $SE = 2.31$, $CI_{95} = [-5.57, 3.52]$, $t(336) = 0.443$, $p = 0.658$. [H1_A] Similar mixed-effects models were created for participants exposed to the generic summarisation format, with pairwise comparisons of estimated marginal means (EMMs) conducted to assess pre-post changes within each generic summary condition (i.e., explicit and implicit). Results indicated that perceptions of source trust significantly decreased from pre ($M = 47.1$, $SE = 2.59$, $CI_{95} = [42.0, 52.2]$) to post ($M = 34.9$, $SE = 2.59$, $CI_{95} = [29.8, 40.0]$) in the explicit revelation condition, with an estimated mean difference of -12.24 , $SE = 2.27$, $CI_{95} = [-16.70, -7.77]$, $t(336) = -5.389$, $p < 0.001$. Perceptions of source trust also significantly decreased from pre ($M = 47.9$, $SE = 2.58$, $CI_{95} = [42.8, 52.9]$) to post ($M = 35.3$, $SE = 2.58$, $CI_{95} = [30.2, 40.4]$) in the implicit revelation condition, with an estimated mean difference of -12.59 , $SE = 2.26$, $CI_{95} = [-17.02, -8.15]$, $t(336) = -5.575$, $p < .001$. [H1_B] See Table 8 for descriptive statistics.

Table 8

Descriptive Statistics for Perceptions of Source Trust


| Summarisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|-------------------------|----------------|------------------------|-------------------------|----------------|
| | Implicit | | | Explicit | | |
| | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change |
| Generic | 47.88 (22.66) | 35.28 (22.35) | -12.60 | 47.14 (23.36) | 34.90 (22.46) | -12.24 |
| Mean-Only | 46.50 (21.30) | 48.53 (24.32) | +2.03 | 49.80 (24.21) | 43.29 (25.03) | -6.51 |
| PCS-Only | 42.10 (27.15) | 41.07 (27.13) | -1.03 | 43.75 (23.53) | 50.57 (24.54) | +6.82 |
| Combination | 49.45 (25.84) | 55.94 (26.33) | +6.49 | 52.02 (24.01) | 59.87 (23.41) | +7.85 |

Exploratory t-tests were performed to confirm the patterns observed above. These analyses confirmed that there were no significant pre-post changes in perceptions of source trust among participants who were exposed to the PCS summarisation format in the implicit alternative choice revelation condition ($M = -1.03$, $SD = 20.78$; $t(166) = 0.244$, $p = .807$, $d = 0.04$), but also suggested that no significant pre-post changes took place in the explicit alternative choice revelation condition ($M = 6.82$, $SD = 17.14$; $t(160) = 1.804$, $p = .073$, $d = 0.28$). Similar t-tests confirmed that there were significant pre-post changes in perceptions of source trust among participants exposed to the generic summarisation format in both the implicit alternative choice

revelation condition ($M = -12.60, SD = 21.36; t(174) = 3.711, p < .001, d = 0.56$) as well as the explicit alternative choice revelation condition ($M = -12.24, SD = 24.54; t(172) = 3.524, p < .001, d = 0.53$; Figure 8).

Figure 8⁶³

Pre-Post Changes in Perceptions of Source Trust for Each Condition



Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .013$, ** $p < .003$, *** $p < .0003$

Similar mixed-effects modelling and pairwise comparisons of estimated marginal means determined that levels of audience appreciation of source did not significantly increase from pre ($M = 48.2, SE = 2.96, CI_{95} = [42.4, 54.0]$) to post ($M = 52.3, SE = 2.96, CI_{95} = [46.5, 58.1]$) in the explicit revelation condition, with an estimated mean difference of 4.15, $SE = 2.68, CI_{95} = [-1.12, 9.42], t(336) = 1.548, p = 0.123$. Contrary to our hypothesis, the analyses suggested that levels of audience appreciation significantly decreased from pre ($M = 51.6, SE = 2.90, CI_{95} = [45.9, 57.3]$) to post ($M = 44.7, SE = 2.90, CI_{95} = [39.0, 50.4]$) in the implicit revelation condition, with an estimated mean difference of $-6.92, SE = 2.63, CI_{95} = [-12.08, -1.76], t(336) = -2.68, p = 0.009$. [H2_A] Similar mixed-effects models were created for participants exposed to the generic summarisation format, with pairwise comparisons of estimated marginal means (EMMs) conducted to assess pre-post changes within each generic summary condition (i.e., explicit and implicit). Results indicated

⁶³ Please note that for all visualisations which depict pre-post changes, we've opted to truncate the y-axis. Due to the fact that the graphics depict change on a 0 to 100 scale, the full range (i.e., which would encompass all possible responses) would extend from -100 (for a participant who changed from a response of 100 to a response of 0) to 100 (for a participant who changed from a response of 0 to a response of 100). A 200-point range on the y-axis dramatically compresses the bar charts, rendering significant and non-significant intra-condition changes nearly indistinguishable from one another and, arguably, making the graphic virtually useless as a tool for comparison and comprehension. For example, there are several instances where a pre-post shift of less than 10 points constitutes a significant change. When depicted via a full-range (i.e., 200 point) y-axis, the bar is compressed so substantially that it looks nearly identical to an adjacent bar depicting non-significant change. Thus, while we acknowledge that the truncation introduces the risk that individuals will come to overestimate the degree of intra-condition change (or inter-condition difference), we believe that, conversely, refusing to truncate introduces the risk that an individual walks away with an equally-skewed impression that pre-post change (or difference in change) was insignificant. Please see *General Discussion* for a slightly more detailed discussion.

that levels of audience appreciation significantly decreased from pre ($M = 48.9$, $SE = 2.85$, $CI_{95} = [43.3, 54.5]$) to post ($M = 32.5$, $SE = 2.85$, $CI_{95} = [26.9, 38.1]$) in the explicit revelation condition, with an estimated mean difference of -16.36 , $SE = 2.59$, $CI_{95} = [-21.46, -11.26]$, $t(336) = -6.325$, $p < 0.001$. Levels of audience appreciation also significantly decreased from pre ($M = 52.0$, $SE = 2.84$, $CI_{95} = [46.4, 57.6]$) to post ($M = 36.1$, $SE = 2.84$, $CI_{95} = [30.5, 41.7]$) in the implicit revelation condition, with an estimated mean difference of -15.88 , $SE = 2.57$, $CI_{95} = [-20.94, -10.82]$, $t(336) = -6.174$, $p < .001$. [H2_B] See Table 9 for descriptive statistics.

Table 9


Descriptive Statistics for Level of Audience Appreciation

| Summarisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|-------------------------|----------------|------------------------|-------------------------|----------------|
| | Implicit | | | Explicit | | |
| | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change |
| Generic | 52.00 (25.65) | 36.13 (25.61) | -15.87 | 48.90 (25.42) | 32.54 (22.37) | -16.36 |
| Mean-Only | 55.42 (23.05) | 52.12 (28.00) | -3.30 | 55.58 (24.95) | 42.55 (26.87) | -13.03 |
| PCS-Only | 51.63 (28.77) | 44.71 (29.47) | -6.92 | 48.19 (28.06) | 52.33 (27.31) | +4.14 |
| Combination | 57.42 (26.81) | 60.87 (27.77) | +3.45 | 55.48 (26.65) | 64.45 (25.27) | +8.97 |

Exploratory t-tests were performed to confirm the patterns observed above. These analyses suggested that there were no significant pre-post changes in levels of audience appreciation of the summarisation format among participants who were exposed to PCS data in the implicit alternative choice revelation condition ($M = -6.92$, $SD = 23.42$; $t(166) = 1.539$, $p = .126$, $d = 0.24$) or the explicit alternative choice revelation condition ($M = 4.14$, $SD = 22.98$; $t(160) = 0.953$, $p = .342$, $d = 0.15$). T-tests also confirmed that there were significant pre-post changes in levels of audience appreciation of the summarisation format among participants who were exposed to generic data in both the implicit alternative choice revelation condition ($M = -15.87$, $SD = 24.43$; $t(174) = 4.109$, $p < .001$, $d = 0.62$) as well as the explicit alternative choice revelation condition ($M = -16.36$, $SD = 25.46$; $t(172) = 4.505$, $p < .001$, $d = 0.68$). Exploratory t-tests also determined statistically-significant pre-post changes (post-Bonferroni corrections to alpha thresholds) in levels of appreciation of summarisation format for sources that chose to summarise data using group means in the explicit alternative choice revelation condition ($M = -13.03$, $SD = 25.43$; $t(176) = 3.353$, $p < .001$, $d = 0.50$; Figure 9).

Figure 9

Pre-Post Changes in Level of Appreciation of Summarisation Format for Each Condition



Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .013$, ** $p < .003$, *** $p < .0003$

Mixed-effects modelling and pairwise comparisons of estimated marginal means determined that perceptions of source intent did not significantly change from pre ($M = 60.4$, $SE = 2.04$, $CI_{95} = [56.3, 64.4]$) to post ($M = 58.8$, $SE = 2.04$, $CI_{95} = [54.8, 62.8]$) in the explicit revelation condition, with an estimated mean difference of -1.54 , $SE = 1.91$, $CI_{95} = [-5.29, 2.21]$, $t(336) = -0.809$, $p = 0.419$. However, the analyses did suggest that perceptions of source intent significantly changed (towards “unite”) from pre ($M = 59.6$, $SE = 2.01$, $CI_{95} = [55.6, 63.5]$) to post ($M = 63.9$, $SE = 2.01$, $CI_{95} = [59.9, 67.8]$) in the implicit revelation condition, with an estimated mean difference of 4.27 , $SE = 1.87$, $CI_{95} = [0.58, 7.96]$, $t(336) = 2.282$, $p = 0.023$. [H3_A] Similar mixed-effects models were created for participants exposed to the generic summarisation format, with pairwise comparisons of estimated marginal means (EMMs) conducted to assess pre-post changes within each generic summary condition (i.e., explicit and implicit). Results indicated that perceptions of source intent did not significantly change from pre ($M = 40.1$, $SE = 1.97$, $CI_{95} = [36.2, 44.0]$) to post ($M = 37.3$, $SE = 1.97$, $CI_{95} = [33.5, 41.2]$) in the explicit revelation condition, with an estimated mean difference of -2.74 , $SE = 1.84$, $CI_{95} = [-6.36, 0.88]$, $t(336) = -1.486$, $p = 0.138$. However, perceptions of source intent significantly decreased from pre ($M = 40.1$, $SE = 1.96$, $CI_{95} = [36.3, 44.0]$) to post ($M = 34.9$, $SE = 1.96$, $CI_{95} = [31.0, 38.7]$)

in the implicit revelation condition, with an estimated mean difference of -5.25 , $SE = 1.83$, $CI_{95} = [-8.84, -1.66]$, $t(336) = -2.869$, $p = .004$. [H3_B] See Table 10 for descriptive statistics.

Table 10


Descriptive Statistics for Perceptions of Source Intent

| Summarisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|-------------------------|----------------|------------------------|-------------------------|----------------|
| | Implicit | | | Explicit | | |
| | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change |
| Generic | 40.15 (15.74) | 34.90 (19.96) | -5.25 | 40.08 (17.41) | 37.34 (18.67) | -2.74 |
| Mean-Only | 47.16 (14.21) | 46.06 (16.70) | -1.10 | 44.02 (16.54) | 39.75 (19.54) | -4.27 |
| PCS-Only | 59.58 (19.19) | 63.86 (18.59) | +4.28 | 60.36 (16.66) | 58.81 (20.44) | -1.55 |
| Combination | 47.40 (19.96) | 47.84 (20.26) | +0.44 | 54.31 (17.82) | 53.53 (16.73) | -0.78 |

Exploratory t-tests determined that there were no significant pre-post changes in perceptions of source intent among participants who were exposed to the PCS summarisation format in either the implicit alternative choice revelation condition ($M = 4.28$, $SD = 15.05$; $t(166) = 1.466$, $p = .145$, $d = 0.23$) or the explicit alternative choice revelation condition ($M = -1.55$, $SD = 19.20$; $t(160) = 0.527$, $p = .599$, $d = 0.08$). Similarly, t-tests also determined that there were no significant pre-post changes in perceptions of source intent among participants who were exposed to the generic summarisation format in the implicit alternative choice revelation condition ($M = -5.25$, $SD = 16.18$; $t(165.04) = 1.938$, $p = .054$, $d = 0.29$) or the explicit alternative choice revelation condition ($M = -2.74$, $SD = 18.03$; $t(172) = 1.000$, $p = .319$, $d = 0.15$; Figure 10).

Figure 10

Pre-Post Changes in Perceptions of Source Intent for Each Condition



Mixed-effects modelling and pairwise comparisons of estimated marginal means determined that desire to see more data summarised in this manner did not significantly change from pre ($M = 46.6$, $SE = 2.98$, $CI_{95} = [40.7, 52.4]$) to post ($M = 48.4$, $SE = 2.98$, $CI_{95} = [42.6, 54.3]$) in the explicit revelation condition, with an

estimated mean difference of 1.88, $SE = 2.70$, $CI_{95} = [-3.43, 7.19]$, $t(336) = 0.694$, $p = 0.488$. However, the analyses did suggest that desire to see more data summarised in this format did decrease from pre ($M = 48.2$, $SE = 2.93$, $CI_{95} = [42.4, 53.9]$) to post ($M = 41.8$, $SE = 2.93$, $CI_{95} = [36.0, 47.5]$) in the implicit revelation condition, with an estimated mean difference of -6.39 , $SE = 2.65$, $CI_{95} = [-11.60, -1.18]$, $t(336) = -2.409$, $p = 0.017$. [H4_A] Similar mixed-effects models were created for participants exposed to the generic summarisation format, with pairwise comparisons of estimated marginal means (EMMs) conducted to assess pre-post changes within each generic summary condition (i.e., explicit and implicit). Results indicated that desire to see more data summarised in this manner significantly reduced from pre ($M = 48.4$, $SE = 2.88$, $CI_{95} = [42.8, 54.1]$) to post ($M = 31.1$, $SE = 2.88$, $CI_{95} = [25.5, 36.8]$) in the explicit revelation condition, with an estimated mean difference of -17.33 , $SE = 2.61$, $CI_{95} = [-22.47, -12.19]$, $t(336) = -6.648$, $p < 0.001$. and they also significantly decreased from pre ($M = 49.4$, $SE = 2.86$, $CI_{95} = [43.7, 55.0]$) to post ($M = 32.2$, $SE = 2.86$, $CI_{95} = [26.6, 37.8]$) in the implicit revelation condition, with an estimated mean difference of -17.19 , $SE = 2.59$, $CI_{95} = [-22.28, -12.08]$, $t(336) = -6.627$, $p < .001$. [H4_B] See Table 11 for descriptive statistics.

Table 11

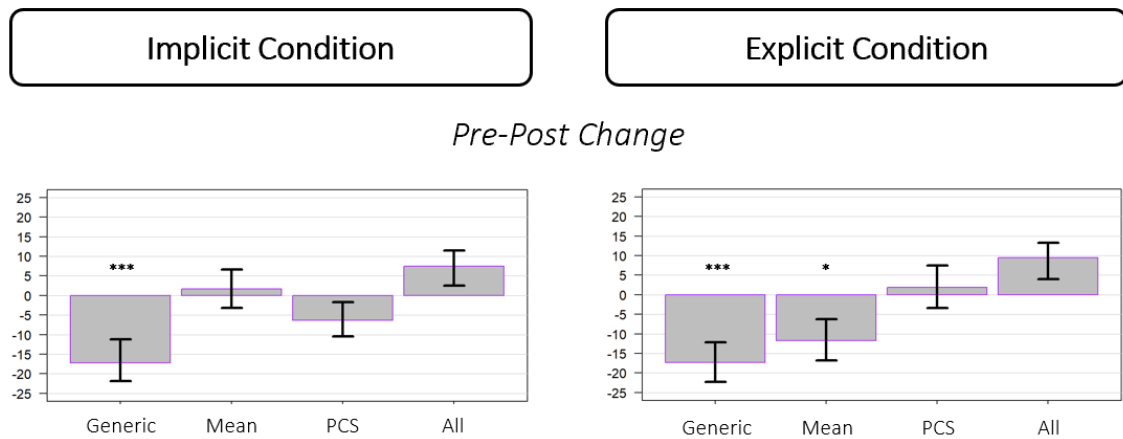
Descriptive Statistics for Desire to See More Data Summarised in this Manner

| Summarisation Condition | Alternative Revelation Condition | | | | | |
|-------------------------|----------------------------------|-------------------------|----------------|------------------------|-------------------------|----------------|
| | Implicit | | | Explicit | | |
| | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change | Pre-Exposure M (SD) | Post-Exposure M (SD) | Mean Change |
| Generic | 49.38 (25.26) | 32.19 (26.35) | -17.19 | 48.45 (25.55) | 31.11 (24.01) | -17.34 |
| Mean-Only | 47.01 (26.01) | 48.66 (28.62) | +1.65 | 50.45 (27.31) | 38.73 (25.42) | -11.72 |
| PCS-Only | 48.17 (28.60) | 41.77 (29.30) | -6.40 | 46.57 (29.34) | 48.44 (26.25) | +1.87 |
| Combination | 51.54 (27.63) | 58.96 (26.07) | +7.42 | 51.86 (27.31) | 61.29 (27.95) | +9.43 |

Exploratory t-tests were unable to confirm significant pre-post changes in desire to see more data summarised in this manner among participants who were exposed to the PCS summarisation format in the implicit alternative choice revelation condition ($M = -6.40$, $SD = 21.11$; $t(166) = 1.431$, $p = .154$, $d = 0.22$), but corroborated that there was no significant change in the explicit alternative choice revelation condition ($M = 1.87$, $SD = 25.09$; $t(160) = 0.429$, $p = .669$, $d = 0.07$). T-tests did confirm that there were significant pre-post changes in desire to see more data summarised in this manner among participants who were exposed to the generic summarisation format in both the implicit alternative choice revelation condition ($M = -17.19$, $SD = 24.38$; $t(174) = 4.416$, $p < .001$, $d = 0.67$) as well as the explicit alternative choice revelation condition ($M = -17.34$, $SD = 26.34$; $t(172) = 4.611$, $p < .001$, $d = 0.70$). Exploratory t-tests also determined statistically-significant pre-post changes (post-Bonferroni corrections to alpha thresholds) in levels of desire to see more data summarised in this manner for sources that chose to summarise data using group means in the explicit alternative choice revelation condition ($M = -11.72$, $SD = 25.57$; $t(176) = 2.963$, $p = .003$, $d = 0.44$; Figure 11).

Figure 11

Pre-Post Changes in Desire to See More Data Summarised in this Manner for Each Condition



Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .013$, ** $p < .003$, *** $p < .0003$

Finally, exploratory analyses of the engagement with content dependent variables (i.e., “liking,” commenting, and sharing) failed to find any evidence to suggest that differences in summarisation format yielded significant asymmetries in engagement tendencies. Specifically, a one-way ANOVA on likelihood of sharing the article found no significant main effect of summarisation format ($F(3, 681) = 2.530, p = .056, \eta_p^2 = 0.011$), and chi-square tests of independence (chosen due to the variables being categorical and response variables being discrete) determined no significant association between visualisation format and liking ($\chi^2 = 10.319, df = 6, p = 0.112$), commenting ($\chi^2 = 3.726, df = 3, p = 0.293$), or clickthrough behaviour ($\chi^2 = 11.525, df = 6, p = 0.073$).

Discussion

Across the two studies, the results suggest that A) mode of presentation (be it visualisation type or summarisation format) fails to play a meaningful role in audience engagement with content (i.e., significant main effects of mode of presentation were not observed, suggesting presentation method is unlikely to alter audience propensities across several traditional engagement metrics) and B) it does exert a significant impact on how audiences come to perceive, think about, and intend to engage with a source. In regards to the former finding, while it would have been ideal to have found that modes which have proven capable of exerting positive impacts on political perception (e.g., greater accuracy, lower perceptions of polarisation,

etc., such as full-range bar charts and icon array histograms) also positively impacted audience engagement, we would argue that null results here are still a net-positive finding. Considering how many “engagement bait” tactics rely on sensationalist claims and anger (e.g., “rage bait”) to drive engagement, results very well could have indicated that formats such as truncated bar charts (which might be conceptualised as “sensationalised” depictions of partisan attitudes relative to full-range bar charts) were significantly better at spurring user engagement, leading to a greater motivation for sources to employ them. Thus, we would argue that the fact that mode (whether it be of visualisation or summarisation) fails to significantly impact engagement with content should encourage sources to turn to the other metrics featured in the study to determine how to present data, many of which *do* suggest heterogeneous participant responses.

The results demonstrated that visualisation format plays a critical role in shaping an audiences’ perception of source trust and credibility. Specifically, participants exposed to the full-range bar chart and icon array histogram under explicit revelation conditions exhibited significant increases in trust, credibility, source appreciation, and desire to see more data depicted in this manner compared to those exposed to the truncated bar chart. However, in implicit revelation conditions, such changes were generally less pronounced, particularly for the full-range bar chart relative to the truncated bar chart, suggesting that explicitly guiding audiences to consider the merits of alternative visualisations amplifies the impact of these formats. The icon array histogram, however, often continued to catalyse significant pre-post improvements in source perceptions relative to the truncated bar chart even in the implicit revelation conditions, suggesting that such a visualisation format might benefit source reputations even if audiences are not explicitly made aware of alternative visualisation modalities.

In the summarisation study, similar patterns emerged. Summarisation formats emphasising overlap between groups, such as PCS (percentage of common scores), significantly enhanced trust of source and appreciation of source under the explicit revelation condition. This effect may stem from the PCS format’s ability to emphasise differences *as well as* commonalities, all while providing an audience with sufficient detail to support perceptions of credibility. Conversely, generic summarisations resulted in significant decreases in trust and appreciation under both revelation conditions. This suggests that vague or overly simplistic descriptions may fail to engage audiences or, worse, reinforce perceptions of bias, apathy, or lack of information. Notably, PCS summaries were associated with perceptions of unifying intent, particularly under implicit revelation, highlighting their potential to mitigate perceptions of a source contributing to widening polarisation. Generic summaries, on the other hand, were consistently associated with reductions in unifying intent, further underscoring their limitations in fostering positive audience evaluations of source.

The results also revealed that audience desire to see more data summarised in a similar manner was significantly influenced by the interaction of summarisation format and revelation condition. PCS formats

generally maintained or slightly increased audience desires to see more data summarised in this manner under explicit revelation conditions, indicating that when audiences are guided to understand the rationale behind such summaries, their receptiveness increases. In contrast, generic summaries consistently led to significant decreases in audience desire across both revelation conditions, which may underscore a broader dissatisfaction with generic summaries while also highlighting the importance of providing audiences with clear, detailed, and meaningful data breakdowns.

These findings contribute to the literature on strategic communication and media perceptions by demonstrating the importance of media *choices* in fostering audience trust, appreciation, and engagement. The differential effects of explicit and implicit revelation further emphasise the importance of audience education and awareness in driving these responses. Explicit revelation amplifies the benefits of effective visualisations and summaries but may also exacerbate the limitations of weaker formats like generic summarisations. These findings suggest that communicators need to carefully consider both the content and the framing of their data presentations to achieve desired audience outcomes.

Practically, these results offer valuable insights for organisations, policymakers, and media practitioners. Visualisations such as icon array histograms and full-range bar charts are likely to foster greater trust and appreciation of sources, particularly when paired with explicit explanations of their benefits. Similarly, PCS summarisations that emphasise group overlap hold promise for not only reducing perceptions of polarisation, but also for fostering positive opinions relating to source trust, appreciation, and intent. Moreover, the risks associated with generic summarisation formats cannot be overstated; such summaries consistently failed to engage audiences and, in many cases, actively undermined perceptions of trust and appreciation. This highlights the need for communicators to adopt clear and detailed data summarisation strategies when presenting information to audiences.

Limitations

While providing some useful insights, the current studies could be critiqued in a number of ways. Arguably the most sensitive design element related to the way in which the relative advantages and disadvantages of visualisation types and summarisation techniques were explained to participants in the explicit alternative revelation conditions. Our priority was to discuss the pros and cons of each presentation format in a way that was accurate and, ideally, neutral in tone. Moreover, with an eye towards neutrality, we sought to standardise the content of the descriptions so that each one featured at least one advantage and at least one disadvantage of the format. And when formats featured more than one advantage, such as the icon array histogram (i.e., “they show the full distribution of each group's responses and provide a more balanced depiction of both group disagreement and group agreement”), efforts were made to counterbalance this

with an equal number of disadvantages (i.e., “[they] are a bit more visually complicated than bar charts and can be harder for some to interpret”).

However, to avoid what we deemed to be the risk of creating false equivalences in the minds of participants (e.g., “it seems that truncated bar charts have just as many pros and cons as full-range bar charts – they must be just as good, just different”), the standardisation procedure was not overly rigid. For example, when it came to describing the full-range bar chart we noted that they are “easy to understand, but critics argue that they might visually accentuate the differences between two groups while failing to highlight group similarities” (i.e., one pro and one con). For the truncated bar chart, however, we included the same pro and con as above but also added that the source chose “to truncate (or shorten) the y-axis on their chart. This has the effect of making the difference between the two groups appear to be greater; making their average responses seem ‘farther apart’ than if they had kept the y-axis its full size.” Thus, this format did not have a perfectly counterbalanced number of advantages and disadvantages. While we believed this trade-off was a prudent one, future researchers may consider altering the manner in which the formats are described to determine whether it yields results which deviate from the patterns documented in the present set of studies.

Additionally, the way in which we collected data about participant engagement with content was very rudimentary (i.e., by simply asking them whether they would like or not like the content, comment or not comment on it, etc. if they saw it on their social media). This means of measuring online behaviour is admittedly contrived and perhaps unlikely to accurately capture the true behavioural proclivities likely to be displayed in a more realistic online environment. Future research should consider devising a more sophisticated means of capturing online engagement, perhaps by embedding the dependent variables (and perhaps the study as a whole) in environments that more faithfully replicate online ecosystems.

Finally, examining cultural and contextual factors, such as pre-existing biases or inter-party attitudes, could shed light on how these strategies interact with individual differences to either reinforce or challenge views of media sources depending on the data presentation strategies they employ.

Conclusion

Overall, the results of the two studies demonstrate that audiences may come to significantly shift their perceptions of, opinions about, and intentions to engage with sources based on the data presentation formats sources choose to employ. While these results appear to be most powerful when individuals are both exposed to alternative data presentation options *and* taught about their relative advantages and disadvantages, the results suggest that, for some metrics, differential audience evaluations of sources can be triggered simply by *seeing* how other sources have chosen to present identical data. Although this set of findings is an encouraging step in dissuading sources from engaging in “questionable presentation practices,”

it would be naïve to believe that news outlets will immediately shift course in light of such results. These entities are often driven primarily by factors such as audience engagement (with the platform as a whole as well as with content on the site) and revenue. While our studies provide preliminary evidence that the former *could* be compromised by questionable data presentation, consequences for the latter remain untested and speculative, leaving immediate large-scale change quite unlikely. A more measured interpretation of the impact of the present results on sources would be that, as audiences become more familiar with the various means of presenting identical data, sources that become more deliberative in how they choose to visualise and summarise data will likely stand to benefit more than their less deliberative peers, who may be penalised (via reduced perceptions of trust, reductions in intentions to engage, etc.) should they present data in unfavourable manners. However, when (and if so, at what rate and to what degree) audiences will develop a greater familiarity with data presentation options remains an open question.

CHAPTER 5

EXPLORING AN INDIVIDUAL-FOCUSED INTERVENTION: TESTING THE PROSPECT OF INOCULATING AGAINST TRUNCATION-INDUCED POLARISATION EFFECTS⁶⁴

Abstract

Psychological inoculation has proven to be one of the most effective interventions for combatting the proliferation of online misinformation, with many interventions informing participants of how polarisation is a common technique wielded by agents of misinformation. The present study sought to determine whether individuals could be inoculated against a specific polarisation tactic: the truncation of the y-axis when depicting partisan data. Our results (n = 394) suggest that a short (i.e., approximately five minute) intervention can render individuals more capable of differentiating between misleading and non-misleading visualisations and ultimately significantly less likely to succumb to exaggerated perceptions of polarisation catalysed by inadvisable truncations of the y-axis (i.e., “truncation-induced polarisation effects”). Moreover, individuals exposed to the intervention also came to rate sources as more or less reliable depending on their use (or lack thereof) of y-axis truncation when depicting intergroup data.

Due to the perverse incentives that characterise pockets of the modern media landscape, when it comes to fostering more responsible data presentation practices, it’s idealistic (and perhaps naïve) to believe that media outlets will voluntarily police themselves, adopting standards which prioritise the public good over maximising engagement and profits. Consequently, it’s critical to investigate the possibility of deploying interventions which directly target *individuals*, seeking to make them less vulnerable to the misleading and manipulative techniques which will likely continue to pervade some mainstream and social media platforms. One of the most promising lines of research suitable to achieving these ends seems to be *psychological inoculation*.

While currently being heavily featured in the misinformation and disinformation literature as a means of combating the spread and impact of online falsehoods (e.g., Lewandowsky & van der Linden, 2021), psychological inoculation was originally developed as a technique designed to confer a more general

⁶⁴ University of Cambridge ethics identification number 8431.231

“resistance to persuasion” (e.g., McGuire, 1961). The technique draws inspiration from the medical field, as evidenced by the title of an early paper (authored by one of the originators of the theory, William J. McGuire) “Vaccine for Brainwash” (McGuire, 1970). In essence, psychological inoculation operates on the same core principles as biological inoculation. However, whereas the latter seeks to artificially induce immunity to specific pathogens, the former seeks to induce resistance to specific persuasive (and often untrue) appeals (McGuire, 1964). Dynamically, the two forms of inoculation operate in similar fashions: just as biological inoculation may expose a host to a weakened strain of a virus, permitting the body the opportunity to successfully “fight it off,” and thereby preparing it to more capably defend itself against a stronger strain, psychological inoculation operates by exposing an individual to a weakened form of a persuasive message, teaching them how to “defend” against it, and thereby “immunising” that individual against more potent forms of similar persuasive tactics in the future.

Broadly, there are two critical components which comprise an inoculation message: threat and refutational pre-emption (sometimes colloquially referred to as “pre-bunking”; Traberger, Roozenbeek, and van der Linden, 2022). Early work on psychological inoculation showed that pre-emptive exposure to (and refutation of) weakened forms of expected counterarguments not only made a belief more resistant to stronger forms of the counterarguments refuted, but even conferred a more general resistance to novel counterarguments (Papageorgis & McGuire, 1961). Since then, more than half a century of investigation into the theory has proven that psychological inoculation is a viable intervention across a multitude of contexts and domains, with some researchers claiming it to be “the most consistent and reliable method for conferring resistance to persuasion” (Miller et al., 2013, p. 127). More specifically, inoculation interventions have proven to successfully confer resistance to persuasive appeals related to social media misinformation (e.g., Roozenbeek, van der Linden, Goldberg, Rathje, & Lewandowsky, 2022), aggressive political campaigning (e.g., Pfau & Burgoon, 1988), and maladaptive responses to health messages (e.g., Richards & Banas, 2015) just to name a few. Moreover, meta-analyses on the efficacy of inoculation interventions have found them to be superior to alternative techniques when it comes to producing attitudinal resistance (Banas & Rains, 2010; Lu, Hu, Li, Bi, & Ju, 2023).

However, while psychological inoculation has proven to be a potent means by which one can defend against undue influence, there remains conflicted findings about its efficacy in certain domains as well as questions concerning the mechanisms via which it operates. In regards to the former, recent research has called into question the ability of inoculation interventions to counter some of the world’s most pressing problems, such as climate change and its subsequent denial. While prior research has demonstrated the ability of inoculation interventions to succeed within this context (e.g., van der Linden, Leiserowitz, Rosenthal, & Maibach, 2017), a recent study conducted across 12 countries found that, beyond a resistance to an initial misinformation encounter, there was little evidence to support the notion that inoculation interventions

provided extensive protection against subsequent exposures to climate disinformation statements (Spampatti, Hahnel, Trutnevyte, & Brosch, 2024). In regards to the latter, when resistance *is* conferred, scholars debate the underlying mechanisms through which it operates. While many inoculation interventions succeed in amplifying one’s ability to *discern*, for instance, truthful statements from falsehoods (e.g., Lewandowsky & van der Linden, 2021; Lu et al., 2023; Roozenbeek, Traberg, & van der Linden, 2022), others – while not necessarily rendering the intervention is ineffective – present evidence which make unclear whether the intervention is sufficiently *targeted* in what it addresses. For example, some results have suggested that inoculation interventions may not only promote scepticism towards *warranted* targets (i.e., sources of misinformation), but may foster a more indiscriminate sense of *overall* scepticism, which may result in “collateral damage” to unintended targets (i.e., sources of true information; e.g., Buczel, Szyszka, Siwiak, Szpitalak, & Polczyk, 2022). However, other research maintains that inoculation interventions do not increase scepticism towards “real” (as opposed to “fake”) news (Roozenbeek, Maertens, McClanahan, & van der Linden, 2021).

Some of the most widely-used and successful inoculation interventions have broadly addressed how stoking intergroup polarisation is a commonly used misinformation *technique* (see van der Linden & Roozenbeek, 2020), and recent research has offered promising evidence to suggest that “technique-based” inoculations⁶⁵ (i.e., where participants are taught about the particular techniques employed by agents of misinformation) can promote a significant reduction in the perceived reliability of content which aims to polarise (Harrop, Roozenbeek, Madsen, & van der Linden, 2022). However, there currently exists little empirical research into the feasibility of using psychological inoculation as a way to combat specific polarisation strategies, such as the manipulation of depictions of political data.

For the purposes of the present study, we sought to investigate the possibility of inoculating individuals against a response consistently observed throughout the studies covered in Chapter 1 and Chapter 2: namely, the tendency for individuals to dramatically underestimate levels of intergroup agreement and significantly overestimate levels of polarisation when presented with bar charts that have a truncated y-axis. Troublingly, this effect was observed not only when bar charts were “maximally-truncated” (i.e., truncated to the greatest extent possible while still ensuring all essential chart elements are visible) but also when they were truncated to a degree which has been shown, in different (i.e., non-political) contexts, to promote appropriate sensitivity to effect size (i.e., approximately 1.5 SD of the grand mean; see Witt, 2019). As bar charts – including truncated bar charts – remain a ubiquitous means of depicting data across the sciences (see Hanel, Maio, & Manstead, 2019) and news media, we sought to develop and test the effectiveness of

⁶⁵ Such inoculation approaches are especially valuable in dynamic information environments as they “prepare individuals to resist being persuaded by messages that may be different in content, but use the same underlying persuasion strategy” (Roozenbeek, Traberg, & van der Linden, 2022, p.2).

an intervention which would A) warn of the threat of being manipulated by media sources who can benefit from disseminating content which paints a more divisive, sensationalised political picture than may actually exist (e.g., Somer, McCoy, & Luke, 2023) and B) teach participants how to spot when the manipulative technique (i.e., truncation of the y-axis) was being employed. In doing so, we hoped to both improve individuals' ability to recognise misleading visualisations as well as to inoculate them against “truncation-induced” polarisation effects.

Though speculative, we suspect a number of factors could account for variability in both overall *susceptibility* to truncation-induced polarisation effects and – should the intervention prove successful – heterogeneity in the overall efficacy of the intervention across the sample. Firstly, graphical literacy (e.g., Durand, Yen, O'Malley, Elwyn, & Mancini, 2020) may play a role. It's possible, for instance, that individuals who are low in graphical literacy are more likely to take data visualisations “at face value” – being less capable (and, consequently, perhaps less likely) to critically interrogate the composition of the graphical depiction of the data. Conversely, those high in graph literacy may display a greater ability to recognise when manipulative techniques are being deployed, providing them with an increased opportunity to reevaluate how the data should be interpreted. Secondly, it's possible that one's propensities concerning reflectiveness and deliberation could play a role. For example, recent research has demonstrated that one's degree of cognitive reflection – assessed via metrics such as the “Cognitive Reflection Test” (Frederick, 2005) – significantly correlates with both the types of accounts they follow on X (formerly Twitter) as well as, critically, the *reliability* of the news they tended to share on the platform (Mosleh, Pennycook, Arechar, & Rand, 2021). While those low on graphical literacy may have no *choice* but to take data visualisations at face value, it's possible that those lower in cognitive reflection may – *regardless* of their level of graphical literacy – simply be more *prone* to accepting data as it is presented without meaningfully engaging with it. Finally, should individuals be particularly hostile towards those with differing political views, it's possible that they may be more prepared and/or willing to accept cross-party data being depicted via truncated y-axis – which will accentuate the degree of difference between the two groups – due to a *desire* to see themselves as distinct from their political opponents (see Abramowitz & Webster, 2018).⁶⁶

Method

Participants

⁶⁶ This is by no means an exhaustive list of potential moderators. For instance, strong cases could also be made for measures such as intellectual humility or conspiracy mentality. Those high on the former, for instance, might be more willing to believe they are capable of being deceived, and thus more likely to engage with the intervention, whereas those high on the latter may be more likely to distrust and scrutinise political data overall, which might render them less susceptible to the truncation-induced polarisation effects.

A sample of American participants were recruited via the Prolific Academic platform and paid an approximate rate of £7.50/hour. Following exclusions for non-consent, attention check failure⁶⁷, or incomplete data, the final sample was 394.⁶⁸ Table 1 summarises the demographic composition of the sample.

Table 1

Composition of Sample for Inoculation Study

| Identity Characteristic | Sample Composition |
|-------------------------|---------------------------------------------------------------------------------------------------------|
| Sex | Female: 62.94% ($n = 248$), Male: 34.26% ($n = 135$), Other: 2.79% ($n = 11$) |
| Age | $M = 40.94$, Range: 19-78 |
| Education | 50.00% ($n = 197$) had attained a bachelor's degree or higher |
| Political Affiliation | Democrat: 24.37% ($n = 96$), Republican: 22.59% ($n = 89$), Independent/Other: 53.05% ($n = 209$) |

Among participants who openly identified as Democrat or Republican, 62.50% ($n = 60$) categorised themselves as “strong” Democrats and 42.70% ($n = 38$) categorised themselves as “strong” Republicans. Among those who openly identified as Independent/Other, when forced to choose which of the two major US political parties they considered themselves closest to, 62.68% of this group chose Democrat ($n = 131$) and 37.32% ($n = 78$) chose Republican. Ideologically, the sample skewed slightly liberal ($M = 43.46$ on a 0 to 100 ideological self-placement scale).

Design

The study employed a two-group, between-subjects design (see Figure 1) where participants were randomly assigned to one of two possible conditions:

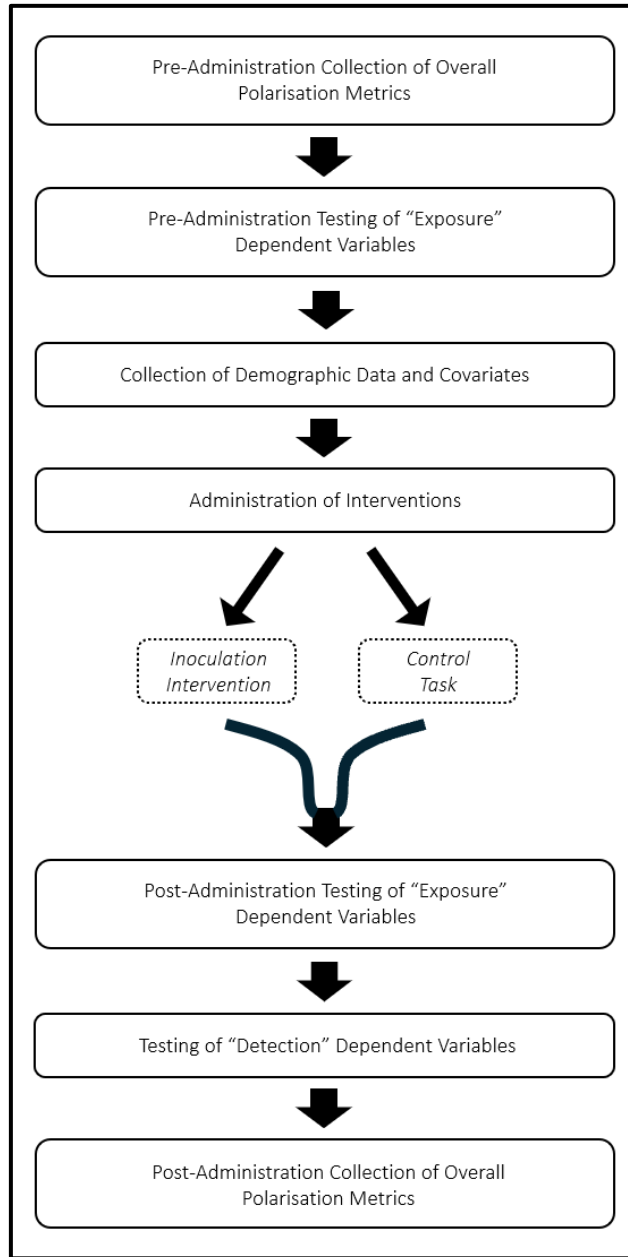
1. Control condition
2. Inoculation condition

⁶⁷ The study featured two separate attention checks, and failure to successfully complete both of which resulted in dismissal from the study. The first, amidst a series of four other questions utilising a 101-point sliding scale, told participants explicitly to slide the scale to precisely 72 if they were paying attention. The second began as an opinion question about politics being a “survival-of-the-fittest” competition, but explicitly stated in the second (of three) sentences for the participant to select “Neither Agree nor Disagree” if they were paying attention. A total of five response options were presented. In total, 91.84% of all participants successfully passed both attention checks.

⁶⁸ A priori power analyses were conducted to determine the sample size required to achieve sufficient statistical power to detect a small-to-medium effect size (i.e., Cohen’s d of 0.35) with a power level of 0.90 and an alpha threshold of 0.05. The analysis indicated a sample size of approximately 173 participants per group would be required. The effect size was chosen based on typical effect sizes observed across similar inoculation studies. Prior to exclusions, the total sample size was 429. For those who provided data, the sample composition prior to exclusions had a mean age of 40.85, was comprised of 261 females, 143 males, and 12 other (with 13 non-responses), and had 103 self-reported Democrats, 95 self-reported Republicans, and 218 self-reported other (with 13 non-responses).

Figure 1

Design Outline for Inoculation Study



Procedure

Participants first completed a pre-intervention round of the “exposure” task, wherein they were presented with eight bar charts depicting intergroup data and asked to share their opinions concerning perceptions of ideological polarisation between the groups as well as perceptions of source reliability (see *Measures*). Next, participants were required to answer a series of demographic questions (including age, sex, highest level of education attained, political affiliation, political identity strength, and ideological leaning) before completing

a set of scales designed to serve as possible covariates when analysing A) overall susceptibility to truncation-induced polarisation effects and B) efficacy of the inoculation intervention. These scales were:

- The Short Graph Literacy Scale (Okan, Janssen, Galesic, & Waters, 2019)⁶⁹
- The Revised Cognitive Reflection Test (Thomson & Oppenheimer, 2016)⁷⁰
- The Moral Disengagement Sub-Scale (taken from their “Partisan Hostility Scale”; Kalmoe & Mason, 2019)⁷¹

Participants were then exposed to their condition-specific intervention (i.e., inoculation or control). Participants in the inoculation condition received a short (i.e., approximately five minute) intervention which taught them about the practice of y-axis truncation, its impact on data visualisations, and how to spot when the technique is being used (see Appendix G for full inoculation script). Control condition participants were administered a neutral intervention exercise (i.e., a word completion task).

Following administration of their respective interventions, participants in each condition then completed the post-intervention round of the exposure task before moving on to complete the “detection” task (see *Measures*).

Additionally, at the very beginning and the very end of the study, participants were asked to respond to a collection of items concerning their perceptions of overall ideological polarisation between Republicans and Democrats, general affective polarisation between the groups, and beliefs about the level of intergroup discreteness (a component of essentialism; e.g., see Haslam & Whelan, 2008).

Measures

The study featured three primary dependent variables situated within two separate tasks.

Exposure Task: The exposure task took place at two different timepoints: once prior to the administration of the intervention and once following the administration. At each timepoint, participants were shown intergroup data depicted via eight bar charts. Half of the bar charts attributed the data to the opinions of political groups (i.e., Republicans and Democrats) while the other half attributed the data to the opinions of unspecified non-political groups (e.g., Group A and Group B). Additionally, half of the charts were untruncated (i.e., “full-range”) while half were truncated.

The four charts used in the political and non-political subgroups were divided into two separate *pairs* of charts. Each pair was said to depict group opinion data on a polarising issue (e.g., abortion, climate change,

⁶⁹ Which had a Cronbach’s alpha of 0.53, but which the developers of the scale justify as a figure that “should be expected from a 4-item scale that purposively varied the type of graph and graph comprehension skills required” (Okan, Janssen, Galesic, & Waters, 2019, p. 187).


⁷⁰ Cronbach’s alpha of 0.874.

⁷¹ Cronbach’s alpha calculated to be 0.703.

etc.) and featured both an untruncated version of the data as well as a truncated version of the data (which was truncated to exactly ½ the range of the untruncated version). Additionally, as a means of examining truncation effects across group types, the political and non-political groups each had *identical* chart pairs (i.e., four charts with the exact same data and dimensions), the only differences being A) the group attribution (and colouring) and B) the issue the data was said to represent (see Figure 2).

Figure 2

Example of Stimuli Used in Exposure Task



For each chart, participants had to indicate their responses to two questions (each on a 0 to 100 scale): 1) How polarised do the groups seem to be on this issue, and 2) Based on the way this data was presented, how reliable do you believe this source to be?

Detection Task: The detection task took place after the respective interventions (i.e., inoculation or control). It used an identical number of charts (i.e., eight) and division of stimuli (i.e., half political groups, half non-political; half truncated bar charts, half untruncated) but featured different chart data and issue attributions. For each chart, participants were simply asked to indicate their response (on a 0 to 100 scale) to the following question: How misleading (if at all) would you say this depiction of the data is?

Finally, the set of pre-post items (i.e., administered at both the very beginning and very end of the study) included:

Measure of Perceived Overall Ideological Polarisation: Participants were asked to indicate, on a scale from 0 (“Not at All Polarised”) to 100 (“Extremely Polarised”), how polarised they believed Republicans and Democrats were overall on key issues.

Measure of General Affective Polarisation: Participants were asked to indicate, on a scale from 0 (“Extremely Cold/Negative”) to 100 (“Extremely Warm/Positive”), how they generally felt towards members of each political party (i.e., Republicans and Democrats).

Measure of Perceived Intergroup Discreteness: Participants were asked to indicate how much they agree or disagree, on a scale from 0 (“Completely Disagree”) to 100 (“Completely Agree”), with the following statement: Republican voters and Democratic voters just seem to be two fundamentally different types of people.

Hypotheses

Table 2 depicts the pre-registered (OSF URL: <https://osf.io/fsq3t>) hypotheses for the study:

Table 2

Pre-Registered Hypotheses for the Study

| Label | Hypothesis |
|-------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H1 | Participants assigned to the inoculation condition will evaluate the groups being depicted via a truncated bar chart as being significantly less polarised in the post-assessment of the exposure paradigm than in the pre-assessment. No significant differences will be observed for participants in the control condition. |
| H2 | Participants assigned to the inoculation condition will evaluate the sources depicting data via a truncated bar chart as being significantly less reliable in the post-assessment of the exposure paradigm than in the pre-assessment. No significant differences will be observed for participants in the control condition. |
| H3 | Participants in the inoculation condition will perform significantly better in the detection paradigm than participants in the control condition. |
| H4 | Participants assigned to the inoculation condition will evaluate the groups being depicted via a truncated bar chart as being significantly less polarised than participants assigned to the control condition in the post-assessment of the exposure paradigm. |
| H5 | Participants assigned to the inoculation condition will evaluate the sources depicting group data via a truncated bar chart as being significantly less reliable than participants assigned to the control condition in the post-assessment of the exposure paradigm. |
| H6 | Significant main effects of condition will be observed in the pre-post differences of perceptions of overall ideological polarisation. More specifically, participants assigned to the inoculation condition will evaluate Republican voters and Democratic voters as being significantly less polarised in the post-assessment than in the pre-assessment. ⁷² |

Results


While t-tests revealed that participants in the control condition ($M = 69.52$, $SD = 19.52$) and participants in the inoculation condition ($M = 67.48$, $SD = 19.71$) did not significantly differ in their perceptions of polarisation for groups depicted via a bar chart with a truncated y-axis *prior* to the administration of their

⁷² Please note that while the pre-registration indicated our analysis plan would involve an ANOVA, we ultimately decided a t-test would be more straightforward.

respective interventions ($t(392) = 1.03, p = .30, d = 0.10$), a statistically significant difference was observed in their post-intervention perceptions of polarisation for groups depicted via a bar chart with a truncated y-axis ($M_{CONT} = 68.17, SD_{CONT} = 20.78, M_{INOC} = 51.27, SD_{INOC} = 26.53, t(367.34) = 7.03, p < .001, d = 0.71$). Although participants in the control condition did not display a significant change in their pre-post perceptions of polarisation for groups being depicted via a truncated bar chart ($M = -1.34, SD = 11.54; t(396) = 0.66, p = .51, d = 0.07$), t-tests revealed that participants in the inoculation condition did exhibit a significant post-intervention reduction in their perceptions of polarisation for groups depicted via a truncated bar chart ($M = -16.22, SD = 27.96; t(358.19) = 6.85, p < .001, d = 0.69$; see Figure 3).

Figure 3⁷³

Pre-Post Change in Perceptions of Intergroup Polarisation by Chart Type



* $p < .05$, ** $p < .01$, *** $p < .001$

Exploratory analyses revealed that, in addition to the inoculation intervention eliciting significant reductions in pre-post perceptions of polarisation for groups depicted via truncated bar charts, it also produced statistically significant reductions in perceptions of polarisation for groups depicted via *full-range* bar charts ($M = -12.22, SD = 18.68; t(388) = 6.45, p < .001, d = 0.65$). Critically, however, although the latter effect was unintentional (as the intervention sought to primarily address excessive perceptions of polarisation engendered by truncated bar charts), analyses suggest that the intervention still appears to have effectively reduced the polarisation perception *gap* between groups depicted by truncated bar charts and full-range bar charts for participants in the inoculation condition, as evidenced by the fact that the effect size difference in perceptions of polarisation between the two chart types was reduced from $d = 0.90$ pre-intervention to $d = 0.55$ post-intervention.


Regarding perceived source reliability, participants in the control condition exhibited no significant differences in their pre-post perceptions of reliability for sources depicting groups either via truncated bar

⁷³ Please note that Figure 3 and Figure 4 feature a truncated y-axis. This has been used because the theoretical full-range of the y-axis would be 200 (i.e., accommodating for up to a 100 point perception change in either direction). Using such a large y-axis compresses the bars so substantially that it is challenging for a reader to be able to meaningfully examine the chart, including either the inter- or intra-condition differences. While we believe this moderate degree of truncation depicts the conditions (and the aforementioned differences) in a reasonable manner, it's important to recognize that truncation can exaggerate these features, and thus readers should ensure their visual perceptions are grounded in the statistical results presented.

chart ($M = -0.39, SD = 10.48; t(396) = 0.18, p = .86, d = 0.02$) or full-range bar chart ($M = 0.76, SD = 10.09; t(396) = 0.38, p = .71, d = 0.04$). Conversely, participants in the inoculation condition *did* exhibit statistically significant differences in their pre-post perceptions of source reliability for sources depicting groups via both truncated bar chart ($M = -28.01, SD = 26.07; t(388) = 12.33, p < .001, d = 1.25$) and via full-range bar chart ($M = 28.84, SD = 28.68; t(388) = 13.32, p < .001, d = 1.35$; see Figure 4). More specifically, for participants in the inoculation condition, pre-post perceptions of source reliability *decreased* for sources choosing to depict groups via a truncated bar chart (from $M_{PRE} = 48.56$ to $M_{POST} = 20.56$) and *increased* for sources choosing to depict groups via a full-range bar chart (from $M_{PRE} = 45.43$ to $M_{POST} = 74.27$). Ultimately, although no significant pre-intervention differences existed between participants in the control condition and those in the inoculation condition concerning perceptions of source reliability for sources depicting groups either via a truncated bar chart or a full-range bar chart, participants across conditions exhibited significantly different post-intervention perceptions of source reliability for both sources depicting groups via a truncated bar chart ($M_{CONT} = 52.22, SD_{CONT} = 21.54, M_{INOC} = 20.56, SD_{INOC} = 22.92; t(392) = 14.13, p < .001, d = 1.42$) as well as sources depicting groups via a full-range bar chart ($M_{CONT} = 48.09, SD_{CONT} = 20.56, M_{INOC} = 74.27, SD_{INOC} = 21.79; t(392) = 12.27, p < .001, d = 1.24$).

Figure 4

Pre-Post Change in Perceptions of Source Reliability by Chart Type




* $p < .05$, ** $p < .01$, *** $p < .001$

T-tests also confirmed statistically significant differences emerged between control condition participants and inoculation condition participants in their sensitivity to the detection of “misleading” and “non-misleading” visualisations (i.e., visualisations utilizing a truncated y-axis and visualisations utilizing a full-range y-axis, respectively). Both the control group ($M_{TRUNC} = 60.43, SD_{TRUNC} = 19.74, M_{FULL} = 52.85, SD_{FULL} = 21.61; t(396) = 3.65, p < .001$) and the inoculation group ($M_{TRUNC} = 79.09, SD_{TRUNC} = 17.08, M_{FULL} = 22.11, SD_{FULL} = 24.42; t(347.11) = 26.70, p < .001$) offered assessments of misleadingness which differed significantly between chart types (i.e., truncated versus full-range), and critically, both groups exhibited the appropriate *directionality* of divergence in their assessment (i.e., rating truncated charts as being *more* misleading than full-range charts). However, the effect sizes which accompanied these differences were very different for each set of participants: control group participants provided assessments of misleadingness for truncated

and full-range bar charts which had an effect size difference of $d = 0.37$, while inoculation group participants provided assessments of misleadingness whose effect size, as measured by the assessment differences between the two chart types, was $d = 2.70$. Overall, inoculation group participants assessed truncated bar charts to be significantly *more* misleading than control group participants ($t(386.05) = 10.04, p < .001, d = 1.01$) and assessed full-range bar charts to be significantly *less* misleading than control group participants ($t(392) = 13.24, p < .001, d = 1.33$; see Figure 5).

Figure 5

Perceptions of Visualisation Misleadingness by Chart Type



* $p < .05$, ** $p < .01$, *** $p < .001$

Additionally, there is evidence to suggest that the inoculation intervention also promoted reductions in both perceptions of overall ideological polarisation between Republicans and Democrats as well as a component of perceived intergroup essentialism (i.e., discreteness) for the two groups. In regard to the former, t-tests confirmed the emergence of significant pre-post differences in perceptions of overall ideological polarisation between participants who underwent the inoculation intervention and those who underwent the control intervention ($t(341.51) = 9.02, p < .001, d = 0.91$), with participants in the inoculation condition exhibiting a significantly greater reduction in perceived overall ideological polarisation ($M = -19.96, SD = 24.32$) than those in the control condition ($M = -1.01, SD = 16.58$). In regard to the latter, exploratory t-tests revealed significant pre-post differences in perceptions of intergroup discreteness between the two participant groups ($t(378.10) = 5.92, p < .001, d = 0.60$), with inoculation group participants exhibiting a significantly greater reduction in perceived intergroup discreteness ($M = -14.97, SD = 19.81$) than control group participants ($M = -4.04, SD = 16.65$).

Exploratory analyses attempted to examine factors which may A) contribute to susceptibility to “truncation-driven polarisation effects” and B) explain variance in intervention efficacy among inoculation group participants. To assess the former, we first determined participants susceptibility to changes in chart type by subtracting their perceptions of polarisation for the full-range version of each chart (from the pre-exposure paradigm) from their perceptions of polarisation for the corresponding truncated version. This showed how much more polarised participants perceived (identical) groups to be simply by virtue of the y-

axis being truncated. Overall, across all participants, the act of truncation led to an average increase in perceived polarisation of 17.81 points. Using this metric as the dependent variable, linear regression models were constructed to determine which variables might operate as significant predictors to account for variance in “susceptibility.” Neither demographic variables (e.g., age, sex, education) nor other hypothesised covariates (e.g., graph literacy, cognitive reflection) proved to be significant predictors of variance when plugged into single-predictor regression models (i.e., models constructed with a single explanatory variable). Using degree of moral disengagement (a sub-scale of partisan hostility) as a sole explanatory variable did yield a model which was significantly predictive of variance in susceptibility ($R^2 = .01$, $F(1, 392) = 5.782$, $p = .02$), however this model fails to retain significance once Bonferroni corrections are applied to adjust the alpha threshold. To determine whether any variables were capable of explaining variance in the efficacy of the intervention, a metric was created which measured the differences in truncation-driven polarisation susceptibility pre- and post-intervention. Using this metric as a dependent variable (and data from inoculation condition participants only), single explanatory variable models were tested, first on demographic variables and then on aforementioned scale measures. While no demographic variables accounted for variance in intervention efficacy in a statistically significant manner, two of the scales (i.e., cognitive reflection and moral disengagement), when used in sole independent variables in the linear model, yielded statistical significance (i.e., $\beta = -4.924$, $F(1, 193) = 9.437$, $p = .002$, $R^2 = .04$ and $\beta = 0.695$, $F(1, 193) = 7.675$, $p = .006$, $R^2 = .03$, respectively) even after Bonferroni corrections were applied. However, when the model inputs both scale measures as explanatory variables, only moral disengagement remains as a significant predictor of variance ($\beta = 1.065$, $p = .04$).

Discussion

The results suggest that a brief (i.e., approximately five minute) intervention in which participants are A) made aware of the practice of y-axis truncation (including why media outlets might choose to engage in such a practice), B) warned of its potential impact on intergroup perception, and C) taught how to spot when the practice was being employed can effectively mitigate truncation-driven polarisation effects relative to participants who took part in a control intervention. More specifically, the inoculation intervention seems to be effective in promoting a number of separate yet interrelated “defensive” responses.

First of all, it significantly reduces perceptions of ideological polarisation for groups depicted via truncated bar chart relative to individuals who have not undergone the intervention. This is an incredibly important feature of the treatment, as evidence from Chapter 1 and Chapter 2 shows how individuals exposed to intergroup data via truncated bar charts consistently perceive groups as being significantly more polarised than when identical data is depicted via full-range bar charts. However, the intervention did have the unintended effect of also reducing perceptions of polarisation for groups depicted via *untruncated* bar charts.

On the one hand, it's possible that such an effect could materialise as a result of the intervention fostering an *overall sense of data scepticism* as opposed to a more targeted discernment of misleading visualisations. Similar results have been reported within other inoculation paradigms, where participants seem to become sceptical not only of information which *warrants* scepticism (i.e., misinformation), but seemingly of *all* types of information (including that which contains no misinformation; e.g., Buczel et al., 2022). While priming certain forms of scepticism can be useful in motivating a healthy resistance to undue influence (e.g., Li, 2023), heightened levels of indiscriminate scepticism could feasibly have deleterious effects, causing individuals to question accurate, and thus useful, information. On the other hand, however, in this particular context, it's possible that this inadvertent byproduct of the treatment could be a net-positive for two reasons. First, considering the fact that multiple sets of longitudinal data document a decades-long pattern of overexaggerated perceptions of polarisation within the United States (e.g., Westfall, Van Boven, Chambers, & Judd, 2015), it's likely that a reduction in perceived polarisation across *both* graph types is warranted, especially in light of the fact that participants exposed to full-range bar charts – although performing significantly better than their counterparts exposed to truncated bar charts – may *still* significantly underestimate *actual* levels of intergroup opinion overlap (e.g., see Chapter 1 and Chapter 2). Secondly, the reason truncated bar charts are an effective weapon of misinformation is that they skew perceptions relative to full-range bar charts. As they are depicting identical data, in an ideal world, participants exposed to each format would report the same perceptions of and draw the same inferences from the data. While not catalysing that degree of parity, the intervention did manage to *reduce the difference* between perceptions of polarisation produced by truncated bar charts and perceptions of polarisation produced by untruncated bar charts – effectively reducing the asymmetry in their impact.

Secondly, participants exposed to the inoculation intervention appear to have become more discerning when assessing source reliability, evaluating sources that have chosen to truncate the y-axis significantly less favourably than those that have chosen to depict the data using a full-range y-axis. This feature of the treatment is critical as its effects may continue to exert influence on the media ecosystem beyond the initial exposure. For example, should an individual notice a particular source truncating bar charts, the corresponding reduction in perceived reliability of the source may compel them to search elsewhere for information. Extrapolating this behaviour on a large enough scale, a shifting of demand-side characteristics driven by more data-literate audiences may drive migrations from less reliable sources (i.e., as determined by the proxy signal of y-axis truncation) to more reliable sources, and reductions in metrics like website traffic and audience engagement (as discussed in Chapter 4) are precisely the type of levers which would motivate media outlets to adjust their content strategies.

Finally, as evidenced by participant responses in the detection task, the intervention significantly alters perceptions of visualisation misleadingness. Such an alteration may either be driven by A) participants

becoming more *capable* of identifying misleading visualisations and/or B) participants developing a heightened sensitivity to the practice of y-axis truncation when rendering judgments of misleadingness (i.e., the presence or absence of this practice *informs* their judgments to a larger degree). In terms of the former, it is worth reiterating that participants in the control condition, to their credit, *also* perceived data depicted via truncated bar charts to be significantly more misleading than data depicted via untruncated bar charts, however the size of the effect (i.e., $d = 0.37$) was dwarfed by the size of the effect in the inoculation group ($d = 2.70$), perhaps lending credence to the latter explanation.

Beyond the aforementioned effects, the inoculation intervention also seems to have played a role in reducing perceptions of *overall* ideological polarisation as well as perceptions of intergroup discreteness. This represents a pair of valuable residual effects as it suggests that not only can an inoculation intervention deployed in this manner reduce issue-specific perceptions of truncation-driven polarisation, but the intervention – perhaps due to its heavy emphasis on the media’s role in sustaining and, in worst-case scenarios, manufacturing heightened perceptions of polarisation – also seems to leave participants believing that Republicans and Democrats are both less divided overall as well as less “fundamentally different” from one another – a result with important implications for dehumanisation (e.g., Haslam, 2006).

Moreover, in examining the efficacy of the intervention, it appears that moral disengagement was the only statistically significant predictor. The relationship, however, appears to be that individuals with higher scores on moral disengagement show less “improvement” (as measured by the reduction in how different one perceives identical data presented via different chart types) following the inoculation. While one can speculate as to why such a finding might emerge (e.g., perhaps individuals high on this measure were less motivated by the narrative within which the intervention was couched; namely, that Republicans and Democrats are “not always as different as the media might make them seem”), it warrants a more formal examination in the future to determine the precise underlying mechanisms.

While the inoculation intervention was effective across a variety of metrics, “success” within this study was almost solely measured as the performance of the inoculation group participants relative to the control group participants. More specifically, it was determined that the intervention would be judged to be successful if it was able to reduce perceptions of polarisation (for data depicted via a truncated bar chart) more than the control condition, if it promoted a more discerning evaluation of source reliability (i.e., becoming more likely to evaluate sources that chose to use a full-range y-axis as “reliable” while deeming those choosing to use a truncated y-axis as “less reliable”), and if it led to a greater sensitivity in one’s ability to detect misleadingness within this data visualisation format. However, while the consequences of the intervention are likely all net-positive contributions to depolarisation efforts, it’s important to be aware that it’s possible for interventions of this nature to go “too far.”

While it is fair for a depolarisation agent to aim to reduce the negative *effects* of polarisation (e.g., anti-democratic behaviour, partisan hostility, etc.) “as much as possible,” one should not seek to blindly reduce *perceptions* of polarisation as much as possible. If two groups disagree on an issue, then some degree of polarisation *does* exist, and thus the intention of the responsible practitioner should not be to convince the two sides that there is no disagreement, but rather to ensure that they are gauging the extent of the disagreement appropriately. Similarly, the concern underlying the truncation-induced polarisation effects observed throughout Chapter 1 and Chapter 2 was not that truncated bar charts were merely reducing perceptions of polarisation relative to alternate data visualisation techniques (e.g., full-range bar charts), but that they were also *compromising the accuracy* of intergroup perceptions (as measured by estimates of PCS). Consequently, a limitation of the current study is that it was devoid of more objective measures of *accuracy* against which we could benchmark the treatment’s effects. For example, while it is *likely* good that participants in the inoculation condition came to see the groups as less polarised, there will inevitably be a limit wherein further reductions of perceived polarisation will once again start to *contribute to* perception inaccuracies, at which point the intervention bypasses its intention of “correcting” perceptions and effectively crosses into “manipulating” them.

Thus, we believe that researchers interested in building upon the current study could do so in at least five ways. First, in regard to the discussion above, future research should aim to incorporate accuracy benchmarks into the designs of similar studies, enabling researchers to determine the degree to which the inoculation is correcting (versus over-correcting) participant perceptions. Secondly, future researchers should try to understand how important a role the media narrative within which the inoculation script was couched plays in the overall efficacy of the intervention. More specifically, would a similar intervention devoid of such a narrative yield similar impact or did the specificity of the narrative substantially improve motivation to avoid being deceived? Thirdly, future researchers should aim to determine the longitudinal effects and rates of decay of the intervention, as has been explored in other inoculation investigations (e.g., Maertens, Roozenbeek, Basol, & van der Linden, 2021). Fourthly, considering that the intervention also significantly shifted “non-targeted” perceptions (such as those concerning intergroup discreteness), future research should also seek to investigate whether such an intervention may also yield “cross-protection,” or resistance to other persuasion techniques not directly inoculated against (see Parker, Rains, & Ivanov, 2016). Finally, the current study featured a control condition which had participants complete a timed word-construction task. While the design of this activity was similar to the types of control interventions featured in prior studies on inoculation (e.g., word sorting tasks and word puzzles; Maertens, Anseel, & van der Linden, 2020), future research should investigate how alternative control conditions may impact the degree of divergence observed between inoculation and control participants across the study’s dependent variables. For example, instead of using a “true neutral” control intervention, perhaps future researchers could

investigate how motivating control participants to heighten their vigilance to manipulation (without ever inoculating them in any specific way) might impact the between-condition differences.

The results of the study have important practical implications. Perhaps not coincidentally, America has witnessed a simultaneous rise of both levels of polarisation as well as levels of political disinformation (e.g., Tucker et al., 2018). While, in an ideal world, institutions such as major news outlets, whose level of reach and influence capital cannot be matched by individual practitioners, would play a more active role in depolarisation efforts, to count on such institutional-driven change would arguably be a triumph of hope over experience. Consequently, to address these alarming trends, academics have begun spearheading large-scale searches for interventions which might prove useful in stemming the advance of escalating levels of anti-democratic attitudes and partisan animosity (e.g., the “Strengthening Democracy Challenge”; see Voelkel et al., 2022). The current study represents a small but meaningful contribution to both the inoculation literature as well as to the growing body of depolarisation tactics. In regard to the former, it provides evidence to suggest that technique-based inoculation can operate as a mechanism to counter misleading visualisation tactics – an area which, to the best of our knowledge, has not been empirically tested in the extant inoculation literature. In regard to the latter, it speaks to the viability of administering inoculation-based interventions as a concise and potent tool in the fight against accelerating polarisation. In just five minutes, it appears such interventions can significantly reduce perceptions of polarisation (while simultaneously affecting the manner in which source reliability and visualisation misleadingness is evaluated), making it a valuable addition to the depolarisation repertoire.

GENERAL DISCUSSION

THE IMPACT OF DATA PRESENTATION ON POLITICAL PERCEPTION AND INTERGROUP DYNAMICS

A Summary of Research Objectives

The research featured throughout this thesis sought to determine the role data presentation plays in the moderation of political perception, affect, and action as well as to investigate potential solutions to address what we might refer to as “questionable presentation practices” (similar to “questionable research practices”; e.g., Xie, Wang, & Kong, 2021). To achieve this, we devised a research agenda which first aimed to systematically test the impacts of two distinct formats of data presentation (i.e., data visualisation and data summarisation), with the ultimate goal of elucidating the consequences of data presentation choices on political perception and behaviour and ideally determining the boundary conditions of these effects. Upon reviewing the results obtained from our first six studies pertaining to the impact of data presentation on metrics such as perceptions of ideological polarisation and intergroup opinion overlap, it became clear that there were “better” and “worse” ways to present data (particularly when it came to visualisation choice). Consequently, the latter half of our research agenda came to be dedicated to conducting early-stage investigations into two possible intervention approaches (one systemically-focused, the individually-focused), with the hope of providing preliminary evidence regarding the feasibility of possible intervention strategies to mitigate the undesirable effects of questionable presentation practices.

In Chapter 1, drawing inspiration from the work of Hanel, Maio, and Manstead (2019), we attempted to extend (and partially replicate) their findings concerning “mode of presentation,” but aimed to do so using data from within a polarised political context (i.e., the American political system). More specifically, we sought to understand if (and, if so, to what degree) visualising identical data in different manners might affect political perception and/or behaviour. Arguably the most experimentally rigorous of our research undertakings, the first four studies attempted to disentangle the effects of data visualisation at varying levels of artificiality and abstraction.

In Chapter 2, we attempted to replicate our work from Chapter 1 using two large-scale, international samples from separate geopolitical contexts (i.e., the United States and the United Kingdom).

In Chapter 3, we tested whether long-term, repeated exposure to certain types of summarisation techniques (e.g., only sharing group means, only sharing degree of overlap/percentage of common scores, etc.) or

summarisation-visualisation combinations (e.g., sharing group means and a corresponding bar chart) would yield significant impacts on partisan perception and/or behaviour.

In Chapter 4, in a pair of methodologically-identical studies, we investigated whether the different presentation formats (i.e., visualisation types and summarisation techniques) sources *choose* to use would impact A) audience engagement with the content, B) perceptions of the sources (e.g., trustworthiness, credibility, etc.), and/or C) future intent to engage with the source.

Finally, in Chapter 5, we tested the viability of creating and deploying a short “inoculation” intervention aimed at reducing the effect y-axis truncation has on perceptions of ideological polarisation (as evidenced by the results obtained throughout Chapter 1 and Chapter 2). The study not only endeavoured to determine whether this type of intervention – which, to our knowledge, had never been tested in this fashion – could mitigate truncation-induced polarisation effects, but also whether it could make participants more discerning about how sources *choose* to display data, ultimately leading to a heightened sensitivity surrounding y-axis manipulation when rendering judgements about source reliability.

Synthesising Key Findings

Across five chapters and nine separate studies, several key insights can be gleaned which we believe have applicability for both researchers and practitioners alike. Below, I attempt to synthesise what we consider to be the most robust and consequential findings from our research.

Visualisation Choice Exerts a Significant Impact on Perceptions of Ideological Polarisation

While Hanel et al. (2019) had obtained results which suggested that mode of presentation A) was capable of significantly impacting the manner in which identical underlying data was perceived and interpreted and B) that certain modes (e.g., superimposed normal distributions) could promote improved intergroup inferences (e.g., more positive attitudes towards outgroups), it was not clear whether such patterns would replicate when the groups being depicted were opponents in a highly-polarised political environment. However, as our findings throughout Chapter 1 and Chapter 2 made clear, exposure to the same data, depicted via different visualisation formats, can substantially alter perceptions of ideological polarisation even under these circumstances.

This effect, while first observed in rather contrived experimental conditions (e.g., with artificial data on unspecified issues attributed to non-political groups), proved robust to experimental manipulations which incrementally increased the external validity by reducing artificiality and experimental control; from using non-political group labels to political group labels, from providing no issue specificity to ascribing data to divisive issues, from moving from the relative sterility of artificial data to the messiness of real world statistics,

the effect held. Moreover, the effect was replicated (albeit with a reduced effect size⁷⁴) in large, representative samples in two separate geopolitical contexts.

As depicted in Table 1, the effect sizes observed as a result of simply altering the visualisation format were routinely pronounced, often exceeding Cohen’s *d* of 0.80 (i.e., the threshold at which we typically interpret such an effect to be considered “large”).

Table 1

Range of Effect Sizes Between Visualisation Formats on Issue-Specific Perceptions of Ideological Polarisation

| Study | Effect Sizes (Cohen’s <i>d</i>) | | | |
|--------------------|--------------------------------------------------|-----------------------------------------------------|--------------------------------------------------|-----------------------------------------------------|
| | <i>No Visualisation and Full-Range Bar Chart</i> | <i>Truncated Bar Chart and Full-Range Bar Chart</i> | <i>No Visualisation and Icon Array Histogram</i> | <i>Truncated Bar Chart and Icon Array Histogram</i> |
| Chapter 1, Study 1 | NA | 0.93 ¹ | NA | 0.64 ¹ |
| Chapter 1, Study 2 | NA | 1.03 | NA | 0.78 |
| Chapter 1, Study 3 | 1.88 | 1.21 | NA | NA |
| Chapter 1, Study 4 | NA | 1.08/0.89 ² | NA | 0.80/0.45 ² |
| Chapter 2 | 0.22/0.27 ³ | 0.23/0.21 ³ | 0.18/0.20 ³ | 0.18/0.13 ³ |

*Note*₁. While perceived polarisation was measured in Study 1, an issue had not yet been specified or attributed to the data, and both political as well as non-political group labels were used.

*Note*₂. The two statistics per cell for Study 4 represent perceptions of ideological polarisation for the immigration issue and the assault rifle issue, respectively.

*Note*₃. The two statistics per cell for Chapter 2 represent the UK data and US data, respectively.

Critically, these changes in visualisation choice were not merely shifting participant perceptions on obscure or inconsequential issues – these interventions were consistently catalysing significant shifts in perception on highly-divisive issues for which partisans likely held strong, historically-entrenched attitudinal prior attitudes. From immigration to gun control, visualisation format proved capable of substantially altering beliefs concerning how polarised groups were on some of today’s most vitriol-laden, hotly-debated topics. Such a finding is especially important as it suggests that these interventions may not only be capable of shifting perceptions on nascent or emerging political issues, but may also be able to alter how issues which

⁷⁴ While the overall pattern of results between Chapter 1 and Chapter 2 remained broadly consistent, the effect sizes obtained in the latter were often noticeably smaller than those observed in the former (as evidenced by the data in Table 1). While we hypothesise that the incongruity is likely due to differences in sample composition and/or sampling methodology (as samples for all Chapter 1 studies came from Prolific’s online participant recruitment platform and the sample for the Chapter 2 replication came from Ipsos MORI’s proprietary recruitment and sampling methodology), preliminary analyses of possible explanatory factors like educational attainment of the samples did not reveal any pronounced heterogeneity between the two sets of samples. Consequently, as of this writing, it remains an open question as to what might be driving these differences in effect sizes.

have an extensive track record of partisan discord – issues which may otherwise be described as intractable – are viewed and, possibly, approached.

Finally, across several studies, it was found that the effect that visualisation choice exerted on *issue-specific* perceptions of ideological polarisation also extended to yield a corresponding impact on *overall* perceptions of intergroup polarisation. The fact that perception changes originating in a particular policy domain generalise to affect perceptions of overall division across the broader policy landscape is an especially intriguing byproduct of this intervention, as it implies that the application of more conscientious data visualisation practices on a specific issue can potentially unleash a cascade of perception alterations on separate, otherwise unrelated issues. However, as the current collection of studies did not measure the second-level effects of the intervention on any *specific* political issue, we cannot conjecture as to the power (or lack thereof) it may have on more tangible downstream perceptions.

More “Balanced” Visualisations Catalyse More Accurate Perceptions – But Still Significantly Underestimate Actual Levels of Intergroup Agreement

Another feature of Hanel et al.’s (2019) work which the current research sought to replicate within an explicitly political context concerned the disparate intergroup perceptions catalysed by more “typical” modes of presentation (e.g., bar charts, especially those with a truncated y-axis) and those catalysed by more novel, “balanced” modes of presentation (e.g., superimposed normal distributions, icon array histograms, etc.). Hanel et al. (2019) originally found that visualisations which provided more “similarity” information (e.g., showing degree of overlap as opposed to merely divergence) produced more accurate perceptions.

Measuring how different modes of presentation affect the *accuracy* of perceptions is especially important in this type of work. While our underlying goal was to determine whether certain data presentation methodologies might be capable of reducing perceptions of polarisation, the circumstances under which these reductions might occur were of critical importance. Specifically, if certain modes managed to do so by *also reducing the accuracy* of intergroup perception (i.e., causing participants to believe opposing political groups agreed *more* than they actually do), this would constitute manipulation in the same way that cultivating inaccuracies in the opposite direction would (i.e., influencing participants to believe political groups agree *less* than they actually do). Consequently, while the end goal of our research was to try to identify a possible way to curtail runaway polarisation (as is the goal of many other polarisation researchers), we must not ignore the *means* by which this curtailing might be achieved (and whether such means are desirable, or even ethical). Fortunately, our findings not only suggested that more “balanced” visualisations often led to lower perceptions of polarisation but also tended to support Hanel et al.’s (2019) notion that these modes of presentation also frequently promoted more *accurate* perceptions of the data (as measured

by the relative degree to which they approached a PCS or PCR benchmark). However, this finding has two important caveats.

Firstly, we must be careful in how we define “balanced.” While balanced can certainly pertain to comparatively *novel* visualisation formats which show levels of group disagreement *as well as* levels of group agreement (such as superimposed normal distributions and icon array histograms), we would argue that balanced could also mean – at least in the context of a polarised information environment – visualisations which do not *unnecessarily accentuate* group differences (such as a full-range bar chart). Although a full-range bar chart can accurately be categorised as a “typical” means of depicting data, it is likely an *atypical* depiction within hostile online ecosystems and partisan echo chambers, where participants are more likely to opt for data visualisations which most dramatically amplify the appearance of group differences (e.g., truncated bar charts). Thus, full-range bar charts, relative to their truncated counterparts, represent a *comparatively* more balanced depiction of data – one which our results show reliably catalyses greater reductions in perceived polarisation as well as more accurate perceptions of intergroup agreement.

Secondly, although more balanced depictions of data led to estimates of intergroup agreement which *more closely approached* the accuracy benchmarks, they still frequently produced estimates which fell significantly below these measures. The implication, of course, is that even by utilising data presentation interventions which leverage the most promising visualisation techniques featured in our research (e.g., full-range bar charts and icon array histograms, which consistently promoted some of the highest levels of intergroup perception accuracy), individuals will often continue to dramatically underestimate the degree to which voters on either side of the aisle agree with one another. Additionally, when sub-optimal visualisation formats are used or – sometimes worse – when participants are provided with no visualisation at all, the degree of underestimation becomes especially pronounced. This represents an important limitation of the work that practitioners should heed: while “better” visualisation methods can *improve* perception accuracy, they often cannot fully correct pervasive cross-party misperceptions.

In the Context of Reducing Polarised Perceptions, Exposure to Any Form of Data Often Seems Preferable to Exposure to No Data at All

Within Chapters 1, 2, and 3, a pattern emerged which seems to suggest that, in many circumstances, providing individuals with *any* form of (accurate) data may yield preferable results to providing them with no data whatsoever.

In fact, while the truncated bar chart and the no visualisation condition produced responses that generally did not significantly differ from one another in Chapter 2, in both Study 3 and Study 4 of Chapter 1, even the truncated bar chart – a decidedly sub-optimal visualisation format for the purposes of political data – often produced *more positive* perceptions than the corresponding no visualisation condition. Thus, though it was

perhaps unsurprising that “better” visualisations (e.g., full-range bar charts and icon array histograms) proved vastly superior than no visualisation across virtually all metrics, what was surprising is that, evidently, it’s often better to provide an audience with data visualised sub-optimally (e.g., via truncated bar chart) than to provide them with no data at all (and consequently leave them to make judgments based on their pre-existing beliefs).

Relatedly, while the longitudinal examination of the role of repeated exposure to various summarisation techniques on perceptions of polarisation (Chapter 3) generally failed to find consistent significant differences *between* experimental conditions, it did present compelling evidence to suggest that repeated exposure to *any* form of data summarisation seemed to be superior to participants receiving no data at all on metrics such as perceived ideological polarisation, perceived intergroup agreement, and perceived ideological distance of average partisan voters. More specifically, in all instances where this pattern was observed, the no visualisation control was generating perceptions which were “worse” (e.g., higher perceptions of polarisation, lower perceptions of intergroup agreement, etc.) than the experimental conditions, implying that participants were more likely to walk away with a more positive intergroup outlook after eight weeks had they been exposed to some form of consistent data summarisation than had they been left to assess intergroup dynamics without the aid of data.

As briefly discussed in Chapter 3, we believe this is an important finding as it provides evidence for a pressing question: in polarised political contexts, is presenting partisans with data on divisive issues more likely to provoke or placate intergroup tension? In the absence of evidence, it wouldn’t be unreasonable to assume that weekly reminders of where partisans stand on contentious policy areas might inflame an already heightened partisan animosity, but the results of our studies seem to suggest otherwise: namely, that such exposure actually has a mollifying effect on beliefs about the extent of political disagreement.

We believe such an effect might largely be a byproduct of the pervasive misperceptions that characterise the political landscape of the United States (e.g., Westfall, Van Boven, Chambers, & Judd, 2015). One could conjecture that, if individuals within this political context are naturally prone to *believing* that political groups are further apart on key issues than they actually are, then making assessments based on nothing but prior (i.e., likely skewed) knowledge is bound to lead to exaggerated perceptions of difference. However, as is evidenced by the responses to the “Level of Intergroup Agreement Relative to Expectations” metric from both Study 2 of Chapter 1 as well as from Chapter 2, most individuals seem to *expect* partisan data to depict more pronounced divergences between the two parties than are actually warranted; they were most likely to say that the truncated bar chart – a mode of presentation which serves to visually accentuate group differences – depicted levels of agreement that met their expectations, while visualisations like the full-range bar chart and icon array histogram were relatively more likely to violate these expectations (insofar as they

presented data in a fashion that gave the impression the groups agreed marginally *more* than they would have otherwise expected).

Ultimately, we believe that the widespread inaccuracies that permeate the American political landscape and (mis)inform the cross-party perceptions of the average citizen make it inadvisable for agents interested in sparking depolarisation efforts to do so without the help of exposure to accurate intergroup data. In fact, as suggested by some of our results, under certain circumstances, even exposure to suboptimal data visualisations can be beneficial relative to the alternative of providing no exposure whatsoever.

Data Presentation Choices Impact How Individuals Perceive and Intend to Interact with Sources

Although the pair of studies detailed in Chapter 4 revealed that data presentation does not appear to play a significant role in one's *engagement* with content (i.e., "liking" behaviour, desire to share, click-through rate, etc.), it also made clear that mode of presentation – whether it be visualisation type or summarisation technique – exerts a strong influence on the way individuals come to view and intend to interact with sources.

Critically, the study featured separate conditions designed to simulate the distinct ways in which individuals may become acquainted with the alternative methods a source could use to visualise or summarise data. The *implicit* alternative choice revelation condition simply exposed participants to the other ways data could be presented whereas the *explicit* alternative choice revelation condition both exposed participants to alternative presentation modalities *and* expounded upon the relative advantages and disadvantages of each method. The former condition was meant to simulate how an individual might react to the data presentation choices of a source after simply becoming aware that alternatives exist; the latter condition was meant to simulate how they might react were they to become aware *and learn* about the comparative pros and cons of the available choices.

Our analyses showed that, in the visualisation study, the vast majority of significant results were obtained in the explicit alternative revelation condition. Under these conditions, where participants were both exposed to alternatives *and taught* about each's pros and cons, we saw significant shifts in (among other things) perceptions of source trust, perceptions of source credibility, and audience intent to return to source. In the summarisation study, we not only saw a myriad of significant results in the explicit alternative revelation condition, but also in the *implicit* one, as well. Here, the simple act of becoming aware of alternative summarisation techniques led to significant shifts in (among other things) perceptions of trust, desire to see more data summarised in this fashion, and audience intent to return to source.

These results have two important implications. First, regardless of whether it's choosing between visualisation formats or summarisation techniques, the choice always seems to carry significant consequences for a source if audiences *come to learn* about the advantages and disadvantages of options

within the available choice set. In both studies, significant main effects of presentation style emerged across a variety of outcome variables for participants in the explicit condition, meaning that participants come to actively differentiate between sources based on their use of these options, both forming impressions and rendering judgments on this basis.

Secondly, audiences seem to be even more sensitive to differentiating between data *summarisation* choices than data visualisation choices. Even when participants were not explicitly taught about the pros and cons of each summarisation technique, and were instead merely *exposed* to alternative options, significant main effects of summarisation format emerged across several metrics. This was particularly the case for sources that chose to use generic statements, as they experienced statistically significant reductions in perceptions and future intended action on five out of six outcome variables. Consequently, when it comes to summarisation choices, sources must recognise that audiences will likely come to judge them based on the decisions they make (as long as they have an awareness of the other summarisation options available).

While such findings were encouraging, especially as significant main effects of data presentation were observed on measures like trust and credibility, which have routinely been implicated as factors crucial to the maintenance of a source's reputation and influence (e.g., Baum & Gussin, 2007; Stroud & Lee, 2013), it remains to be seen whether such a set of results could precipitate market pressures capable of motivating change (or at least greater conscientiousness) within major outlets. Overall, we remain sceptical that such results will inspire widespread change without an additional accelerant such as an indignant and sweeping public outcry against questionable presentation practices or the prospect of swift and severe financial repercussions for their continuation.

Individuals Can Be Inoculated Against Truncation-Induced Polarisation Effects

One of the more robust findings across the studies within Chapter 1 and Chapter 2 concerned the deleterious impact *truncated* bar charts had on perception accuracy relative to other modes of presentation. While a more systemically- or institutionally-focused intervention would be a preferable means by which to deter outlets from engaging in such suboptimal data presentation practices (for the reasons elucidated in Chapter 4; e.g., superior scope and ease of scalability), it would be naïve to believe that such outlets would be so easily persuaded to alter their strategies (especially if more “responsible” data presentation comes at the expense of engagement or revenue).

Fortunately, Chapter 5 demonstrated the feasibility of inoculating individuals against the type of truncation-induced polarisation effects observed in Chapters 1 and 2. More specifically, a brief, approximately five-minute intervention proved to be an effective means of moderating perceptions of polarisation for (critically) groups depicted via truncated bar chart as well as (unexpectedly) groups depicted via full-range bar chart. The intervention also had the effect of generating divergent perceptions of reliability for sources which

chose to use one chart type over the other, comparatively “punishing” sources (via diminished reliability ratings) that chose to engage in data disseminations via truncated bar charts.

This is an important set of results for two reasons. Firstly, while recent work has shown that technique-based inoculation can reduce the perceived reliability of content which aims to polarise (Harrop, Roozenbeek, Madsen, & van der Linden, 2022), to the best of our knowledge, ours is the first intervention which has sought to directly test whether technique-based inoculation can operate as a mechanism to counter one’s tendency to polarise as a response to misleading data visualisation tactics. Considering the success of this proof-of-concept, it both provides practitioners with another tool with which to combat misinformation as well as opens the door for future researchers to investigate how inoculation interventions may provide similar “protection” against questionable presentation practices. Secondly, it provides those interested in mitigating the negative impact of questionable presentation practices with a viable “safety-net” option. While our hope would be that the results obtained in Chapter 4 would motivate some media outlets (and other news sources) to adopt more responsible data presentation practices, such a hope is admittedly idealistic. It’s clear that some sources rely on the strategic use of sensationalist claims and division narratives to drive audience retention and thus are unlikely to be swayed by the implications of our work. However, considering the fact that the inoculation precipitated a marked decline in perceptions of reliability for sources choosing to use truncated bar charts, it’s possible (though admittedly speculative) that – if the intervention were ever deployed widely enough – exposure could organically promote a decline in readership/viewership of those engaging in such manipulative visualisation tactics, creating a shift in market dynamics which could force them to reevaluate the continued use of these methods.

Ultimately, in the absence of the voluntary cooperation of influential sources (who could help facilitate a systemic solution), it’s important to have an individual-focused alternative which can help to promote a more informed, discerning, and ultimately less vulnerable public. We believe this inoculation intervention – while by no means perfect – represents a promising first step towards achieving that end.

Shifting Perceptions of Ideological Polarisation Might Not Exert Meaningful Influence on Levels of Affective Polarisation or Interparty Behaviour

While the results obtained throughout Chapter 1 and Chapter 2 demonstrated that type of visualisation was a variable consistently capable of shifting perceptions of ideological polarisation, its effect on affective polarisation and a variety of intergroup behavioural measures was often non-significant. Additionally, when the effect did promote statistically-significant differences in these outcome variables, it did so with less consistency and with considerably smaller effect sizes.

Such a finding is useful, but also perhaps intuitively unsurprising. The data visualisations employed throughout the first chapter of studies ultimately provided different perspectives through which to view the

ideological positions of political partisans on specific policy issues. While these alterations in the way ideological disagreements are depicted may naturally motivate a participant to reconsider just how divided opposing groups are, it may not necessarily inspire a reconsideration of *how they feel* about the other party (or how they'll choose to behave towards them). This is perhaps because the relationship between ideological and affective polarisation, while not necessarily completely independent from one another (e.g., Abramowitz & Webster, 2016), has been proposed to be tenuously related at best (e.g., Iyengar, Lelkes, Levendusky, Malhotra, & Westwood, 2019).

The general dynamic described above – wherein a *belief* is substantially shifted but where the underlying *feeling* and intended *action* seem to remain the same – is reminiscent of work conducted by researchers like Swire, Berinsky, Lewandowsky, and Ecker (2017), who demonstrated that participant levels of belief in factually incorrect statements attributed to Donald Trump could be reduced via interventions which identified the claims as false and provided corrections. However, in spite of the shifted *beliefs* (which would presumably undermine Trump's credibility), the researchers found that Trump's supporters indicated no change in their *feelings* towards Trump or their *intention to vote* for him. This pattern of results was broadly replicated (this time utilising a bipartisan design) by Swire-Thompson, Ecker, Lewandowsky, and Berinsky (2020) in a paper revealingly titled *They Might be a Liar, but They're My Liar*. Using a sample which included both Trump supporters as well as supporters of Bernie Sanders, the authors found that while participants were amenable to belief corrections which called into question the veracity of their preferred candidate, these corrections generally did not precipitate a reduction in support.

General findings within the misinformation literature – which seem to mirror the results observed throughout our studies – were nicely summarised by Ullrich Ecker, who wrote “people's misconceptions can be reduced by fact checking, but it's harder to make people change their feelings towards the candidate or to change their voting intentions” (Ecker, 2017, p. 85). Similarly, our work tends to corroborate this hierarchy of malleability, with several of our interventions proving capable of substantially shifting participant *perceptions* and *beliefs*, but typically struggling to elicit meaningful movement on partisan affect or intergroup behaviour.

Ultimately, it may simply be that ideological and affective polarisation are, at least to a degree, distinct in their development and subsistence and subsequently distinct in the mechanisms through which they might be moderated, and perhaps our interventions have only managed to effectively target one of these domains reliably. For example, recent research suggests that partisan *identity* is a more reliable determinant of affective polarisation and interparty animus than divergence on policy positions (Dias & Lelkes, 2022). Although our interventions may be able to shift perceptions of *how far apart* the groups are from one

another, it likely cannot meaningfully alter the foundations of one's own political identity nor their beliefs concerning the identity of political "others."

Directions for Future Research

At the end of each study, we've attempted to provide thoughtful critiques of factors which may have impacted our outcomes, such as the specifics of our experimental approaches or the choice of which dependent variables to use. These critiques were posited with an eye towards suggesting ways in which future researchers may constructively build upon our work, both via improving designs as well as addressing existing limitations and shortcomings. While we hope these sets of observations and suggestions will broaden one's perspectives concerning the theoretical and methodological considerations germane to this work, we also wish to suggest a set of unresolved questions which we believe can both motivate and guide future investigations.

Does Shifting Perceptions on Specific Types of Interparty Data Reduce Affective Polarisation?

It's important to acknowledge that the minimal movement observed on measures of affective polarisation following manipulations of data presentation could be an inherent limitation of these types of data presentation interventions. For example, it may simply be that shifting perceptions of issue-specific ideological polarisation may not meaningfully contribute to the moderation of affective polarisation. Such a finding would broadly align with recent research which suggests that correcting misperceptions of political opponents has little effect on reducing partisan animosity or support for partisan violence (Dias et al., 2024). While there are critical ways in which affective and ideological polarisation may interact (Abramowitz & Webster, 2016), such a finding would be broadly compatible with much of the extant literature, as many researchers see the two forms of polarisation as distinct; both unique in their origins and independent in their changes over time (e.g., Iyengar et al., 2019). In fact, certain research has even shown that extreme divergence on policy issues is not a necessary condition for affective polarisation (e.g., Mason, 2015) and that levels of affective polarisation may *rise* even as ideological divisions fall (e.g., Levendusky & Malhotra, 2016).

However, should future researchers seek to better understand the circumstances under which data presentation interventions *might* elicit significant reductions in levels of affective polarisation, we would recommend investigating two possible pathways via which change might occur.

The first involves better understanding the potential mediating role perceptions of cooperation and compromise might have on levels of affective polarisation. Recent work by Bersoff (2024) suggests that an underexplored mechanism by which we might reduce affective polarisation is by increasing cross-party confidence in the possibility of dispute resolution. Bersoff hypothesises that it is the perceived *intractability*

of political disagreement (and the corresponding despondency it promotes) which may facilitate heightened levels of partisan animosity. However, he suggests that changing this perception, and inspiring renewed hope in our ability to constructively work through political challenges, may yield reductions in affective polarisation. Such a conjecture is promising for our purposes as several studies reviewed within this thesis have demonstrated how mode of presentation can significantly impact beliefs about the likelihood of compromise. While shifting these beliefs alone did not consistently result in corresponding reductions of affective polarisation, if we were to accept that perceptions of compromise likelihood are a likely *precursor* to beliefs about the possibility of dispute resolution, then data presentation interventions, due to their ability to increase optimism about the prospect of intergroup compromise, may be able to significantly impact a potentially key mediator of affective polarisation.

The second involves determining whether the *type* of data being visualised or summarised might play a pivotal role in which facets of polarisation are affected. Throughout the current work, due to our primary focus being on trying to shift perceptions of ideological polarisation, we've chosen to visualise and summarise data on *political issues*. However, considering the inconsistent relationship between affective and ideological polarisation elucidated above, learning about where one's political group and the opposing group stand on policies – even if it's closer than they originally believed – may not make much of a difference in how they ultimately *feel* about their political opponents. Affective polarisation is more than just a byproduct of political disagreement; it's a manifestation of how one *feels towards* members of the opposing party and what they *believe about* their character (e.g., Iyengar, Sood, & Lelkes, 2012; Levendusky & Malhotra, 2016). Thus, if it is possible to change affective polarisation via data presentation interventions, it might be that we need to consider different *types* of data to summarise or depict. For example, while learning that members of the opposing party are not as far apart from you on *policy* issues may not fundamentally shift how you feel about them, learning that you are closer on personal values or life priorities may.⁷⁵ Future researchers should consider this possibility and devise studies to test it accordingly.

Does Shifting Perceptions of Ideological Polarisation Lead to Reliable Shifts in Any Interparty Behaviours?

Across the vast majority of our research, we found minimal evidence to support the belief that using data presentation techniques to shift perceptions of ideological polarisation can lead to corresponding shifts in interparty behaviour. Although glimpses of hope for such a relationship could be seen in results such as how the valence of sentiment in response to an optimistic political statement was significantly altered depending

⁷⁵ Researchers should be aware, however, that while there are likely personal values for which significant cross-party similarity can be found, discrepancies in certain types of values between opposing party members may indeed be even *more* pronounced than their policy positions. For example, measures of authoritarian values – a preference for strict adherence to ingroup norms and punishment for those who violate them (Altemeyer, 1996) – were found to be exceptionally different among those who voted “Leave” versus “Remain” in the 2016 British EU Referendum (i.e., “Brexit”; Undzenas, Dunn, & Spaiser, 2022).

on the data visualisation method to which participants were exposed (see Chapter 2), the overwhelming majority of behavioural dependent variables tested failed to reject the null hypothesis.

Similar to the direction proposed in the preceding section, should future researchers wish to continue to explore the possibility of behaviourally-relevant repercussions of these interventions, we once again hypothesise intergroup compromise and cooperation to be the most promising behavioural domain to examine in future experimental iterations. As discussed, shifting the manner in which intergroup data was visualised demonstrated the ability to significantly impact how likely groups believed cross-party compromise to be (see Table 2 below).

Table 2

Range of Effect Sizes Between Visualisation Formats on Beliefs about Likelihood of Compromise

| Study | Effect Sizes (Cohen's <i>d</i>) | | | |
|--------------------|--------------------------------------------------|-----------------------------------------------------|--------------------------------------------------|-----------------------------------------------------|
| | <i>No Visualisation and Full-Range Bar Chart</i> | <i>Truncated Bar Chart and Full-Range Bar Chart</i> | <i>No Visualisation and Icon Array Histogram</i> | <i>Truncated Bar Chart and Icon Array Histogram</i> |
| Chapter 1, Study 1 | NA | 0.93 ¹ | NA | 0.99 ¹ |
| Chapter 1, Study 2 | NA | 1.13 | NA | 1.06 |
| Chapter 1, Study 3 | 1.21 | 0.96 | NA | NA |
| Chapter 1, Study 4 | NA | NA | NA | NA |
| Chapter 2 | 0.36/0.39 ² | 0.39/0.40 ² | 0.34/0.35 ² | 0.37/0.35 ² |

*Note*₁. Two forms of truncated bar charts were used in Study 1. The data above refers to the 1.5 SD bar chart (as this was the standard used in subsequent studies, rendering it a more appropriate cross-party comparison metric).

*Note*₂. The two statistics per cell for Chapter 2 represent the UK data and US data, respectively.

Consequently, if we were to operate under the assumption that increases in *perceptions* of compromise likelihood *should* result in corresponding increases in compromise *attempts* (a hypothesis which, of course, would need to be empirically validated before being accepted), we would strongly suggest that researchers attempt to test behavioural dependent variables which are situated at the intersection of intergroup compromise and cooperation. We believe that, if these types of interventions are indeed capable of exerting meaningful effects on behaviour, these are the two most likely domains in which such changes would be observed.

Will Similar Effects Emerge in Dissimilar Samples?

The present collection of research, like many academic investigations before it (e.g., Gurven, 2018; Apicella, Norenzayan, & Henrich, 2020), likely suffered from a dearth of sample diversity. While a portion of our research did feature a large-scale international replication (i.e., Chapter 2), using demographically-

representative samples from both the US and the UK, the vast majority of our work was derived from samples collected via platforms like Prolific, leaving the generalisability of our results beyond these populations uncertain.

While Prolific's participant pool may be more diverse and offer higher-quality responses than its competitors (e.g., Amazon MTurk; see Peer, Brandimarte, Samat, & Acquisti, 2017), it still likely suffers from self-selection biases, such as the fact that participants who sign up *want* to participate in research. Moreover, even if we were to assume such platforms could offer perfectly representative samples devoid of any such self-selection biases, our work still remained largely limited from a geographic standpoint, as the vast majority of our studies recruited American participants only. When the research extended beyond the confines of the US (as it did in Chapter 2), it still only sampled from one other country, the United Kingdom, and it's fair to wonder just how different these samples are considering that they both are likely to predominantly fall under the WEIRD umbrella (e.g., Apicella et al., 2020).

However, while we hope it does not discourage future researchers to pursue such ends, it is worth noting that extending this research into substantively different cultural, geographic, and sociopolitical contexts comes with notable challenges. For example, we initially hoped to carry out a portion of the replication in India, but following a conversation with experts in that region familiar with this type of undertaking, we were advised that recruiting a suitable participant pool might be difficult, as Indian media outlets might not be disseminating data on where voters stand on polarising political issues as frequently or as transparently as outlets in the West. While we did manage to secure a convenience sample from the country, we worried it skewed too heavily towards educated portions of the population (not to mention suffered from the aforementioned self-selection biases) to be useful.

Should future researchers more capable, familiar with, or well-connected than ourselves manage to secure usable samples, it will be interesting to see not only how different cultural environments may moderate data presentation effects, but also how factors like differences in political systems and non-dichotomous party structures might influence the results.

Concluding Thoughts

Political polarisation remains one of the most pressing geopolitical issues of our time. In the United States, scholars have suggested that pervasive and longstanding misperceptions within the political arena might contribute to the continuation (and perhaps exacerbation) of these political divisions. Our research posed a

simple question: could the manner in which we present data about political partisans play a role in moderating political perceptions?

Our work has demonstrated that, far from being trivial design or dissemination features, the way in which we depict and describe political data can have significant impacts on how individuals come to perceive the state of political polarisation.

Such a finding is simultaneously both empowering and concerning. It is empowering because this means a simple, powerful, and highly-scalable depolarisation tactic is already available to us: being more deliberate and conscientious about the way we present data. Sometimes, something as simple as shifting from depicting data via a truncated bar chart to instead opting to depict the same data via a full-range bar chart or icon array histogram can be enough to significantly shift how polarised individuals perceive an issue to be – even for issues with a protracted history of partisan hostility. But it is concerning because it means that agents and agencies may either inadvertently (due to obliviousness about the consequences of data presentation) or intentionally (due to apathy, conflicting priorities, or ill-intent) continue to bolster political misperceptions should they continue the use of questionable presentation practices.

However, we want to make clear that the “questionable presentation practices” we’ve referenced throughout this thesis refer specifically to the data presentation choices made when depicting *political rivals in a highly-polarised political ecosystem*. This is not to say that truncated bar charts, for example, cannot be considered a questionable presentation practice outside of this context, but all we are confident declaring definitively as a result of our particular line of research is that, when alternate visualisation options exist, truncated bar charts appear to usually be an inappropriate choice within it.⁷⁶

Ultimately, we believe that equipping a bar chart with a full-range y-axis is a best practice to which practitioners should widely adhere. It is a prudent default and, in terms of producing technically-sound visualisations, is usually the safest choice. However, beyond perhaps the circumstances cited above (i.e., depicting data on political rivals in a highly-polarised political ecosystem), we do not consider ourselves staunch anti-truncation absolutists; we leave ourselves open to the possibility that, in certain scenarios (especially ones where pervasive misperceptions do not already exist), truncating a y-axis might be the correct (or at least an ethically-defensible) design choice.

For example, we could imagine a hypothetical scenario where a highly-contagious and deadly virus is infecting large swaths of a population. A drug manufacturer creates a drug which, while not extraordinarily effective, is significantly more likely to result in an individual surviving the virus once contracted than if they had not taken the drug. However, “significantly” more effective, especially across a large enough sample,

⁷⁶ However, as we saw in the latter studies of Chapter 1, even a truncated bar chart, while perhaps sub-optimal relative to other modes, may be a preferable means of disseminating political data to another alternative: providing individuals with *no visualization at all*.

could be mere fractions of a percentage point – a difference that is meaningful to (and well understood by) those with a certain degree of statistical knowledge, but which is likely to be underappreciated by a less statistically-savvy audience (which is arguably the majority of many populations; see Gaissmaier & Gigerenzer, 2008). Consequently, for groups low in numeracy or graphical literacy, depicting the survival probability data for those who take the drug and those who do not as a full-range bar chart would produce a visualisation that would likely create the impression that the two options are “virtually the same,” undermining motivation to acquire the potentially life-saving medicine. In such a situation, a degree of truncation might be warranted as it could assist the public in developing an appropriate sensitivity to the impact that such seemingly small differences, once aggregated, might have on health outcomes.⁷⁷ Ultimately, this is similar to the point being made by Witt (2019). Although the recommendations put forth in that article (e.g., truncating a y-axis to 1.5-2.0 SD of the grand mean) proved to be damaging (relative to alternative visualisation choices) in the domain of political perceptions, we have reason to believe that they are entirely appropriate under different circumstances.

Although truncating the y-axis has proven to be a decidedly questionable presentation practice relative to other visualisation alternatives when depicting data on political rivals in a highly-polarised political ecosystem, the good news is that our research has also provided preliminary evidence for two interventions – one systemic, one individual – which show promise in disincentivising the use of these practices (for the former) and mitigating their harmful impact (for the latter).

Political data presentation is a hammer: a tool in the hands of some, a weapon in the hands of others. We sincerely hope our work may serve as a blueprint for those who wish to build the foundations for a less polarised world as opposed to those who wish to chip away at its crumbling pillars.

⁷⁷ One could make the argument that, in such a circumstance, perhaps it might be more appropriate to use a different visualisation format rather than truncate the y-axis. We do not dispute this contention. However, if a media outlet has decided that it will use a bar chart as the visualisation vehicle (for instance, because it believes it is the format for which the majority of the general public will most easily comprehend), then we stand by the assertion that truncation could be justified.

APPENDICES

Appendix A: The Construction and Usage of Icon Array Histograms (Introduction, Chapter 1, and Chapter 2)

Icon array histograms are essentially superimposed histograms constructed by using icons as the building blocks. Our icon array histograms relied primarily on “paired icons,” which are two “restroom icons⁷⁸,” joined side-by-side to represent the opinion of two separate respondents. For each bin, we first make as many cross-party pairs as possible. For example, if ten members of Group A gave a response of “1” on a Likert scale, and six members of Group B gave that same response, then we would be able to make six cross-party pairs. This would leave us with a surplus of four Group A respondents in that response bin, who would then become two *same-party* pairs. Thus, that response bin would contain six cross-party pairs and two same-party pairs, representing all 16 individuals who offered that response. This matching pairs system permits icon array histograms to highlight “overlap” (i.e., where members of two groups are providing the same response) more effectively than traditional modes of presentation (e.g., bar charts), which often offer minimal distributional information, while also representing the true distribution of responses (which may otherwise be skewed by relying on normal distributions) in a faithful and accurate manner. Moreover, by reducing the total number of respondents depicted in each group to 100, the total number of “overlapping” pairs (i.e., cross-party pairs) will also correspond precisely to the PCS (or “percentage of common scores”) of the two groups, making this mode an efficient and intuitive way of calculating group agreement with minimal graphical or statistical sophistication, as well.

While icon arrays, due to their comparative novelty, *do* pose the risk of being more difficult to comprehend or interpret than more traditional data visualisation formats (e.g., bar charts), pilot testing indicated that individuals seemed to have a fairly good intuitive sense of what they were depicting (e.g., a representative, modal response to the question “What is this visualisation showing?” was “how much the groups agree on the issue and how much they disagree”). Moreover, considering the fact that it routinely promoted political perceptions whose *accuracy* was among the best of the formats studied, it seems that it is providing a visual cue that is, at the very least, *useful* in correcting perceptions. However, the mechanisms by which it accomplishes this – both in isolation as well as relative to other formats – remain unclear, so should the format continued to be used as an alternative means of data dissemination, it would be prudent for future research to focus on better understanding how individuals perceive and relate to the visualisation as it is presented. We believe specific attention should be paid to better understanding if (and how) individuals

⁷⁸ One of the more niche findings within the icon array literature concerns the type of icon used. Zikmund-Fisher et al. (2014) tested a number of icon types (e.g., blocks, ovals, etc.) to determine whether the choice of symbol used plays a role in how the information is processed and retained. Overall, icon type was found to play a significant role in perception, retention, and recall of risk, with “restroom icons” – the minimalist, gendered silhouettes traditionally used to differentiate men’s restrooms from women’s – to be especially promising.

comprehend what precisely the visualisation depicts (e.g., do they have a strong grasp of what each of the “paired” icons – particularly the purple ones – represent?) as well as determining *how* individuals interpret and use the mode to inform their decisions (e.g., is it used in a deliberative way, such as counting “agreement pairs,” or used more to get a general “gist” of the overall level of agreement and disagreement?).

Our approach throughout the work was to take a minimalist approach to the presentation of the icon array format to participants (i.e., not formally explaining what it was or providing anything beyond the legend for information as to what it depicts). We did this deliberately, with an eye towards examining how the visualisation would fare “in the real world,” where publications rarely have time (or space, in the case of print publications) to explain how one should properly interpret a visualisation prior to deploying it. However, should researchers look to refine and optimise the format in the future, we suggest a few possible areas to explore. First, one should consider whether the legend, as it currently stands, is optimal. Ultimately, the legend needs to be concise while also containing enough information so that most individuals can properly decipher what they’re seeing. We did our best to strike a balance between succinctness and informativeness with our phrasing (e.g., “Republican and Democrat with the same opinion”), but we’re not confident that all individuals are interpreting this phrase in the same manner. Moreover, while we believe the visual “proportions” of red to blue to purple icons makes interpreting the *relative* frequency of one group to another an intuitive exercise for many, it’s not clear whether participants truly understand what each of the building blocks of the visualisation (i.e., the “paired” icons) mean. Consequently, further investigation is needed to understand participant thought processes following exposure here. Additionally, the impact of the graphic element choices of the visualisation format, such as colour and icon type, should be examined. For the former, while we do believe using colours traditionally associated with each political party to be appropriate, it remains to be seen whether making the “agreement pairs” purple is the best methodology. For example, one could also imagine a format wherein the agreement pairs were constructed of an individual red and blue icon (which might be a more natural depiction of agreement for some). Additionally, considerations should be made for colour blind participants, attempting to choose colours between which this sub-sample can easily distinguish. For the latter, while the “restroom icons” have been found to be a useful tool for icon array construction, we have used only “male” icons in our constructions. Future research should make efforts to understand how utilising different combinations of restroom icons (e.g., female only, mixed) might moderate perception.

Appendix B: The Appropriateness of Different Measures of Similarity for Different Data Visualisations (Chapter 1)

While PCR and PCS are very highly correlated (i.e., $r_s \geq .96$; see Hanel et al., 2019), and thus can typically serve as interchangeable measures of similarity depending on practitioner preference, we believe there are specific circumstances in which the use of one versus the other is more or less appropriate.

PCR is calculated by using Cohen's d and relies on assumptions of normality in the distribution of responses. For example, two normally distributed groups with a Cohen's d of 0.80 would have a PCR of 69%. However, in spite of their high correlation, there are times when the PCR and PCS may substantially diverge. One such occasion just happened to be in the artificial datasets constructed for use throughout Chapter 1. Each dataset was deliberately built in such a way so that they would be normally distributed and have a Cohen's d of exactly 0.80, however, in spite of this, the PCS was 59% (10 percentage points lower than the PCR).

Because PCR is essentially calculating the degree to which the areas underneath two normally-distributed curves overlap (which will depend on how "far apart" their means are from one another; i.e., the Cohen's d), it is almost always an *estimate* of actual overlap. This is because – unless respondents are given an infinite number of response options – the PCR is "transforming" discrete data into continuous data, "smoothing" what would otherwise be separate bins of responses. Conversely, PCS is calculated by going "bin-by-bin" measuring overlap, meaning A) it can produce a more precise measure of similarity for any two datasets that have used discrete response options (i.e., it is not "estimating" in the manner PCR does as described above), and B) critically, it can account for *non*-normally distributed data. Extremes can be important, perhaps especially when examining political compositions, and normal distributions do not show this faithfully.

However, most relevant to the current work is certain measures of similarity may be more or less *appropriate* accuracy benchmarks for different types of data visualisations. For example, in Study 1 and Study 2 from Chapter 1, what the icon array histograms "show" is the PCS – the number of overlapping pairs is exactly 59 (out of 100 total pairs) – whereas what the superimposed normal distributions "show" is the PCR (i.e., the overlapping area is 69% of the total. Consequently, it would be inappropriate to use PCS as a benchmark of accuracy for superimposed normal distributions or PCR as a benchmark of accuracy for icon array histograms, as each visualisation depicts a different measure of similarity.

Appendix C: Chapter 1, Study 1 Additional Hypotheses and Results

While Hanel et al. (2019) primarily investigated how mode of presentation affects perceptions of intergroup similarity, we were interested in determining whether such effects would also extend to more politically-relevant perceptions and dynamics (e.g., perceptions of polarisation, beliefs about compromise likelihood, etc.). Although primarily exploratory in nature, the literature is not devoid of precedent to suggest that perceptions of similarity may ultimately exert downstream impacts on political preference. For example, studies have demonstrated a relationship between perceptions of similarity and one's support for wealth redistribution policies (e.g., Senik, Stichnoth, & Van der Straeten, 2009; Ordabayeva & Fernandes, 2017).

Consequently, it may not be unreasonable to assume that shifts in perceptions of similarity, driven by mode of presentation interventions, may also lead to corresponding shifts in other politically-salient items. Specifically, based on the results from Study 3 and Study 5 in Hanel et al. (2019) – and mirroring the rationale described above – we hypothesised that the two comparatively “novel” modes (i.e., superimposed normal distributions and icon array histograms), which depicted the full range of intergroup responses (including overlap) would promote more “positive” political perceptions and expectations than identical data depicted using a maximally-truncated bar chart.

Measures (cont.)

Perceived Issue Polarisation: Participants were asked to indicate their response (via a 101-point scale that ranged from “Not at All Polarising” to “Extremely Polarising”) to the item “Based on the graph, how polarising do you believe this issue is for the two groups?”

Perceived Affective Polarisation: Participants were asked to indicate their response (via a 101-point scale that ranged from “Extremely Cold/Negative” to “Extremely Warm/Positive”) to the item “Based on the graph, how warm (i.e., positive) or cold (i.e., negative) do you think members of each group would feel towards members of the other group on this issue?”

Willingness to Cooperate: Participants were asked to indicate their response (via a 101-point scale that ranged from “Not at All Willing” to “Extremely Willing”) to the item “Based on the graph, how willing do you think members of each group would be to try to cooperate with one another on this issue?”

Optimism about Compromise: Participants were asked to indicate their response (via a 101-point scale that ranged from “Not at All Optimistic” to “Extremely Optimistic”) to the item “Based on the graph, how optimistic are you that the two groups could find ways to compromise on the issue?”

Trust about Fairness: Participants were asked to indicate their response (via a 101-point scale that ranged from “Can’t Trust at All” to “Can Trust Completely”) to the item “Imagine that you were a member of one of the groups, and the other group was given the ability to make decisions on this issue. Based on the graph, how much do you think you could trust the other group to make decisions that would be fair to both sides?”

Hypotheses (cont.)

Table 1

Additional Hypotheses for Study 1

| Label | Hypothesis |
|-------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H4 | The modes of presentation which display the full range of responses (i.e., superimposed normal distributions and icon array histograms) will elicit more “positive” intergroup perceptions across the five secondary DVs (i.e., higher perceptions of willingness to cooperate, greater optimism about compromise, lower perceived issue polarisation, lower affective polarisation, and higher trust about fairness) than the maximally-truncated bar chart. |

Results (cont.)

A series of t-tests determined, as hypothesised [H4], that the modes of presentation which displayed the full range of group responses (i.e., superimposed normal distributions and icon array histograms) did promote responses which were significantly “more positive” than the maximally-truncated bar chart across the five remaining dependent variables (see Table 2).

Table 2

Results of Pairwise T-Tests between Maximally-Truncated Bar Charts and the “Full-Range” Modes Across Five Dependent Variables

| Dependent Variable | Results of Pairwise T-Tests | |
|----------------------------------|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| | <i>Maximally-Truncated Bar Chart versus Superimposed Normal Distribution</i> | <i>Maximally-Truncated Bar Chart versus Icon Array Histogram</i> |
| Willingness to Cooperate | $M_{TBC} = 42.31, SD_{TBC} = 25.42$ | $M_{TBC} = 42.31, SD_{TBC} = 25.42$ |
| | $M_{SND} = 66.83, SD_{SND} = 18.35$ $t(246.81) = 9.220^{***}, p < .001, d = 1.11$ | $M_{IAH} = 56.90, SD_{IAH} = 18.92$ $t(233.98) = 5.739^{***}, p < .001, d = 0.67$ |
| Optimism about Compromise | $M_{TBC} = 43.36, SD_{TBC} = 25.88$ | $M_{TBC} = 43.36, SD_{TBC} = 25.88$ |
| | $M_{SND} = 71.17, SD_{SND} = 20.29$ $t(257.78) = 9.981^{***}, p < .001, d = 1.20$ | $M_{IAH} = 61.54, SD_{IAH} = 18.67$ $t(228.68) = 7.084^{***}, p < .001, d = 0.83$ |
| Perceived Issue Polarisation | $M_{TBC} = 55.84, SD_{TBC} = 26.48$ | $M_{TBC} = 55.84, SD_{TBC} = 26.48$ |
| | $M_{SND} = 39.82, SD_{SND} = 21.95$ $t(264.28) = 5.500^{***}, p < .001, d = 0.66$ | $M_{IAH} = 45.02, SD_{IAH} = 18.91$ $t(226.98) = 4.129^{***}, p < .001, d = 0.49$ |
| Perceived Affective Polarisation | $M_{TBC} = 60.72, SD_{TBC} = 21.12$ | $M_{TBC} = 60.72, SD_{TBC} = 21.12$ |
| | $M_{SND} = 40.08, SD_{SND} = 18.15$ $t(278) = 8.781^{***}, p < .001, d = 1.05$ | $M_{IAH} = 48.10, SD_{IAH} = 15.47$ $t(231.15) = 6.000^{***}, p < .001, d = 0.70$ |
| Level of Trust about Fairness | $M_{TBC} = 36.36, SD_{TBC} = 22.42$ | $M_{TBC} = 36.36, SD_{TBC} = 22.42$ |
| | $M_{SND} = 59.59, SD_{SND} = 19.04$ $t(278) = 9.359^{***}, p < .001, d = 1.12$ | $M_{IAH} = 50.26, SD_{IAH} = 20.26$ $t(340) = 5.956^{***}, p < .001, d = 0.66$ |

* $p < .005$, ** $p < .001$, *** $p < .0001$

Note. Bonferroni corrections have been applied to significance thresholds.

Appendix D: Chapter 1, Study 2 Additional Hypotheses and Results

Table 1

Additional Pre-Registered Hypotheses for Study 2⁷⁹

| Label | Hypothesis |
|-------|------------|
|-------|------------|

⁷⁹ Although these hypotheses were not explicitly labeled as exploratory in the pre-registration, they were intended to be treated as such – providing us with an opportunity to test how some peripherally-related political psychology phenomena (such as conservative loyalty, deference to authority, and disapproval of social norm violation; e.g., Haidt & Graham, 2007; Feldman & Weber, 2016).

| | |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H6 _A | For Republican participants, the effect size difference in responses to the compromise variable for those in the full-range bar chart condition and those in the truncated bar chart condition will be greater when the data is attributed to elites than when attributed to citizens. |
| H6 _B | For Republican participants, the effect size difference in responses to the compromise variable for those in the superimposed normal distributions condition and those in the truncated bar chart condition will be greater when the data is attributed to elites than when attributed to citizens. |

Results (cont.)

For Republican respondents, differences in effect size between responses elicited by the truncated bar chart and the full-range bar chart on the compromise variable were not greater when the data was attributed to elites ($d = 0.83$) than to citizens ($d = 1.24$) [H6_A]. Similarly, for Republican respondents, differences in effect size between responses elicited by the truncated bar chart and the superimposed normal distributions on the compromise variable were not greater when the data was attributed to elites ($d = 0.82$) than to citizens ($d = 0.93$) [H6_B].

Appendix E: Chapter 1, Study 3 Additional Hypotheses and Results

Table 1

Additional Pre-Registered Hypotheses for Study 3

| Label | Hypothesis |
|-----------------|--------------------------------------------------------------------------------------------------------------------------|
| H6 | Main effects of mode of presentation will be observed on the compromise measure. |
| H6 _A | The truncated bar chart will elicit lower perceptions of compromise optimism than the full-range bar chart. |
| H6 _B | The truncated bar chart will elicit lower perceptions of compromise optimism than the superimposed normal distributions. |


Results (cont.)

Exploratory pairwise t-tests were conducted between the no visualisation control and each of the three experimental mode of presentation conditions on perceived optimism about compromise. These analyses determined that participants assigned to the no visualisation control exhibited significantly lower perceptions of compromise optimism ($M = 29.95$, $SD = 22.09$) post-Bonferroni corrections than those assigned to the truncated bar chart condition ($M = 35.74$, $SD = 23.48$; $t(610) = 2.707$, $p = .007$, $d = 0.25$), the full-range bar chart condition ($M = 58.72$, $SD = 24.23$; $t(602) = 13.089$, $p < .001$, $d = 1.21$), and the superimposed normal distribution condition ($M = 60.29$, $SD = 24.89$; $t(630) = 13.606$, $p < .001$, $d = 1.25$). A 3x3 ANOVA was conducted to assess the main and interaction effects of mode of presentation and source on compromise optimism across all experimental conditions, which indicated a significant main effect of mode [H6] ($F(2, 1368) = 147.420$, $p < .001$, $\eta_p^2 = 0.177$; see Figure 1) but no significant main effect of source ($F(2, 1368) = 0.361$, $p = .697$, $\eta_p^2 = 0.001$) nor a significant interaction between the two variables ($F(4, 1368) = 1.198$, $p = .310$, $\eta_p^2 = 0.003$). As hypothesised, truncated bar charts elicited significantly lower perceptions

of compromise optimism than full-range bar charts [H6_A] ($t(900) = 14.464, p < .001, d = 0.96$) or superimposed normal distributions [H6_B] ($t(928) = 15.460, p < .001, d = 1.01$).

Figure 1

Perceptions of the Likelihood of Intergroup Compromise by Mode




* $p < .05$, ** $p < .01$, *** $p < .001$

Note. Only significant differences between the no visualisation condition and each of the three experimental conditions has been plotted.

An exploratory 3x3 (mode by source) ANOVA was conducted to assess main and interaction effects of mode of presentation and source on level of intergroup agreement relative to expectations (see Figure 2). The ANOVA revealed a significant main effect of mode (truncated bar chart ($M = 3.21, SD = 0.98$), full-range bar chart ($M = 4.01, SD = 0.93$), superimposed normal distributions ($M = 4.10, SD = 0.89$); $F(2, 1368) = 125.942, p < .001, \eta_p^2 = 0.155$; see Figure 18) but no significant main effect of source (Fox News ($M = 3.77, SD = 1.01$), CNN ($M = 3.78, SD = 1.00$), no source ($M = 3.78, SD = 1.04$); $F(2, 1368) = 0.138, p = .871, \eta_p^2 = 0.000$) nor a significant interaction between the two variables ($F(4, 1368) = 0.198, p = .939, \eta_p^2 = 0.001$).

Figure 2

Perceptions of Level of Intergroup Agreement Relative to Expectations by Mode




* $p < .05$, ** $p < .01$, *** $p < .001$

Note. Significance levels featured are based on the results obtained from Tukey HSD post-hoc tests.

An exploratory 3x3 (mode by source) ANOVA was conducted to assess main and interaction effects of mode of presentation and source on perceptions of group discreteness (see Figure 3). The ANOVA revealed a significant main effect of mode (truncated bar chart ($M = 4.33$, $SD = 1.24$), full-range bar chart ($M = 2.97$, $SD = 1.13$), superimposed normal distributions ($M = 2.98$, $SD = 1.16$); $F(2, 1368) = 198.599$, $p < .001$, $\eta_p^2 = 0.225$; see Figure 19) but no significant main effect of source (Fox News ($M = 3.42$, $SD = 1.35$), CNN ($M = 3.46$, $SD = 1.33$), no source ($M = 3.40$, $SD = 1.33$); $F(2, 1368) = 0.435$, $p = .647$, $\eta_p^2 = 0.001$) nor a significant interaction between the two variables ($F(4, 1368) = 0.095$, $p = .984$, $\eta_p^2 = 0.000$).

Figure 3

Perceptions of Group Discreteness by Mode




* $p < .05$, ** $p < .01$, *** $p < .001$

Note. Significance levels featured are based on the results obtained from Tukey HSD post-hoc tests.

Exploratory pairwise t-tests were conducted between the no visualisation control and each of the three experimental mode of presentation conditions on level of trust concerning out-party fairness (see Figure 4). These analyses determined that participants assigned to the no visualisation control exhibited significantly lower levels of trust ($M = 22.92$, $SD = 20.01$) post-Bonferroni corrections than those assigned to the truncated bar chart condition ($M = 28.37$, $SD = 21.54$; $t(610) = 2.785$, $p = .006$, $d = 0.26$), the full-range bar chart condition ($M = 44.86$, $SD = 22.76$; $t(602) = 10.712$, $p < .001$, $d = 0.99$), and the superimposed normal distribution condition ($M = 45.87$, $SD = 23.92$; $t(314.98) = 11.846$, $p < .001$, $d = 1.00$). An exploratory 3x3 ANOVA was conducted to assess the main and interaction effects of mode of presentation and source on level of trust concerning out-party fairness across all experimental conditions, which indicated a significant main effect of mode ($F(2, 1368) = 85.432$, $p < .001$, $\eta_p^2 = 0.111$) but no significant main effect of source (Fox News ($M = 38.86$, $SD = 23.23$), CNN ($M = 40.76$, $SD = 24.82$), no source ($M = 39.66$, $SD = 24.33$); $F(2, 1368) = 1.013$, $p = .364$, $\eta_p^2 = 0.001$) nor a significant interaction between the two variables ($F(4, 1368) = 2.083$, $p = .081$, $\eta_p^2 = 0.006$).

Figure 4

Level of Trust Regarding Other-Party Fairness by Mode



* $p < .05$, ** $p < .01$, *** $p < .001$

*Note*₁. Only levels of statistical significance between experimental conditions and the no visualisation control are depicted.

*Note*₂. Significance levels featured are based on the results obtained from Tukey HSD post-hoc tests.

Appendix F: Chapter 1, Study 4 Additional Hypotheses and Results

Table 1

Additional Pre-Registered Hypotheses for Study 4

| Label | Hypothesis |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| H4 _A | Participants exposed to the full-range bar chart will offer donations that are significantly higher than those offered by participants exposed to either the truncated bar chart or the no visualisation conditions. |
| H4 _B | Participants exposed to the icon array histogram condition will offer donations that are significantly higher than those offered by participants exposed to either the truncated bar chart or the no visualisation conditions. |

- H4_C ANCOVA analyses will show mode of presentation to be a statistically significant predictor of donation amount even when controlling for other variables.
- H5_A Participants exposed to the full-range bar chart will show a significantly higher willingness to participate in civil cross-party discussion than those exposed to either the truncated bar chart or the no visualisation conditions.
- H5_B Participants exposed to the icon array histogram will show a significantly higher willingness to participate in civil cross-party discussion than those exposed to either the truncated bar chart or the no visualisation conditions.
- H5_C ANCOVA analyses will show mode of presentation to be a statistically significant predictor of willingness to engage in civil cross-party discussions even when controlling for other variables.
- H6_A Participants exposed to the full-range bar chart will produce responses to the optimistic political statement whose sentiment will be significantly more positive than sentiments of participants exposed to either the truncated bar chart or the no visualisation condition.
- H6_B Participants exposed to the icon array histogram will produce responses to the optimistic political statement whose sentiment will be significantly more positive than sentiments of participants exposed to either the truncated bar chart or the no visualisation condition.
- H7_A For participants allocated to the booster condition of the Trust Game, those exposed to the full-range bar chart will send significantly more to other-party partners than those exposed to the truncated bar chart.
- H7_B For participants allocated to the booster condition of the Trust Game, those exposed to the icon array histogram will send significantly more to other-party partners than those exposed to the truncated bar chart.
- H9 Significant main effects of booster condition will be observed on the amount sent in the Trust Game.
-

Results (cont.)

Participants exposed to the full-range bar chart did not offer donation amounts ($M = 2.58$, $SD = 3.82$) [H4_A] that differed significantly from those exposed to the truncated bar chart ($M = 3.13$, $SD = 4.04$; $t(264) = 1.138$, $p = .256$, $d = 0.14$) but did offer *significantly less* than participants in the no visualisation condition ($M = 3.71$, $SD = 4.34$; $t(258) = 2.230$, $p = .027$, $d = 0.28$). Participants exposed to the icon array histogram did not offer donation amounts ($M = 3.57$, $SD = 4.26$) [H4_B] that differed significantly from those exposed to the truncated bar chart ($t(267) = 0.880$, $p = .380$, $d = 0.11$) or those in the no visualisation condition ($t(261) = 0.255$, $p = .799$, $d = 0.03$). An ANCOVA analysis (controlling for possible covariates of age, sex, education, and political party) also determined no significant main effect of mode on donation generosity [H4_C] ($F(3, 517) = 2.228$, $p = .084$).

Participants exposed to the full-range bar chart did not offer to participate in a civil conversation ($M = 2.32$, $SD = 0.76$) [H5_A] at a rate that differed significantly from those exposed to the truncated bar chart ($M = 2.26$, $SD = 0.79$; $t(264) = 0.632$, $p = .528$, $d = 0.08$) or participants in the no visualisation condition ($M = 2.23$, $SD = 0.75$; $t(258) = 1.012$, $p = .312$, $d = 0.13$). Participants exposed to the icon array histogram also did not offer to participate ($M = 2.27$, $SD = 0.76$) [H5_B] at a rate that differed significantly from those exposed to the truncated bar chart ($t(267) = 0.094$, $p = .925$, $d = 0.01$) or participants in the no visualisation condition ($t(261) = 0.468$, $p = .640$, $d = 0.06$). An ANCOVA analysis (controlling for possible covariates of age, sex, education, and political party) also determined no significant main effect of mode on willingness to participate [H5_C] ($F(3, 517) = 0.240$, $p = .869$).


Participants exposed to the full-range bar chart did not produce response sentiments ($M = 3.80$, $SD = 1.34$) that differed significantly [H6_A] from those exposed to the truncated bar chart ($M = 3.82$, $SD = 1.47$; $t(264) = 0.109$, $p = .913$, $d = 0.01$) or participants in the no visualisation condition ($M = 3.64$, $SD = 1.59$; $t(258) =$

0.873, $p = .383$, $d = 0.11$). Participants exposed to the icon array histogram also did not produce response sentiments ($M = 4.06$, $SD = 1.25$) that differed significantly [H6_B] from those exposed to the truncated bar chart ($t(267) = 1.462$, $p = .145$, $d = 0.18$) but *did* produce responses that contained a significantly more positive sentiment than participants in the no visualisation condition ($t(238.81) = 2.375$, $p = .018$, $d = 0.30$).

Contrary to what was hypothesised, participants exposed to the full-range bar chart in the booster condition did not send more to opposing-party partners ($M = 7.29$, $SD = 3.17$) than those exposed to the truncated bar chart ($M = 6.47$, $SD = 3.03$) in the booster condition [H7_A] ($t(47) = 0.304$, $p = .762$, $d = 0.09$). Similarly, participants exposed to the icon array histogram in the booster condition also did not send more ($M = 5.95$, $SD = 3.69$) to opposing-party partners than those exposed to the truncated bar chart in the booster condition [H7_B] ($t(48) = 0.164$, $p = .870$, $d = 0.05$). Pairwise t-tests for amount sent (i.e., “trusted”) by participants revealed no statistically significant differences between participants assigned to the no visualization control ($M = 6.06$, $SD = 3.49$) versus those in either the truncated bar chart condition ($M = 6.38$, $SD = 3.35$; $t(190) = 0.652$, $p = .515$, $d = 0.09$), the full-range bar chart condition ($M = 6.71$, $SD = 3.34$; $t(188) = 1.324$, $p = .187$, $d = 0.19$), or the icon array histogram condition ($M = 6.37$, $SD = 3.47$; $t(197) = 0.639$, $p = .524$, $d = 0.09$). A three-way (mode by booster condition by partner party congruence) ANOVA across all experimental conditions (i.e., excluding the no visualisation control condition) also indicated no significant main effect of mode on amount sent ($F(2, 302) = 0.354$, $p = .702$, $\eta_p^2 = 0.002$), no significant main effect of booster condition [H9] on amount sent ($F(1, 302) = 0.112$, $p = .738$, $\eta_p^2 = 0.000$) but did reveal a significant main effect of partner party congruence (i.e., whether an individual was told they were playing with someone from their own party or the other party) on amount sent ($F(1, 302) = 10.841$, $p = .001$, $\eta_p^2 = 0.035$; see Figure 1). No significant interaction effects between any of the three variables were observed.

Figure 1

Amount Sent to Partner in Trust Game by Mode, Party Congruence, and Booster Condition



Appendix G: Rationale for Selecting Immigration as the Issue Presented to Participants (Chapter 2)

The particular items used in the study were selected for a number of reasons. Firstly, our original series of studies began by using artificial datasets that depicted intergroup data with an effect size (i.e., Cohen’s d) of precisely 0.80 (i.e., the commonly-accepted threshold for an effect to be categorized as “large”). As we moved from artificial data to actual data in our earlier studies, we sought to identify a policy area that typically yielded intergroup (i.e., between Republicans and Democrats) effect sizes that hovered around the 0.80 range. Based on prior ANES data, immigration seemed to be a satisfactory fit for that criterion. Similarly, the data from the US immigration item used in the present study has a Cohen’s d of 0.72. While the UK immigration item has a slightly higher Cohen’s d (i.e., 0.90 when dividing respondents into “Leavers” and “Remainers”), we felt the difference was justifiable, especially since utilizing nearly-identical items for both countries would make cross-cultural comparisons more useful.

Appendix H: The Interaction Between Education and Perceived Issue Polarisation (Chapter 2)

Though not originally planned, an exploratory ANOVA for perceived ideological polarisation which included both mode *and* educational attainment as independent variables also found the latter to exert significant main effects for both the UK sample and the US sample, such that greater educational attainment corresponded with higher levels of perceived ideological polarisation (see Figure 1 below).




Figure 1: Perceptions of ideological polarisation by educational attainment for the UK sample (left) and the US sample (right).

Appendix I: Stimuli Accompanying Summary Statistics for Participants in Visualisation Conditions (Chapter 3)

Below are the depictions of data that were presented to participants whose summarisation condition included a corresponding visualisation. For each figure, the (full-range) bar chart is on the left while the icon array histogram is on the right.

Figure 1

Opinions on the Legality of Abortion (Week 1)




Figure 2

Opinions on Ideal Levels of Legal Immigration (Week 2)




Figure 3

Opinions on the Correct Level of Support the US Should Provide to Ukraine (Week 3)




Figure 4

Opinions of Level of Sympathy with the Israeli People (Week 4)




Figure 5

Opinions on Quality of the US Healthcare System Relative to Other Developed Nations (Week 5)




Figure 6

Opinions on Correct Level of Strictness for Gun Laws (Week 6)




Figure 7

Opinions on the Morality of the Death Penalty for Murder (Week 7)





Figure 8

Beliefs about Climate Change (Week 8)



Appendix J: Pairwise T-Tests between Experimental Summary Statistic Conditions (Chapter 3)

Below are three tables summarising the pairwise t-tests between each of the experimental conditions on the three measures of interest: pre-post changes in perceptions of overall ideological polarisation (Table 1), pre-post changes in perceptions of intergroup agreement (Table 2), and pre-post changes in perceptions of ideological distance between average partisan voters (Table 3).

Table 1*Pairwise T-Tests between Each Experimental Condition on Pre-Post Changes in Perceived Overall Ideological Polarisation*

| Summary Condition | Comparison Condition | Pairwise T-Test Result |
|----------------------------|----------------------------|------------------------------------------|
| Mean-Only | Mean with Visualisation | $t(155.12) = 0.279, p = 0.780, d = 0.04$ |
| | Overlap-Only | $t(162) = 0.187, p = 0.852, d = 0.03$ |
| | Overlap with Visualisation | $t(168) = 0.829, p = 0.408, d = 0.13$ |
| | Both (No Visualisation) | $t(181) = 0.980, p = 0.329, d = 0.14$ |
| | Both with Visualisation | $t(178) = 1.316, p = 0.190, d = 0.20$ |
| Mean with Visualisation | Overlap-Only | $t(148.83) = 0.438, p = 0.662, d = 0.07$ |
| | Overlap with Visualisation | $t(152.68) = 0.985, p = 0.326, d = 0.15$ |
| | Both (No Visualisation) | $t(174) = 1.125, p = 0.262, d = 0.17$ |
| | Both with Visualisation | $t(171) = 1.411, p = 0.160, d = 0.21$ |
| Overlap-Only | Overlap with Visualisation | $t(152) = 0.636, p = 0.526, d = 0.10$ |
| | Both (No Visualisation) | $t(165) = 0.775, p = 0.439, d = 0.12$ |
| | Both with Visualisation | $t(162) = 1.107, p = 0.270, d = 0.17$ |
| Overlap with Visualisation | Both (No Visualisation) | $t(171) = 0.151, p = 0.880, d = 0.02$ |
| | Both with Visualisation | $t(168) = 0.489, p = 0.626, d = 0.08$ |
| Both (No Visualisation) | Both with Visualisation | $t(181) = 0.339, p = 0.735, d = 0.05$ |

Table 2*Pairwise T-Tests between Each Experimental Condition on Pre-Post Changes in Perceived Intergroup Agreement*

| Summary Condition | Comparison Condition | Pairwise T-Test Result |
|----------------------------|----------------------------|------------------------------------------|
| Mean-Only | Mean with Visualisation | $t(171) = 0.279, p = 0.754, d = 0.11$ |
| | Overlap-Only | $t(137.06) = 0.087, p = 0.931, d = 0.01$ |
| | Overlap with Visualisation | $t(168) = 1.758, p = 0.081, d = 0.27$ |
| | Both (No Visualisation) | $t(170.65) = 0.761, p = 0.448, d = 0.11$ |
| | Both with Visualisation | $t(178) = 1.625, p = 0.106, d = 0.24$ |
| Mean with Visualisation | Overlap-Only | $t(155) = 0.562, p = 0.575, d = 0.09$ |
| | Overlap with Visualisation | $t(161) = 2.229, p = 0.027, d = 0.35$ |
| | Both (No Visualisation) | $t(174) = 0.020, p = 0.984, d = 0.00$ |
| | Both with Visualisation | $t(171) = 2.137, p = 0.034, d = 0.33$ |
| Overlap-Only | Overlap with Visualisation | $t(152) = 1.569, p = 0.119, d = 0.25$ |
| | Both (No Visualisation) | $t(165) = 0.572, p = 0.568, d = 0.09$ |
| | Both with Visualisation | $t(162) = 1.466, p = 0.145, d = 0.23$ |
| Overlap with Visualisation | Both (No Visualisation) | $t(171) = 2.211, p = 0.029, d = 0.34$ |
| | Both with Visualisation | $t(168) = 0.150, p = 0.881, d = 0.02$ |
| Both (No Visualisation) | Both with Visualisation | $t(181) = 2.131, p = 0.034, d = 0.32$ |

Table 3

Pairwise T-Tests between Each Experimental Condition on Pre-Post Changes in Perceived Ideological Distance between Average Partisan Voters

| Summary Condition | Comparison Condition | Pairwise T-Test Result |
|----------------------------|----------------------------|------------------------------------------|
| Mean-Only | Mean with Visualisation | $t(144.46) = 1.310, p = 0.192, d = 0.20$ |
| | Overlap-Only | $t(162) = 1.733, p = 0.085, d = 0.27$ |
| | Overlap with Visualisation | $t(136.05) = 2.598, p = 0.010, d = 0.41$ |
| | Both (No Visualisation) | $t(169.43) = 1.783, p = 0.076, d = 0.26$ |
| | Both with Visualisation | $t(155.3) = 2.178, p = 0.031, d = 0.32$ |
| Mean with Visualisation | Overlap-Only | $t(155) = 0.225, p = 0.822, d = 0.04$ |
| | Overlap with Visualisation | $t(161) = 1.138, p = 0.257, d = 0.18$ |
| | Both (No Visualisation) | $t(174) = 0.300, p = 0.764, d = 0.05$ |
| | Both with Visualisation | $t(171) = 0.731, p = 0.466, d = 0.11$ |
| Overlap-Only | Overlap with Visualisation | $t(152) = 0.974, p = 0.331, d = 0.16$ |
| | Both (No Visualisation) | $t(165) = 0.076, p = 0.939, d = 0.01$ |
| | Both with Visualisation | $t(162) = 0.538, p = 0.592, d = 0.08$ |
| Overlap with Visualisation | Both (No Visualisation) | $t(171) = 0.916, p = 0.361, d = 0.14$ |
| | Both with Visualisation | $t(168) = 0.425, p = 0.672, d = 0.07$ |
| Both (No Visualisation) | Both with Visualisation | $t(181) = 0.476, p = 0.634, d = 0.07$ |

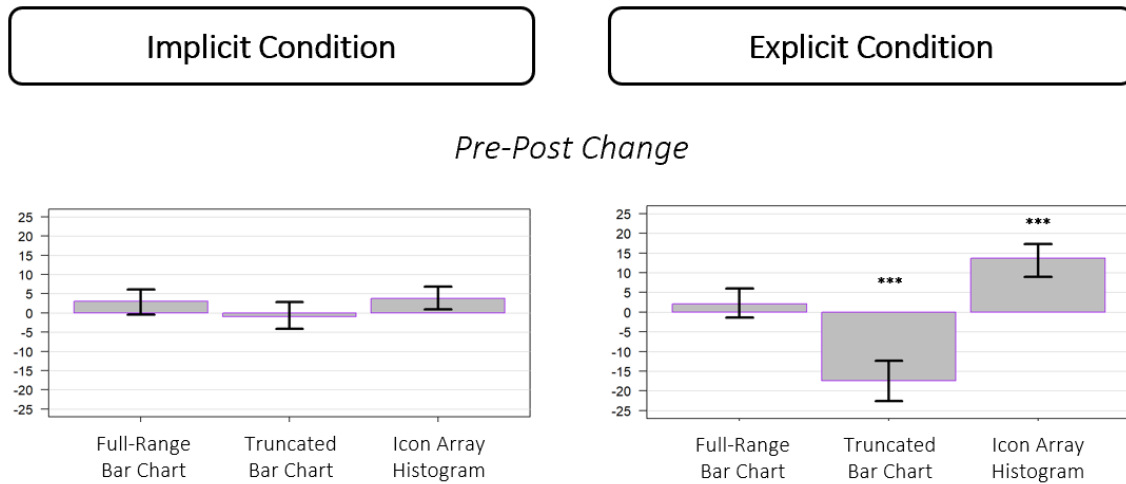
Appendix K: Additional Analyses, Results, and Discussion (Chapter 4)

Visualisation Study

For participants in the explicit choice revelation condition, exploratory t-tests determined statistically-significant differences (post-Bonferroni corrections to alpha thresholds) in pre-post changes of perceptions of trust for sources that chose to depict data via a truncated bar chart (i.e., perceptions of trust significantly *declined*; $M = -17.45, SD = 24.44; t(162) = 4.522, p < .001, d = 0.71$) and sources that chose to depict data via an icon array histogram (i.e., perceptions of trust significantly *increased*; $M = 13.66, SD = 18.49; t(164) = 4.033, p < .001, d = 0.63$; Figure 1).

Figure 1⁸⁰

Pre-Post Changes in Perceptions of Source Trust for Each Condition



Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).


* $p < .017$, ** $p < .003$, *** $p < .0003$

For participants in the explicit choice revelation condition, exploratory t-tests determined statistically-significant differences (post-Bonferroni corrections to alpha thresholds) in pre-post changes of perceptions of credibility for sources that chose to depict data via a truncated bar chart (i.e., perceptions of credibility significantly *declined*; $M = -13.12$, $SD = 22.04$; $t(162) = 3.343$, $p = .001$, $d = 0.52$) and sources that chose to depict data via an icon array histogram (i.e., perceptions of credibility significantly *increased*; $M = 13.49$, $SD = 21.14$; $t(164) = 4.020$, $p < .001$, $d = 0.62$; Figure 2).

⁸⁰ Please note that for all visualisations which depict pre-post changes, we've opted to truncate the y-axis. Due to the fact that the graphics depict *change* on a 0 to 100 scale, the full range (i.e., which would encompass all possible responses) would extend from -100 (for a participant who changed from a response of 100 to a response of 0) to 100 (for a participant who changed from a response of 0 to a response of 100). A 200-point range on the y-axis dramatically compresses the bar charts, rendering significant and non-significant intra-condition changes nearly indistinguishable from one another and, arguably, making the graphic virtually useless as a tool for comparison and comprehension. For example, there are several instances where a pre-post shift of less than 10 points constitutes a significant change. When depicted via a full-range (i.e., 200 point) y-axis, the bar is compressed so substantially that it looks nearly identical to an adjacent bar depicting non-significant change. Thus, while we acknowledge that the truncation introduces the risk that individuals will come to *overestimate* the degree of intra-condition change (or inter-condition difference), we believe that, conversely, refusing to truncate introduces the risk that an individual walks away with an equally-skewed impression that pre-post change (or difference in change) was insignificant. Please see *General Discussion* for a slightly more detailed discussion.

Figure 2

Pre-Post Changes in Perceptions of Source Credibility for Each Condition




Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .017$, ** $p < .003$, *** $p < .0003$

For participants in the explicit choice revelation condition, exploratory t-tests determined statistically-significant differences (post-Bonferroni corrections to alpha thresholds) in pre-post changes of perceptions of intent for sources that chose to depict data via a truncated bar chart (i.e., perceptions of intent significantly shifted towards “divide”; $M = -8.77$, $SD = 18.51$; $t(162) = 2.759$, $p = .006$, $d = 0.43$; Figure 3).

Figure 3

Pre-Post Changes in Perceptions of Source Intent for Each Condition




Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .017$, ** $p < .003$, *** $p < .0003$

For participants in the explicit choice revelation condition, exploratory t-tests determined statistically-significant differences (post-Bonferroni corrections to alpha thresholds) in pre-post desire to see data depicted in such a manner for sources that chose to depict data via a truncated bar chart (i.e., desire significantly *declined*; $M = -20.28, SD = 26.32; t(162) = 4.828, p < .001, d = 0.75$) and sources that chose to depict data via an icon array histogram (i.e., desire significantly *increased*; $M = 11.90, SD = 23.33; t(164) = 3.292, p = .001, d = 0.51$; Figure 4).

Figure 4

Pre-Post Changes in Desire to See More Data Visualised in this Manner for Each Condition




Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .017$, ** $p < .003$, *** $p < .0003$

For participants in the explicit choice revelation condition, exploratory t-tests determined statistically-significant differences (post-Bonferroni corrections to alpha thresholds) in pre-post changes of level of audience appreciation for sources that chose to depict data via a truncated bar chart (i.e., levels of appreciation significantly *decreased*; $M = -21.64, SD = 25.64; t(162) = 5.219, p < .001, d = 0.82$) and sources that chose to depict data via an icon array histogram (i.e., levels of appreciation significantly *increased*; $M = 13.73, SD = 25.88; t(164) = 3.502, p < .001, d = 0.54$; Figure 5).

Figure 5

Pre-Post Changes in Level of Appreciation of Visualisation Format for Each Condition



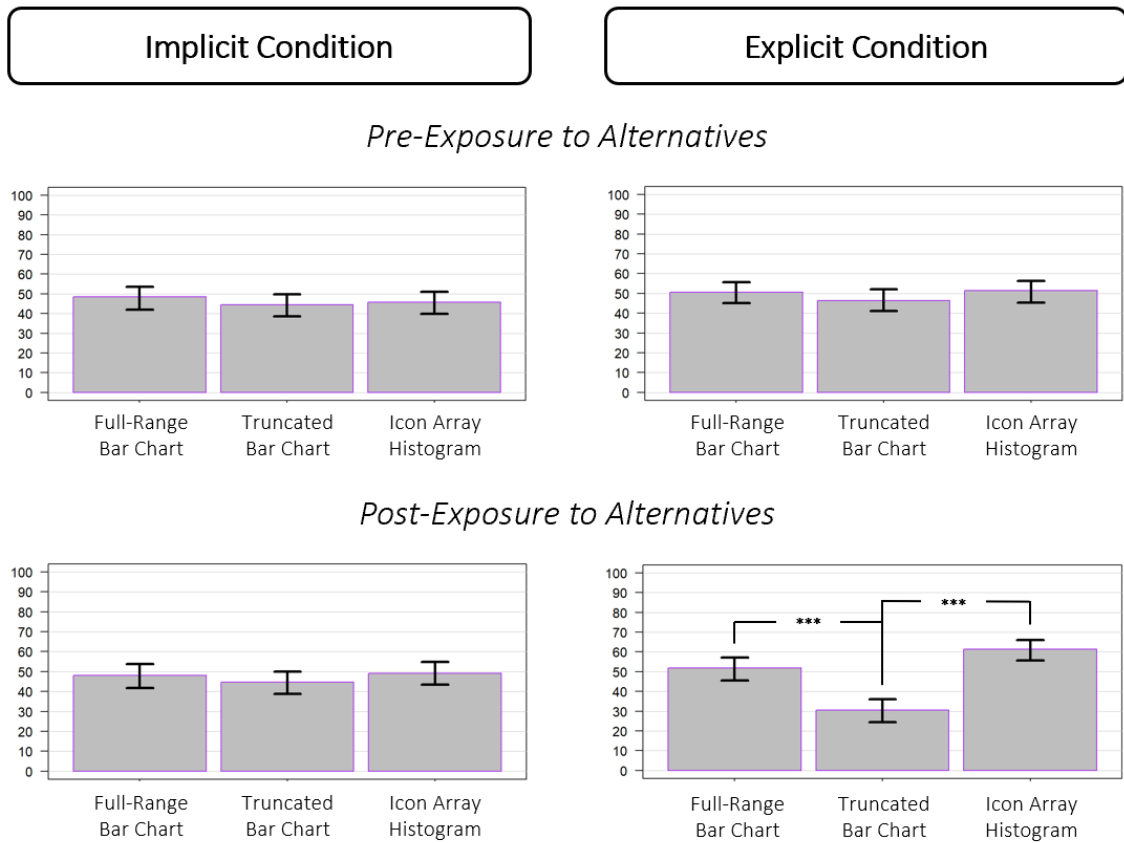
Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .017$, ** $p < .003$, *** $p < .0003$

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of visualisation type on intent to return to source pre-exposure to alternatives (truncated bar chart ($M = 44.42$, $SD = 25.56$), full-range bar chart ($M = 48.52$, $SD = 25.87$), icon array histogram ($M = 45.67$, $SD = 24.93$); $F(2, 245) = 0.564$, $p = .570$, $\eta_p^2 = 0.005$) or post-exposure (truncated bar chart ($M = 44.76$, $SD = 25.83$), full-range bar chart ($M = 48.08$, $SD = 27.26$), icon array histogram ($M = 49.18$, $SD = 26.60$); $F(2, 245) = 0.621$, $p = .538$, $\eta_p^2 = 0.005$; see Figure 6). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of visualisation type on intent to return to source pre-exposure to alternatives (truncated bar chart ($M = 46.30$, $SD = 24.40$), full-range bar chart ($M = 50.63$, $SD = 23.80$), icon array histogram ($M = 51.48$, $SD = 25.43$); $F(2, 244) = 1.052$, $p = .351$, $\eta_p^2 = 0.009$) but there was a significant main effect post-exposure (truncated bar chart ($M = 30.48$, $SD = 26.74$), full-range bar chart ($M = 51.91$, $SD = 26.69$), icon array histogram ($M = 61.29$, $SD = 24.53$); $F(2, 244) = 30.390$, $p < .001$, $\eta_p^2 = 0.199$).

Figure 6

Pre- and Post-Exposure Measures of Intent to Return to Source for Each Condition




Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

For participants in the explicit choice revelation condition, exploratory t-tests determined statistically-significant differences (post-Bonferroni corrections to alpha thresholds) in pre-post changes of intent to return for sources that chose to depict data via a truncated bar chart (i.e., intent significantly *decreased*; $M = -15.83$, $SD = 26.51$; $t(162) = 3.960$, $p < .001$, $d = 0.62$) and sources that chose to depict data via an icon array histogram (i.e., intent significantly *increased*; $M = 9.81$, $SD = 24.51$; $t(164) = 2.528$, $p < .012$, $d = 0.39$; Figure 7).

Figure 7

Pre-Post Changes in Intent to Return to Source for Each Condition



Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

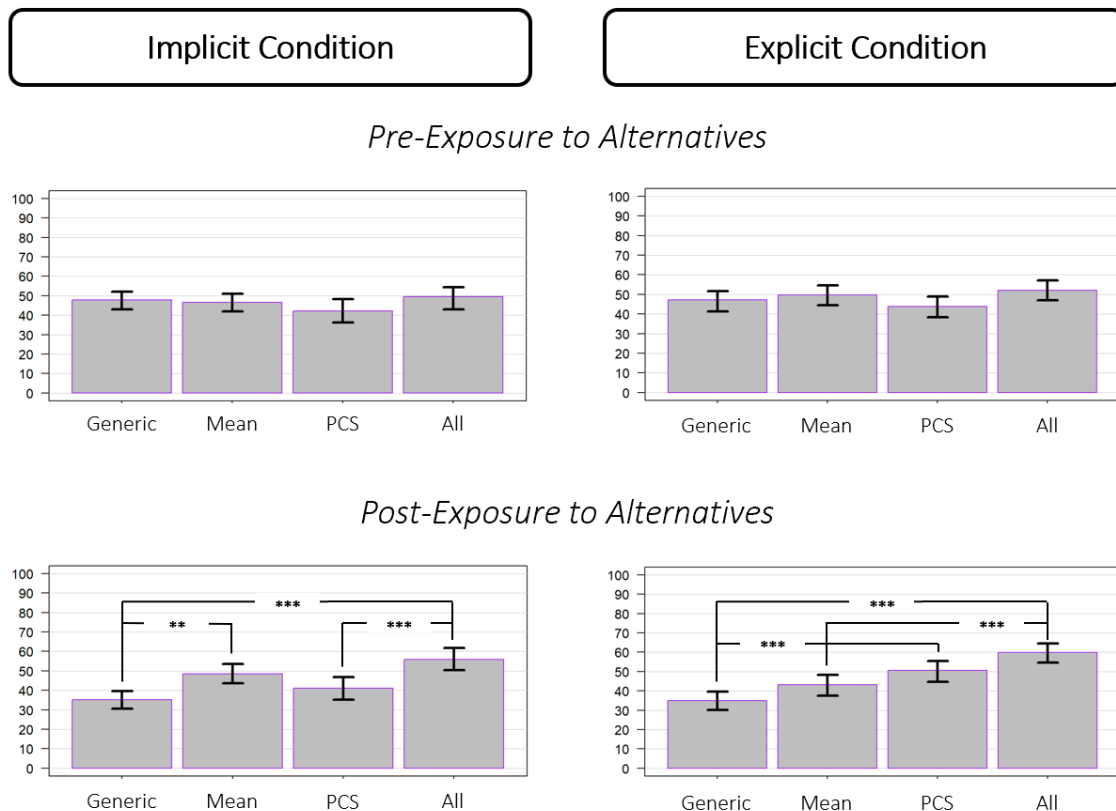
* $p < .017$, ** $p < .003$, *** $p < .0003$

Summarisation Study

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of summarisation format on perceptions of source trust pre-exposure to alternatives (generic ($M = 47.88$, $SD = 22.66$), mean ($M = 46.50$, $SD = 21.30$), PCS ($M = 42.10$, $SD = 27.15$), all ($M = 49.45$, $SD = 25.84$); $F(3, 339) = 1.428$, $p = .234$, $\eta_p^2 = 0.012$) but a significant main effect post-exposure (generic ($M = 35.28$, $SD = 22.35$), mean ($M = 48.53$, $SD = 24.32$), PCS ($M = 41.07$, $SD = 27.13$), all ($M = 55.94$, $SD = 26.33$); $F(3, 339) = 11.060$, $p < .001$, $\eta_p^2 = 0.089$; see Figure 8). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of summarisation format on perceptions of source trust pre-exposure to alternatives (generic ($M = 47.14$, $SD = 23.36$), mean ($M = 49.80$, $SD = 24.21$), PCS ($M = 43.75$, $SD = 23.53$), all ($M = 52.02$, $SD = 24.01$); $F(3, 338) = 1.867$, $p = .135$, $\eta_p^2 = 0.016$) but there was a significant main effect post-exposure (generic ($M = 34.90$, $SD = 22.46$), mean ($M = 43.29$, $SD = 25.03$), PCS ($M = 50.57$, $SD = 24.54$), all ($M = 59.87$, $SD = 23.41$); $F(3, 338) = 17.000$, $p < .001$, $\eta_p^2 = 0.131$).

Figure 8

Pre- and Post-Exposure Measures of Perceptions of Source Trust for Each Condition



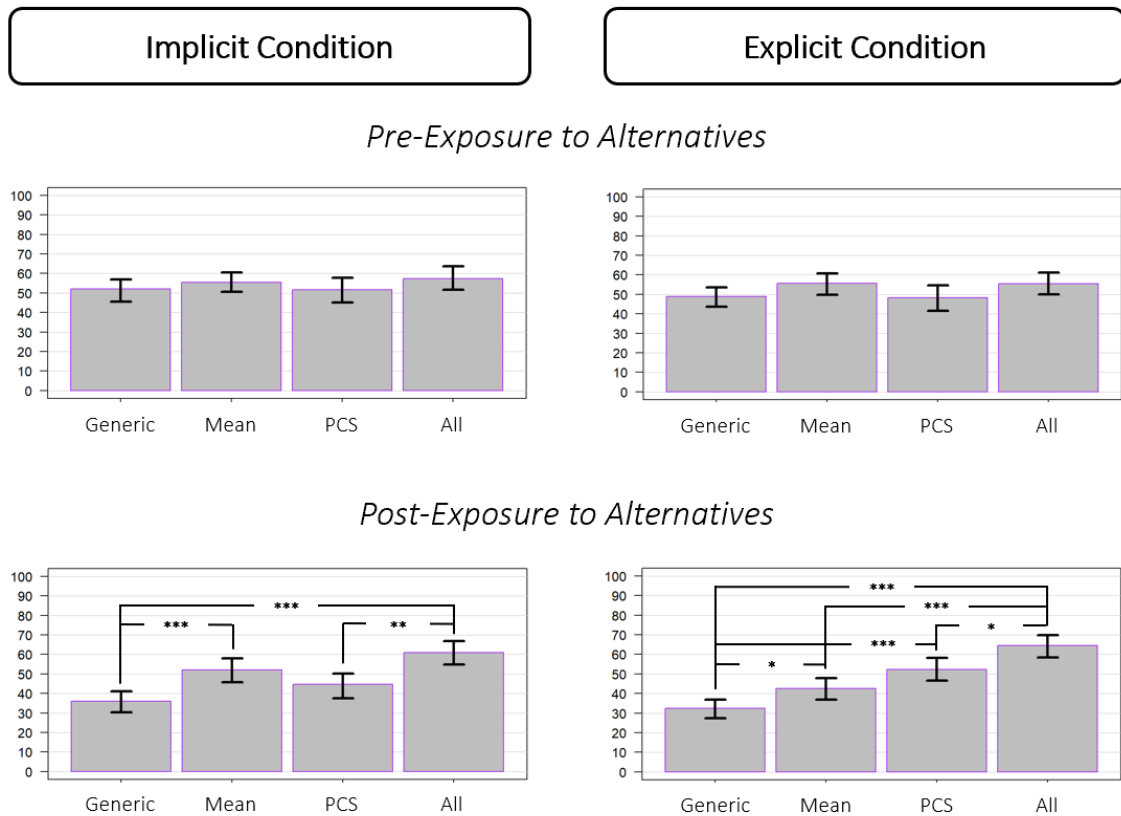
Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of summarisation format on levels of appreciation of the summarisation format pre-exposure to alternatives (generic ($M = 52.00$, $SD = 25.65$), mean ($M = 55.42$, $SD = 23.05$), PCS ($M = 51.63$, $SD = 28.77$), all ($M = 57.42$, $SD = 26.81$); $F(3, 339) = 0.971$, $p = .407$, $\eta_p^2 = 0.009$) but a significant main effect post-exposure (generic ($M = 36.13$, $SD = 25.61$), mean ($M = 52.12$, $SD = 28.00$), PCS ($M = 44.71$, $SD = 29.47$), all ($M = 60.87$, $SD = 27.77$); $F(3, 339) = 12.490$, $p < .001$, $\eta_p^2 = 0.100$; see Figure 9). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of summarisation format on levels of appreciation of the summarisation format pre-exposure to alternatives (generic ($M = 48.90$, $SD = 25.42$), mean ($M = 55.58$, $SD = 24.95$), PCS ($M = 48.19$, $SD = 28.06$), all ($M = 55.48$, $SD = 26.65$); $F(3, 338) = 2.025$, $p = .110$, $\eta_p^2 = 0.018$) but there was a significant main effect post-exposure (generic ($M = 32.54$, $SD = 22.37$), mean ($M = 42.55$, $SD = 26.87$), PCS ($M = 52.33$, $SD = 27.31$), all ($M = 64.45$, $SD = 25.27$); $F(3, 338) = 24.550$, $p < .001$, $\eta_p^2 = 0.179$).

Figure 9

Pre- and Post-Exposure Measures of Level of Appreciation of Summarisation Format for Each Condition



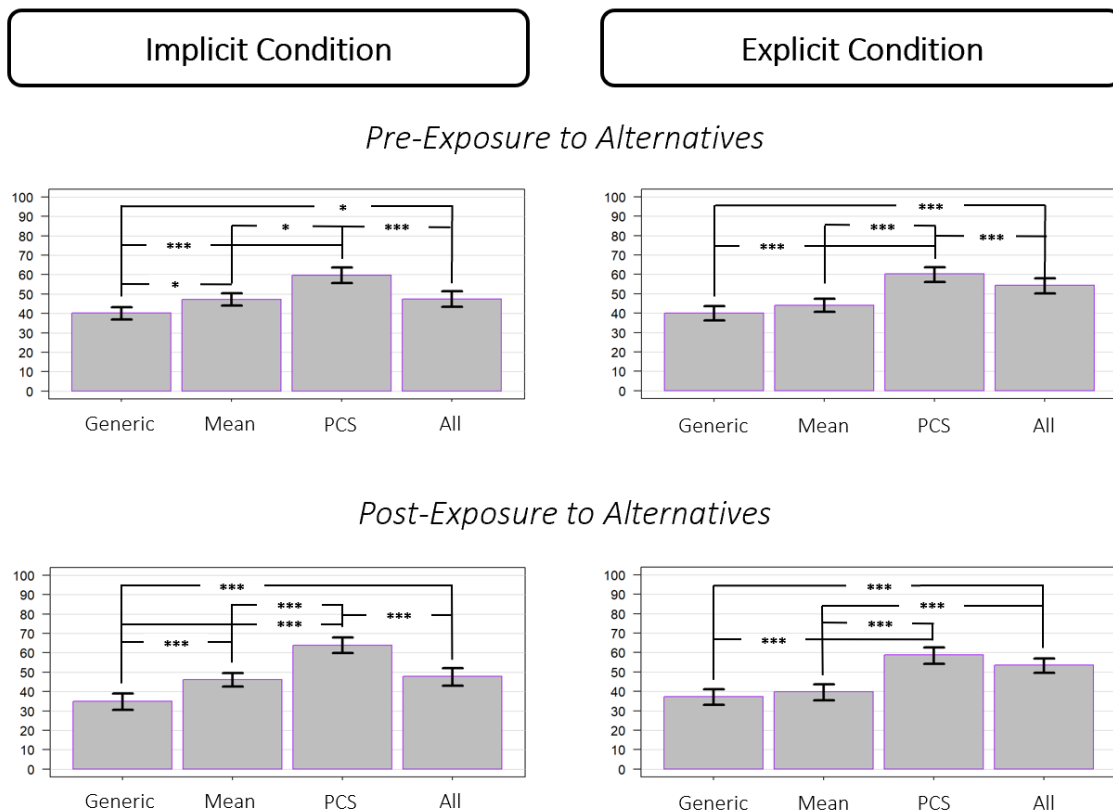
Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there were significant main effects of summarisation format on perceptions of source intent pre-exposure to alternatives (generic ($M = 40.15$, $SD = 15.74$), mean ($M = 47.16$, $SD = 14.21$), PCS ($M = 59.58$, $SD = 19.19$), all ($M = 47.40$, $SD = 19.96$); $F(3, 339) = 18.390$, $p < .001$, $\eta_p^2 = 0.140$) as well as significant main effects post-exposure (generic ($M = 34.90$, $SD = 19.96$), mean ($M = 46.06$, $SD = 16.70$), PCS ($M = 63.86$, $SD = 18.59$), all ($M = 47.84$, $SD = 20.26$); $F(3, 339) = 34.000$, $p < .001$, $\eta_p^2 = 0.231$; see Figure 10). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there were significant main effects of summarisation format on perceptions of source intent pre-exposure to alternatives (generic ($M = 40.08$, $SD = 17.41$), mean ($M = 44.02$, $SD = 16.54$), PCS ($M = 60.36$, $SD = 16.66$), all ($M = 54.31$, $SD = 17.82$); $F(3, 338) = 24.930$, $p < .001$, $\eta_p^2 = 0.181$) as well as significant main effects post-exposure (generic ($M = 37.34$, $SD = 18.67$), mean ($M = 39.75$, $SD = 19.54$), PCS ($M = 58.81$, $SD = 20.44$), all ($M = 53.53$, $SD = 16.73$); $F(3, 338) = 25.910$, $p < .001$, $\eta_p^2 = 0.187$).

Figure 10

Pre- and Post-Exposure Measures of Perceptions of Source Intent for Each Condition



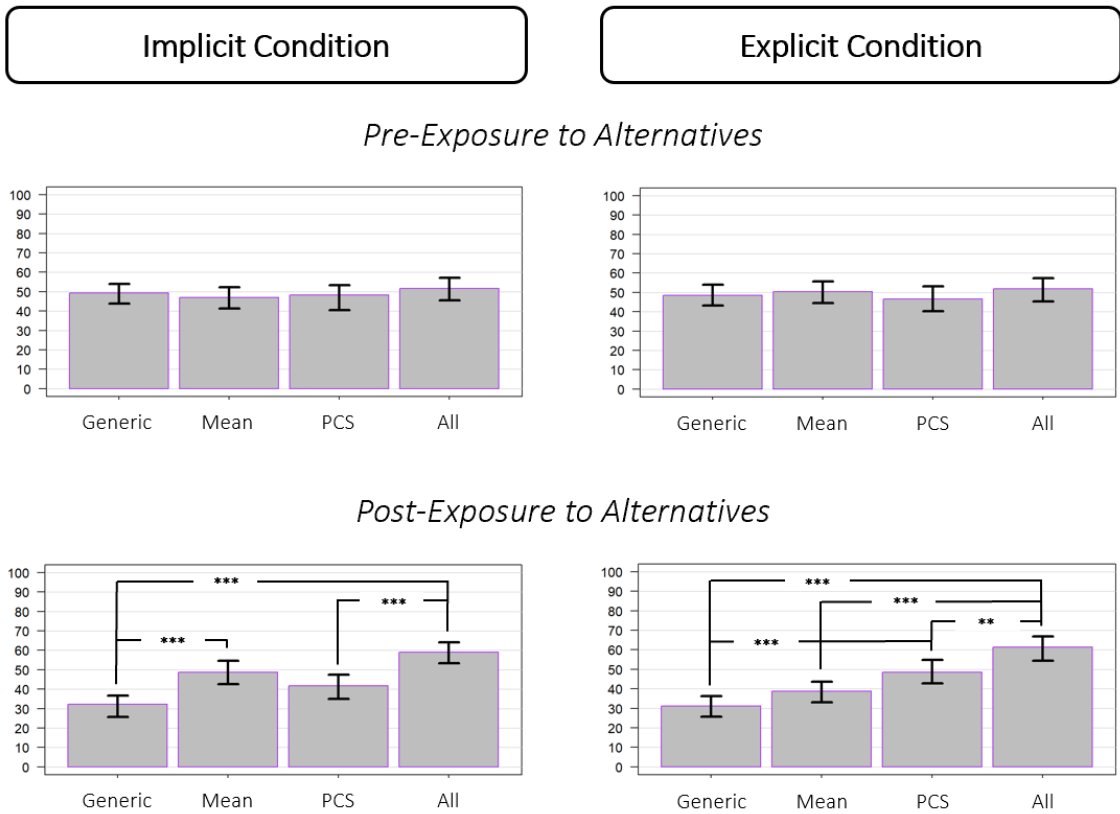
Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of summarisation format on desire to see more data summarised in this manner pre-exposure to alternatives (generic ($M = 49.38$, $SD = 25.26$), mean ($M = 47.01$, $SD = 26.01$), PCS ($M = 48.17$, $SD = 28.60$), all ($M = 51.54$, $SD = 27.63$); $F(3, 339) = 0.442$, $p = .723$, $\eta_p^2 = 0.004$) but a significant main effect post-exposure (generic ($M = 32.19$, $SD = 26.35$), mean ($M = 48.66$, $SD = 28.62$), PCS ($M = 41.77$, $SD = 29.30$), all ($M = 58.96$, $SD = 26.07$); $F(3, 339) = 14.440$, $p < .001$, $\eta_p^2 = 0.113$; see Figure 11). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of summarisation format on desire to see more data summarised in this manner pre-exposure to alternatives (generic ($M = 48.45$, $SD = 25.55$), mean ($M = 50.45$, $SD = 27.31$), PCS ($M = 46.57$, $SD = 29.34$), all ($M = 51.86$, $SD = 27.31$); $F(3, 338) = 0.596$, $p = .618$, $\eta_p^2 = 0.005$) but there was a significant main effect post-exposure (generic ($M = 31.11$, $SD = 24.01$), mean ($M = 38.73$, $SD = 25.42$), PCS ($M = 48.44$, $SD = 26.25$), all ($M = 61.29$, $SD = 27.95$); $F(3, 338) = 21.710$, $p < .001$, $\eta_p^2 = 0.162$).

Figure 11

Pre- and Post-Exposure Measures of Desire to See More Data Summarised in this Manner for Each Condition




Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

Exploratory t-tests determined statistically-significant pre-post changes (post-Bonferroni corrections to alpha thresholds) in perceptions of credibility for sources that chose to summarise data using a generic statement in both the implicit ($M = -15.31$, $SD = 22.81$; $t(174) = 4.587$, $p < .001$, $d = 0.69$) and explicit ($M = -15.28$, $SD = 22.82$; $t(172) = 4.360$, $p < .001$, $d = 0.66$) alternative choice revelation conditions (Figure 12).

Figure 12

Pre-Post Changes in Perceptions of Source Credibility for Each Condition



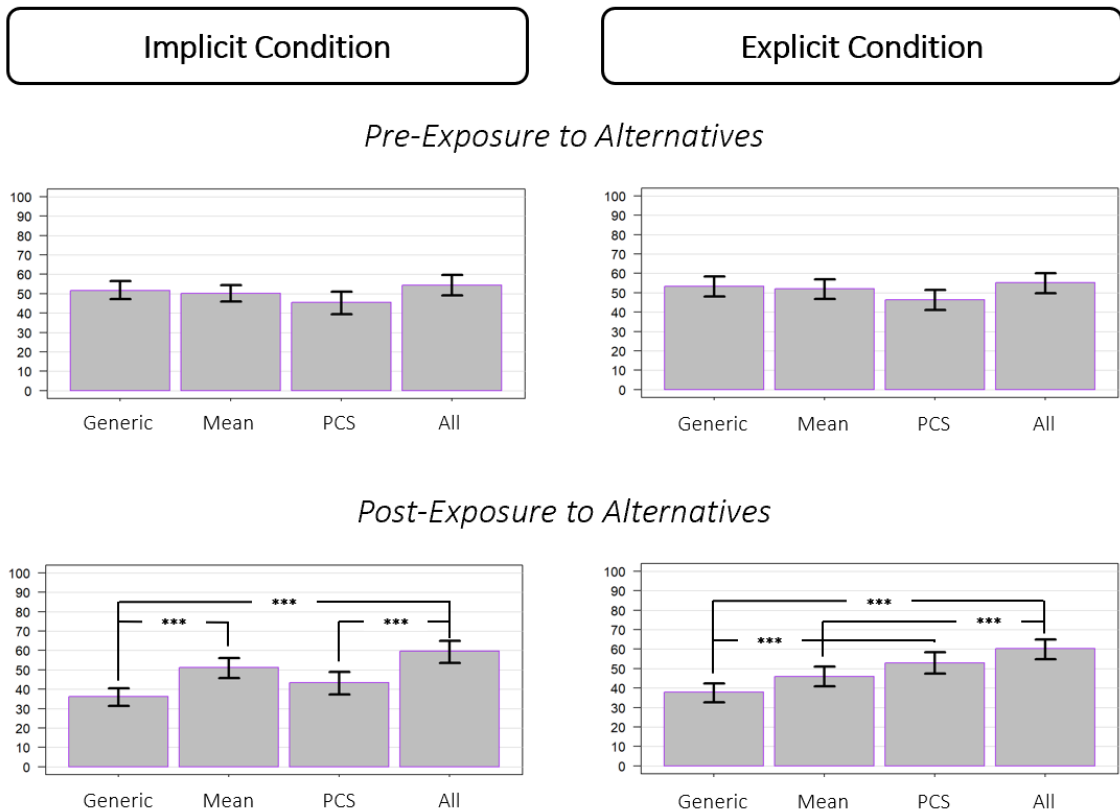
Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .013$, ** $p < .003$, *** $p < .0003$

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of summarisation format on perceptions of source credibility pre-exposure to alternatives (generic ($M = 51.61$, $SD = 21.68$), mean ($M = 50.08$, $SD = 19.64$), PCS ($M = 45.56$, $SD = 26.76$), all ($M = 54.45$, $SD = 24.37$); $F(3, 339) = 2.160$, $p = .093$, $\eta_p^2 = 0.019$) but a significant main effect post-exposure (generic ($M = 36.31$, $SD = 22.58$), mean ($M = 51.15$, $SD = 24.72$), PCS ($M = 43.42$, $SD = 27.05$), all ($M = 59.68$, $SD = 25.55$); $F(3, 339) = 13.970$, $p < .001$, $\eta_p^2 = 0.110$; see Figure 13). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of summarisation format on perceptions of source credibility pre-exposure to alternatives (generic ($M = 53.24$, $SD = 23.45$), mean ($M = 52.15$, $SD = 24.20$), PCS ($M = 46.40$, $SD = 23.95$), all ($M = 55.27$, $SD = 23.54$); $F(3, 338) = 2.109$, $p = .099$, $\eta_p^2 = 0.018$) but there was a significant main effect post-exposure (generic ($M = 37.97$, $SD = 22.76$), mean ($M = 46.03$, $SD = 25.14$), PCS ($M = 52.95$, $SD = 25.62$), all ($M = 60.35$, $SD = 24.04$); $F(3, 338) = 13.210$, $p < .001$, $\eta_p^2 = 0.105$).

Figure 13

Pre- and Post-Exposure Measures of Perceptions of Source Credibility for Each Condition




Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

Exploratory t-tests determined statistically-significant pre-post changes (post-Bonferroni corrections to alpha thresholds) in intent to return to source for sources that chose to summarise data using a generic statement in both the implicit ($M = -14.56$, $SD = 24.96$; $t(174) = 3.909$, $p < .001$, $d = 0.59$) and explicit ($M = -13.92$, $SD = 26.17$; $t(172) = 3.600$, $p < .001$, $d = 0.55$) alternative choice revelation conditions (Figure 14).

Figure 14

Pre-Post Changes in Intent to Return to Source for Each Condition



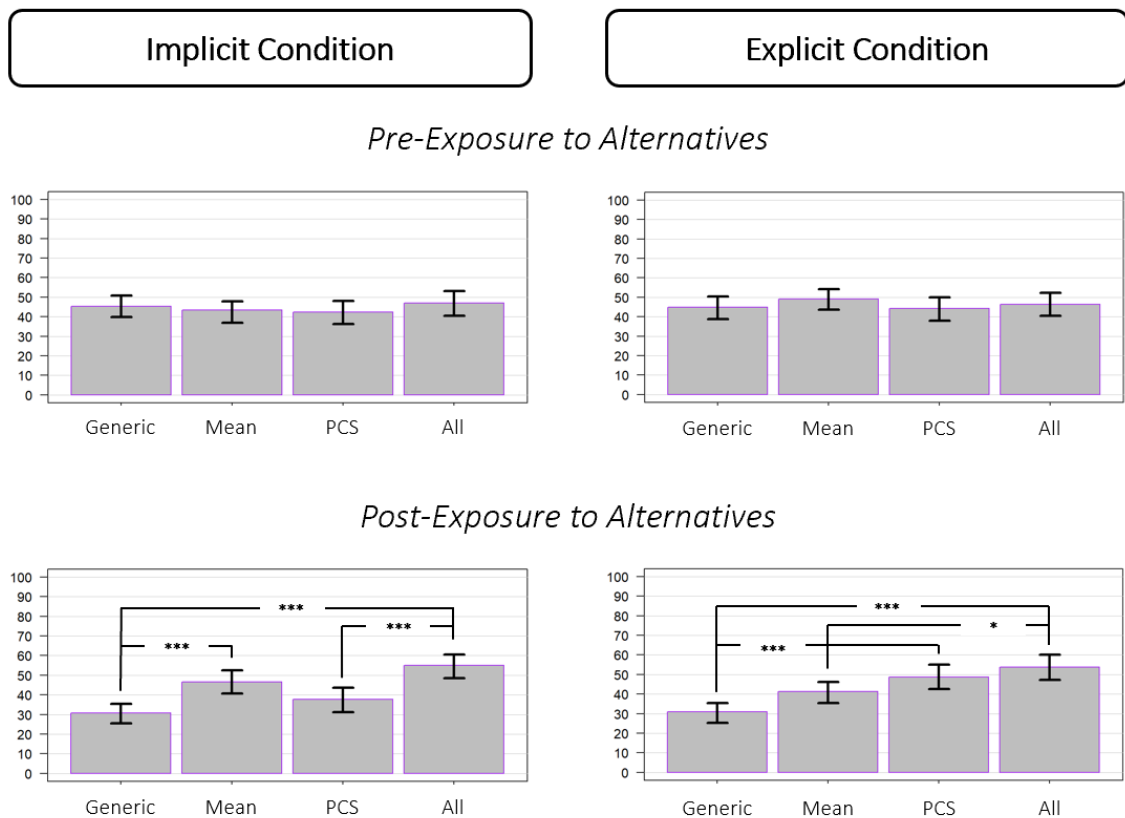
Note. Asterisks denote levels of significance for pre-post changes within-condition (following appropriate Bonferroni corrections to alpha thresholds).

* $p < .013$, ** $p < .003$, *** $p < .0003$

Exploratory one-way ANOVAs determined that, for participants in the implicit revelation condition, there was no significant main effect of summarisation format on intent to return to source pre-exposure to alternatives (generic ($M = 45.24$, $SD = 25.32$), mean ($M = 43.34$, $SD = 25.92$), PCS ($M = 42.40$, $SD = 28.85$), all ($M = 47.05$, $SD = 28.79$); $F(3, 339) = 0.487$, $p = .691$, $\eta_p^2 = 0.004$) but a significant main effect post-exposure (generic ($M = 30.68$, $SD = 24.07$), mean ($M = 46.67$, $SD = 27.71$), PCS ($M = 37.80$, $SD = 29.65$), all ($M = 54.92$, $SD = 28.00$); $F(3, 339) = 12.770$, $p < .001$, $\eta_p^2 = 0.102$; see Figure 15). Identical one-way ANOVAs determined that, for participants in the explicit revelation condition, there was no significant main effect of summarisation format on intent to return to source pre-exposure to alternatives (generic ($M = 44.95$, $SD = 27.18$), mean ($M = 49.17$, $SD = 26.01$), PCS ($M = 44.20$, $SD = 27.07$), all ($M = 46.45$, $SD = 27.88$); $F(3, 338) = 0.570$, $p = .635$, $\eta_p^2 = 0.005$) but there was a significant main effect post-exposure (generic ($M = 31.03$, $SD = 23.71$), mean ($M = 41.35$, $SD = 25.16$), PCS ($M = 48.79$, $SD = 28.95$), all ($M = 53.80$, $SD = 29.66$); $F(3, 338) = 11.590$, $p < .001$, $\eta_p^2 = 0.093$).

Figure 15

Pre- and Post-Exposure Measures of Intent to Return to Source for Each Condition



Note. Asterisks denote levels of significance for between-condition differences as determined by Tukey HSD post-hoc tests.

* $p < .05$, ** $p < .01$, *** $p < .001$

In regards to how disparate data presentation methods impact perceptions of, opinions about, and intentions to engage with sources, four patterns of results warrant a more extensive discussion: 1) main effects of data presentation pre-exposure to alternatives, 2) shifts in perception, opinion, and intended action catalysed by the two alternative choice revelation conditions, 3) emergent preference hierarchies post-exposure to alternatives, and 4) presentation methods which resulted in consistent shifts in perception, opinion, and intended action across studies.

Main Effects of Data Presentation Pre-Exposure to Alternatives

Prior to examining the role that exposure to *alternative* presentation methods may have on source perceptions, opinions, and audience intentions, it's critical to determine whether presentation methods *in isolation* (i.e., devoid of alternatives against which comparisons can be made) yield significant effects on such metrics.

Within both the visualisation study and the summarisation study, main effects of presentation method (i.e., visualisation type or summarisation format) were only observed pre-exposure for one variable: source intent. This implies that, devoid of exposure to any alternative visualisation or summarisation practices, the data presentation methods that sources choose to use do not appear to be a differentiating factor in the way audiences evaluate variables such as source trust, credibility, and appreciation. However, the fact that we *do* see pre-exposure differences on the source intent variable across both studies suggests that audiences use both the visualisation and summarisation choices of a source to differentially infer their intent (e.g., unite, divide, etc.), without the need for alternative visualisations to serve as comparison points. More specifically, in the visualisation study, when participants are exposed to sources depicting data via a truncated bar chart, they tend to evaluate that source's intent as being more divisive than if they had depicted identical data using a full-range bar chart or icon array histogram. Similarly, in the summarisation study, sources summarising data using generic statements tend to be seen as having an intent that is significantly more divisive than those that have chosen to use a combination of all three formats, while sources that have opted to use PCS are seen as having an intent geared significantly more towards uniting than those that have chosen to use group means.

Shifts In Perception, Opinion, and Intended Action Catalysed by the Two Alternative Choice Revelation Conditions

The inclusion of two separate alternative data presentation revelation conditions within each study were designed to serve as a way to compare how perceptions might change if audiences were to simply *become aware* of alternative means of visualising or summarising the data (i.e., the implicit condition) versus if they were both made aware and also *taught* the relative advantages and disadvantages of each practice (i.e., the explicit condition).

In the visualisation study, main effects of visualisation type were only observed post-exposure in the implicit condition on two of the six perception metrics (i.e., source intent and appreciation of visualisation used). For all other metrics (e.g., perceived trust, credibility, intent to return, etc.), exposure to the implicit alternative choice revelation format (i.e., simply seeing how other sources had chosen to visualise the same data) did not yield significant main effects of visualisation type. Conversely, for participants assigned to the explicit alternative revelation condition, significant main effects of visualisation were observed post-exposure for all six outcome variables. This suggests that once audiences are both made aware of alternative visualisation methods *and* taught the pros and cons of each, the visualisation choices sources have made become critical differentiators in how they evaluate source trust, credibility, level of appreciation, intent to return, etc.

In the summarisation study, however, the gap between the efficacy of the implicit and explicit alternative choice revelation conditions was reduced, with significant main effects of summarisation format being

observed across all dependent variables for *both* conditions. Such a finding is important as it suggests that while only being exposed to alternative choices *and* taught about their relative pros and cons is a reliable intervention for catalysing main effects of visualisation type, simple exposure (without accompanying teaching) is enough to catalyse such effects for summarisation format. For sources, the takeaway is that, as audiences become aware of alternative summarisation techniques being used, their perceptions of sources as well as their intentions to engage with them are likely to be altered based on whether they deem their choices to be more or less appropriate.

Emergent Preference Hierarchies Post-Exposure to Alternatives

As participants were exposed to the two intervention conditions, post-exposure evaluations began to exhibit certain discernible patterns of choice preference.

In the visualisation study, participants in the implicit condition often did not display results which revealed any apparent hierarchies of visualisation type post-exposure. However, when they did (i.e., for perceptions of source intent and appreciation of visualisation format), the pattern always indicated a preference for icon array histograms over truncated bar charts (i.e., significant between-condition differences were observed for these visualisation types, with icon array histograms always promoting more favourable results) and one time indicated a preference for full-range bar charts over truncated bar charts (i.e., in the appreciation metric). Participants in the explicit condition, however, exhibited consistent, recognisable patterns of preference for visualisation types across all metrics. In every instance, participants evaluations were significantly more favourable towards sources using an icon array histogram or a full-range bar chart than those using a truncated bar chart. And on one metric (i.e., desire to see more data visualised in this manner), a significant difference also emerged between the full-range bar chart and the icon array histogram (with participants wanting to see more data visualised in the latter's format than the former's).

In the summarisation study, five of the six outcome variables elicited a similar pattern of responses in the implicit condition (the only exception being the source intent variable – explained momentarily). Following simple exposure to other summarisation formats, participants came to view sources employing the generic statement least favourably, followed by the PCS, then the group means, and then the combination technique (viewed most favourably). More specifically, significant differences emerged between sources using the generic statement and those using the group means (with the latter being preferred) as well as between sources using the PCS and those using the combination technique (again, with the latter being preferred). The greatest statistically-significant difference was between the sources that chose to use the generic statement and those that chose to use the combination technique. Overall, these results suggest that, should audiences simply become aware of alternative summarisation practices, sources that choose to summarise data using generic statements or PCS are likely to be penalised relative to sources that choose

to summarise identical data using either group means or a comprehensive combination. For perceived source intent, however, participants in the implicit condition post-exposure come to see sources using PCS as having the greatest intent to “unite,” followed by those using group means and combination techniques, with those employing generic statements being viewed as comparatively most divisive in their intent.

In the explicit condition, post-exposure evaluations of each outcome variable (with the exception of perceived source intent – discussed momentarily) always resembled a “stair step” pattern, with generic statements always generating the least favourable ratings and group means always generating the second least favourable, followed by PCS generating the second most favourable and finally the combination technique generating the most favourable evaluations. Within this pattern, sources that choose to summarise data using generic statements are always viewed significantly less favourably than those that chose to use either PCS or a combination, while sources that chose to use group means are always evaluated less favourably than those that used a combination technique. For perceived source intent, participants in the explicit condition post-exposure come to see sources using PCS or a combination as having a significantly greater intent to “unite” than sources using group means or generic statements.

Presentation Methods Which Elicit Significant Pre-Post Shifts

Finally, across both studies, we see patterns emerging concerning which methods of data presentation, when used by sources, exhibit significant pre-post shifts in perceptions, opinions, and intentions to engage.

In the visualisation study, no significant pre-post shifts were seen in any visualisation type in the implicit condition. In the explicit condition, however, participants routinely exhibited significant shifts in their views of sources that chose to visualise data using either a truncated bar chart or an icon array histogram. Sources that visualised their data using an icon array histogram saw significant pre-post shifts in five of the six outcome variables related to source perception, with each shift occurring in a more “desirable” direction (e.g., sources were perceived as significantly *more* trustworthy and credible, audience intentions to return to the source *increased*, etc.). Conversely, sources that visualised their data using a truncated bar chart saw significant pre-post shifts across all six outcome variables related to source perception, with each shift occurring in a less “desirable” direction (e.g., sources were perceived as significantly *less* trustworthy and credible, audience intentions to return to the source *decreased*, etc.). This suggests that if audiences are exposed to alternative visualisation methods and taught of their relative advantages and disadvantages, meaningful shifts in perceptions can occur – in both directions – if sources opt to use these particular modes of presentation.

In the summarisation study, only one summarisation format exhibited significant pre-post shifts within the implicit condition: generic statements. Across five of the six outcome metrics (the only exception being perceptions of source intent), evaluations of sources that had used generic statements significantly declined,

indicating that the favourability of sources can be adversely affected if they choose to use generic statements, even if audiences are not explicitly taught about the pros and cons of alternative summarisation techniques. Generic statements also elicited significant pre-post declines in the same five metrics in the explicit condition, further solidifying that such a summarisation technique, when employed by sources, risks damaging perceptions of source trust and credibility, while also adversely affecting audience intentions to return. Within the explicit condition, one other summarisation format exhibited significant pre-post changes: group means. Both levels of appreciation for the summarisation format as well as desire to see more data summarised in this manner significantly declined for participants exposed to the group mean summarisation format in the explicit condition, suggesting that audiences come to view this form of data presentation as comparatively unfavourable after learning about the relative advantages and disadvantages of it and several alternative techniques.

Appendix L: Full Script of the Inoculation Intervention (Chapter 5)

Below is the exact text shown to participants who were randomly assigned to the inoculation condition within the inoculation study described in Chapter 5:

Nobody's saying Republicans and Democrats are incredibly similar. They're not. The two groups are very different in many ways, and they disagree on a lot of important issues.

However, although members of Congress and political elites make it seem like compromise is impossible, the truth is that the opinions of average Republicans and Democrats actually aren't that far apart. In fact, they're sometimes closer than some media outlets would like.

You see, some media outlets keep their followers by painting a picture of *extreme* polarisation. They want you to believe that Democrats and Republicans *hate* each other, and that they could never agree on *anything*. Focusing on the beliefs of the most extreme members of each side and pushing a narrative of other-party hatred gets a lot of clicks and views.

But sometimes the truth is more boring than the media would like. While Democrats and Republicans do not hold the same stances on key issues, many members of each party hold pretty moderate views, and are willing to find compromises with the other side.


But this story of compromise, moderation, and shared opinions doesn't sell.

So when Republicans and Democrats seem to be agreeing on key issues more than fits their preferred narrative, some outlets employ subtle manipulation tactics to make the groups seem farther apart than they actually are.

If you aren't aware of these manipulation tactics, you are vulnerable to being fooled and misled. Fortunately, once you know what to look out for, it's easy to detect these manipulations and protect yourself from misinformation.

Here's an example: imagine two groups were asked how they felt about a certain issue on a scale from 1 to 7 (1 being "Strongly Disagree" and 7 being "Strongly Agree"). The average response given by members of Group 1 was 3.1 on a 1 to 7 scale; the average response given by members of Group 2 was 3.9 on a 1 to 7 scale.


Below is a graph that *accurately* depicts the data:



Notice how the y-axis starts at 1 and goes all the way to 7 – just like it does in the original survey. This is an accurate way to show the data.


You must pay close attention to the *vertical* axis on the left side (otherwise known as the y-axis). In the graph above, the source included the *full* y-axis (i.e., it went all the way from 1 to 7, just as the original survey had responses that ranged from 1 to 7). But what some sources will do in an effort to manipulate individuals into believing groups are more divided than they actually are is to *shorten* (or "truncate") the y-axis. So instead of going from 1 to 7, they might only show it from 3 to 4 (or something similar). This has the effect of making the two groups seem "farther apart" than they actually are.

Here's an example of what that tactic looks like, using the same data as the last graph:



Here you can see that the y-axis has been *shortened*; it only goes from about 3 to 4 (instead of from 1 to 7). This makes the groups seem farther apart than they actually are.

Now look at the two graphs side-by-side:




When you see the two graphs side-by-side, you get an idea of how much of an impact shortening the y-axis makes. The groups seem fairly similar in their opinions in Graph 1, but once the y-axis is shortened (like it is in Graph 2), it suddenly makes it seem like they are very different.

This is one of the ways agents on social and traditional media may try to trick you into thinking groups are more polarized than they are.

Fortunately, spotting this manipulation trick is easy once you know how to look out for it. You always want to pay close attention to the full range of responses used in the original survey (e.g., 1 to 7) and make sure that the vertical-axis (or y-axis) of the graph goes all the way from the lowest possible response (i.e., 1) to the highest possible response (i.e., 7). If it doesn't show the full range, then the groups will appear further apart than they actually are. And the more the y-axis is truncated (i.e., the more they reduce the original length), the more distorted the graph will be.

Don't be fooled! Let's practice spotting deceptive tactics:

An organization asked 1,000 men and women their opinion on a new law. The responses ranged from 1 (“Strongly Disagree”) to 9 (“Strongly Agree”). Here’s how the organization depicted the responses in their latest publication:




Was this visualization deceptive:

- Yes
- No

Good! Because the organization used a full-range y-axis (i.e., one that went from 1 to 9, just like the original survey), the visualization was not deceptive.

Let’s try one more.

Members of two political groups, the Reds and the Blues, were recently asked to give their opinions on a key housing issue. The survey asked how much they supported the issue on a scale from 1 (“Not at All”) to 7 (“Completely”). Here is how a local newspaper reported the results:



Was this visualization deceptive:

- Yes
- No

Exactly! Since the newspaper opted to shorten (or “truncate”) the y-axis from 1 to 7 (i.e., the original range of responses) to now only going from 2 to 5, it makes the “gap” between the two groups seem larger than it actually is. Hence, it is a deceptive visualization.

Great job! Just remember: if you don’t want to be manipulated by questionable media practices, be sure to always check whether the y-axis of the graph matches the length of responses in the original survey.

REFERENCES

- Abramowitz, A., & McCoy, J. (2019). United States: Racial resentment, negative partisanship, and polarisation in Trump's America. *The ANNALS of the American Academy of Political and Social Science*, 681(1), 137-156.
- Abramowitz, A. I., & Saunders, K. L. (2008). Is polarisation a myth?. *The Journal of Politics*, 70(2), 542-555.
- Abramowitz, A. I., & Webster, S. (2016). The rise of negative partisanship and the nationalization of US elections in the 21st century. *Electoral Studies*, 41, 12-22.
- Abramowitz, A. I., & Webster, S. W. (2018). Negative partisanship: Why Americans dislike parties but behave like rabid partisans. *Political Psychology*, 39, 119-135.
- Ahler, D. J. (2014). Self-fulfilling misperceptions of public polarisation. *The Journal of Politics*, 76(3), 607-620.
- Ahler, D. J., & Sood, G. (2018). The parties in our heads: Misperceptions about party composition and their consequences. *The Journal of Politics*, 80(3), 964-981.
- Alieva, I. (2023). How American media framed 2016 presidential election using data visualization: The case study of the New York times and the Washington post. *Journalism Practice*, 17(4), 814-840.
- Allport, G. W., Clark, K., & Pettigrew, T. (1954). *The nature of prejudice*.
- Altemeyer, B. (1996). *The authoritarian specter*. Harvard University Press.
- American National Election Studies (2021). *ANES 2020 Time Series Study Full Release* [dataset and documentation]. February 10, 2022 version. www.electionstudies.org
- American Political Science Association. (1950). Toward a more responsible two-party system: a report of the Committee on Political Parties, American Political Science Association. (*No Title*).
- Anscombe, F. J. (1973). Graphs in statistical analysis. *The american statistician*, 27(1), 17-21.

- Apicella, C., Norenzayan, A., & Henrich, J. (2020). Beyond WEIRD: A review of the last decade and a look ahead to the global laboratory of the future. *Evolution and Human Behavior*, 41(5), 319-329.
- Arceneaux, K., Johnson, M., & Murphy, C. (2012). Polarized political communication, oppositional media hostility, and selective exposure. *The Journal of Politics*, 74(1), 174-186.
- Armaly, M. T., & Enders, A. M. (2021). The role of affective orientations in promoting perceived polarisation. *Political Science Research and Methods*, 9(3), 615-626.
- Arnheim, R. (1980). A Plea for Visual Thinking. *Critical Inquiry*, 6(3), 489–497.
- Bafumi, J., & Herron, M. C. (2010). Leapfrog representation and extremism: A study of American voters and their members in Congress. *American Political Science Review*, 104(3), 519-542.
- Baldassarri, D., & Page, S. E. (2021). The emergence and perils of polarisation. *Proceedings of the national academy of sciences*, 118(50), e21116863118.
- Banas, J. A., & Rains, S. A. (2010). A meta-analysis of research on inoculation theory. *Communication Monographs*, 77(3), 281-311.
- Bastian, B., & Haslam, N. (2006). Psychological essentialism and stereotype endorsement. *Journal of Experimental Social Psychology*, 42(2), 228-235.
- Battle-Baptiste, W., & Rusert, B. (Eds.). (2018). *WEB Du Bois's data portraits: Visualizing black America*. Chronicle Books.
- Baum, M. A., & Groeling, T. (2009). Shot by the messenger: Partisan cues and public opinion regarding national security and war. *Political Behavior*, 31(2), 157-186.
- Baum, M. A., & Gussin, P. (2008). In the eye of the beholder: How information shortcuts shape individual perceptions of bias in the media. *Quarterly Journal of political science*, 3(1), 1-31.
- Becker, S. O., Fetzer, T., & Novy, D. (2017). Who voted for Brexit? A comprehensive district-level analysis. *Economic Policy*, 32(92), 601-650.
- Berinsky, A. J., & Kinder, D. R. (2006). Making sense of issues through media frames: Understanding the

- Kosovo crisis. *The Journal of Politics*, 68(3), 640-656.
- Bersoff, D. M. (2024). Overcoming political polarisation: Strategies for diminishing intransigence and reducing intergroup animus. *Behavioral Science & Policy*, 0(0). <https://doi.org/10.1177/23794607241238037>
- Binder, S. A. (2004). *Stalemate: Causes and consequences of legislative gridlock*. Rowman & Littlefield.
- Boldry, J. G., Gaertner, L., & Quinn, J. (2007). Measuring the measures: A meta-analytic investigation of the measures of outgroup homogeneity. *Group Processes & Intergroup Relations*, 10(2), 157-178.
- Bolsen, T., & Shapiro, M. A. (2017). Strategic framing and persuasive messaging to influence climate change perceptions and decisions. In *Oxford research encyclopedia of climate science*.
- Boxell, L., Gentzkow, M., & Shapiro, J. M. (2022). Cross-country trends in affective polarisation. *Review of Economics and Statistics*, 1-60.
- Branscombe, N. R., Ellemers, N., Spears, R., & Doosje, B. (1999). The context and content of social identity threat. *Social identity: Context, commitment, content*, 35-58.
- Branscombe, N. R., & Wann, D. L. (1994). Collective self-esteem consequences of outgroup derogation when a valued social identity is on trial. *European journal of social psychology*, 24(6), 641-657.
- Bresciani, S., & Eppler, M. J. (2009). The risks of visualization. *Identität und Vielfalt der Kommunikationswissenschaft*, 165-178.
- Brick, C., Hood, B., Ekroll, V., & de-Wit, L. (2020). Illusory essences: A bias holding back theorizing in psychological science.
- Brownstein, R. (2022). America is growing apart, possibly for good. *The Atlantic*. <https://www.theatlantic.com/politics/archive/2022/06/red-and-blue-state-divide-is-growing-michael-podhorzer-newsletter/661377/>
- Buczel, K. A., Szyszka, P. D., Siwiak, A., Szpitalak, M., & Polczyk, R. (2022). Vaccination against

- misinformation: The inoculation technique reduces the continued influence effect. *Plos one*, 17(4), e0267463.
- Bürigi, P., & Roos, J. (2003). Images of strategy. *European Management Journal*, 21(1), 69-78.
- Cairo, A. (2020). The dawn of a philosophy of visualization. *Data visualization in society*, 17-18.
- Campbell, A. (1980). *The American voter*. University of Chicago Press.
- Campbell, D. E., Green, J. C., & Layman, G. C. (2011). The party faithful: Partisan images, candidate religion, and the electoral impact of party identification. *American Journal of Political Science*, 55(1), 42-58.
- Carlin, R. E., & Love, G. J. (2018). Political competition, partisanship and interpersonal trust in electoral democracies. *British Journal of Political Science*, 48(1), 115-139.
- Cassese, E. C. (2021). Partisan dehumanization in American politics. *Political Behavior*, 43, 29-50.
- Chambers, J. R., Baron, R. S., & Inman, M. L. (2006). Misperceptions in intergroup conflict: Disagreeing about what we disagree about. *Psychological science*, 17(1), 38-45.
- Chater, N., & Loewenstein, G. (2023). The i-frame and the s-frame: How focusing on individual-level solutions has led behavioral public policy astray. *Behavioral and Brain Sciences*, 46, e147.
- Chong, D., & Druckman, J. N. (2007). Framing theory. *Annual Review of Political Science*, 10, 103-126.
- Curbal (2016, February 16). *The "Bill Gates Effect" in Google Analytics and how to get rid of it with Power BI*. Retrieved June 30, 2024 from <https://curbal.com/the-bill-gates-effect-in-google-analytics-and-how-to-get-rid-of-it-with-power-bi>
- D'Amore, C., van Zomeren, M., & Koudenburg, N. (2024). How perceived polarization predicts attitude moralization (and vice versa): A four-wave longitudinal study during the 2020 US election. *Journal of personality and social psychology*.
- de Benedictis-Kessner, J., Baum, M. A., & Berinsky, A. J. (2019). Polarisation and Media Usage: Disentangling Causality. In Suhay, E., Grofman, B., & Trechsel, A. H. (Eds.). (2020). *The Oxford handbook of electoral persuasion*. Oxford University Press, USA.

- Deppe, K. D., Gonzalez, F. J., Neiman, J., Pahlke, J., Smith, K., & Hibbing, J. R. (2015). Reflective liberals and intuitive conservatives: A look at the Cognitive Reflection Test and ideology.
- De Langhe, B., & Fernbach, P. (2019). The dangers of categorical thinking: we're hardwired to sort information into buckets-and that can hamper our ability to make good decisions. *Harvard Business Review*, 97(5), 80-92.
- Dias, N. C., Aarslew, L. F., Frederiksen, K. V. S., Lelkes, Y., Pradella, L., & Westwood, S. J. (2024). Correcting misperceptions of partisan opponents is not effective at treating democratic ills. *PNAS nexus*, 3(8), 304.
- Dias, N., & Lelkes, Y. (2022). The nature of affective polarization: Disentangling policy disagreement from partisan identity. *American Journal of Political Science*, 66(3), 775-790.
- Dimant, E. (2024). Hate trumps love: The impact of political polarization on social preferences. *Management Science*, 70(1), 1-31.
- Dimara, E., Franconeri, S., Plaisant, C., Bezerianos, A., & Dragicevic, P. (2018). A task-based taxonomy of cognitive biases for information visualization. *IEEE transactions on visualization and computer graphics*, 26(2), 1413-1432.
- Druckman, J. N. (2001). Evaluating framing effects. *Journal of economic psychology*, 22(1), 91-101.
- Druckman, J. N. (2001). The implications of framing effects for citizen competence. *Political Behavior*, 23(3), 225-256.
- Druckman, J. N. (2011). What's it all about? Framing in political science. *Perspectives on framing*, 279, 282-296.
- Druckman, J. N., & Levendusky, M. S. (2019). What do we measure when we measure affective polarisation?. *Public Opinion Quarterly*, 83(1), 114-122.
- Druckman, J. N., Levendusky, M. S., & McLain, A. (2018). No need to watch: How the effects of partisan media can spread via interpersonal discussions. *American Journal of Political Science*, 62(1), 99-112.

- Druckman, J. N., & Levy, J. (2022). Affective polarisation in the American public. In T. J. Rudolph (Ed.), *Handbook on politics and public opinion* (p. 257-269). Edward Elgar Publishing.
- Druckman, J. N., Peterson, E., & Slothuus, R. (2013). How elite partisan polarisation affects public opinion formation. *American Political Science Review*, *107*(1), 57-79.
- Durand, M. A., Yen, R. W., O'Malley, J., Elwyn, G., & Mancini, J. (2020). Graph literacy matters: Examining the association between graph literacy, health literacy, and numeracy in a Medicaid eligible population. *PloS one*, *15*(11), e0241844.
- Dweck, C. S. (2013). *Self-theories: Their role in motivation, personality, and development*. Psychology Press.
- Ecker, U. K. (2017). Why rebuttals may not work: the psychology of misinformation. *Media Asia*, *44*(2), 79-87.
- El Haffar, G., Durif, F., & Dubé, L. (2020). Towards closing the attitude-intention-behavior gap in green consumption: A narrative review of the literature and an overview of future research directions. *Journal of cleaner production*, *275*, 122556.
- Enders, A. M., & Armaly, M. T. (2019). The differential effects of actual and perceived polarisation. *Political Behavior*, *41*, 815-839.
- Esteban, J. M., & Ray, D. (1994). On the measurement of polarisation. *Econometrica: Journal of the Econometric Society*, 819-851.
- Evergreen, S. D. (2019). *Effective data visualization: The right chart for the right data*. SAGE publications.
- Fanelli, D. (2010). "Positive" results increase down the hierarchy of the sciences. *PloS one*, *5*(4), e10068.
- Fang, X., Singh, S., & Ahluwalia, R. (2007). An examination of different explanations for the mere exposure effect. *Journal of consumer research*, *34*(1), 97-103.
- Fansher, M., Adkins, T. J., Lalwani, P., Boduroglu, A., Carlson, M., Quirk, M., ... & Jonides, J. (2022). Icon arrays reduce concern over COVID-19 vaccine side effects: a randomized control study. *Cognitive Research: Principles and Implications*, *7*(1), 1-9.

- Fehr, E., & Gächter, S. (2000). Cooperation and punishment in public goods experiments. *American Economic Review*, 90(4), 980-994.
- Feinberg, M., & Willer, R. (2015). From gulf to bridge: When do moral arguments facilitate political influence?. *Personality and Social Psychology Bulletin*, 41(12), 1665-1681.
- Feldman, S., & Weber, C. (2016). Understanding Social Conservatism: The Effects of Moral Values and Authoritarianism.
- Feldman-Stewart, D., Brundage, M. D., & Zotov, V. (2007). Further insight into the perception of quantitative information: judgments of gist in treatment decisions. *Medical Decision Making*, 27(1), 34-43.
- Fernbach, P. M., & Van Boven, L. (2022). False polarisation: Cognitive mechanisms and potential solutions. *Current opinion in Psychology*, 43, 1-6.
- Finkel, E. J., Bail, C. A., Cikara, M., Ditto, P. H., Iyengar, S., Klar, S., ... & Druckman, J. N. (2020). Political sectarianism in America. *Science*, 370(6516), 533-536.
- Fiorina, M. P., & Abrams, S. J. (2008). Political polarisation in the American public. *Annu. Rev. Polit. Sci.*, 11, 563-588.
- Fiorina, M. P., Abrams, S. J., & Pope, J. C. (2011). *Culture war?*. Longman.
- Fiske, S. T., Cuddy, A. J., & Glick, P. (2007). Universal dimensions of social cognition: Warmth and competence. *Trends in cognitive sciences*, 11(2), 77-83.
- Flaxman, S., Goel, S., & Rao, J. M. (2016). Filter bubbles, echo chambers, and online news consumption. *Public opinion quarterly*, 80(S1), 298-320.
- Friendly, M. (2008). The golden age of statistical graphics. *Statistical Science*, 502-535.
- Friendly, M., & Denis, D. J. (2001). Milestones in the history of thematic cartography, statistical graphics, and data visualization. URL <http://www.datavis.ca/milestones>, 32, 13.
- Frederick, S. (2005). Cognitive reflection and decision making. *Journal of Economic perspectives*, 19(4), 25-42.

- Gaissmaier, W., & Gigerenzer, G. (2008). Statistical illiteracy undermines informed shared decision making. *Zeitschrift für Evidenz, Fortbildung und Qualität im Gesundheitswesen*, *102*(7), 411-413.
- Galesic, M., Garcia-Retamero, R., & Gigerenzer, G. (2009). Using icon arrays to communicate medical risks: overcoming low numeracy. *Health Psychology*, *28*(2), 210.
- Gallagher, K. M., & Updegraff, J. A. (2012). Health message framing effects on attitudes, intentions, and behavior: a meta-analytic review. *Annals of behavioral medicine*, *43*(1), 101-116.
- Garcia-Retamero, R., Galesic, M., & Dhami, M. K. (2013). Reducing denominator neglect. *Transparent Communication of Health Risks: Overcoming Cultural Differences*, 145-164.
- Garcia-Retamero, R., Galesic, M., & Gigerenzer, G. (2010). Do icon arrays help reduce denominator neglect?. *Medical Decision Making*, *30*(6), 672-684.
- Garzia, D., Ferreira da Silva, F., & Maye, S. (2023). Affective Polarisation in Comparative and Longitudinal Perspective. *Public Opinion Quarterly*, *87*(1), 219-231.
- Gelman, S. A., & Rhodes, M. (2012). "Two-thousand years of stasis": How psychological essentialism impedes evolutionary understanding. In K. S. Rosengren, S. K. Brem, E. M. Evans, & G. M. Sinatra (Eds.), *Evolution challenges: Integrating research and practice in teaching and learning about evolution* (pp. 3–21). Oxford University Press.
- Gerend, M. A., & Cullen, M. (2008). Effects of message framing and temporal context on college student drinking behavior. *Journal of Experimental Social Psychology*, *44*(4), 1167-1173.
- Gervais, B. T. (2015). Incivility online: Affective and behavioral reactions to uncivil political posts in a web-based experiment. *Journal of Information Technology & Politics*, *12*(2), 167-185.
- Gervais, B. T. (2019). Rousing the partisan combatant: Elite incivility, anger, and antideliberative attitudes. *Political Psychology*, *40*(3), 637-655.
- Glare, P., Fridman, I., & Ashton-James, C. E. (2018). Choose your words wisely: the impact of message framing on patients' responses to treatment advice. *International review of neurobiology*, *139*, 159-190.

- Graham, J., Haidt, J., & Nosek, B. A. (2009). Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology, 96*(5), 1029.
- Graham, J., Nosek, B. A., & Haidt, J. (2012). The moral stereotypes of liberals and conservatives: Exaggeration of differences across the political spectrum. *PloS one, 7*(12), e50092.
- Guess, A. M., Barberá, P., Munzert, S., & Yang, J. (2021). The consequences of online partisan media. *Proceedings of the National Academy of Sciences, 118*(14), e2013464118.
- Guess, A. M., & Lyons, B. A. (2020). Misinformation, disinformation, and online propaganda. *Social media and democracy: The state of the field, prospects for reform, 10*.
- Gurven, M. D. (2018). Broadening horizons: Sample diversity and socioecological theory are essential to the future of psychological science. *Proceedings of the National Academy of Sciences, 115*(45), 11420-11427.
- Haidt, J., & Graham, J. (2007). When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social justice research, 20*(1), 98-116.
- Hanel, P. H., Maio, G. R., & Manstead, A. S. (2019). A new way to look at the data: Similarities between groups of people are large and important. *Journal of personality and social psychology, 116*(4), 541.
- Hanel, P. H., & Wolf, L. J. (2020). Leavers and Remainers after the Brexit referendum: More united than divided after all?. *British Journal of Social Psychology, 59*(2), 470-493.
- Harper, C. A., & Hogue, T. E. (2019). The role of intuitive moral foundations in Britain's vote on EU membership. *Journal of Community & Applied Social Psychology, 29*(2), 90-103.
- Harrop, I., Roozenbeek, J., Madsen, J., & van der Linden, S. (2022). Inoculation can reduce the perceived reliability of polarizing social media content. *International Journal of Communication, 16*, 1–24.
- Haslam, N. (2006). Dehumanization: An integrative review. *Personality and social psychology review, 10*(3), 252-264.
- Haslam, N., Bastian, B., Bain, P., & Kashima, Y. (2006). Psychological essentialism, implicit theories,

- and intergroup relations. *Group Processes & Intergroup Relations*, 9(1), 63-76.
- Haslam, N., Bastian, B., & Bissett, M. (2004). Essentialist beliefs about personality and their implications. *Personality and Social Psychology Bulletin*, 30(12), 1661-1673.
- Haslam, N., & Whelan, J. (2008). Human natures: Psychological essentialism in thinking about differences between people. *Social and Personality Psychology Compass*, 2(3), 1297-1312.
- Hetherington, M., & Weiler, J. (2018). *Prius or pickup?: How the answers to four simple questions explain America's great divide*. Houghton Mifflin.
- Hobolt, S. B. (2016). The Brexit vote: a divided nation, a divided continent. *Journal of European Public Policy*, 23(9), 1259-1277.
- Homar, A. R., & Cvelbar, L. K. (2021). The effects of framing on environmental decisions: A systematic literature review. *Ecological Economics*, 183, 106950.
- Horgan, J. (1991). Profile: Physicist John A. Wheeler: Questioning the "It from Bit." *Scientific American*, 264(6), 36-41. <http://www.jstor.org/stable/24936938>
- Hradziushka, A. A., Argylov, N. A., & Babuk, P. A. (2024). The Digital Transition of Local Newspapers: New Communication Strategies for Audience Retention. In *2024 Communication Strategies in Digital Society Seminar (ComSDS)* (pp. 60-64). IEEE.
- Huber, M., Van Boven, L., Park, B., & Pizzi, W. T. (2015). Seeing red: Anger increases how much Republican identification predicts partisan attitudes and perceived polarisation. *PLoS one*, 10(9), e0139193.
- Huff, D. (1954). *How to Lie with Statistics*. W.W. Norton & Company.
- Humpherson, D. (2023, February 27). Ed Humpherson to David Pares: Treasury Inflation Infographic. Office for Statistics Regulation. <https://osr.statisticsauthority.gov.uk/correspondence/ed-humpherson-to-david-pares-treasury-inflation-infographic/>
- Hux, J. E., & Naylor, C. D. (1995). Communicating the benefits of chronic preventive therapy: does the format of efficacy data determine patients' acceptance of treatment?. *Medical Decision*

Making, 15(2), 152-157.

Huxley, A. (1978). *The human situation*. Chatto and Windus.

Inman, H. F., & Bradley, E. L., Jr. (1989). The overlapping coefficient as a measure of agreement between probability distributions and point estimation of the overlap of two normal densities. *Communications in Statistics Theory and Methods*, 18, 3851–3874.

Iyengar, S., Lelkes, Y., Levendusky, M., Malhotra, N., & Westwood, S. J. (2019). The origins and consequences of affective polarisation in the United States. *Annual Review of Political Science*, 22, 129-146.

Iyengar, S., Sood, G., & Lelkes, Y. (2012). Affect, not ideology social identity perspective on polarisation. *Public opinion quarterly*, 76(3), 405-431.

Jerit, J., & Zhao, Y. (2020). Political misinformation. *Annual Review of Political Science*, 23, 77-94.

Judd, C. M., Ryan, C. S., & Park, B. (1991). Accuracy in the judgment of in-group and out-group variability. *Journal of personality and social psychology*, 61(3), 366.

Kahneman, D. (2011). *Thinking, fast and slow*. Macmillan.

Kalmoe, N. P., & Mason, L. (2019, January). Lethal mass partisanship: Prevalence, correlates, and electoral contingencies. In *National Capital Area Political Science Association American Politics Meeting* (pp. 1-41).

Keltner, D., & Robinson, R. J. (1993). Imagined ideological differences in conflict escalation and resolution. *International Journal of Conflict Management*.

Kennedy, H., & Engebretsen, M. (2020). Introduction: The relationships between graphs, charts, maps and meanings, feelings, engagements. *Data visualization in society*, 19-32.

Kennedy, H., Hill, R. L., Aiello, G., & Allen, W. (2016). The work that visualisation conventions do. *Information, Communication & Society*, 19(6), 715-735.

Klein, E. (2020). *Why we're polarized*. Simon and Schuster.

Komorita, S. S., & Parks, C. D. (1995). Interpersonal relations: Mixed-motive interaction. *Annual review*

of psychology, 46(1), 183-207.

Kteily, N., Bruneau, E., Waytz, A., & Cotterill, S. (2015). The ascent of man: Theoretical and empirical evidence for blatant dehumanization. *Journal of personality and social psychology*, 109(5), 901.

Kugler, K. G., & Coleman, P. T. (2020). Get complicated: The effects of complexity on conversations over potentially intractable moral conflicts. *Negotiation and Conflict Management Research*, 13(3), 211-230.

Landry, A. P., Ihm, E., Kwit, S., & Schooler, J. W. (2021). Metadehumanization erodes democratic norms during the 2020 presidential election. *Analyses of Social Issues and Public Policy*, 21(1), 51-63.

Landry, A. P., Schooler, J. W., Willer, R., & Seli, P. (2023). Reducing explicit blatant dehumanization by correcting exaggerated meta-perceptions. *Social Psychological and Personality Science*, 14(4), 407-418.

Lane, D. M., & Sándor, A. (2009). Designing better graphs by including distributional information and integrating words, numbers, and images. *Psychological methods*, 14(3), 239.

Lauer, C., & O'Brien, S. (2020). How people are influenced by deceptive tactics in everyday charts and graphs. *IEEE Transactions on Professional Communication*, 63(4), 327-340.

Laustsen, L., & Petersen, M. B. (2016). Winning faces vary by ideology: How nonverbal source cues influence election and communication success in politics. *Political Communication*, 33(2), 188-211.

Lecheler, S., & de Vreese, C. H. (2013). What a difference a day makes? The effects of repetitive and competitive news framing over time. *Communication research*, 40(2), 147-175.

Lecheler, S., & De Vreese, C. H. (2016). How long do news framing effects last? A systematic review of longitudinal studies. *Annals of the International Communication Association*, 40(1), 3-30.

Lecheler, S., Keer, M., Schuck, A. R., & Hänggli, R. (2015). The effects of repetitive news framing on

- political opinions over time. *Communication Monographs*, 82(3), 339-358.
- Lees, J., & Cikara, M. (2020). Inaccurate group meta-perceptions drive negative out-group attributions in competitive contexts. *Nature Human Behaviour*, 4(3), 279-286.
- Lees, J., & Cikara, M. (2021). Understanding and combating misperceived polarisation. *Philosophical Transactions of the Royal Society B*, 376(1822), 20200143.
- Levendusky, M. S. (2013). Why do partisan media polarize viewers?. *American journal of political science*, 57(3), 611-623.
- Levendusky, M., & Malhotra, N. (2016). Does media coverage of partisan polarisation affect political attitudes?. *Political Communication*, 33(2), 283-301.
- Levendusky, M. S., & Malhotra, N. (2016). (Mis) perceptions of partisan polarisation in the American public. *Public Opinion Quarterly*, 80(S1), 378-391.
- Levendusky, M. S., & Pope, J. C. (2011). Red states vs. blue states: going beyond the mean. *Public Opinion Quarterly*, 75(2), 227-248.
- Lewandowsky, S., & van der Linden, S. (2021). Countering misinformation and fake news through inoculation and prebunking. *European Review of Social Psychology*, 32(2), 348-384.
- Li, J. (2023). Not all skepticism is "healthy" skepticism: Theorizing accuracy-and identity-motivated skepticism toward social media misinformation. *new media & society*, 14614448231179941.
- Linde, S. (2020). The politicization of risk: Party cues, polarisation, and public perceptions of climate change risk. *Risk Analysis*, 40(10), 2002-2018.
- Lisnic, M., Lex, A., & Kogan, M. (2024). "Yeah, this graph doesn't show that": Analysis of Online Engagement with Misleading Data Visualizations. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (pp. 1-14).
- Lo, L. Y. H., Gupta, A., Shigyo, K., Wu, A., Bertini, E., & Qu, H. (2022). Misinformed by visualization: What do we learn from misinformative visualizations?. In *Computer Graphics Forum* (Vol. 41, No. 3, pp. 515-525).

- Lotto, L., Tasso, A., Gavaruzzi, T., Carnaghi, A., & Rumiati, R. (2014). How framing and numerical information affect people's judgments when they read a newspaper story. *New Ideas in Psychology, 32*, 33-41.
- Lu, C., Hu, B., Li, Q., Bi, C., & Ju, X. D. (2023). Psychological inoculation for credibility assessment, sharing intention, and discernment of misinformation: Systematic review and meta-analysis. *Journal of Medical Internet Research, 25*, e49255.
- Lu, Y., & Pan, J. (2021). Capturing clicks: How the Chinese government uses clickbait to compete for visibility. *Political Communication, 38*(1-2), 23-54.
- Maertens, R., Anseel, F., & van der Linden, S. (2020). Combatting climate change misinformation: Evidence for longevity of inoculation and consensus messaging effects. *Journal of Environmental Psychology, 70*, 101455.
- Maertens, R., Roozenbeek, J., Basol, M., & van der Linden, S. (2021). Long-term effectiveness of inoculation against misinformation: Three longitudinal experiments. *Journal of Experimental Psychology: Applied, 27*(1), 1.
- Maier, M., Adam, S., & Maier, J. (2017). Does the messenger matter? A comparison of the effects of Eurosceptic messages communicated by mainstream and radical right-wing parties on citizens' EU attitudes. *Journal of Elections, Public Opinion and Parties, 27*(3), 330-349.
- Mansky, J. (2018, November 15). W.E.B. Du Bois' Visionary Infographics Come Together for the First Time in Full Color. *Smithsonian Magazine*. <https://www.smithsonianmag.com/history/first-time-together-and-color-book-displays-web-du-bois-visionary-infographics-180970826/>
- Markant, D. B., Rogha, M., Karduni, A., Wesslen, R., & Dou, W. (2022). Can Data Visualizations Change Minds? Identifying Mechanisms of Elaborative Thinking and Persuasion. In *2022 IEEE Workshop on Visualization for Social Good (VIS4Good)* (pp. 1-5). IEEE.
- Mason, L. (2015). "I disrespectfully agree": The differential effects of partisan sorting on social and issue polarisation. *American Journal of Political Science, 59*(1), 128-145.

- Mason, L. (2018). *Uncivil agreement: How politics became our identity*. University of Chicago Press.
- Mayhew, D. R. (2005). *Divided we govern: Party control, lawmaking and investigations, 1946-2002*. Yale university press.
- McCaffrey, D., & Keys, J. (2000). Competitive framing processes in the abortion debate: Polarisation-vilification, frame saving, and frame debunking. *Sociological Quarterly*, *41*(1), 41-61.
- McCarty, N. (2016). Polarisation, Congressional Dysfunction, and Constitutional Change. *Ind. L. Rev.*, *50*, 223.
- McCarty, N., Poole, K. T., & Rosenthal, H. (2016). *Polarized America: The dance of ideology and unequal riches*. MIT Press.
- McCoy, J., Rahman, T., & Somer, M. (2018). Polarisation and the global crisis of democracy: Common patterns, dynamics, and pernicious consequences for democratic polities. *American Behavioral Scientist*, *62*(1), 16-42.
- McGuire, W. J. (1961). Resistance to persuasion conferred by active and passive prior refutation of the same and alternative counterarguments. *The Journal of Abnormal and Social Psychology*, *63*(2), 326.
- McGuire, W. J. (1964). Some contemporary approaches. In *Advances in experimental social psychology* (Vol. 1, pp. 191-229). Academic Press. [https://doi.org/10.1016/S0065-2601\(08\)60052-0](https://doi.org/10.1016/S0065-2601(08)60052-0).
- McGuire, W. J. (1970). Vaccine for brainwash. *Psychology today*, *3*(9), 36-64
- McNutt, A., Kindlmann, G., & Correll, M. (2020). Surfacing visualization mirages. In *Proceedings of the 2020 CHI Conference on human factors in computing systems* (pp. 1-16).
- Mehlhoff, I. D. (2023). A Group-Based Approach to Measuring Polarisation. *American Political Science Review*, 1-9.
- Mena, P. (2023). Reducing misperceptions through news stories with data visualization: The role of readers' prior knowledge and prior beliefs. *Journalism*, *24*(4), 729-748.
- Mernyk, J. S., Pink, S. L., Druckman, J. N., & Willer, R. (2022). Correcting inaccurate metaperceptions

- reduces Americans' support for partisan violence. *Proceedings of the National Academy of Sciences*, 119(16), e2116851119.
- Mertens, S., Herberz, M., Hahnel, U. J., & Brosch, T. (2022). The effectiveness of nudging: A meta-analysis of choice architecture interventions across behavioral domains. *Proceedings of the National Academy of Sciences*, 119(1), e2107346118.
- Miller, C. H., Ivanov, B., Sims, J., Compton, J., Harrison, K. J., Parker, K. A., ... & Averbek, J. M. (2013). Boosting the potency of resistance: Combining the motivational forces of inoculation and psychological reactance. *Human communication research*, 39(1), 127-155.
- Miller, D. T., & Prentice, D. A. (2016). Changing norms to change behavior. *Annual review of psychology*, 67, 339-361.
- Milosh, M., Painter, M., Sonin, K., Van Dijke, D., & Wright, A. L. (2021). Unmasking partisanship: Polarisation undermines public response to collective risk. *Journal of Public Economics*, 204, 104538.
- Mitchell, V. W., Bakewell, C., Jackson, P., & Heslin, C. (2015). How message framing affects consumer attitudes in food crises. *British Food Journal*, 117(8), 2200-2211.
- Moore-Berg, S. L., Ankori-Karlinsky, L. O., Hameiri, B., & Bruneau, E. (2020). Exaggerated meta-perceptions predict intergroup hostility between American political partisans. *Proceedings of the National Academy of Sciences*, 117(26), 14864-14872.
- Mosleh, M., Pennycook, G., Arechar, A. A., & Rand, D. G. (2021). Cognitive reflection correlates with behavior on Twitter. *Nature communications*, 12(1), 921.
- Mullainathan, S., & Shleifer, A. (2005). The market for news. *American economic review*, 95(4), 1031-1053.
- Mutz, D. C. (1998). *Impersonal influence: How perceptions of mass collectives affect political attitudes*. Cambridge University Press.
- Nassar, M. R., Wilson, R. C., Heasley, B., & Gold, J. I. (2010). An approximately Bayesian delta-rule model

- explains the dynamics of belief updating in a changing environment. *Journal of Neuroscience*, 30(37), 12366-12378.
- Nelson, T. E., Clawson, R. A., & Oxley, Z. M. (1997). Media framing of a civil liberties conflict and its effect on tolerance. *American Political Science Review*, 567-583.
- Newell, B. R., Vigouroux, S., & Greenwell, H. (2023). Expectations, opportunities, and awareness: A case for combining i-and s-frame interventions. *Behavioral & Brain Sciences*, 46.
- Nguyen, V. T., Jung, K., & Gupta, V. (2021). Examining data visualization pitfalls in scientific publications. *Visual Computing for Industry, Biomedicine, and Art*, 4(1), 1-15.
- Nguyen, H. V., Nguyen, C. H., & Hoang, T. T. B. (2019). Green consumption: Closing the intention-behavior gap. *Sustainable Development*, 27(1), 118-129.
- Nicolini, D. (2007). Studying visual practices in construction. *Building Research & Information*, 35(5), 576-580.
- Nielsen, R. K., Schulz, A., & Fletcher, R. (2020). The BBC is under scrutiny. Here's what research tells about its role in the UK. *Reuters Institute for the Study of Journalism*.
- Okan, Y., Janssen, E., Galesic, M., & Waters, E. A. (2019). Using the short graph literacy scale to predict precursors of health behavior change. *Medical Decision Making*, 39(3), 183-195.
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349, aac4716.
- Ordabayeva, N., & Fernandes, D. (2017). Similarity focus and support for redistribution. *Journal of Experimental Social Psychology*, 72, 67-74.
- Ostrom, T. M., & Sedikides, C. (1992). Out-group homogeneity effects in natural and minimal groups. *Psychological bulletin*, 112(3), 536.
- Pandey, A. V., Manivannan, A., Nov, O., Satterthwaite, M., & Bertini, E. (2014). The persuasive power of data visualization. *IEEE transactions on visualization and computer graphics*, 20(12), 2211-2220.

- Pandey, A. V., Rall, K., Satterthwaite, M. L., Nov, O., & Bertini, E. (2015). How deceptive are deceptive visualizations? An empirical analysis of common distortion techniques. In *Proceedings of the 33rd annual acm conference on human factors in computing systems* (pp. 1469-1478).
- Papageorgis, D., & McGuire, W. J. (1961). The generality of immunity to persuasion produced by pre-exposure to weakened counterarguments. *The Journal of Abnormal and Social Psychology*, 62(3),475.
- Park, B. (2018). "How are we apart? Continuity and change in the structure of ideological disagreement in the American public, 1980– 2012." *Social Forces* 96, 1757–1784.
- Park, H. J., & Lin, L. M. (2020). Exploring attitude–behavior gap in sustainable consumption: Comparison of recycled and upcycled fashion products. *Journal of business research*, 117, 623-628.
- Parker, K. A., Rains, S. A., & Ivanov, B. (2016). Examining the “blanket of protection” conferred by inoculation: The effects of inoculation messages on the cross-protection of related attitudes. *Communication Monographs*, 83(1), 49-68.
- Peer, E., Brandimarte, L., Samat, S., & Acquisti, A. (2017). Beyond the Turk: Alternative platforms for crowdsourcing behavioral research. *Journal of experimental social psychology*, 70, 153-163.
- Pew Research Center (2016). Partisanship and political animosity in 2016. *Politics and policy*.
- Pew Research Center (2017). The partisan divide on political values grows even wider. *Trust, facts and democracy*. <https://www.pewresearch.org/politics/2017/10/05/the-partisan-divide-on-political-values-grows-even-wider/>
- Pew Research Center (2020). *Americans Who Mainly Get Their News on Social Media Are Less Engaged, Less Knowledgeable*. <https://www.pewresearch.org/journalism/2020/07/30/americans-who-mainly-get-their-news-on-social-media-are-less-engaged-less-knowledgeable/>
- Pfau, M., & Burgoon, M. (1988). Inoculation in political campaign communication. *Human Communication Research*, 15(1), 91-111.

- Pinker, S. (2015, August 23). *On my radar: Steven Pinker's cultural highlights*. The Guardian. Retrieved from <https://www.theguardian.com/culture/2015/aug/23/on-my-radar-steven-pinker-psychologist-author>
- Playfair, W. (1786), *The Commercial and Political Atlas; Representing, by Means of Stained Copper-Plate Charts, the Exports, Imports, and General Trade of England, at a Single View. To Which are Added, Charts of the Revenue and Debts of Ireland, Done in the Same Manner by James Corry*, London: Debrett; Robinson; and Sewell.
- Prior, M. (2013). Media and political polarization. *Annual review of political science*, 16(1), 101-127.
- Rhodes, M., & Mandalaywala, T. M. (2017). The development and developmental consequences of social essentialism. *Wiley Interdisciplinary Reviews: Cognitive Science*, 8(4), e1437.
- Richards, A. S., & Banas, J. A. (2015). Inoculating against reactance to persuasive health messages. *Health communication*, 30(5), 451-460.
- Robinson, R. J., Keltner, D., Ward, A., & Ross, L. (1995). Actual versus assumed differences in construal: "Naive realism" in intergroup perception and conflict. *Journal of personality and social psychology*, 68(3), 404.
- Roozenbeek, J., Maertens, R., McClanahan, W., & van der Linden, S. (2021). Disentangling item and testing effects in inoculation research on online misinformation: Solomon revisited. *Educational and Psychological Measurement*, 81(2), 340-362.
- Roozenbeek, J., Traberg, C. S., & van der Linden, S. (2022). Technique-based inoculation against real-world misinformation. *Royal Society open science*, 9(5), 211719.
- Roozenbeek, J., van der Linden, S., Goldberg, B., Rathje, S., & Lewandowsky, S. (2022). Psychological inoculation improves resilience against misinformation on social media. *Science advances*, 8(34), eabo6254.
- Rosch, E. H. (1973). Natural categories. *Cognitive Psychology*, 4(3), 328-350.
- Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological bulletin*, 86(3),

638.

Rothbart, M., & Taylor, M. (1992). Category labels and social reality: Do we view social categories as natural kinds? In G. R. Semin & K. Fiedler (Eds.), *Language and social cognition* (pp. 11-36).

London: Sage.

Rothermund, P., & Deutsch, R. (2024). Exaggerating differences back and forth: Two levels of intergroup accentuation. *British Journal of Social Psychology*, *63*(2), 708-722.

Ruggeri, K., Večkalov, B., Bojanić, L., Andersen, T. L., Ashcroft-Jones, S., Ayacaxli, N., ... & Folke, T. (2021). The general fault in our fault lines. *Nature Human Behaviour*, *5*(10), 1369-1380.

Rutchick, A. M., Smyth, J. M., & Konrath, S. (2009). Seeing red (and blue): Effects of electoral college depictions on political group perception. *Analyses of Social Issues and Public Policy*, *9*(1), 269-282.

Santos, C. Q., da Cunha, H. S., Teixeira, C., de Souza, D. R., Tietzmann, R., Manssour, I., ... & Barros, R. C. (2017). Media professionals? opinions about interactive visualizations of political polarisation during Brazilian presidential campaigns on twitter. In *Proceedings of the 50th Hawaii International Conference on System Sciences, 2017, United States*.

Schaffner, B. F., & Sellers, P. J. (Eds.). (2009). *Winning with words: the origins and impact of political framing*. Routledge.

Schultz, P. W., Nolan, J. M., Cialdini, R. B., Goldstein, N. J., & Griskevicius, V. (2007). The constructive, destructive, and reconstructive power of social norms. *Psychological science*, *18*(5), 429-434.

Schulz, A., Levy, D. A., & Nielsen, R. (2019). Old, educated, and politically diverse: The audience of public service news.

Schwartz, S. H., Melech, G., Lehmann, A., Burgess, S., Harris, M., & Owens, V. (2001). Extending the cross-cultural validity of the theory of basic human values with a different method of measurement. *Journal of cross-cultural psychology*, *32*(5), 519-542.

Seely, N. (2018). Virtual vitriol: A comparative analysis of incivility within political news discussion

- forums. *Electronic News*, 12(1), 42-61.
- Senik, C., Stichnoth, H., & Van der Straeten, K. (2009). Immigration and natives' attitudes towards the welfare state: evidence from the European Social Survey. *Social indicators research*, 91, 345-370.
- Seyle, D. C., & Newman, M. L. (2006). A house divided? The psychology of red and blue America. *American Psychologist*, 61(6), 571.
- Slooman, S., Slooman, S. A., & Fernbach, P. (2018). *The knowledge illusion: Why we never think alone*. Penguin.
- Slothuus, R. (2008). More than weighting cognitive importance: A dual-process model of issue framing effects. *Political Psychology*, 29(1), 1-28.
- Slothuus, R., & Bisgaard, M. (2021). How political parties shape public opinion in the real world. *American Journal of Political Science*, 65(4), 896-911.
- Somer, M., McCoy, J. L., & Luke, R. E. (2023). Pernicious polarization, autocratization and opposition strategies. In *Resilience of Democracy* (pp. 61-80). Routledge.
- Snow, J. (1855). *On the mode of communication of cholera*. John Churchill.
- Spampatti, T., Hahnel, U. J., Trutnevyte, E., & Brosch, T. (2024). Psychological inoculation strategies to fight climate disinformation across 12 countries. *Nature Human Behaviour*, 8(2), 380-398.
- Spence, I. (2006). William Playfair and the psychology of graphs. In *Proceedings of the American Statistical Association, Section on Statistical Graphics* (pp. 2426-2436).
- Spiegelhalter, D. (2017). Risk and uncertainty communication. *Annual Review of Statistics and Its Application*, 4, 31-60.
- Stinson, D. A., Cameron, J. J., Wood, J. V., Gaucher, D., & Holmes, J. G. (2009). Deconstructing the "reign of error": Interpersonal warmth explains the self-fulfilling prophecy of anticipated acceptance. *Personality and Social Psychology Bulletin*, 35(9), 1165-1178.
- Stroud, N. J. (2010). Polarisation and partisan selective exposure. *Journal of communication*, 60(3), 556-576.

- Stroud, N. J., & Lee, J. K. (2013). Perceptions of cable news credibility. *Mass Communication and Society, 16*(1), 67-88.
- Swami, V., Barron, D., Weis, L., & Furnham, A. (2018). To Brexit or not to Brexit: The roles of Islamophobia, conspiracist beliefs, and integrated threat in voting intentions for the United Kingdom European Union membership referendum. *British Journal of Psychology, 109*(1), 156-179.
- Swire, B., Berinsky, A. J., Lewandowsky, S., & Ecker, U. K. (2017). Processing political misinformation: Comprehending the Trump phenomenon. *Royal Society open science, 4*(3), 160802.
- Swire-Thompson, B., Ecker, U. K., Lewandowsky, S., & Berinsky, A. J. (2020). They might be a liar but they're my liar: Source evaluation and the prevalence of misinformation. *Political psychology, 41*(1), 21-34.
- Syropoulos, S., & Leidner, B. (2023). Emphasizing Similarities Between Politically Opposed Groups and Their Influence in Perceptions of the Political Opposition: Evidence From Five Experiments. *Personality and Social Psychology Bulletin, 01461672231192384*.
- Tajfel, H. (1969). Cognitive aspects of prejudice. *Journal of biosocial science, 1*(S1), 173-191.
- Tajfel, H. (1981). *Human groups and social categories: Studies in social psychology*. Cup Archive.
- Tajfel, H., & Turner, J. C. (1986). The social identity theory of intergroup behavior. In S. Worchel & W. Austin (Eds.), *Psychology of intergroup relations* (pp. 7–24). Chicago, IL: Nelson Hall.
- Taylor, D. G. (1982). Pluralistic ignorance and the spiral of silence: A formal analysis. *Public opinion quarterly, 46*(3), 311-335.
- Terren, L., & Borge-Bravo, R. (2021). Echo chambers on social media: A systematic review of the literature. *Review of Communication Research, 9*, 99-118.
- Thaler, R. H. (2015). *Misbehaving: The making of behavioral economics*. WW Norton & Company.
- Thaler, R. H., & Sunstein, C. R. (2021). *Nudge: The final edition*. Yale University Press.
- Thibodeau, P. H., & Boroditsky, L. (2011). Metaphors we think with: The role of metaphor in

- reasoning. *PloS one*, 6(2), e16782.
- Thomson, K. S., & Oppenheimer, D. M. (2016). Investigating an alternate form of the cognitive reflection test. *Judgment and Decision making*, 11(1), 99-113.
- Todd, P. M., & Gigerenzer, G. (2012). *Ecological rationality: Intelligence in the world*. OUP USA.
- Todorov, A., Mandisodza, A. N., Goren, A., & Hall, C. C. (2005). Inferences of competence from faces predict election outcomes. *Science*, 308(5728), 1623-1626.
- Traberg, C. S., Roozenbeek, J., & van der Linden, S. (2022). Psychological inoculation against misinformation: Current evidence and future directions. *The ANNALS of the American Academy of Political and Social Science*, 700(1), 136-151.
- Tucker, J. A., Guess, A., Barberá, P., Vaccari, C., Siegel, A., Sanovich, S., ... & Nyhan, B. (2018). *Social media, political polarisation, and political disinformation: A review of the scientific literature*.
- Tufte, E. R. (2001). *The visual display of quantitative information* (Vol. 2, p. 9). Cheshire, CT: Graphics press.
- Tufte, E. R., & Robins, D. (1997). *Visual explanations* (p. 52). Cheshire, CT: Graphics.
- Tukey, J. W. (1977). *Exploratory data analysis* (Vol. 2, pp. 131-160).
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124-1131.
- Tversky, A., & Kahneman, D. (1989). Rational choice and the framing of decisions. In *Multiple criteria decision making and risk analysis using microcomputers* (pp. 81-126). Springer, Berlin, Heidelberg.
- Tyler, M., & Iyengar, S. (2023). Testing the Robustness of the ANES Feeling Thermometer Indicators of Affective Polarisation. *American Political Science Review*, 1-7.
- Uleman, J. S., & Kressel, L. M. (2013). A brief history of theory and research on impression formation. *Oxford handbook of social cognition*, 53-73.
- Undzenas, D., Dunn, K., & Spaiser, V. (2022). Re-examining the EU Referendum vote: right-wing

- authoritarianism and social dominance orientation as indirect trait-level motivation. *Journal of Elections, Public Opinion and Parties*, 32(4), 938-959.
- van der Linden, S., Leiserowitz, A., Rosenthal, S., & Maibach, E. (2017). Inoculating the public against misinformation about climate change. *Global challenges*, 1(2), 1600008.
- van der Linden, S., & Roozenbeek, J. (2020). Psychological inoculation against fake news. In R. Greifeneder, M. Jaffé, E. Newman, & N. Schwarz (Eds.), *The psychology of fake news: Accepting, sharing, and correcting misinformation* (pp. 147–170). New York, NY: Psychology Press.
doi:10.4324/9780429295379-11
- Van Doorn, J., Lemon, K. N., Mittal, V., Nass, S., Pick, D., Pirner, P., & Verhoef, P. C. (2010). Customer engagement behavior: Theoretical foundations and research directions. *Journal of service research*, 13(3), 253-266.
- Van Wijk, J. J. (2006). Views on visualization. *IEEE transactions on visualization and computer graphics*, 12(4), 421-432.
- Vettehen, P. H., Zhou, S., Kleemans, M., d'Haenens, L., & Lin, T. T. (2012). Competitive pressure and arousing television news: A cross-cultural study. *Asian Journal of Communication*, 22(2), 179-196.
- Voelkel, J. G., Stagnaro, M., Chu, J., Pink, S., Mernyk, J., Redekopp, C., ... & Willer, R. (2022). Megastudy identifying successful interventions to strengthen Americans' democratic attitudes.
- Wainer, H. (1984). How to display data badly. *The American Statistician*, 38(2), 137-147.
- Walkup, L. E. (1965). Creativity in science through visualization. *Perceptual and Motor Skills*, 21(1), 35-41.
- Walter, N., Cohen, J., Holbert, R. L., & Morag, Y. (2020). Fact-checking: A meta-analysis of what works and for whom. *Political Communication*, 37(3), 350-375.
- Weber, W. (2020). Exploring narrativity in data visualization in journalism. *Data visualization in society*, 295-311.

- Westfall, J., Van Boven, L., Chambers, J. R., & Judd, C. M. (2015). Perceiving political polarisation in the United States: Party identity strength and attitude extremity exacerbate the perceived partisan divide. *Perspectives on Psychological Science, 10*(2), 145-158.
- Whitt, S., Yanus, A. B., McDonald, B., Graeber, J., Setzler, M., Ballingrud, G., & Kifer, M. (2021). Tribalism in America: Behavioral experiments on affective polarisation in the trump era. *Journal of Experimental Political Science, 8*(3), 247-259.
- Williams, M. J., & Eberhardt, J. L. (2008). Biological conceptions of race and the motivation to cross racial boundaries. *Journal of personality and social psychology, 94*(6), 1033.
- Willis, J., & Todorov, A. (2006). First impressions: Making up your mind after a 100-ms exposure to a face. *Psychological science, 17*(7), 592-598.
- Wilson, A. E., Parker, V. A., & Feinberg, M. (2020). Polarisation in the contemporary political and media landscape. *Current Opinion in Behavioral Sciences, 34*, 223-228.
- Wilson, E. J., & Sherrell, D. L. (1993). Source effects in communication and persuasion research: A meta-analysis of effect size. *Journal of the academy of marketing science, 21*, 101-112.
- Witt, J. K. (2019). Graph construction: An empirical investigation on setting the range of the y-axis. *Meta-psychology*.
- Wojcieszak, M., de Leeuw, S., Menchen-Trevino, E., Lee, S., Huang-Isherwood, K. M., & Weeks, B. (2023). No polarisation from partisan news: Over-time evidence from trace data. *The International Journal of Press/Politics, 28*(3), 601-626.
- Wolak, J. (2020). *Compromise in an age of party polarization*. Oxford University Press.
- Wolf, M. R., Strachan, J. C., & Shea, D. M. (2012). Forget the good of the game: Political incivility and lack of compromise as a second layer of party polarization. *American Behavioral Scientist, 56*(12), 1677-1695.
- Xie, Y., Wang, K., & Kong, Y. (2021). Prevalence of research misconduct and questionable research practices: A systematic review and meta-analysis. *Science and engineering ethics, 27*(4), 41.

Yudkin, D., Hawkins, S., & Dixon, T. (2019). The perception gap: How false impressions are pulling Americans apart.

Zagefka, H., Nigbur, D., Gonzalez, R., & Tip, L. (2013). Why does ingroup essentialism increase prejudice against minority members?. *International Journal of Psychology, 48*(1), 60-68.

Zaller, J. (1992). *The nature and origins of mass opinion*. Cambridge University Press.

Zikmund-Fisher, B. J., Witteman, H. O., Dickson, M., Fuhrel-Forbis, A., Kahn, V. C., Exe, N. L., ... & Fagerlin, A. (2014). Blocks, ovals, or people? Icon type affects risk perceptions and recall of pictographs. *Medical decision making, 34*(4), 443-453.