Manners cost nothing?

Civility on Twitter and political engagement, trust, and affective polarisation in Britain

A thesis submitted in partial fulfilment of the requirements for the Degree of Master of Philosophy in Politics (Comparative Government).

Candidate ID: 1030857
Word Count (excluding End Matter): 29,907

University of Oxford
Trinity 2020

Abstract

The precise mechanisms through which social media affects democratic political discourse are academically underexplored, but pertinent given evidence that the online information environment is characterised by increasing polarisation, factual relativity, incivility, and disengagement. This thesis contributes to scholarly understanding of the implications of social media as a forum for public debate in advanced democracies, and tests the mediating role of issue salience on these effects. To do this, it employs two identically-structured experiments on issues of contrasting current salience in Britain (Brexit and poverty) using Twitter, a prominent social media platform for public debate. The experiments seek to measure the attitudinal effects of exposure to civil and uncivil debate on Twitter for political engagement and trust, focussing on the contrasting reactions of partisans and passive observers. Results indicate that civility in online political debate can significantly improve levels of trust and engagement, particularly for those with low political interest who are unsure of their political orientation. Interestingly, these findings are most consistent in the poverty-related condition and corresponding negative impacts of incivility are not as prominent in either experiment. Concerningly, this implies that incivility is becoming normalised on Twitter and that polarised issue contexts inhibit the ameliorating value of civility. Implications of these findings are discussed with reference to existing literature, and the thesis concludes that attempts to diffuse polarised topics offline while cultivating norms of civility online could help to recast social media as platforms for the inclusion of currently disengaged voices and a force for democratic stability.

Acknowledgements

I am immensely grateful to my supervisor, my family and friends for their enduring support and guidance throughout.

Contents

Part I: The Topic

Chapter 1: Introduction

The online information environment: social media and news consumption

Social media has become part of daily life across the developed world, now embedded within the communications structure of industrialised societies. Defined as "forms of electronic communication through which users create online communities to share information, ideas, personal messages, and other content" (Merriam-Webster, 2019), social media is used for a variety of purposes. Prominent examples are Twitter, Facebook, and Instagram, boasting a combined 4.2 billion global users (Chaffey, 2019). For reasons explained later, this thesis focuses on Twitter in the UK as an example of the multifaceted impacts of social media on politics in modern societies. Given the rapid and exponential growth of these platforms over the last decade (Blank and Lutz 2017:742), academic research like this into different forms of social media in a variety of contexts is a crucial pre-requisite of effective future policy-making.

Twitter began in 2006, and allows users to post ('tweet') and share ('retweet') content of up to 280 characters. In a crucial difference with other platforms, users can 'follow' accounts and interact with others' content without reciprocal acceptance or any prior connection, creating the potential for tweets to reach a large audience very quickly (Duggan and Smith 2016:7). Twitter usage is now extremely widespread, both among the public and elites of various kinds. Actors, sports stars, musicians, journalists and politicians all use Twitter to boost their public profile and interact with the wider public (Khamis et al, 2017). In this context, Twitter's potential advantages for communication are undoubted, bringing together global virtual communities with shared interests (Snowden, 2016), fostering innovation in the workplace (Hanna et al, 2017) and entertainment outside it. The platform also has tangible benefits for healthcare policy exchange (Al-Saleh et al, 2017), combatting loneliness (Allen et al, 2014), facilitating youth activism (Teruelle, 2011), and improving the speed that journalists report information of public importance (Daniels, 2014).

An inevitable product of the free and fast movement of content is news production and consumption on Twitter. Vastly increasing numbers of people now say that social media sites are their primary source of political information. Taking the UK as an example, a 2018 report for communications regulator OFCOM found that 64% of adults use the internet to find news and social media was the most popular source for this (Jigsaw Research 2018:2). Much like other types of content, news on Twitter comes from a myriad of sources and spreads rapidly, often without verification. This challenges established norms around public political socialisation, information diffusion and reliability (Valentino et al 2008; Gil de Zuniga et al, 2014). However, for good or ill, social media appears to a permanent fixture of the communications landscape, and understanding the democratic implications of this new information environment is a crucial endeavour to which this thesis contributes.

Social media and democracy: enlightened public sphere or echo chamber?

The question of whether the internet enhances or threatens the stable functioning of democracy has been extensively debated. Owing to ubiquitous usage and the connectivity advantages above, many scholars have argued that social media (especially Twitter) has the potential to broaden public political conversation on an unprecedented scale (Bode 2016:26). Literature suggests that Twitter facilitates passive news exposure, increasing the socialisation of disengaged demographics into politics. Boczkowski et al show that "consuming news

incidentally on social media has moved from the periphery to the centre of the contemporary repertoire of online information practices, in particular among young people" (2018:3524). This could help bring previously underrepresented voices into the public conversation, and if a healthy democracy should reflect the widest possible array of interests and opinions (Habermas, 1989) such exposure could enrich political discourse. Relatedly, because users can access information and debate quickly with relative ease, hopes have arisen that social networks could precipitate a more politically informed electorate (Papacharissi 2004:260). Indeed, "for many, the internet held out hopes of reinvigorating democracy by encouraging discourse among those of opposing views, one where the status of participants is less important, and where ideas sink or swim by virtue of the strength of their arguments" (Wojcieszak and Mutz 2009:40). Few could question the principle that a diverse exchange of ideas supported by readily accessible evidence is in the public interest (Dahl, 1998). A third potential benefit is increased accountability, as Twitter directly connects the public to politicians, fostering incentives for responsiveness among leaders and weakening detachment between political elites and the public (LaMarre and Suzuki-Lambrecht 2013:361). To exemplify, Robert Cramer, a Swiss politician in his 60s, reflected that "I managed to not be on Facebook; it's my parliamentary assistant who reads my email messages, but Twitter seems to be inevitable!" (quoted in Rauchfleisch and Metag 2015:2413). If citizen access to elites matters for civic life, the advantages of Twitter are clear.

However, the reality of political discourse online is often fragmented, polarised and uncivil. This largely emanates from social media's facilitation of an unprecedented degree of personal choice over information consumption (Sunstein 2018:88). As mentioned above, Twitter allows an individual to build a bespoke network, with consequences for the type of information received. This network selectivity has cultivated the development of niche 'issue publics'; communities that coalesce over a specific viewpoint and become further immersed from interaction with like-minded others (Lawrence et al, 2010). Such 'echo chambers' can have significant impacts, with evidence that "selective information use contributes to increases in extremity and issue voting patterns" (Young 2009: 254). Public debate then becomes easily fragmented, as users with opposing viewpoints consume information that predominantly accords with pre-existing views, which can then become more extreme because of group psychology (Moscovici and Zavalloni, 1969). This phenomenon is exacerbated by the disengagement of those uninterested in politics, a tendency reinforced by the impression that political debate is polarised whenever incidental exposure does occur. This is a concerning trend given that voices without strong partisan passions online tend to be exposed to a greater variety of news sources, and could thus be moderating influences in debate (Cardenal et al 2019:372). While contextually important to note research suggesting that offline news audiences are no less balkanised (Fletcher and Neilsen 2017:476), it is evident that the very features of online discourse that fostered hopes for inclusivity and diversity are increasingly feeding the opposite effects of fragmentation and disengagement.

Echo chambers also threaten information reliability and shared acceptance of facts, two key pillars of stable democracy. Much evidence suggests that the online space is increasingly characterised by informational relativity and denigration of expertise discordant with existing views (Prochazka et al, 2018; Searles et al, 2018). The features of Twitter which allow information and opinion to spread quickly without assessment of its veracity (Hermida 2012:659), combined with the reinforcement of existing views through the aforementioned network homophily, encourages the development of many simultaneous 'realities' as motivated reasoning guides reception of evidence, often leading to dismissal of facts which challenge existing opinions. As Van Aelst et al conclude, "we are witnessing increasing relativism towards facts, evidence and empirical knowledge; a development in which factual

information more and more comes to be seen as a matter of opinion, in which evidence is neglected, and in which misinformation, rumours and conspiracy theories increasingly permeate public discourse" (2017:14). While inaccurate to hold social media solely accountable for these trends, network selectivity online is fuelling echo chambers which in-turn contribute to a dangerous intolerance of opposing opinion and evidence.

Polarisation and incivility

In this online information environment, users with different political viewpoints often have little common ground and interactions become easily polarised. For observers who are not inclined to participate or confident of their views, it fosters the perception that political conversation is divisively hostile and reinforces disengagement (Anderson et al, 2014). Consequently, often only small factions with atypically passionate and hardened opinions participate in debate, resulting in the amplification of unrepresentative minorities. Indeed, recent research in America found that "most news consumers are not partisan in the sources they use to learn about politics. Those that are — a relatively small number of people — drive a vastly disproportionate amount of traffic to partisan outlets and also appear to participate more in politics" (Guess 2019:31). This hyper-partisan polarisation is then worsened by the dynamics of group psychology. A range of evidence shows that members of a broadly like-minded deliberating group usually end the discussion with a more extreme position in the same direction as their prior inclinations (Yardi and Boyd 2010:318). Sunstein identifies three central mechanisms. The first is a limited argument pool, causing heavy imbalance towards ideas that support the group's initial inclination and making opposing points both fewer and weaker. Secondly, individuals frequently adjust their position to match the dominant viewpoint in a group because of social desirability bias. Finally, learning that others share your views exerts powerful reinforcement effects, and such corroboration increases confidence in existing inclinations (2017:101). Linking this to the online information environment, he concludes that "if the public is balkanised, and if different groups are designing their own preferred communications packages, the consequence will be not merely the same but still more balkanisation, as group members move one another toward more extreme points. At the same time, different deliberating groups will be driven increasingly far apart, simply because most of their discussions are with one another" (Sunstein 2017:102).

Polarisation can naturally induce incivility, owing to the lack of unifying influences between opposing camps and the perception that critics have disingenuous or sinister motives (Capellini 2019:435). As discussed above, the denigration of expertise coupled with the proliferation of contradictory information sources online inhibits conflict de-escalation, especially because attempts to correct misleading claims can result in an even stronger commitment to the original view (Nyhan and Riefler, 2010). The structural affordances of Twitter provide particularly fertile ground for incivility, largely owing to the anonymity that users have. Recent American research found that 84% of respondents agreed that "social media are places where people say things while discussing politics that they would never say in person", while "roughly half of users feel the political conversations they see on social media are angrier than in other areas of life" (Duggan and Smith 2016:4). The lack of in-person consequences can give individuals the impression that "someone looking for them would be searching for a needle in a haystack" (McNealy 2011:6), emboldening disproportionate incivility during disagreements. Further, echo chambers and desire for popularity online can glorify incivility as peers often find it entertaining (Sydnor, 2017), not aided by 'troll' accounts that provoke argument and abuse online (Ward and McLoughlin, 2017). Finally, Twitter's 280 character limit creates a need for that bluntness inhibits nuance and exacerbates tensions (Benjamin, 2011).

The prevalence of echo chambers and polarisation create incentives for political elites to adopt corresponding rhetoric to retain support and generate attention online (Iyengar and Westwood, 2015). As an obvious example, US president Donald Trump's use of Twitter to insult opponents in a deliberately provocative fashion means audiences "don't have to like it to get hooked, and the result is to keep the whole country, and much of the world, entranced" (Wu, 2017). Through incendiary rhetoric online, political elites can increase media coverage and boost their public profile, while observers and the public are given signals that normalise incivility, embolden hyper-partisans, and reinforce the perception of division among observers. A multitude of evidence delineates the negative emotional effects of incivility as inducing anger - inhibiting rational evaluation of opinion challenging new information (Massaro and Stryker, 2012) – and aversion – disincentivising participation and precipitating apathy (Gervais 2015; Druckman et al 2017). The potential implications of the online information environment for political socialisation and democratic discourse are thus clear, as is the pressing relevance of research like this into its specificities.

This research: aims and purpose

This research aims to isolate the effect of varying civility in political debate on Twitter for observers. As explained fully later, incivility is conceptualised as a threefold behavioural phenomenon, marked by insults, distortion of opposing positions and negative emotionality (Gervais 2017:389). This is explored in Britain, a fertile ground to investigate online discourse for two main reasons. Firstly, Twitter usage is high: almost 14 million Britons accessed it regularly in 2018, the fourth highest globally (Statista, 2019). Secondly, the political environment is at time of writing strongly polarised over Britain's exit from the EU (Brexit), a divisive issue in which social media has played an active role (see Chapter 2). Britain is thus well-suited to examining the effects of Twitter as a public political outlet.

There is a rich and expanding social science literature around online discourse and its interaction with echo chambers, factual relativity, polarisation, misinformation, and disengagement. This thesis should be seen as complementary to this, while also addressing important gaps. Firstly, comparatively little work has studied Twitter in the UK or the role of political context in its impacts. Given Brexit's evolution into a febrile ongoing debate with sharply balkanised opposing camps and the invocation of identity politics, the UK is currently especially suitable for testing the impact of Twitter during moments of major political salience. Secondly, much research in the areas listed above has inevitably focussed directly on participants in online political discourse, whether elites or their publicly visible supporters and detractors. This work is undoubtedly useful, but risks missing an important aspect of social media's interaction with democracy: the hidden scale and heterogeneity of audiences to this content, many of whom will not participate but will have their opinions and attitudes shaped by what they see, which will affect their subsequent political behaviours. Accordingly, subgroup analysis is the predominant focus of this study, to specifically identify what and who is driving the phenomena described above. In so doing, this thesis aims to contribute to theoretical understanding of the diversity of information consumers on Twitter, the conditions under which their respective attitudes are shaped, and the potential pathways for improvement of the current problems.

In brief, this study employs two identically structured online experiments, with an overall sample size of 1790 participants. Each experiment began with exposure to a tweet from a fictional expert on a political topic and a question assessing their level of agreement with it. The first treatment group then saw connected public debate that was civil and respectful in tone. Conversely, the second treatment group saw debate that was uncivil and highly polarised, while the control group saw no debate. All participants were then asked questions

assessing their a) trust in the public b) trust in experts c) trust in politicians d) likelihood to vote in the next general election. The two treatment groups were also asked a fifth question reassessing their agreement with the expert's tweet. To test whether events that cut across partisan divides and left-right political orientation such as Brexit evoke particular reactions, one experiment focussed on Brexit while the other debated the causes of poverty, symbolic of more traditional partisan divides. A full explanation of the research design follows in Chapter 4.

As discussed at length in chapters 5 to 7, the findings of this study suggest that the extant literature's theoretical focus on the negative impacts of incivility may need to be partially revised to fully capture current realities on social media. Results indicate that civility in online political debate can significantly improve levels of trust and engagement, particularly for non-voters in previous elections, those of low political interest and those unsure of their political orientation. Interestingly, these findings are strongest in the poverty-related condition and corresponding negative impacts of incivility are not as prominent in either experiment. Concerningly, this both implies that incivility is becoming normalised on Twitter and that polarised issue contexts are likely to inhibit the ameliorating value of civility.

Structure and roadmap

This thesis proceeds in three parts. These sections will focus on what is being studied, how it is being studied, and evaluation of the findings. The following chapter will close the opening section and explain the case selection rationale, exploring Brexit and the connected role of social media. Section two (chapters 3, 4 and 5) is method-based, focussing on the study of social media in political science, the research design of this thesis, and the results of the experiments. Finally, section three (chapters 6 and 7) situates the findings within academic context. It explores the study's contribution to existing literature and its implications for future policy-making, with a view to realising social media's potential as forums for deliberative democracy to thrive.

Chapter 2: Why Britain?

Old certainties and new realities: the atypical salience of Brexit

Britain's vote to leave the European Union by popular referendum in June 2016 was a momentous event in the country's modern history. The result surprised pundits, politicians and bookmakers alike, the majority of whom anticipated a Remain victory (Hobolt 2016:1260). The close margin (51.9% to 48.1%) combined with a divisive campaign period and the complexities of extrication from the EU have served as a recipe for polarisation since, and it is commonly agreed that Brexit has had a transformative impact on British politics. Traditionally, post-war political dynamics have been dominated by two competing parties of power, The Conservatives and Labour, competing on a predominantly left/right economic distribution axis (Webb, 2000; Maor, 1997). Discourse had sometimes produced inter-party polarisation and sometimes broad consensus, emanating from Labour's vacillation between left and centre-left policy programs or the varying extent to which the Conservatives endorsed unfettered free-market economics. British politics before 2016 had thus been far from static, but the basis of party competition had proven resilient to social change and shifting impressions of the left-right political divide (Webb 2004:34).

Conservative victory in the 2015 General Election triggered a chain of events which have threatened these established patterns. For various reasons, then Prime Minister David Cameron decided to offer a referendum on EU membership as part of his 2015 election manifesto (Green and Prosser 2016:1300). Notwithstanding eurosceptic prominence, it is noteworthy that there was nothing inevitable about a referendum; EU membership was not among the most salient issues among the public prior to the election (Curtice 2017:25) and the Conservative majority that ensued surprised many analysts. However, the result meant that a nationwide referendum was held on 23 June 2016.

Social media and the EU Referendum: elite and popular influences

As contended in Chapter 1, the profound influence of the referendum on British society provides a fruitful context to examine the increasingly central role that the social media information environment plays in democratic discourse. Indeed, there is already some interesting work in this area (Hanska and Bauchowitz, 2017; Khatua and Khatua, 2016). Given that the implications of the referendum are still unfolding at time of writing, it is reasonable to expect more scholars to follow. This is because it is now widely acknowledged on both sides of the referendum argument that social media had a considerable impact in the campaign; used to target voters, reinforce key messages and spread content with a speed and reach unparalleled by traditional media (Gorodnichenko et al 2018:3).

While perhaps unwarranted to claim that social media decided the referendum, it is clear that Leave supporters had the larger online presence throughout, both in official campaigns and popular discourse (Llewellyn, 2016). There were clear early indications that Leave was ahead online, perhaps showing that that these forums should be taken more seriously in future elections. When it is also considered that "the predominance of Euroscepticism on social media mirrored its dominance in the press" (Hanska and Bauchowitz 2017:27), the referendum result was perhaps less surprising than many analysts thought. This also demonstrates that social media's interaction with other forms of news consumption is an important contextual factor when assessing its influence because it is often a gateway to other information sources (Boczkowsi et al 2018:3528).

The vital importance of online marketing was exemplified by both campaigns spending vast marketing funds on social media advertising. The online space thus became a highly

consequential battleground, such that “the 2016 EU Referendum could be characterised as the first ‘digital referendum’”, with “the internet, social media and new political communication technologies utilised for the purposes of voter registration; fundraising; intelligence gathering; and message dissemination” (Mullen, 2016). Vote Leave (the official pro-Brexit campaign) spent over £2.7 million on Facebook adverts alone, crafting 1433 different messages (BBC, 2018). Through alleged links to consultancy companies Cambridge Analytica and AggregateIQ using data harvesting techniques that have since been the subject of extensive legal and ethical enquiry, the campaign profiled thousands of adults based on their online activity, many of whom would never usually vote. From here, they designed bespoke content for different demographics, which most other people would never see and could not challenge (Cadwalladr, 2017). While such practices encourage voter mobilisation, they also actively contribute to the development of separated publics online and exacerbate polarisation. Morality aside, the widely acknowledged centrality of online campaigning to Vote Leave’s strategy is further evidence that social media is rapidly transforming the electoral domain.

Importantly for this thesis, study of the online debate among members of the public with no formal connection to the campaigns is arguably as important. Owing to social media’s bypassing of the traditional gatekeepers of information dissemination (Usherwood, 2016) and the strength of feeling on both sides, the referendum attracted high levels of mobilisation on Twitter, evident in the popularity of various referendum-related hashtags (Khatua and Khatua, 2016). This type of discussion was not subject to regulation or fact-checking and had potentially vast reach given the ubiquity of incidental news exposure online (Boczkowski, et al, 2018). Unsurprisingly, there is evidence that misinformation about important issues proliferated with ease and that numerous competing claims left many voters unsure of the truth (Polonski, 2016a). Such problems were worsened by ‘bots’ (automated accounts designed to provoke incivility) which were prominent on Twitter during the campaign (Howard and Kollanyi, 2016; Flynn, 2017). These accounts convincingly appear as real people and feed misperceptions about opposing viewpoints, further underlining the difficulties stemming from anonymity online (Gorodnichenko et al 2018:4).

These issues were exacerbated by Leave and Remain supporters forming echo chambers. As perhaps expected for an online debate taking two binary opposing sides, there is evidence of substantial partisan segregation and network selectivity across multiple platforms. Analysing referendum related news on Facebook, Del Vicario et al find that “consumption patterns elicit the emergence of two distinct communities of news outlets” (2017:6). Likewise, Polonski shows that on Twitter “semantic analysis of hashtags reveals an extremely polarised networked public characterised by clusters of hashtags that are rarely used in combination”. He concludes that this had very important consequences, fuelling yet more polarisation because “pro-EU posts are mostly addressing the remain camp, while anti-EU posts are mostly addressing the leave camp” (2016b). Congruent with the theoretical expectations of the previous chapter, research shows that this lack of exposure to alternative arguments hardened opinion on both sides and induced motivated reasoning which inhibited rational evaluation of new information (Carl et al 2019:90). Reflecting this, Bonnachi et al write that many used Twitter during the referendum campaign as a medium to “express their political identities and craft ‘hoped for’ political futures” (2018:175). Combining these insights, referendum discourse was likely to be populated with highly fragmented information distribution and emotionally charged opinion, separated from that of the regulated official campaigns but actively shaping opinions nonetheless (Llewellyn, 2016). This raises wider questions about the suitability of the online information environment for facilitating considered political choices.

Leave and Remain: the new divide

When the campaign had finished, the referendum results revealed stark cleavages within British society. Menon and Salter conclude that "the referendum represented a turning-point in British politics. Debates about it polarised the country in the weeks before 23 June, and on the day itself, a high turnout testified to the mobilisation that had been achieved. Yet the outcome revealed a country profoundly divided by class, by wealth, by education and by geography" (2016:1297). Further work has also demonstrated the importance of age, with over 65s more than twice as likely as under 25s to have voted Leave (Curtice 2017:34). The geographical schisms had two main dimensions, firstly between the four nations of the United Kingdom. England and Wales voted comfortably in favour of leaving, while the reverse was true for Scotland and Northern Ireland (Young 2017:763). Secondly, there was a strong divide between urban (Remain) and rural (Leave) areas. This reflects the "diverging trajectories of economic development and politics taken by locations that have prospered in a globalised knowledge economy - predominantly cities - contrasted with places on the periphery, in towns and rural areas" (Jennings and Stoker 2018:155). Interestingly, there was no apparent connection between EU exposure (as measured by immigration and trade interaction) and likelihood of voting Leave (Becker et al 2017:605), suggesting that campaign arguments and rhetoric exerted a profound influence on the result. Hobolt, Leaper and Tilley neatly summarise the overall political effects of the referendum, writing that "these new identities reflect pre-existing, but less-politicised, social divisions, like age and education, which were mobilised in the context of the referendum and have consolidated into the newly salient identities: 'Leave' and 'Remain'" (2020:3).

Importantly, these identities cut across ideological and party lines with significant division within both left and right-wing voters. Britain's main political parties, formed along a left-right axis of competition, subsequently saw their vote share split on Europe. The majority of Conservatives voters chose Leave (61% to 39%). Labour voters (65%) and Liberal Democrats (68%) largely supported Remain but significant minorities backed Leave (YouGov, 2016). This has created conflicting allegiances for voters and dilemmas for parties, with tentative indications at both the 2017 and 2019 General Elections that Brexit divisions are strong enough to trigger a long-term realignment in party politics (Vaccari et al, 2020; Mellon et al, 2018; Sabbagh, 2019).

Since 2016: the hardening of identities and coarsening of debate

Given the number and clarity of dividing lines and the antagonistic nature of the campaigns, it is not surprising that Brexit-related polarisation has endured long after the referendum result (Usherwood, 2016). The issue has become fertile ground for rising incivility and abuse across elite and popular political discourse, a trend particularly apparent online owing to the previously discussed structural affordances of social media and their concurrent consequences for polarisation (Ward and McLaughlin, 2017).

Events since the referendum have undoubtedly coarsened debate, antagonised divisions and eroded trust in politics, expertise and fellow members of the public (Basillakis et al 2018:29). Trust (or lack thereof) was central to the Leave campaign's arguments, through the portrayal of Remain-supporting politicians as a self-serving political class divorced from the opinions and experiences of the wider public. This populist anti-elite rhetoric proved highly resonant given that the vast majority of MPs supported Remain, exploiting existing mistrust of politicians as a legacy of the 2008 economic crisis (Menon and Salter 2016:1297). Indeed, research shows those with the lowest trust in political leaders were among the most likely to support Leave (Iakhnis et al 2018:1). The campaign's scepticism of politicians was accompanied by a

denigration of expertise and the dismissal of warnings about the economic consequences of leaving (Hobolt 2016:1262). This was neatly encapsulated by prominent Leave-supporting MP Michael Gove's proclamation that the British public has "had enough of experts" (Mance, 2016). Such sentiments actively complemented existing trends towards informational relativity and polarisation online, which have in-turn been exacerbated by post-Brexit political instability.

Indeed, the three and a half years that followed the EU Referendum were a tumultuous period in British politics, marked by two unscheduled General Elections, the creation of new political parties and an unprecedented number of government defeats in parliament (Intal and Yasseri, 2019). In July 2019 Boris Johnson became Britain's third prime minister in as many years, after both David Cameron and Theresa May had resigned for Brexit-related reasons. The complexity and significance of EU extrication have augmented the visceral strength of feeling on both sides and escalated the tone of debate, both among politicians and the public (Elliot 2017:45). Echoing the theoretical impacts of polarisation outlined previously, partisans on both sides have become increasingly extreme in their positions. Prior to Britain's formal exit from the EU in January 2020, many Remain supporters argued that no form of Brexit could be countenanced and campaigned for a second referendum or the revocation of Article 50. At the same time, frustration with the withdrawal negotiations led many Leavers to advocate a 'no-deal Brexit', that is to leave the EU without agreement on future trade or security arrangements. Both of these positions were originally in the extreme fringes of opinion, but became increasingly commonplace with time (Maher et al 2018:205). This is a probable consequence of continued interaction with like-minded others, congruent with echo chambers and group psychology (Sunstein 2018:99). The wider implications of this extreme polarisation clearly undermined the potential for compromise and civil deliberation in the Brexit debate, creating self-reinforcing antagonism and entrenching divisions (Browning 2019:223).

As an illustration of this, there is emerging evidence that Leave and Remain identities are now stronger than party allegiance. Duffy et al write that "the available evidence suggests that we are seeing a fragmentation of political support, alongside affective polarisation related to Brexit identities – and these identities are superseding people's weakening party-political ones" (2019:8). Recent research also found that in 2018, two years after the result, only 6% of people did not identify with either Leave or Remain. When this is compared with the 22% of people without party political attachment, the authors conclude that "Brexit has quickly and dramatically replaced the traditional party allegiances in the hearts and minds of voters" (Evans and Schaffner, 2019). To the extent that majoritarian systems like Britain's depend on some consistency in voting allegiance over time (Golder and Ferland, 2018), Brexit divisions may precipitate electoral volatility long into the future.

Of equal if not greater importance to British democracy is whether Brexit polarisation is so entrenched as to adversely affect public perceptions of others and norms of civility in social interaction. Growing evidence suggests that post-Brexit political divisions are also negatively influencing non-political judgements. In fascinating experimental research, Murray et al find that "EU voting behaviour has a significant moderating effect on pro-social behaviour towards competing partisans" (2017:28) and that Brexit allegiances affect trust in non-political social situations. They show that Leavers and Remainers are consistently more suspicious of an opposing Brexit partisan. Similarly, recent survey work has found that "Remainers and Leavers both describe each other as 'hypocritical', 'selfish' and 'closed-minded', and their own group as 'honest', 'intelligent' and "open-minded'. There is a remarkable symmetry in responses, reflecting the animus that people feel towards the other group". The authors conclude that this "may compromise people's willingness to talk across the political divide and come to mutual understandings" (Hobolt et al 2018:20). Although not an inevitability, with

formal negotiations over Britain's future relationship with the EU expected to take several years, there is concerning scope for the salience of these new political identities to endure and division to continue. Furthermore, if polarised and conflictual political environments induce disengagement among moderates as posited previously, then sources of common ground are likely to be further reduced (Prior, 2007; Pingree, 2011). Indeed, the phenomenon of 'Brexit fatigue' has become increasingly cited, reflecting growing apathy among some citizens (McGee, 2019). This risks weakening elite accountability and public deliberation, potentially threatening the social fabric of civility and acceptance of difference that is fundamental to democratic culture.

Brexit and the normalisation of incivility: uniquely salient or the new status quo?

This chapter leads to three central propositions underpinning the choice of Britain as the context for this research. Firstly, social media played an influential role in the EU Referendum campaign and continues to be prominent in post-Brexit discourse. Secondly, the Brexit debate is profoundly divided and characterised by echo chambers, polarisation, incivility between opposing partisans and mistrust of expertise or opinion-challenging information. Finally, these phenomena are highly consequential and have the potential to cause prolonged instability. If these points are accepted, then post-Brexit Britain is a clearly fruitful context for research exploring the interplay between social media and democracy.

However, it is important not to over-extrapolate from what could be a uniquely polarised period. To test whether the Brexit debate generates abnormal levels of emotion or whether it is entrenched across political discourse, this thesis will conduct two identically structured experiments, one focussing on Brexit and the other on the traditional political cleavage surrounding the causes of poverty. This design will probe whether the impacts of online political debate are mediated by the salience of the issue being discussed, or whether division is the new status quo in the digital age.

The results of the 2017 election, which saw a resurgence of two party politics following a clear choice between left and right-wing economics, demonstrate that traditional fault-lines have not been destroyed by Brexit (Vaccari et al, 2020). Instead, Brexit has introduced a new issue dimension alongside, and the nature of the interaction between them will be a crucial determinant of political outcomes in coming years (Surridge 2019:8). This research design explores the impact of social media on these uncertainties, and the next section focusses on the methodological aspects of this endeavour.

Part II: Methodology and Results

Chapter 3: Studying Social Media

Social media as a site of research

Perhaps unsurprisingly given its increasing prominence in everyday life, social media is attracting interest from scholars across a range of disciplines. In recent years, it has been the focus of attempts to study everything from the mobilisation of social movements (Zhu, 2017) to the persuasive power of brand marketing (Hennig-Thaurau et al, 2015). On top of its popularity with citizens, social media's capacity to track reactions to live events all over the world as well as preserve this content indefinitely is potentially very advantageous for researchers. The methodological leverage of this as a comparative tool is arguably unprecedented, and will only increase as more people join social media (Bruns and Stieglitz, 2013). As McCormick et al write, "a researcher conducting offline observational work is left with only his or her perceptions (and notes) of what transpired. Interaction on social media, however, is preserved and can be reviewed multiple times and passed to other interested researchers" (2015:392). Social media's potential to inspire more rigorous and systematic data analysis is thus unambiguous.

Added to these advantages is the ability of social media to facilitate observation of multi-layered discourses, tailored to specific research interests. It can be argued that "the ways in which people both access and share information about opinions, attitudes, and behaviours have gone through perhaps a greater transformation in the past decade than in any previous point in history" (Japec et al 2015:855). Social media streams acutely reflect this shift. Further, the search and refinement capacities of sites like Twitter enable issues to be studied from an international, national, local or individual standpoint. The availability of large-scale social media data also allows unprecedented network analyses of who communicates with whom, about what, and to what degree (Capella 2017:547). In short, social networks are cacophonous hubs of activity that can be harnessed to capture a diverse range of social phenomena. All of this can be done with comparative ease and avoids the costly recruitment of in-person participants, the difficulties around construction of artificial investigative scenarios, or the need for the fortune of a natural experiment. Social media platforms thus drastically reduce the entry costs of research while simultaneously increasing its hypothetical scope (Gil de Zuniga and Diehl 2017:5). Subsequently, studies utilising social media will undoubtedly become a more prominent aspect of future research landscapes.

Approaches in context: big data, observational studies, and experimental research

Considering its potential uses in such a variety of contexts, it should be expected that scholars employing social media will take differing approaches. Methodological diversity is particularly inevitable because it is impossible to conceptualise social media as a singular entity. Different platforms have varying audiences, capacities and shortcomings which render them suitable for some research purposes and not others. While far from an exhaustive summary, the three central approaches to social media enquiry in political science are big data methods, observational studies and experimental designs (Bruns and Stieglitz 2013:103). This section will briefly outline the principles, usages, strengths and weaknesses of each.

Big data methods are characterised by four key features: the use of large, complex datasets; the frequent involvement of electronic databases; the employment of computational or algorithmic solutions to generate inferences from these data; and the application of

conclusions to social theory in varied domains (Shah et al 2015:7). Datasets can include millions of datapoints, meaning that studies can be executed with sample sizes of hitherto unseen scale (Boyd and Crawford 2012:663). Machine learning and content recognition techniques are used to identify patterns and establish conclusions at a general level of abstraction, minimising the need for often uncertain judgements about the scalability of conclusions emanating from smaller studies. This has encouraged proclamations that the emergence of big data holds the promise of "opening up a new era in the social and behavioural sciences" (Golder and Macy 2012:7), and that "the convergence of big data, social media, and computational social science allows researchers to rethink interpersonal and mass communication" (Capella 2017:545). Political scientists employing big data methods have certainly reflected this level of ambition; attempting to predict electoral outcomes (Tumasjan et al, 2010), quantify the public popularity of individual political candidates (Ceron et al, 2013), and study the interaction between social media and societal prejudices (Brandtzaeg, 2017).

However, big data studies often face criticism for over-reliance on numerical analysis without an accompanying theoretical framework. Gil de Zuniga and Diehl argue that big data advocates often make the flawed assumption that "all that is necessary is to data mine relationships. On the contrary, in order to better understand the fluid, high-speed world of social media and politics, we still rely on basic principles of social science: validity of constructs, theory that stresses context, generalisability, and ethical concerns" (2017:5). A related critique is that such large datasets inevitably filter out sample nuances, giving misleading impressions of simplicity based on broad generalisations. Indeed, Boyd and Crawford criticise a tendency towards "sweeping dismissal of all other theories" which means "other methods for ascertaining why people do things, write things, or make things are lost in the sheer volume of numbers" (2012:666). A variety of other methodological shortcomings of big data social media studies are commonly articulated. These include: the questionable validity of population-wide assertions given social media's particular demographics (Ruiz-Soler, 2017); the lack of identifying information about individual users/datapoints (McCormick et al, 2015); data retrieval complexities (Golder and Macy, 2012), and privacy concerns (Boyd and Crawford, 2012).

These issues mean scholars are continuing to employ conventional methods of causal inference to study social media. One such approach is observational studies. Applied to online discourse, these utilise network analysis techniques to investigate content focussed on particular events or users. Social media's aforementioned search and refinement capacities allow researchers to study particular areas of online discourse in great detail, by collating posts about a particular topic or from a specific group of users (Bruns and Stieglitz 2013:92). From here, patterns and recurring themes can be easily identified and explored. Previous uses of observational techniques to study social media have greatly enhanced scholarly understanding of important behavioural areas such as network selectivity (Cappelini et al, 2019) and the evolution of public debate following events of major significance (Yardi and Boyd, 2010).

Notwithstanding the utility of such findings, observational studies often have limited applicability beyond very specific contexts or events, which leads to the generation of "static snapshots while neglecting the network's dynamics" (Bruns and Stieglitz 2013:103). Observational approaches are also restricted by lack of knowledge about the intentions of the users being analysed. As Marwick and Boyd point out "many users consciously use Twitter as a platform to obtain and maintain attention, by targeting tweets towards their perceived audience's interest" (2010:122). This means that it is impossible to establish the true

relationship between content posted online and an individual's actual opinion, an important facet of social media discourse revisited in greater detail later. Furthermore, similarly to big data methods, the lack of intervention by researchers inherent to an observational study renders it unfeasible to gather any demographic information about samples beyond the superficial indications available online (McCormick et al 2017:415). These last two shortcomings mean that findings have to be assessed at face-value, often resulting in conclusions based on uncertain assumptions.

The final common approach to social media enquiry is experimental research, distinguished by an objective to empirically quantify the impact of specific elements different online settings on those participating or observing. This is achieved through devising a treatment stimulus that encapsulates the phenomena under investigation, then randomly sorting a sample of recruited participants into groups designed to be equally balanced on relevant characteristics. In a simple design consisting of only two groups, one group is exposed to treatment and the other is used as a control, before comparing their reactions. The key idea is that by randomly assigning treatment in a sample otherwise balanced on factors which might affect response to the stimulus, any differences between treatment and control can be attributed as the effect of what is being studied (Altman 2016:172). Such designs are particularly useful for studying the online space because the link between social media and real-life opinion is not assumed and can be directly tested (usually through post-treatment surveys), unlike observational or big data studies. Furthermore, because relevant information about participants is collected to ensure equilibrium on important pre-treatment characteristics between groups, subgroup analysis can be carried out on these variables (such as age or education), revealing insights about heterogeneous effects of treatment on different parts of the sample (Rubin 1974:698). These advantages make experimental methods uniquely suited to this research, which requires an ability to closely monitor outcomes across different characteristics. Examples of other experimental research into social media include: examination of emotional reactions to varied content (Valentino et al, 2008), investigation of factors influencing reception and evaluation of information (Nyhan and Riefler, 2010), and studies of echo chambers and group psychology online (Lee, 2007).

Although its advantages are clear, there are two key criticisms of experimental research methodology. The first pertains to inconsistency in treatment strategies across research into similar phenomena. This means that two experiments on the same topic can produce conflicting conclusions because of contrasting treatment operationalisation or theoretical assumptions (De Vaus 2016:205). Subsequently, the research landscape can become fragmented and lack commonly agreed findings, which hampers credibility. The second challenge is the well-documented difficulty of assuming stable and lasting treatment effects when questions are often asked directly after exposure to stimuli (Bamia et al, 2012). This leaves experimental research vulnerable to claims that it cannot provide behavioural insights beyond short-term reactions unless the researcher has the resources and capacity to re-run studies at different times.

Common Twitter challenges: representativity, validity, and prediction

Thus far, this chapter has discussed the current state of social media research in general, without detailing specificities pertaining to particular platforms. The focus now moves to research using Twitter, both because it is the platform used in this research and it is readily employed in studies of all three methodological persuasions above. Twitter is especially suited to political research for a multitude of reasons. Firstly, because of the lack of barriers to interaction with unknown others, including politicians, journalists and experts, Twitter facilitates political conversation on a wider scale than other platforms (Lassen and Brown 2011:421). As

briefly touched upon above, Twitter also has many beneficial technical affordances. As Blank summarises, “Twitter data have many qualities that appeal to researchers. They are extraordinarily easy to collect. Furthermore, they are available in very large quantities; millions of tweets are not unusual. With a simple character text limit and few options (hashtags, retweets, and links), they are easy to analyse. As a result of these attractive qualities, over 1,400 papers have been published using Twitter data” (2017:690). In sum, the combination of frequent political discussion and advantageous capacities for research make Twitter an increasingly popular forum for social scientists. However, there are common challenges to overcome across methodologies, most prominently issues of representativity, validity and prediction.

Concerns over representativity emanate from systematic differences between the demographic characteristics and personality traits of an average Twitter sample compared with the general population. As a relevant example, evidence shows that “British Twitter users are younger, wealthier, and better educated than other Internet users, who in turn are younger, wealthier, and better educated than the offline British population” (Blank 2017:679). Similar patterns are found across different national contexts (Gil de Zuniga and Diehl 2015:5), though some differences between online and offline populations have disappeared over time, such as a significant initial bias towards men in the early phases of social media that has equalised in recent years (Blank 2017:680). Importantly for future research, this suggests that social trends can quickly evolve and Twitter may become more representative as more users join. However, representativity issues are furthered by the vastly differing rates of participation among Twitter users, with analysis showing that 15% of the most active users account for almost 85% of overall content globally (Leetaru et al, 2013). This results in the amplification of vocal minorities, even among the subpopulation of Twitter users, further complicating generalisability of results. This study uses participants as observers rather than content creators to combat this problem, but more widely Twitter’s demographic profile and participation patterns make it difficult to apply conclusions to the wider population, a fallacy often committed by scholars (Ruiz-Soler 2017:21).

A related but distinct risk is the problematic assumption that content posted on Twitter is a valid reflection of an individual’s true opinion. As mentioned above, many Twitter users build an online profile or persona which is potentially divergent from how they would act or speak in-person. Marwick and Boyd discuss the ‘imagined audiences’ to which many tweets are designed to appeal, writing that users often “target tweets to specific audience members, and conceal or reveal information based on who they imagine to be listening. Some construct a sophisticated model of who may be reading their tweets based on linguistic, cultural, and identity markers in their Twitter stream” (2010:130). Subsequently, establishing the validity of information shared online is difficult in the absence of an intervention by the researcher, as in experimental designs. This problem is exacerbated by fake accounts and marketing ploys on Twitter (Gorodnichenko et al, 2018), and thus the assumption that it provides an accurate reflection of public discourse is often highly questionable.

Combining the concerns around representativity and validity clearly illustrates the perils of using Twitter to make predictions about political outcomes (Gayo-Avello, 2012). While some have attempted this (often using big data methods) with varying levels of success (see Ecker, 2017 for a summary), Twitter has clear limits for predictive inference which scholars should be acutely aware of. As Blank concludes, “for substantive issues requiring data representative of a population, Twitter data are unlikely to be appropriate” (2017:694).

The value of parsimonious and theoretically informed research designs

These challenges underline the importance of methodological awareness if the vast potential of research on social media is to be realised. Robust academic enquiry requires a keen appreciation of the limitations of the chosen method and an adjustment of research goals to reflect these realities (Golder and Macy 2012:7). The benefits of social media (especially Twitter) for facilitating accessible research in political science are undeniable. However, such platforms are not a methodological 'silver bullet' that reduces the need for a sound theoretical framework as the basis for political research (Gil de Zuniga and Diehl 2017:5). On the contrary, the complex pitfalls described above actively enhance the need for flexible theories of modern social realities, because theoretical awareness will naturally encourage parsimony and rigour in research designs. Only studies that are cognisant of the above-cited challenges can overcome methodological barriers and harness social media's extraordinary capacity for insight (Capella 2017:545). This thesis attempts to firmly situate experimental enquiry within a theoretical framework (see chapters 1 and 2), and attention now turns to the research design.

Chapter 4: The Research Design

The three chapters thus far have collectively outlined the theoretical and methodological context for this research design. In brief, social media is an increasingly influential forum for political discourse worthy of intensive scholarly enquiry, and experimental methods represent a promising avenue for this. Twitter is a particularly important platform in this endeavour, owing to the proliferation of news consumption on it as well as its structural affordances which facilitate the linked phenomena of network selectivity, echo chambers and polarisation. Such fragmentation provides a fertile context for incivility between opposing partisans, which has in turn been shown to have a string of emotional effects on observers, of which there are many on Twitter. The frequently uncivil nature of online political discourse is thus plausibly impacting public perceptions of politics and affecting trust, engagement, and information reception. While the potential impacts for democratic stability are self-evident, little is known about how the Twitter information environment interacts with specific issue contexts, or the mediating role of salience on cognitive responses to online content. Accordingly, this research design quantifies the respective impacts of incivility and civility in online debate on British Twitter users, using the contemporary salience of Brexit as a potential conditioner of effects.

Conceptualising incivility

Uncivil behaviour encapsulates a potentially vast range of occurrences, varying by situational context and severity. Nevertheless, scholars generally agree that it “describes a disrespectful discourse that silences or derogates alternative views” (Jamieson et al 2017:206). Previous work has explored ‘milder’ forms of incivility that include name-calling, mockery, character assassination and belittling others (Borah, 2012; Sobieraj and Berry, 2011), or ‘heavier’ incivility using abusive discriminatory language based on ascriptive characteristics and social identities (Munger, 2016; Theocharis et al., 2016). Collectively, incivility is conceptualised as part of a spectrum of linguistic behaviours with civil language at one end, impoliteness such as sarcasm and insults in the middle, and highly uncivil language such as racial slurs at the other end (Kosmidis and Theocharis, forthcoming).

To refine this concept and tailor its application to the online space, this study follows Gervais (2017) in employing a framework of incivility as marked by a combination of three linguistic devices. The first is invectives and ridicule, generally inclusive of verbal abuse using insults that add no new information but are intended to condescend. Secondly, important aspects of online confrontation involve hyperbole and purposeful distortion of opposing arguments. This attempts to misrepresent the sentiment expressed, making the source appear more radical or immoral through inflammatory language (Gervais 2017:389). The final category is emotionality and histrionics, encapsulating emotive vocabulary and the use of capital letters and exclamation marks to convey passion, urgency, or disbelief. The treatments in these experiments (Appendix 1) combined these three devices to simulate the characteristics of hostile discourse online.

Experiment structure

The research consisted of an online survey experiment hosted on Qualtrics, with participants from the online research recruitment platform Prolific. Respondents were presented with a single survey containing pre-treatment information collection, integrated randomised treatment stimuli, and post-treatment questions (Appendix 1). To be eligible, respondents had to be

British citizens over 18 who used Twitter, and had tweeted over 20 times in the last year. This ensured that the sample suited the goals of the research, explained in detail later. The survey took around 5 minutes to complete. Participants were paid £0.50 and were only reimbursed if they finished the survey and entered a completion code on Prolific. The study had a total sample of 1882 participants, evenly divided between the two experiments (n = 942). Responses were collected between 26 November and 1 December 2019. The experiments each consisted of two treatment groups and a control. Respondents were randomly assigned to these groups equally (Tables 11 and 12 in Appendix 2).

The treatment structure proceeded as follows: after answering pre-treatment questions uniform across groups, participants were given a tweet about British politics from fictional expert Dr Richard Thompson, and were informed that he was a lecturer at Bristol University. His Twitter handle "@DrRichardPolitics" strongly implied specialism in political science. They then rated their level of agreement with the tweet on a 5-point Likert scale. Those in Experiment 1 were shown a tweet on Brexit, while in Experiment 2 it focussed on government responsibility for poverty amelioration. Participants were informed of the subject of the tweet to prevent misinterpretation, and the difference in topic was the only change between the two experiments.

A central aim of this study was to study the effects of varying civility on trust in expertise, and thus the construction of the expert's profile, tweet, and background were important choices. Research shows that reception of expertise is conditioned by existing societal prejudices and that marginalised groups are consistently viewed as less credible sources of authoritative information (Searles et al, 2018). While undoubtedly a pressing injustice, constructing an expert with any minority characteristics thus risked confounding the treatment effects. Accordingly, the Twitter profile portrayed a white man with an Anglophonic name, sadly reflecting commonly trusted traits in experts (Prochazka et al 2018:67). The brief background information provided described a university lecturer with a doctorate, utilising the synonymous association of a PhD with expertise. Unlike the expert's profile and occupation, the tweets could not be kept constant across the two experiments because they focussed on different topics. As detailed later, UK Twitter users tend to be younger and more liberal than the wider population. In the contexts of dispositions towards Brexit and poverty, evidence shows that younger liberals were likely to oppose Brexit and are also left-wing on economic distribution (Curtice 2017:21). Given that the sample contained only Twitter users, a majority of respondents would likely be pro-Remain in Experiment 1 (Brexit) and favour state intervention in Experiment 2 (poverty). This presented an opportunity to probe the impact of personal opinion on trust in expertise, and accordingly the expert in Experiment 1 favoured leaving the EU (against the likely majority), but supported state intervention in Experiment 2 (congruent with the likely majority).

The treatments

Dr Thompson's tweet in Experiment 1 read: "Our focus on the immediate costs of EU departure risks blinding us to the costs of maintaining the status quo, say the Foundation for Economic Education. Eurozone growth rates slowed by more than 50% in the second quarter of 2019, and EU heavy regulation is consistently a barrier to UK productivity by inhibiting innovation and investment. A good time to get out?".

Dr Thompson's tweet in Experiment 2 read: "Food bank use in UK reaches highest rate on record as benefits fail to cover basic costs, new research shows. The Tory welfare 'strategy' introduced by Cameron lacks any semblance of human compassion, and the inevitable rise in poverty-related deaths will be blood on their hands. Government must do more, and minsters should be held responsible.".

Treatment groups in each experiment then saw simulated public debate, in the form of replies to the tweet. The tone of this debate was the independent variable in this research, expected to exert an effect on participant reactions. One treatment group saw civil debate, engaging respectfully with opposing viewpoints and acknowledging their merits. The other was exposed to incivility, marked by highly aggressive and profane language and a refusal to constructively engage. The control group saw no debate. Treatment groups were then asked to assess the tone of the debate as either civil or uncivil, a manipulation check to ensure treatment worked as intended. Appendix 1 details the content and presentation of treatments.

All Twitter material across both experiments was fictitious and created solely for this research. However, respondents were not informed of this until after the survey, at which point they were comprehensively debriefed about the study.

Pre-treatment questions and heterogeneous treatment effects

The same set of pre-treatment questions were used in both experiments. The information gathered served two simultaneous purposes, first to facilitate multivariate models using appropriate controls (despite randomisation rendering this less important). Secondly, to provide for subgroup analysis of heterogeneous treatment effects (HTEs). These are central to the research, enabling discernment of how specific characteristics mediate responses to treatment.

Following standard research practices, the first pre-treatment questions collected basic demographic indicators, namely age, gender and education. Such is the ubiquity of effects driven by these variables in academic enquiry that Blank and Lutz refer to them as "the classical variables used in digital research" (2017:745). As discussed later, there are substantial differences between Twitter users and the wider UK population across age and education (Sloan et al, 2015). These variables are thus crucial to assessing how closely the sample matches the target population (Twitter users).

Emulating a framework developed by Diana Mutz (2015:147), a battery of questions was aimed at assessing individual conflict orientation. Participants were asked for agreement with four statements probing disposition towards different aspects of conflictual situations. These were: preference for arguments in-person, level of discomfort at others arguing around them, impression of the role of emotions in disagreement, and enjoyment of challenging someone else's opinion. Based on the combination of answers, respondents were categorised as 'conflict averse', 'conflict agreeable' or 'neutral'. The criteria used for classification is detailed in Appendix 1. Conflict orientation is a potential mediator of reactions to incivility for obvious reasons, given that some individuals react differently to hostility. Further, self-efficacy, defined as "an individual's personal assessment of their ability to organise and execute a course of action to achieve their goals" (Khang et al 2014:50), is an important individual-level trait linked to levels of engagement and cognitive capacity to assess competing information (Pingree

2011:28). These are key interests in this research and to the extent that efficacy is related to confidence, which is in turn likely to condition reaction to conflictual situations, conflict orientation may tap into differences in personal efficacy. Information was then collected on the news consumption habits and political profile of participants, accounting for the possible effect of motivated reasoning. As documented previously, reception of information and situational judgement is often biased by pre-existing opinions, resulting in opposing partisans assessing the same information entirely differently (Taber and Lodge 2006:757). Therefore, the existing viewpoints and political habits of respondents could lead to significant heterogeneity in treatment effects. These questions assessed favoured news mediums, level of political interest, ideological orientation, partisan identity, and past voting habits in the 2017 General Election and 2016 EU Referendum. Subjects were then shown one of the expert tweets detailed above, followed by accompanying debate for treatment groups. The final stage asked all respondents the same set of post-treatment questions, which operationalised the study's dependent variables.

The dependent variables

The experiments sought to investigate the relationship between the often uncivil and chaotic information environment on Twitter and levels of political trust and engagement with the democratic process. Five dependent variables were employed to this end.

Political trust is a multi-faceted concept encompassing impressions of individuals, institutions, occupations, and the collective citizenry (Van Aelst et al 2017:15). Accordingly, three dependent variables assessed different aspects of trust by measuring agreement with the following statements: 'I have trust in fellow Britons to make good decisions'; 'I have trust in politicians to make good decisions'; and 'I have trust in experts to make good recommendations'.

Relatedly, a separate dependent variable measured any direct effects of civility or incivility on reception of expertise. To do this, respondents were asked to reaffirm their opinion of the expert's original tweet after exposure to debate to test whether their position had changed (treatment groups only). This speaks to previously documented emotional impacts of incivility, which suggest that incivility can inhibit rational evaluation of information credibility and decrease respect for authoritative sources (Gervais, 2015). Effects are expected to be usefully indicative but minor because instantaneous opinion switching on a repeated question is rare in survey experiments (Mutz 2015:162).

Finally, to operationalise political engagement, a question probed future voting behaviour. Participants were asked "are you likely to vote in the next General Election?" and presented with a binary yes/no choice. This was particularly pertinent at the time given that the experiments inadvertently took place in the campaign period ahead of the snap General Election on 12 December 2019.

Dual experiments: Brexit and poverty

Much previous research using Twitter has investigated reactions to significant events. Studies show that related Twitter discourse often mobilises opposing partisans, connected through shared hashtags (Lakhiwal and Kar 2016:534). Hashtags enable users to 'tag' their tweets to a certain topic, making it visible to anyone who searches for it. They simultaneously facilitate

generalised and selective discussion, because specific groups often create bespoke hashtags to discuss the same phenomena, while others are descriptive (consider the difference between #VoteLeave and #Brexit). Discussion surrounding highly salient events can therefore become sites of uncivil polarisation, as opposing partisans clash after forming their opinions in their respective echo chambers (Yardi and Boyd 2010:316). Thus, they are unlikely to have much common ground due to reliance on conflicting information sources and the often extreme opinions induced by interaction with like-minded others, and are emboldened to be vocal because of the practical anonymity afforded by Twitter. The vast majority of this activity comes from a loud minority of highly interested users, but the public nature of tweets means that this discourse can be observed by a much wider audience (Bruns and Stieglitz 2013:98).

However, as previously discussed, little is known about how these partisan clashes over salient issues effect such observers. The prolonged and tumultuous Brexit process since 2016 certainly matches the theory above, inducing high levels of emotion, polarisation and enduring public salience (Polonski, 2016b). By contrast, debates around economic distribution rarely ignite comparable passion beyond those strongly interested in politics (Green and Prosser 2016:1301). The dual experiments in this design provide the opportunity to test the mediating role of issue salience, which could plausibly have one of two interactive impacts on the workings of civility or incivility. It might exaggerate treatment effects owing to the high passions evoked by Brexit and increased investment in the issue. Contrastingly, the cumulative impacts of issue fatigue, opinion entrenchment, and the normalisation of Brexit incivility may weaken the power of treatment such that aggressive language exerts less effect.

The selection and recruitment of participants: internal and external validity in online research

The methodological justification for experimental research was explained extensively in Chapter 3 and will not be revisited here. While a very fruitful approach, in-person experiments can be both costly and labour-intensive (Clifford et al 2015:2). A promising solution to these barriers lies in online surveys, which can be efficiently distributed to a wide audience using participant recruitment websites. This method “has the potential to facilitate low-cost experiments in political science with a diverse subject pool” (Berinsky et al 2012:351). Following this model, this study used questionnaire hosting website Qualtrics to create and present the survey, and online research recruitment platform Prolific to recruit and filter respondents.

Scholars have questioned the accuracy and generalisability of online experimental research, rightly stressing the need for careful methodology to mitigate these risks (Blank and Lutz 2017:743). Objections often use the language of internal and external validity. Following Berinsky et al, external validity describes “an assessment of whether the causal estimates deduced from experimental research would persist in other settings and with other samples. Internal validity pertains to the question of whether causal estimates appropriately reflect the effects of the experimental manipulation among the participants” (2012:354). This section will detail these respective criticisms of online studies and demonstrate how this design combats each critique.

The central pertinent question of external validity relates to the generalisability of findings using online methods. Twitter users in the UK and beyond have systematic differences compared with the general population, an issue of great concern for studies making society-wide claims. Twitter users deviate substantially on many politically relevant dimensions including vote choice, turnout, age and education. On average, they are younger and better

educated than non-users, and more liberal, left-wing, and politically attentive (Mellon and Prosser 2017:1). In spite of this, researchers persist in using social media data to make inferences about the population, attempting to predict election outcomes (Tumasjan et al, 2010), quantify the public popularity of political candidates (Ceron et al, 2013), and infer public issue salience (Golder and Macy, 2012). Unsurprisingly, such studies have faced criticism for failing to acknowledge that "social media data cannot be used to generalise to any population other than themselves" (Blank and Lutz 2017:741). Appropriately, this research made no attempt to be representative of the UK population, because its aim was to study the effects of political discourse on Twitter for those who regularly observe it. Research using Twitter that employs a whole population sample will not produce accurate conclusions because it will capture reactions among those who have no experience of Twitter. This will inevitably mischaracterise the true nature of Twitter's real-life effects (Scott et al 2017:312).

A useful advantage of Prolific over other subject recruitment tools such as Amazon's Mechanical Turk (MTurk) lies in the ability to select very specific pre-screening conditions. Prolific participants answer a detailed range of questions prior to signing up, allowing for the construction of bespoke samples based on anything from basic demographics to lifestyle choices or political inclinations. Respondents in this study had to be British citizens over 18 years old who had tweeted more than 20 times in the last 12 months. If successful, these criteria would produce a sample with similar characteristics to UK Twitter users above.

Twitter samples may preclude the extension of inferences to society collectively, but this does not inhibit their utility. Praising work that is transparent about its applicability, Ruiz-Soler reflects that "the Twitter world is not identical to the offline world, but it is entirely real. Users who desire status, admiration, social approval, and attention in their relationships elsewhere will bring those desires with them to Twitter. Individuals must navigate many of the same social obstacles online as they do when they seek information or political support offline. Consequently, the argument that Twitter is not widespread enough becomes irrelevant if we assume that is not possible to make general inferences to the entire population" (2017:20). Twitter users provide a highly useful snapshot of important demographics, and understanding its role in their political socialisation will be crucial to understanding their democratic behaviour. Furthermore, given evidence that after controlling for age, gender and education, no statistically significant differences arise between social media users and non-users on political attention, values or political behaviour (Mellon and Prosser 2017:4), samples can be said to reflect these demographics with increasing accuracy. As usage increases, representativity will only rise.

Another noteworthy threat to external validity relates to the reliability of treatment effects obtained using online subject recruitment platforms. The question here is whether the same selection criteria would hypothetically produce similar survey responses if taken in a different setting. In short, while the first concern related to the generalisability of behaviour among Twitter users, this pertains to the generalisability of behaviour among Prolific users. Firstly, there is recent evidence confirming that self-reported survey responses often correlate with observed social media activity (Guess et al 2019:241), increasing this design's robustness to this critique. Another risk is habitual survey respondents. By definition, Prolific subjects have enough interest in answering surveys to make an account, risking that samples are constructed from regular survey takers who become conditioned to responding in ways that might differ from non-habitual respondents (Blank 2017:685). While this is impossible to

disprove completely due to Prolific's relative infancy and the lack of bespoke research about its users, there is good reason to trust the reliability of results obtained using online samples. Analysis of the closely comparable MTurk revealed that routine responding is uncommon, and that treatment effects are comparable between habitual users and others (Berinsky et al 2012:365). Further confidence is given by the results of a recent study which replicated 20 in-person survey experiments using MTurk, finding that "there is not a single instance in which the samples produce significant effects in opposite directions" (Mullinix et al 2015:121).

Moving to internal validity, scholars have discussed the potential for respondents to misunderstand or be inattentive to online survey experiments, such that treatments do not work as intended (Sloan et al 2015:18). Although previous research shows online respondents are often more attentive research subjects (Clifford et al 2015:4), this study included a manipulation check to safeguard against this possibility. Correct perception of the tone of the debate was vital to treatment and thus exposure to the stimulus was immediately followed by a question asking participants to choose whether the tone was best described as civil or uncivil. Across the two experiments, 92 respondents in treatment groups failed this manipulation check and were removed prior to analysis, leaving a healthy total active sample size of 1790.

Lastly, internal validity would be compromised if subjects completed the survey more than once (De Vaus, 2016). Although the minimal reimbursement (£0.50) makes this unlikely, Prolific employs specific strategies to prevent this. Firstly, it is impossible to take the survey twice using the same account. To prevent participants having multiple accounts, they use IP address detection and require profile verification with an email address and phone number. Finally, the survey was configured on Qualtrics to only allow one completion per device. Given these substantial protections and a statistically robust sample size , it is extremely unlikely that duplicate participation exerted any effect.

Collectively, the research design contains strong safeguards against the common methodological pitfalls of online survey research. These measures are intended to preserve the internal and external validity of conclusions.

Potential hypotheses and mechanisms

In the context of the theoretical background outlined in previous chapters, there are a number of potential hypotheses centred on the relationship between the tone of online discourse and trust, political engagement, and polarisation. Previous work shows that tonality can have substantially contrasting emotional impacts on different groups, mediated by a number of factors (Brooks and Geer, 2007). Numerous characteristics were captured by the pre-treatment variables, allowing a large range of possible heterogeneous treatment effects to be easily measured. The section below details the full sample hypotheses generated prior to execution of the study, itemised by dependent variable. These draw upon previous theoretical insights, and precise evaluation of the proposed mechanisms follows in later chapters. Hypotheses apply across both experiments unless stated otherwise.

Dependent Variable 1: Trust in Britons

H1: Average trust in the public should be lowest in the uncivil treatment and highest in the civil treatment, with the control in the middle. The straightforward mechanism here is that witnessing hostile public conversation will induce a negative perception of citizens, while respectful discussion has the opposite effect.

H1a: It is plausible that effects will be conditioned by 2016 EU Referendum allegiance, particularly in the Brexit experiment. This is because popular sovereignty has been central to the arguments of Leave voters since the referendum, with common invocations of the 'will of the people' (Browning 2019:224). Thus, it is hypothesised that Leave supporters will have higher trust in the public, although still mediated by treatment.

Dependent Variable 2: Trust in Experts

H2: Average trust in authority of expertise will be lowest in the uncivil treatment, while the civil treatment will have no discernible impact. This reflects the likelihood that incivility in disagreement will damage the perceived credibility of the expert among observers.

H2a: The power of incivility over trust in expertise will be weaker among those with higher educational qualifications, because of likely socialisation into the importance of credible information sources and respect for pedagogical credentials.

H3: Trust in expertise will be lower across all groups in the Brexit experiment. This reflects a combination of the political sentiment expressed by the expert and expected sample biases. As outlined above, the dataset is expected to contain large numbers of Remain-supporting left-wing liberals. Given that the expert opinion was pro-Brexit in Experiment 1 and pro state-intervention in Experiment 2, the average result should reflect the majority views of the sample. This is because motivated reasoning is expected to influence information evaluation, which in turn affects trust.

H3a: For the same reasons, treatment effects will be highly influenced by EU allegiance in the Brexit experiment and political orientation in the poverty-related experiment. Thus, if motivating reasoning affects trust, Leave supporters are more likely to exhibit higher trust in experts in the Brexit experiment, as are left-wing individuals in the poverty-related sample.

Dependent Variable 3: Expert Tweet Opinion Change

H4: Exposure to incivility could reduce the levels of agreement with the expert, relative to pre-treatment averages. This finding would further vindicate the idea that incivility reduces trust in the credibility of expertise.

H4a: These effects are expected to be moderated by political interest and education. Higher ratings on these two variables make it more probable that an individual will have more security about the reasons for holding an opinion, making it harder for incivility to induce a change of viewpoint.

Dependent Variable 4: Trust in Politicians

H5: Average trust in politicians will be lowest in the uncivil treatment and highest in the civil treatment, with the control group between them. Incivility is expected to exert a negative impact upon perceptions of political discourse, and because politicians are a direct symbolisation of this, trust in them is likely to be reduced. The same mechanism will have the opposing effect for civil treatment.

H5a. Those with a higher level of interest in politics are perhaps more likely to exhibit stronger trust in the political process and the motivations of politicians. Interest is only expected to be a minor mediator because the statement refers to all politicians, and more politically interested respondents are expected to mistrust elites of an opposing ideological disposition, which will offset gains of interest.

Dependent Variable 5: Voting in the Next General Election

H6: Incivility will induce a small negative average impact on likelihood of voting in the next General Election. The theoretical link here is that perceptions of hostility tend to precipitate disengagement among observers. It is plausible that civil debate might also encourage political engagement, because of improved argument quality and accessibility.

H6a: These effects will be strongly contingent upon political interest. There are two possible mechanisms here. Firstly, those who are more politically interested are intuitively more prone to democratic participation, weakening treatment effects. Some studies have also illustrated that incivility can have vastly contrasting emotional impacts on different groups, such that , incivility can actually induce enthusiasm and engagement among those who have strong views on a matter (Gervais, 2018). Conversely, for those who are not politically interested, the disengaging impact of incivility could be enhanced because it is likely to decrease epistemic political efficacy, furthering the idea that participating in politics is ‘not for them.’ For this same group, civil treatment may encourage voting, because of lowered entry costs to understanding and reduced hostility.

H6b: Secondly, much evidence has documented that voting is strongly a matter of habit (see Clarke et al, 2004), so it is predicted that non-voters in previous elections will be less likely to vote in future and vice versa.

General sample-wide hypotheses

H7: Given that Brexit has been a strongly emotive issue for a prolonged period, it is predicted that the magnitude of effects will be consistently larger in the Brexit experiment than the poverty-related one.

H7a: However, it is plausible that the opposite could also be true, owing to issue fatigue, opinion entrenchment, and the normalisation of incivility in Brexit-related discussions. Thus, it might be that the magnitude of effects is reduced on average in the Brexit experiment.

H8: Those classified as ‘conflict averse’ should be particularly influenced by treatment across the dependent variables. Conversely, those who meet the criteria for being ‘conflict agreeable’

should exhibit reduced average treatment effects. These differences would tap into individual-level variation in confidence and personal efficacy.

H9: Across the experiments, uncivil treatment is predicted to produce larger heterogeneous treatment effects, relative to civil treatment. This reflects the well-documented impacts of incivility on those with different opinions. Put simply, it is hypothesised that incivility will increase opinion polarisation between groups.

***

With these hypotheses in mind, the following chapter turns to results and analysis.

Chapter 5: Results and Analysis

This chapter will present the results of the experiments and explain the findings of the research. It will first describe the characteristics of the sample, before explaining the statistical methods used in analysis. Results will then be outlined, firstly in relation to the potential hypotheses of the previous chapter before discussion of other findings of interest. The final two sections evaluate the sample-wide hypotheses and summarise the central conclusions of the study.

Sample analysis, randomisation check, and statistical methods

As hoped, the sample conforms to the demographic profile of UK Twitter users in most areas. The previous chapter outlined that people on Twitter are on average younger and better educated than non-users, and they are more liberal, left-wing, and politically attentive (Mellon and Prosser 2017:1). As shown by Tables 1 to 12 in Appendix 2, the sample matches these characteristics and shows heavy skews towards educational qualifications, youth, left-wing orientation, high political interest and Remain support (a realistic indicator of liberal views - see Curtice, (2017)).

Other points of note include a potentially unrepresentative bias towards women (62.5%). Nevertheless, given evidence that gender differences on social media have virtually disappeared on Twitter in the UK (Blank and Lutz 2017:744) and that women were historically more regular users than men (Anderson, 2015), this should not have exerted much impact. This is supported by the lack of widespread gender-driven heterogeneous treatment effects. The sample also includes a markedly higher than population average number of voters in previous elections (84.4% in the 2017 General Election and 81.9% in the EU Referendum). This is further evidence that the sample matches Twitter users in their inclination towards political interest.

Random assignment is a key advantage of experimental research and lends confidence that any treatment effects found are not due to a high concentration of confounding factors within groups (Rubin 1974:688). To confirm that randomisation was successful, t-tests were conducted on individual pre-treatment variables and multinomial regressions studied them as a group. For concision, only treatment group p-values for all relevant characteristics are presented in Appendix 2, which show that randomisation was roundly robust. With the exception of EU Referendum voting allegiance in some groups, none were statistically significant at the 10% level or lower (Table 13). The t-tests also suggest successful randomisation, reporting comparable results across groups for all tested variables in both experiments.

A variety of statistical methods were employed in the analysis. In the interests of comparability across variables and consistency of interpretation, results reported below are predominantly drawn from OLS and general linear regression models. The coefficients from ordered logit models can be somewhat difficult to interpret because they are scaled in terms of logs, making the relationships between variables somewhat opaque (see Capuano et al, 2016). However, it is acknowledged that ordered logits are the most statistically appropriate estimator for the three ordinal dependent variables relating to trust (Williams 2016:7). As indicated by Tables 14 to 19 in Appendix 2, analysis carried out using this method returned estimates and predicted probabilities with similar directionality and magnitude to the OLS results across the trust measures. There is a substantial methodological literature about the trade-offs associated with different approaches, but many studies show that OLS provides reliable and comparable estimates in properly specified models for research of this type (Moulton, 1986; McNeish, 2014; Huang, 2018; Astin and Denson, 2009). For ease of presentation and interpretation

estimates cited in-text are bivariate results, but these remain roundly robust to controls, further validation of the randomisation process. Appendix 2 contains all referenced tables and figures that are not in-text.

Note that for some subgroups, estimates of heterogenous treatment effects (often denoted by β) are a calculation based on the size of interaction effects relative to others in the model output of the relevant table. For the method used to determine treatment effects from tables of this type, please see the interpretation note found beneath Table 21 in Appendix 2. However, in all cases where using this method is necessary, relevant estimates are listed beneath the table. To avoid the presentation of superfluous information, tables of subgroup effects do not display every category, focussing only on those relevant to the analysis.

Dependent Variable 1: Trust in Britons

This variable was operationalised by asking participants to rate their impression of the statement 'I have trust in Britons to make good decisions' using a 5-point Likert scale of agreement.

Potential hypotheses

Hypothesis 1 posited that incivility should lower the average level of trust in the public to make good decisions relative to control, while civility should raise it. The results strongly vindicate this in both experiments, though effects are larger in the poverty-related treatment. In this condition, participants expressed mild disagreement with the statement on average across all groups. Incivility induced a statistically significant reduction of -0.20 (SE = 0.08) from the control group mean of 2.42, (Table 20 overleaf), while civility exerted an even larger and equally significant positive deviation of +0.27 (SE = 0.08), leaving the overall difference in averages between treatment and control groups at 0.47 (Figure 1). This pattern is mirrored in the Brexit experiment, with incivility -0.10 (SE = 0.09) lower than control, and civility +0.20 (SE = 0.09) higher, a statistically significant overall average difference of 0.30. Results are presented in Table 20 and Figure 2. Subgroup analysis demonstrates that these findings are mirrored across education, age and gender, having the strongest effects on the youngest across both experiments and lowest educated in the Brexit condition (Figures 3, 4 and 5). Collectively, these results substantially support the proposed mechanism that witnessing hostile public conversation erodes trust in citizens, whilst civil discussion has the opposite effect.

Table 20: Full Sample Trust in Britons (Brexit and Poverty)

Trust in Britons						
	Bivariate OLS (Brexit)	*Bivariate GLM (Brexit)*	*Multivariate OLS (Brexit)*	*Bivariate OLS (Poverty)*	*Bivariate GLM (Poverty)*	*Multivariate OLS (Poverty)*
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	2.373***	2.373***	2.243***	2.421***	2.421***	2.384***
	(0.061)	(0.061)	(0.550)	(0.057)	(0.057)	(0.365)
Brexit T1: Civility	0.196**	0.196**	0.207**			
	(0.089)	(0.089)	(0.087)			
Brexit T2: Incivility	-0.101	-0.101	-0.087			
	(0.088)	(0.088)	(0.087)			
Poverty T1: Civility				0.272***	0.272***	0.256***
				(0.081)	(0.081)	(0.081)
Poverty T2: Incivility				-0.196**	-0.196**	-0.185**
				(0.082)	(0.082)	(0.082)
Observations	894	894	894	896	896	896
R^2	0.012		0.071	0.034		0.080
Adjusted R^2	0.010		0.056	0.032		0.065
Multivariate Controls	Age, Education, Gender, Conflict Disposition					
Note:						*p<0.1; **p<0.05; ***p<0.01

Interpretation guide for all full sample tables: *The constant here is the average estimate for the control group, and the coefficients for T1 and T2 represent the difference between this figure and the average in the respective treatment group.*

Note: As can be seen, civility significantly increases trust in Britons in both experiments, robust to different methods and the inclusion of covariates. Incivility has the opposite effect, significantly in the poverty condition.

Figure 1: Full Sample Trust in Britons by Treatment Group (Poverty)

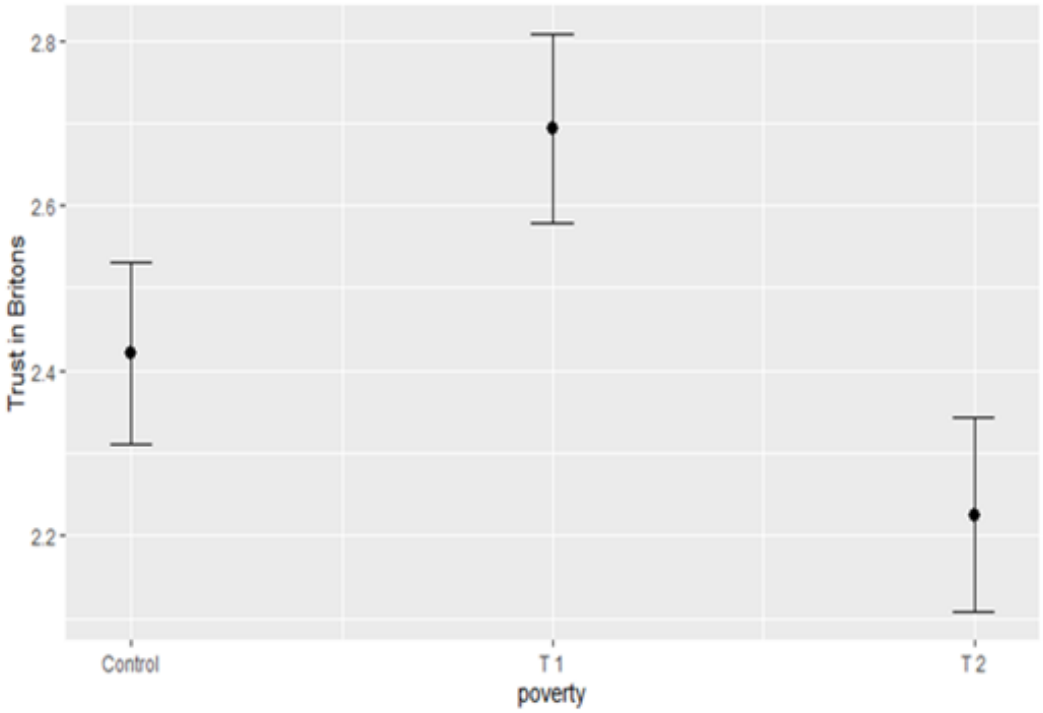

Interpretation guide for this graph and all others: *T1 = Civil treatment, T2 = Uncivil treatment. Centre points represent the average estimate for the group, while the stems show 95% confidence intervals. Method: Bivariate OLS.*

Note: T1: β = 0.27, T2: β = -0.20. Full estimates in Table 20.

Figure 2: Full Sample Trust in Britons by Treatment Group (Brexit)

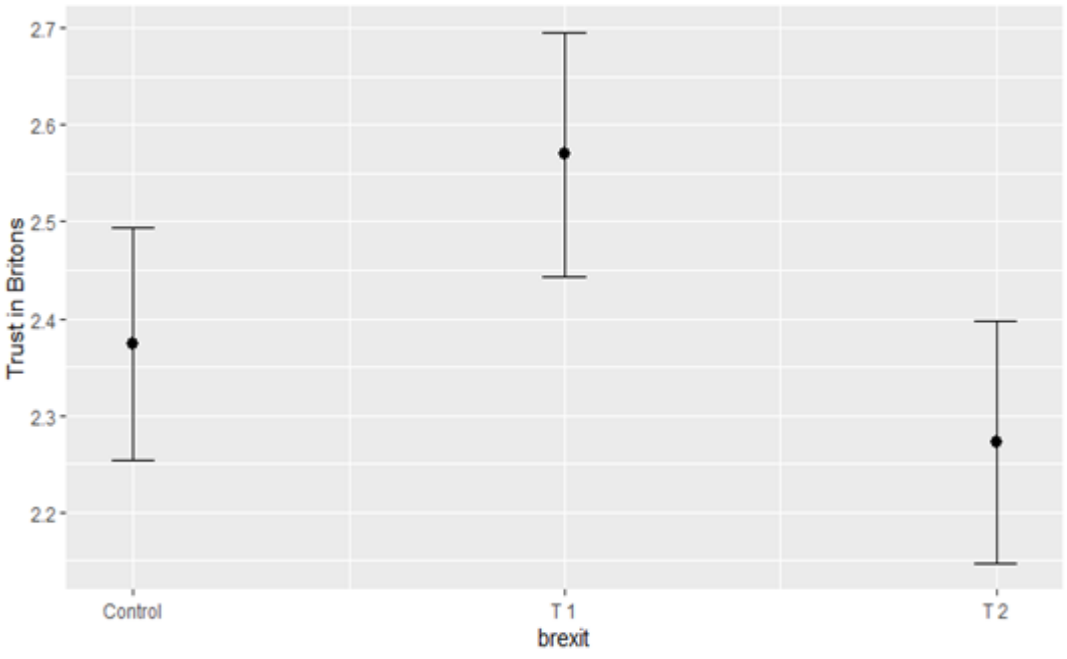

Note: T1 (civility): β = 0.20, T2 (incivility): β = -0.10. Full estimates in Table 20.

It was also predicted that these effects would be conditioned by EU Referendum allegiances in both experiments, particularly in the Brexit condition (Hypothesis 1a). This is again proven correct, those who voted to remain in the EU had markedly lower average trust in the public than Leavers, statistically significant across both experiments but strongest in the Brexit stimulus, as predicted (β = -1.26 in control groups, SE = 0.15; see Tables 21 and 22, Figures 6 and 7). The posited mechanism here reflects the political arguments made by Leave supporters since the referendum, which have commonly invoked the idea that Brexit represents the "will of the people" (Browning, 2019) and notions of popular sovereignty. As such, it follows that advocates of these ideas will have higher trust in the public to make good decisions, just as the majority did in 2016 in their view. These differences between Leavers and Remainers were still substantially conditioned by treatment in both cases, and the larger polarisation observed between camps exposed to Brexit-related stimuli reflects the centrality of ideas around popular sovereignty to this issue.

Figure 6: Trust in Britons by 2016 EU Allegiance (Brexit)

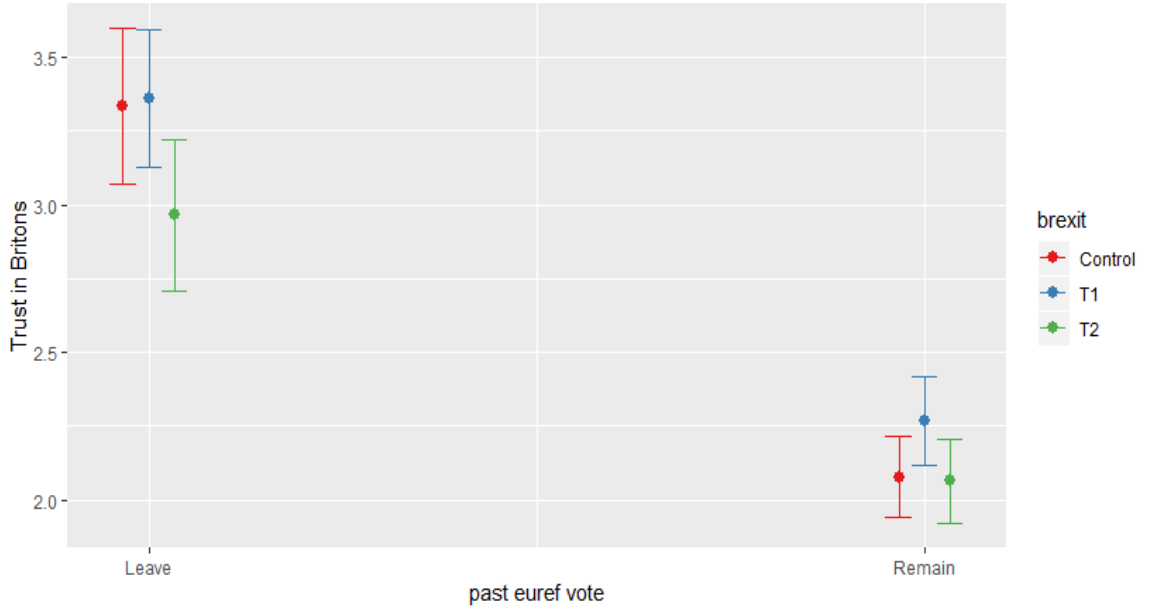

Interpretation guide for all graphs of heterogenous treatment effects: *Red = Control, Blue = T1 (Civil treatment), Green = T2 (Uncivil treatment). The word "brexit" (or "poverty" where appropriate) above the legend denotes the topic of the experiment.*

Note: There are clear and significant differences between Leavers and Remainers in average of trust in Britons across treatment groups (β = -1.255 between control groups, SE = 0.152). Full estimates in Table 21.

Figure 7: Trust in Britons by 2016 EU Allegiance (Poverty)

Note: There are clear and significant differences between Leavers and Remainers in average trust in Britons across treatment groups (β = -0.588 between control groups, SE = 0.148). Full estimates in Table 22.

Other findings

Beyond the predicted hypotheses, it was also found that political orientation and partisan identification are mediators of trust in the public, with right-wingers more likely to have higher faith in Britons (Figure 8). This would perhaps follow a similar mechanism to above, given that those on the right are also more favourably pre-disposed to popular sovereignty (Plattner, 2010). There was also evidence that degree of interest in politics conditions treatment outcomes, with the strongest treatment effects induced for the politically interested in both experiments (Figure 9). Incivility reduced trust by 0.49 (significant) comparative to civility for the most highly interested in the poverty condition (Figure 10). This could indicate that greater investment in political discourse raises importance placed on public trust, perhaps because of increased awareness of the consequences of public decisions for democratic outcomes.

Dependent Variable 2: Trust in Experts

This variable was operationalised by asking participants to rate their impression of the statement 'I have trust in experts to make good recommendations' using a 5-point Likert scale of agreement.

Potential hypotheses

Hypothesis 2 posited that average trust in expertise would be lowest in the uncivil treatments, while the civil treatments would have no discernible impact from control. Some support is found for this, but results are statistically insignificant and coefficients are of limited magnitude (Table 23). Across both the Brexit and poverty-related experiments, averages indicate a consistently lukewarm tendency in favour of trust in experts. Results in the Brexit condition conform with hypothesised directionality, civil debate has indistinguishable impact while

incivility induces a negative effect of almost -0.10 (SE = 0.09) on trust in experts. Interestingly, both treatments exert minor negative impacts in the poverty-related condition, perhaps indicating the social power of opposing arguments in general, but the effect of incivility (β = -0.13, SE = 0.09) is marginally larger and closer to significance. The results provide statistically fragile support for the proposed mechanism that incivility in debate about an idea proposed by an expert will damage the perceived credibility of the source, by inducing uncertainty in evaluation of information.

Moving to related heterogeneous treatment effects, Hypothesis 2a predicted that treatment effects would be weaker among those with higher educational qualifications. This reflected the likelihood that such people will have been socialised to a greater degree into respect for the importance of credible information sources and pedagogical credentials. Figure 11 (below) and Figure 12 illustrate that this was indeed found in both experiments, though the effect was markedly larger in the Brexit condition, reflecting the substantially greater impact of incivility on less educated respondents (β = -0.63, SE = 0.34). This is a large and significant effect (p = < 0.1) and is likely indicative of the mediating role of issue salience, though caution is required due to the small number of less educated respondents.

Figure 11: Trust in Experts by Highest Educational Qualification (Brexit)

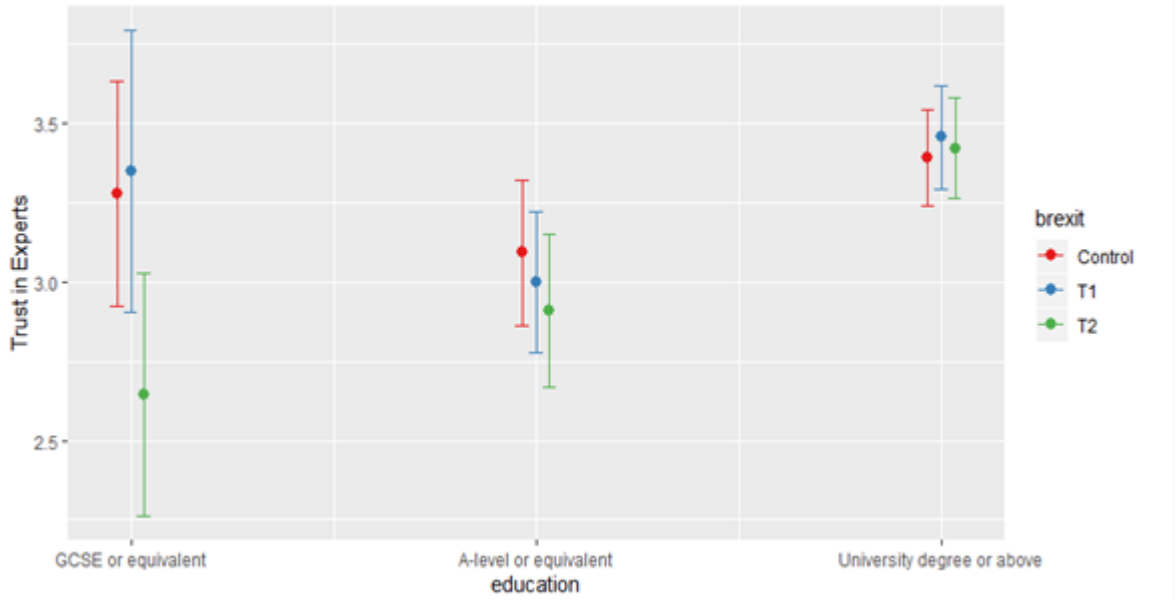

Note: β = -0.63 (SE = 0.34) for the least educated participants in the uncivil treatment (T2).

Hypothesis 3 expected that trust in expertise would be conditioned by the average levels of agreement with the ideological position taken. In the Brexit experiment, the expert favoured leaving the EU, while in the poverty-related experiment he advocated greater state intervention and welfare provision. If the hypothesis was supported, this would be manifested in lower trust in experts across all groups in the Brexit experiment comparative to the poverty-related condition, as a result of UK Twitter users being more likely to be left-wing liberals. Interestingly, little evidence is found to support this. Average estimates across all 6 groups spanning the two experiments are broadly similar, ranging from 3.30 to 3.21 between groups. Somewhat encouragingly for democratic stability, this consistency suggests that aggregate assessments of expert credibility online are not always the product of motivated reasoning.

Owing to the same mechanism of biased information processing according to congruence with ideological priors, Hypothesis 3a posited that treatment effects would be highly influenced by EU allegiance in the Brexit experiment and political orientation in the poverty-related condition. If true, results would show high levels of polarisation between Leave and Remain supporters, and left and right wingers respectively. This is weakly vindicated in the poverty-related experiment, where left-wingers collectively reported marginally higher trust in expertise than those on the right, although mediated insignificantly by incivility. Interestingly, expected effects are not present in the Brexit condition, as average estimates are actually marginally larger across treatment groups among Remainers in spite of their likely disagreement with the expert (Figure 13). This is perhaps indicative of the frequent denigration of expertise among Leave advocates in public discourse (Mance, 2016), and suggests that the marginally higher trust in expertise among Remain supporters may have been markedly greater if the expert had taken a pro-EU position. Thus, there is somewhat contrasting evidence of motivated reasoning at the full-sample and subgroup levels, but collectively it is not overwhelming.

Other findings

However, other mediating variables reveal some interesting trends. Incivility appears to exert differing effects on trust in expertise according to age, gender, and political interest. There is evidence that interest conditions trust in both experiments, and as shown by Table 24 (overleaf) and Figure 14, incivility exerts a statistically significant negative impact for the most interested in politics in the poverty condition ($\beta = 0.48$). This mirrors earlier-cited findings for trust in the public, and is a surprising result that runs contrary to theory suggesting that incivility should have a greater impact on the politically disengaged (Pingree, 2011). Perhaps this needs challenging; it could be that the demobilised already assume that political conversation on Twitter will be hostile, polarised, and to be avoided, such that incivility no longer shapes perceptions. There is also evidence of a significant reduction in expert trust emanating from uncivil treatment among the youngest in the poverty-related condition ($\beta = -0.37$, $SE = 0.17$, see Figure 15), as well as treatment effects varying significantly by gender. Women were both less likely to trust experts on average across all groups, and more impacted by incivility ($\beta = -0.20$, $SE = 0.11$, see Figure 16). Collectively, these trends are potentially alarming. Twitter tends to contain a high proportion of women, youth, and the politically interested, and according to these results such groups trust expertise less when accompanied by incivility in certain contexts, as it often is online. The lesser degree of these effects in the Brexit condition perhaps serves as further evidence that, for some, issue salience dulls the effects of tonality because of anticipated polarisation, an equally consequential conclusion.

Table 24: Trust in Experts by Political Interest (Poverty)

Trust in Experts		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (Control Group: Not at all Interested)	2.429***	2.046***
	(0.399)	(0.400)
T1: Civility (Not at all Interested)	0.460	0.641
	(0.532)	(0.527)
T2: Incivility (Not at all Interested)	0.571	0.671
	(0.520)	(0.514)
T1: Civility (Very Interested treatment effect differential)	-0.854	-1.001*
	(0.555)	(0.550)
T2: Incivility (Very Interested treatment effect differential)	-1.050*	-1.121**
	(0.545)	(0.538)
Control Group: Very Interested	1.307***	1.344***
	(0.413)	(0.409)
Observations	896	875
R^2	0.037	0.079
Adjusted R^2	0.022	0.057
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Statistical output for heterogeneous treatment effects does not produce straightforward estimates for certain subgroups. A full guide for how to interpret tables of this type can be found underneath Table 21 in Appendix 2. Estimates for the Very Interested subgroup in the table above are as follows (shown as a graph in Figure 14):

Control: 3.736 (2.429 + 1.307)

T1: β = -0.394 (0.460 – 0.854)

T2: β = -0.479 (0.571 – 1.050)

Dependent Variable 3: Expert Tweet Opinion Change

This variable applied to treatment groups only and was operationalised by asking participants to reconsider their opinion of the expert tweet stimulus after seeing connected debate, on a 5-point Likert scale.

Potential hypotheses

Hypothesis 4 predicted that incivility would reduce levels of agreement with the expert, relative to pre-treatment averages. Results are presented in Table 25. As expected, there were not huge swings of opinion change, but some support was found for this across both of the experiments, stronger in the Brexit condition. Here, while the civil group average remained constant in practical terms, aggregate levels in the uncivil treatment group declined in a minor but significant way ($\beta = -0.10$, SE = 0.05). In the poverty-related condition, respondents also tended to be indicatively less likely to agree with the expert after uncivil debate ($\beta = -0.05$, SE = 0.04). However, as found with regard to trust in Hypothesis 2, the table shows a similar effect in the civil treatment group ($\beta = -0.07$, SE = 0.03), this time significantly. This shows that the mere presence of debate precipitated a minor decrease in agreement with the expert. These results collectively lend slight support to the idea that incivility can reduce perceptions of expert credibility, as was found in relation to trust above.

Switching to probable heterogeneous treatment effects, Hypothesis 4a held that the likelihood of a respondent updating their opinion would be moderated by political interest and education. Both characteristics are indeed substantial mediators of agreement (Figures 17 and 18). More interested and educated respondents on average more strongly opposed the pro-Brexit opinion and supported the pro-welfare opinion in their respective experiments, and vice versa. However, while education does not appear to exert much impact on probability of opinion changing, political interest is a notable conditioner. The highly interested display largely stable opinions before and after treatment, consistent across civility and incivility. However, rates of opinion change among lowest interest participants are far greater in all four treatment groups, ranging from +0.25 to -0.44 (Tables 26 and 27). Directionality of opinion updates varies between the topics. In the poverty-related condition (Figure 20 and Table 27), both civil and uncivil debate induce a significant movement away from agreement with expertise for the least politically interested, of considerable and similar magnitude ($\beta = -0.44$ and $\beta = -0.30$ respectively). These findings again imply that the mere presence of debate can change initial perceptions of expertise, regardless of tonality. This is congruent with the mechanism proposed, reflecting the decreased likelihood of low interest citizens to be assured of their political opinions due to lack of investment. In the Brexit condition on the other hand (Table 26 overleaf and Figure 19), incivility induces an average effect of 0.25 in favour of agreement with the expert, while civilised debate has the opposite effect ($\beta = -0.40$). This represents a sizable overall differential of 0.65 between treatment groups ($p = < 0.1$), although the power of this result is again hampered by large confidence intervals, reflecting the small number of low interest participants. Nonetheless, the positive directionality of uncivil treatment may indicate that less active observers rely more on the authority of expertise when debate seems conflictual. Alternatively, a more pessimistic opposing mechanism holds that when politically uninterested people encounter hostile debate on a salient issue, it induces motivated reasoning. This could be supported given that politically uninterested respondents were 9.17% more likely to have voted to leave the EU.

Table 26: Expert Tweet Opinion Change by Political Interest (Brexit)

Expert Tweet Opinion Change		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (T1: Civility - Not at all Interested)	-0.400	-0.496*
	(0.262)	(0.276)
T2: Incivility (Not at all Interested)	0.650*	0.714*
	(0.393)	(0.399)
T1: Civility (Very Interested)	0.413	0.556*
	(0.270)	(0.285)
T2: Incivility (Very Interested treatment effect differential)	-0.674*	-0.739*
	(0.403)	(0.411)
Observations	578	562
R^2	0.021	0.044
Adjusted R^2	0.006	0.012
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: There are no control groups in this dependent variable, thus the constant now represents the average level of opinion change in civil treatment (T1) for those in the first listed subgroup (Not at all interested). Effects were considerable for this cohort, unlike estimates in the 'Very interested' subgroup, as seen below:

T1: 0.013 (-0.400 + 0.413)

*T2: -0.011 (**a.** 0.650 – 0.674 = -0.024. **b.** 0.013 - 0.024)*

Other findings

Somewhat predictably, agreement with the expert was heavily conditioned by ideological disposition, though this was partially influenced by treatment. In the Brexit condition, agreement was far higher among Leave supporters in both treatment groups, while the tweet about poverty exhibited far greater support from left-wingers who voted accordingly at the last election (Figures 21, 22 and 23). Interestingly, in the Brexit condition only, age may mediate likelihood to update opinion after treatment, with younger people moving away from agreement with expertise after incivility by -0.24, considerably greater than the impact of civil debate (Table 28). As above with regard to low political interest, this result is tentatively supportive of theories of motivated reasoning impacting assessment of expert credibility under conditions of incivility, given that younger people tended towards supporting Remain. To the extent that political debate on Twitter is often polarised and uncivil, this could imply that objective evaluation of expertise is unlikely to be aided by increasing youth participation on social media.

Figure 21: Expert Tweet Agreement by 2016 EU Allegiance (Brexit)

For graphs of this dependent variable only, Red now represents T1 (civility) while Blue shows estimates for T2 (incivility).

Note: Average levels of agreement with the expert opinion are significantly higher among Leavers than Remainers, regardless of treatment group.

Figure 22: Expert Tweet Agreement by Political Orientation (Poverty)

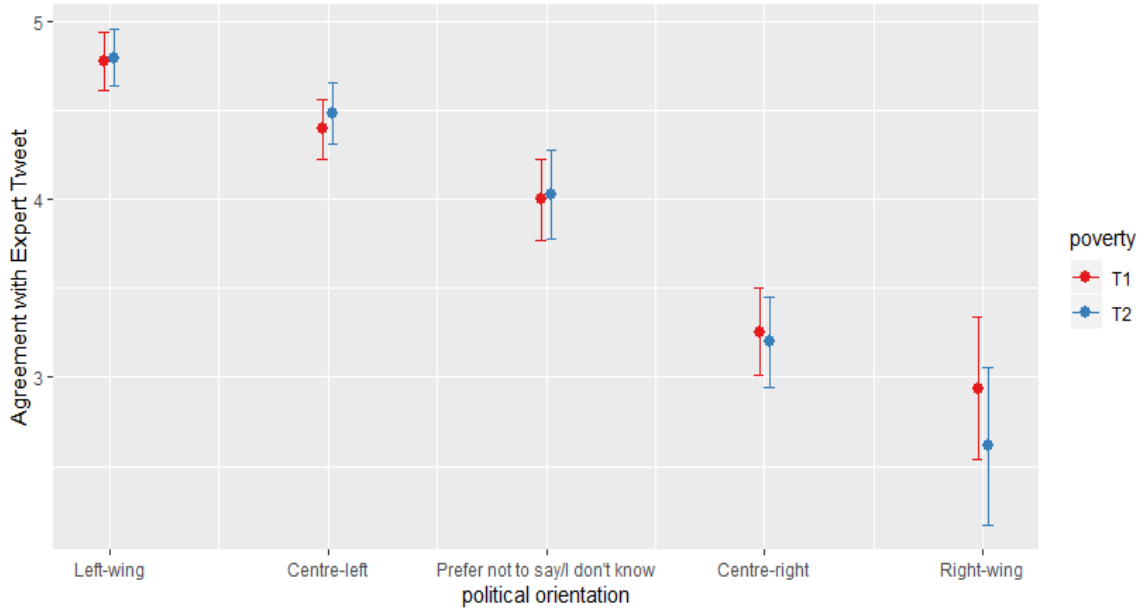

Note: Average levels of agreement with the expert opinion are far higher among the two left-wing subgroups, declining notably in the right-wing categories irrespective of treatment group.

In the only subgroup effects of note found for conflict orientation in the study, there is some evidence that respondents most uncomfortable in hostile situations had lower average agreement with expert opinion after uncivil debate in the poverty condition (β = -0.13) whilst civil debate lowered agreement for the 'conflict agreeable' (β = -0.20, see Figure 24 below and Table 29). Differences are close to significance, and suggest that incivility might be particularly damaging for expert authority online for users with highest aversion to the aggressive conversation regularly found on Twitter, while those that enjoy its spectacle are dissuaded by its absence (civility). However, no such effects were found in the Brexit condition, weakening the weight of this finding.

Figure 24: Expert Tweet Opinion Change by Conflict Disposition (Poverty)

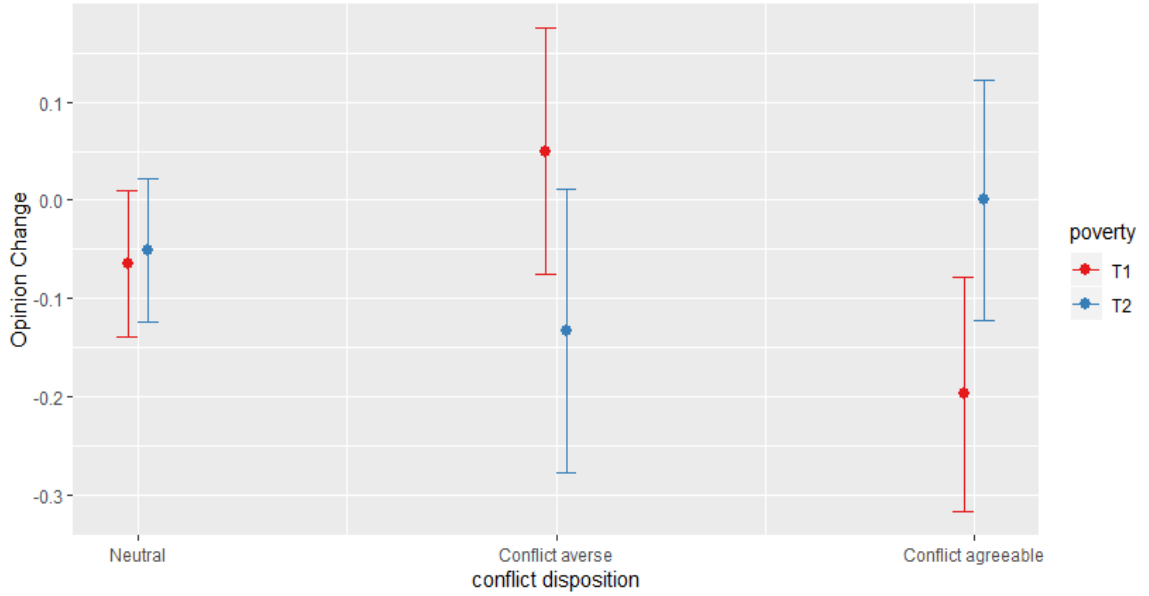

Note: 'Conflict Averse' estimates were lowered by incivility (T2, β = -0.133) and increased by civility (T1, β = 0.050). Conversely, estimates were lowered by civility for the 'Conflict Agreeable' subgroup (β = -0.197). Full estimates in Table 29.

Dependent Variable 4: Trust in Politicians

This variable was operationalised by asking participants to rate their impression of the statement 'I have trust in politicians to make good decisions' using a 5-point Likert scale of agreement.

Potential hypotheses

Hypothesis 5 stated that average trust in politicians would be lowered by exposure to incivility and raised by civility, with control between them. The two experiments display contrasting results. The hypothesis is vindicated for the Brexit condition, though differences between the treatments are strongly driven by the notable positive impacts of civility, with the uncivil estimates largely indistinguishable from control. Indeed, as shown in Table 30 overleaf and Figure 25, civility in the Brexit debate induced a statistically significant increase of 0.14 (SE = 0.08), although it should be noted that trust in politicians still remained low, as it was across all groups. This finding is supportive of the proposed mechanism that civility exerts a positive

impact upon perceptions of political discussion, and that this can increase trust in politicians as the direct symbolisation of that discourse. However, that no distinguishable negative effect is found for incivility (β = -0.02, SE = 0.08), against expectations, is perhaps evidence that its presence in online political discussion about polarised and salient topics is consistent with existing negative perceptions of politics and politicians, such that it exerts no independent influence. Further indication of the mediating effect of conversational context is seen in the poverty-related condition, which showed no positive impact of civility, with both treatment groups registering a very similar small, insignificant reduction in trust (Table 30). The continued lack of notable findings related to incivility support the above-posited explanation, while the nil effect of civility perhaps reflects the expert tweet to which the debate was responding. The poverty-related stimulus attributed a rise in poverty-related deaths directly to government ineptitude and callousness, going as far as to state that ministers will have "blood on their hands". It is thus very plausible that any positive impacts of civility were eroded by the context of the debate.

Table 30: Full Sample Trust in Politicians (Brexit and Poverty)

Trust in Politicians						
	Bivariate OLS (Brexit)	*Bivariate GLM (Brexit)*	*Multivariate OLS (Brexit)*	*Bivariate OLS (Poverty)*	*Bivariate GLM (Poverty)*	*Multivariate OLS (Poverty)*
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	1.902***	1.902***	1.892***	2.019***	2.019***	2.002***
	(0.053)	(0.053)	(0.087)	(0.053)	(0.053)	(0.083)
Brexit T1: Civility	0.143*	0.143*	0.148*			
	(0.077)	(0.077)	(0.076)			
Brexit T2: Incivility	-0.019	-0.019	-0.035			
	(0.076)	(0.076)	(0.076)			
Poverty T1: Civility				-0.053	-0.053	-0.052
				(0.075)	(0.075)	(0.076)
Poverty T2: Incivility				-0.065	-0.065	-0.070
				(0.076)	(0.076)	(0.077)
Observations	894	894	894	896	896	896
R^2	0.006		0.037	0.001		0.005
Adjusted R^2	0.004		0.025	-0.001		-0.007
Multivariate Controls	Age, Education, Gender, Conflict Disposition					
Note:	*p<0.1; **p<0.05; ***p<0.01					

Note: In the Brexit condition, the average in the civil debate treatment (T1) is higher than Control and statistically significant while incivility (T2) shows a null effect. Both treatments exert a minor and insignificant negative impact in the poverty-related condition.

Hypothesis 5a predicted that respondents with a higher level of political interest will have a stronger pre-disposition towards trusting the political process, and thus the motivations of politicians by association. However, because of the negative impression of ideologically opposing politicians induced by partisanship, the aggregate mediating effect was expected to be small. As shown by Figures 26 and 27, patterns across the interest spectrum in both experiments display the predicted relationship, though of minor magnitude and strong statistical insignificance.

Other findings

Further subgroup analysis demonstrates that political orientation was a consistent mediator of trust in politicians in both experiments, with right-wingers tending towards higher trust. These effects are conditioned to varying degrees by treatment, with civility modestly increasing trust across the spectrum in the Brexit condition (Figure 28), while neither treatment had a stable impact in the poverty-related debate. Perhaps relatedly, Leave voters in the control condition exhibited higher values in both Brexit and poverty experiments (β = 0.38, SE = 0.15 and β = 0.17, SE = 0.14 respectively), conditioned insignificantly by treatment (Figures 29 and 30). Gender also had some interesting conditioning impacts, with Brexit-related civility exerting a sizable and significant positive effect upon women (β = 0.25, SE = 0.10, see Figure 31 below and Table 31). There were also notable (but weakly significant) differences between civility and incivility observed for the lowest educated (β = 0.49, Figure 32) and the young (β = 0.39, Figure 33).

Figure 31: Trust in Politicians by Gender (Brexit)

Note: T1 (civility) Female: β = 0.254, SE = 0.095. There is also a significant difference between men and women in control (β = 0.346, SE = 0.112). Full estimates in Table 31.

Dependent Variable 5: Voting in the Next General Election

This variable was operationalised by asking participants to answer 'Yes' or 'No' to the question 'are you likely to vote in the next General Election?'.

Potential Hypotheses

Hypothesis 6 posited that incivility would induce a small reduction in voting intention, while civility may have the opposite effect. The mechanisms proposed here reflect the potential for incivility to precipitate disengagement, arising from the perception that politics is hostile and better avoided. Conversely, civility may encourage democratic participation, because of the greater accessibility of debate and more informative arguments. As fully discussed later, no evidence of any full-sample treatment effects was found across either experiment. Owing to previous work which demonstrating difficulties of experimentally inducing large-scale changes in voting intention (McDonald, 2003), it was expected that any effects found (though important) would be of minor magnitude. This notwithstanding, within each experiment the maximum difference between groups was just 0.1%, and across the whole study the estimates only ranged from 94.0% to 95.9% (Table 32). These percentages also appear to be unrealistically high compared to the reality of turnout in UK elections, even allowing for the politically engaged sample. In all, no full sample effects of note were found.

However, this did not preclude some very interesting findings at the subgroup level. Hypothesis 6a held that likelihood to vote would be heavily conditioned by political interest, with lower interest inducing both a reduced tendency to vote and a stronger mediating impact of treatment. Political interest is indeed a substantial conditioner of declared voting intention, with differences of 71.4% and 60.0% between the lowest and highest interested in the poverty-related and Brexit stimuli respectively. The most interesting results, however, pertain to the contrasting effect of treatment across the two experiments on those with low political interest. Highly attentive participants displayed little variation, very likely to vote regardless of treatment. However, congruent with the same theoretical mechanism outlined for Hypothesis 6 above, civility exerts a strong positive impact on likelihood to vote for low interest participants in the poverty-related condition, vis a vis uncivil treatment (Figure 34 overleaf and Table 33). This result is highly significant for the very least interested, and of substantial magnitude. Indeed, as the graph shows, civility induces a striking increase of 49.2% (SE = 0.11) on voting likelihood comparative to incivility, and 27.8% compared with incivility. This is an important finding which demonstrates that a more civilised discourse could widen democratic participation. This said, the results from the Brexit condition show that these effects are mediated by context in consequential ways. Here, it is the uncivil treatment that appears to have the more rousing impact on in both of the low interest categories comparative to civility (10.0% and 18.8% respectively, see Figure 35 overleaf and Table 34). It was noted in the previous chapter that a potential mechanism behind the conditioning effects of political interest on treatment could be that incivility rouses engagement among the most interested due to strength of feeling and a desire to defend your position from hostile attack. The results from the Brexit condition suggest that this theory may need to be widened to incorporate the mobilising effects of high issue salience on those of low general interest. There can be little doubting that Brexit has raised passions far beyond those who usually follow politics, becoming for many a critical matter of identity. The EU referendum turnout was 72.2%, considerably more than other UK elections, and the continued ferocity of debate since has increased public issue salience yet further (Evans and Schaffner, 2019). The two sets of results thus suggest that tonality of online debate matters substantially for the engagement of those least interested in politics, with the directionality of effects heavily conditioned by issue context. These are undeniably important conclusions, however, as with all conclusions relating

to low political interest in this study, appropriate caution is necessary due to the low number of participants in this category.

Figure 34: Likelihood to Vote in Next General Election by Political Interest (Poverty)

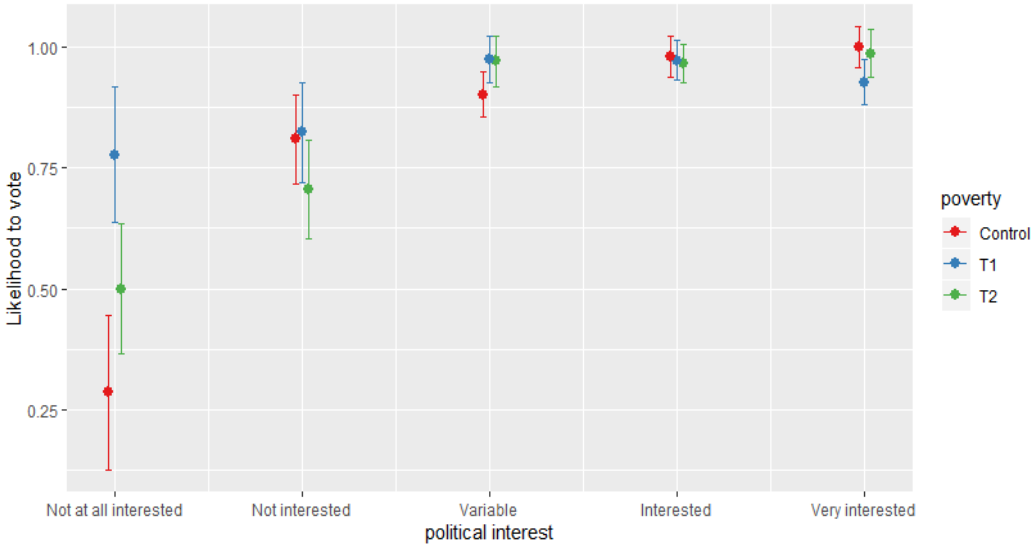

Note: Estimates decline in the two lowest interest categories, but there are major effects in the Not at all interested subgroup. Civil treatment (T1) β = 0.492, 0.278 larger than T2. Full estimates are in Table 33.

Figure 35: Likelihood to Vote in Next General Election by Political Interest (Brexit)

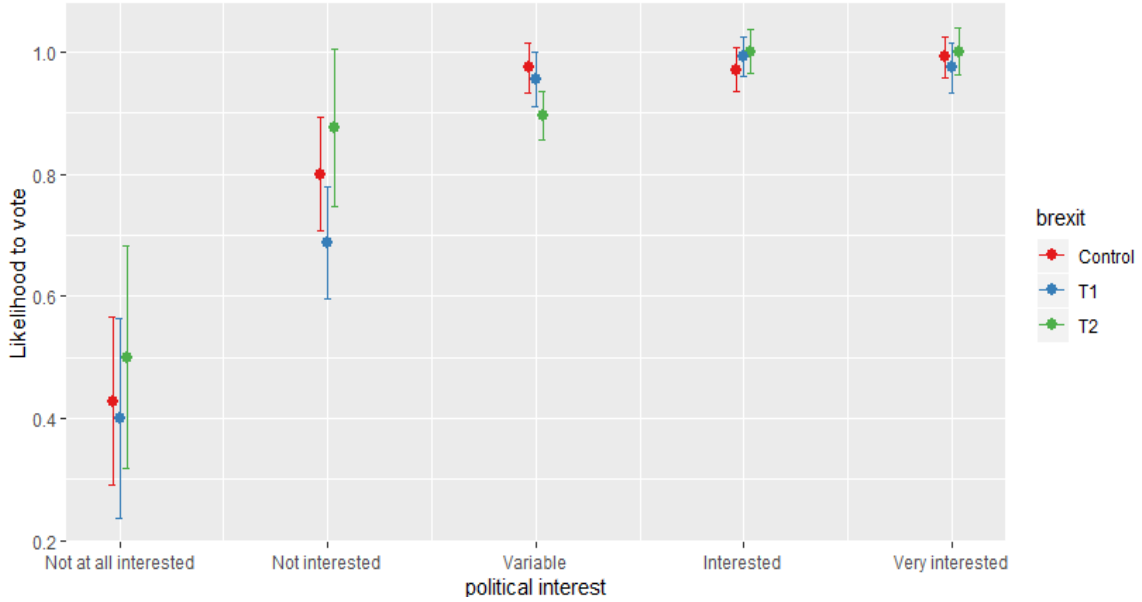

Note: Estimates sharply decline in the two lowest interest categories. Incivility (T2) increased likelihood relative to control, and civility (T1) decreased. Full estimates are in Table 34.

Relatedly, Hypothesis 6b reflected the habitual nature of voting behaviour and predicted that treatment effects would differ between voters and non-voters in previous elections. As expected, past voting heavily influences declared likelihood to vote in future in both experiments (Tables 35 and 36). In strong support of the conclusions for Hypothesis 6a above, in the poverty experiment civility once again exerted a statistically significant increase in likelihood to vote compared with uncivil treatment, this time in relation to non-voters. This demographic was 26.0% more likely to vote than non-voters in the control group (SE = 0.05), and 16.0% higher than those exposed to incivility (see Figure 36 below). Conversely, civility demobilised non-voters in the Brexit condition by 11.5% comparative to voting likelihood in the uncivil treatment and 17.6% from control (SE = 0.05, see Figure 37 overleaf). In addition to the possible engaging effects of issue salience, it could also be true that polarising topics such as Brexit generate expectations of animosity, such that its absence (civility) actually disengages those who do not usually participate. Given the probable correlation between non-voters and low political interest, these results are complementary to the results and mechanisms explained for Hypothesis 6a above.

Figure 36: Likelihood to Vote in Next General Election by Past Voting Turnout (Poverty)

Note: Past voting behaviour substantially conditions anticipated voting, with the estimate for non-voters in Control 0.496 lower than its equivalent for past voters. Treatment effects are positive and substantial for non-voters in civil treatment (T1). Full estimates in Table 35.

Figure 37: Likelihood to Vote in Next General Election by Past Voting Turnout (Brexit)

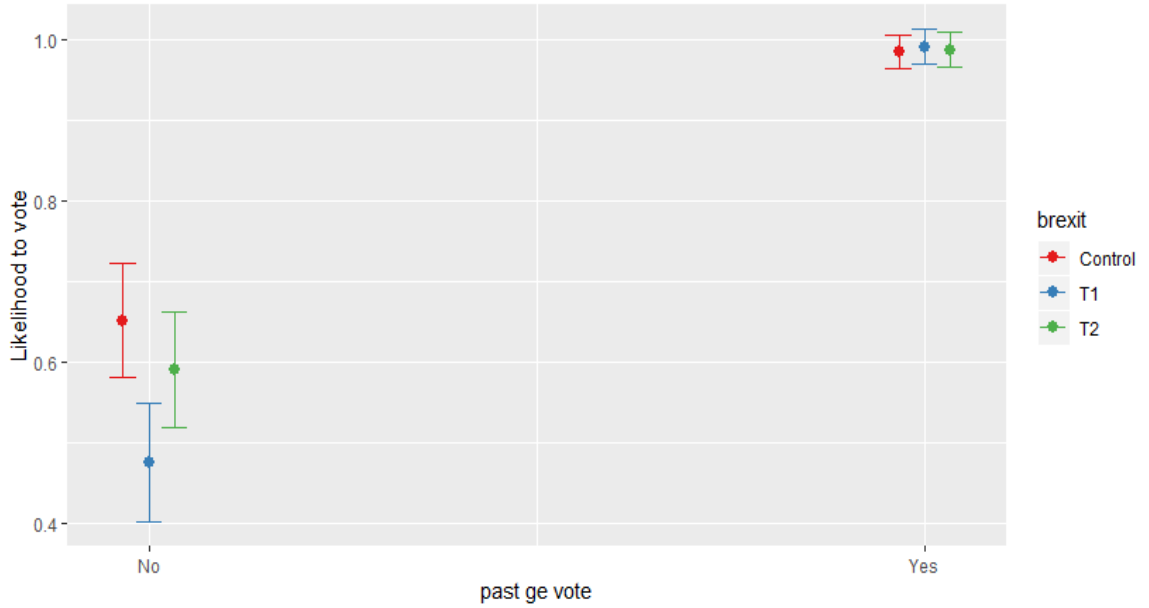

Note: As above, non-voters in the past were substantially less likely to vote, by a magnitude of 0.33 between Control groups. Treatment effects for non-voters are again notable but reversed, with civility reducing likelihood of voting. Full estimates in Table 36.

Other findings

There are additional findings that were not hypothesised but provide further credence to these conclusions. Much like those of low political interest and past non-voters, among respondents unsure of their partisan identity or political orientation in the poverty-related condition, civility exerts statistically significant increases of 23.0% (SE = 0.08) and 17.3% (SE = 0.06) on likelihood of voting respectively, relative to control. As shown in Figures 38 overleaf and Figure 39, civility is also a more powerful mobiliser than incivility for these demographics, by 17.0% for those without a partisan identity (significant) and 9.4% on unknown political orientation (not significant). Combining the insights thus far provides evidence that across four different measures of political expression (interest, voting, partisan identification and ideological orientation) in the poverty-related experiment, civility encourages democratic engagement for the uncertain in a sizable and statistically significant manner. On the last two measures, there is little evidence of any mobilising effects of incivility in the Brexit discussion, but it still holds that civility seems to have little effect when the issue context is salient and polarised (Figure 40).

Figure 38: Likelihood to Vote in Next General Election by Party ID (Poverty)

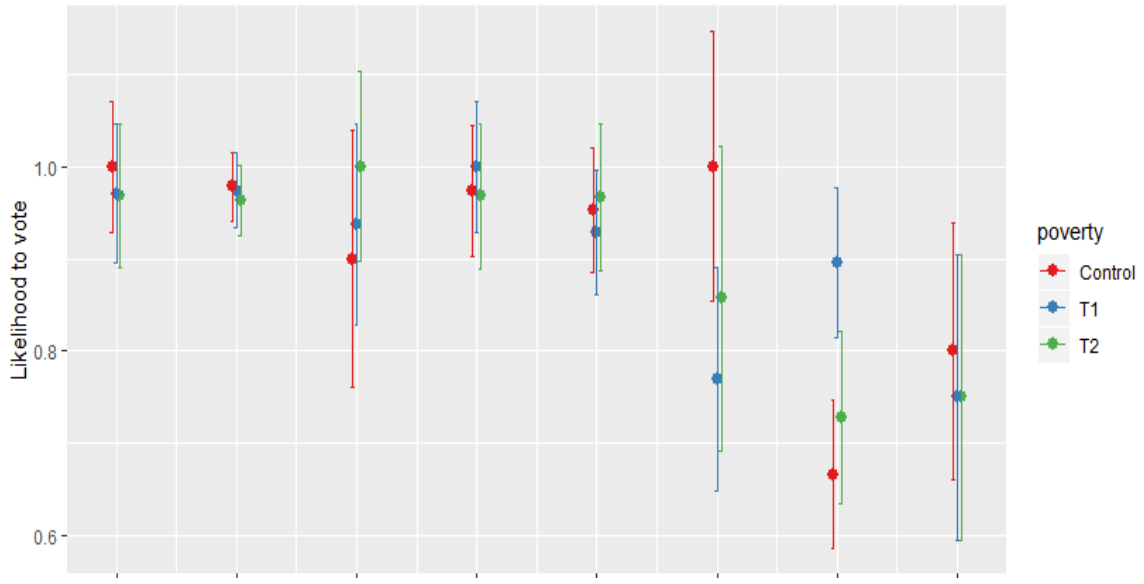

X-axis = Party ID (L to R): *Green Party, Labour Party, Scottish National Party, Liberal Democrats, Conservative Party, Brexit Party,* ***Do not know****, Other*.

Note: The second category from the right shows the clear and significant positive treatment effects of civility (T1) on those without a partisan identity vis-à-vis both Control and incivility (T2) in the poverty experiment.

General sample-wide hypotheses

This section does not relate to specific variables but rather evaluation of the hypothesised sample-wide trends in the previous chapter.

The first two predictions were mutually exclusive and therefore require joint evaluation. Hypothesis 7 held that treatment effects would be consistently larger in the Brexit experiment owing to the prolonged emotive nature of public discourse around Brexit. In this formulation high issue salience was expected to accentuate the impacts of tonality, reflecting the increased average investment that participants would have in the topic. On balance, this is not supported, and with the exception of trust in politicians full sample effects are matched or larger in the poverty experiment. This lends support to the opposing Hypothesis (7a), which projected that high salience would reduce effects compared with a less salient topic. The posited mechanism was a combination of public issue fatigue, opinion entrenchment and normalisation of Brexit incivility, such that polarised animosity comes to be expected and civility cannot override this association. The differences between the experiments are not definitively conclusive, but the directionality of results vindicates Hypothesis 7a to a greater degree.

Hypothesis 8 pertained to individual-level determinants of treatment effect heterogeneity, and expected that participants classified as ‘conflict averse’ would be particularly reactive to treatment. Conversely, ‘conflict agreeable’ respondents would demonstrate more minor responses. Surprisingly, no such patterns were consistently evident and conflict orientation

yielded few results of interest. As discussed below, the failure of this hypothesis to materialise highlights the difficulties of operationalising individual-level factors accurately (Bimber 2001:54), and the lack of effects is not any reflection on the true likely impacts of differences in conflict orientation.

Finally, Hypothesis 9 proposed that treatment effects would be largest in the groups exposed to uncivil debate. Incivility was expected to increase opinion polarisation relative to civil treatment or control because of its previously documented potential to reduce space for compromise between opposing partisans, inhibit rational evaluation of opinion challenging information and increase extremity of opinions. However, as shown above, results varied considerably and this hypothesis is generally not supportable. It was notable that the effects of civil treatment were often larger in important areas than those observed for incivility, such that it could be argued that abusive political debate online has been normalised to the degree that civility now affects the attitudes of observers more.

Caveats and limitations

Together, the results make a useful contribution to academic understanding of the relationship between the level of civility in online discourse and political engagement, trust, and information reception. However, there are three important limitations when assessing the scope of these findings. The first is the negligible size and statistical significance of some results. These remain useful indicators of the phenomena they measure, but certain results need to be studied further due to inconclusive magnitude or large confidence intervals. One of the few shortcomings of a sample based on Twitter users is the unbalanced size of subgroups. Taking education as an example, this study contained 1063 participants with a university degree, compared with 174 whose highest qualification was GCSE level. Discrepancies like these reduce the statistical power of certain potentially interesting results, which need to be verified with a larger sample. Previous work also shows that uncovering the full magnitude of effects on significant and consequential dependent variables such as political trust and voting behaviour is not straightforward (Robison et al 2018:12). This speaks to the habitual nature of political attitudes and behaviour, and the gradual nature of changes that do occur (Spears and Smith 2001:314). As such, temporary and instantaneous experimental treatments of the kind used here are unlikely to capture the full cumulative impact that political socialisation in the Twitter information environment has on users over time. However, designs like this still provide useful conservative estimates, and show how individual characteristics mediate impacts (Davenport et al 2010:18). In order to measure the true magnitude and consistency of effects, however, treatment would have to be repetitious and gradual so as to mirror the lived reality of political behavioural change. Regrettably, such a design lies beyond the financial and temporal capacity of this thesis.

Secondly, it is clear that the full sample averages of voting intention (ranging between 94% and 96%) are likely to be exaggerated estimates of turnout realities. This is perhaps a product of unforeseen and unavoidable circumstances when the study was run. Parliament voted to hold an early General Election on 29 October 2019, held on 12 December. The details and schedule of this research design had already been finalised when the election was announced, meaning the experiments were live (26 November - 1 December) in the immediate campaign period. It is thus likely that a proportion of respondents were actively primed by this political context, over and above treatment effects or whether they actually voted. This issue highlights the inadequacy of voting intention as a sole measure of political engagement, particularly given the problems of desirability bias (McDonald 2003:185). Political enthusiasm is multi-faceted and can be demonstrated in a multitude of ways beyond

the formality of voting (see Teruelle, 2011). If this research were repeated, it would thus contain more varied measures of engagement aimed at depicting its nuanced reality.

A final caveat pertains to the shortcomings of the conflict orientation matrix. Both intuition and a substantial behavioural literature support the expectation that personality traits are likely to condition reactions to civil or uncivil content (Hossain et al, 2018; Anderson et al 2013; Muddiman and Stroud, 2017). However, it would appear that the operationalisation of this variable was ill-suited to capturing its intended phenomena. Previous work shows that different operationalisations of individual-level traits can lead to contrasting results across contexts (Khang et al 2014:51), but given the undoubted importance of personality mediators, reaching a collective standard for their conceptualisation should be a pressing aim for the research field.

Central conclusions

Notwithstanding these caveats, this study contained important findings about the effects of tonality on Twitter for the political attitudes and behaviours of observers. In brief:

- The presence of civility in public debate online exerted a statistically significant and sizable increase in likelihood to vote among both the less politically interested and past non-voters in the poverty-related treatment (49.2% and 26.0% respectively). It had a similarly noteworthy mobilising effect on the voting intention of those unsure about their political orientation or partisan identity. In the clearest indication of the mediating impact of issue salience, little of this holds for the Brexit condition, and there is tentative evidence that incivility actually mobilises those least interested in politics and non-voters, while the lack of 'spectacle' inherent to civilised discussion can have the opposite impact. However, the sample-wide treatment effects on declared intention to vote were negligible and estimates were very high across all groups in both experiments.

- Similarly, civil treatment also induced a clear and statistically significant positive impact on trust in the ability of other citizens to make good judgements. This effect was found in both experiments, largest for the poverty-related debate. Conversely, it was also true that incivility exerted a negative effect on trust in fellow Britons in both the Brexit and poverty condition. These impacts were the most substantial of the three aspects of trust tested, robust across age, gender, and education, but mediated heavily by EU Referendum identities and political orientation.

- Levels of trust in politicians to make good judgements were low across all treatment and control groups, regardless of topic or tone of debate. However, there was further evidence from the Brexit condition that civility can exert a sizable and significant positive impact on these impressions, while the presence of incivility had no distinguishable effect. This is perhaps a reflection of its congruence with existing unfavourable perceptions of political discourse. No such effects were present in the poverty-related debate and minor negative impacts were found across both treatment groups, an indication of the mediating role that particular issue-frames can play, given that politicians were directly criticised in the stimuli. Results differed according to education, age, gender, political orientation, and EU allegiance.

- Trust in the authority of expertise was less significantly affected by treatment, but was lowest among those exposed to conflictual debate in both experiments. There is evidence from both experiments that treatment effects of incivility on perceptions of expertise are moderated significantly in varying directions by political interest, gender, age and education, most often strongest in the poverty-related debate. Interestingly, there was generally weak support for motivated reasoning, and no evidence of increased trust in expertise among Leave voters in spite of the expert's anti-EU stance, a demonstration of entrenched scepticism about expertise in Leave discourse.

- Exposure to debate also marginally reduced levels of agreement with expertise across both civil and uncivil treatments in both experiments. This suggests that the mere presence of debate can weaken credibility assessments. Agreement with the expert's opinion was conditioned by education, political interest, and ideology as hypothesised, though only the politically disengaged updated their views noticeably following debate, dependent on tonality and issue context. Younger respondents and the conflict averse appeared more likely to disagree with expertise after Brexit-related incivility.

- Across the sample collectively, effects appear to be somewhat larger in the poverty-related condition, perhaps evidence that incivility has become normalised in the polarising context of Brexit. Little evidence is found pertaining to the mediating impacts of individual-level conflict orientation. Finally, heterogeneous treatment effects are not noticeably increased by exposure to incivility comparative to civility, suggesting once more that uncivil political discussion may be the default expectation of Twitter users, such that civility stands out. If true, there are perhaps alarming democratic implications given the uncivil nature of the status quo, and a pressing case for the cultivation of a more civilised discourse online.

The following chapter will explore the theoretical implications of these findings and situate them within the context of relevant academia.

Part III: Evaluation

Chapter 6: Theoretical Implications

To establish thematic linkages between the results and explore their democratic implications, there is a need to embed them within extant literature. This chapter will first consider the broader connotations of the flagship findings on political engagement, before analysing the prospects for trust and opinion formation on Twitter. The third section will both compare the impacts of civility and incivility respectively, and the mediating role of issue salience on treatment effects. The final segment summarises the current political information environment on Twitter, and discusses the contribution of this study to scholarly understanding of contemporary developments.

Political engagement: inclination and quality

Much previous work has elucidated Twitter's potential to increase political engagement through its power as an information sharing medium, reducing the costs of following current affairs (Kwak et al 2010:591). However, the striking results from subgroup analysis in this study suggest that social media directly influences political engagement, not merely through information availability, but through its tone, context, presentation and veracity.

Civility exerted a sizable and statistically significant increase in likelihood to vote among both the less politically interested and non-voters in the poverty-related treatment. This is not found for the already interested and existing voters due to a 'ceiling effect'; it is not possible to induce an increase among habitual voters (Bode 2016:29). Unpacking these findings requires understanding of common determinants of turnout decisions among those not routinely engaged or informed. Voting decisions for these demographics are thought to be a product of evaluations of how much information is available, how much skill is needed to obtain it, and the extent that the information reduces uncertainty (Bimber 2001:56). Twitter already facilitates a large amount of readily available information without imposing complex barriers to access. However, the existing information environment does little to reduce uncertainty; the hostile nature of political conversation likely increases it (Valentino et al, 2008). These results suggest that civility could perform this function, possibly unlocking Twitter's potential to widen democratic participation. The same mechanisms likely explain the similarly noteworthy mobilising effect on those unsure of their political orientation or partisan identity, also present in the poverty-related condition. These traits are plausible proxies for low political understanding, highlighting a possibly vital link between tonality and democratic representation. In an era of increasing inequality in political knowledge and interest, civility could be a source of amelioration through inculcation of news-seeking habits online (Prior 2007:134).

Low internal and external efficacy (belief in personal capacity to affect change and the system to deliver it) are closely linked with apathy and abstention from the political process, a conclusion extensively supported by previous research (see Ardevol-Abreu et al 2019:83). If civility promotes greater accessibility in political conversation and reduces perceived barriers to understanding, it could plausibly increase engagement for the politically unsure by improving their efficacy, just as incivility reduces it (Pingree 2011:23). Given the prevalence of incidental news exposure on Twitter, these ideas complement findings that inadvertent encounters with political content can reduce the engagement gap between high and low interest citizens (Valeriani and Vaccari 2016:1857). The poverty-related experiment strongly

suggests that such equalisation is again likely to be heavily contingent upon the degree of civility in this news exposure.

It is intriguing that little of this holds for the Brexit condition. In the clearest indication of the mediating impact of issue salience, there is tentative evidence of the opposite effect in that incivility actually mobilises those least interested in politics and non-voters. Although not of the same magnitude or significance of the civility results in the poverty-related treatment, this finding is potentially concerning. Given that a combination of emotional arousal and topical relevance can increase likelihood of retransmitting information online (Capella et al:411), incivility could have a wider reach in salient contexts, undermining evidence-based discussion and prospects for tension diffusion. Such content been readily evident in Twitter discourse surrounding Brexit (Polonski, 2016a). These findings underline that political engagement is not an unconditional democratic good; electors need access to a fair impression of circumstances to make voting choices based upon rational evaluation of competing information (Almond and Verba, 1963). This becomes less likely when arguments are presented with hostility (Massaro and Stryker 2012:413), and as such these results suggest that Twitter could be promoting democratic participation among the politically unsure based upon adversarial emotions over evaluation of complex issues.

Findings from the Brexit treatment also suggest that the lack of 'spectacle' inherent to civilised discussion can actively demotivate potential voters in a salient issue context. As explained below, political discussion on Twitter forms part of a complex ecosystem of content in which politically disinterested users most often seek entertainment (Kim et al 2013:2607). According to recent research, UK Twitter users are 38% more likely than non-users to search the internet for entertainment (Blank 2017:688). Relatedly, negative online adverts have been shown to have strong turnout effects on the habitually disinterested in salient election campaigns, who are roused by incivility (Hopp and Vargo 2017:370). It is also known that abusive language can boost observer attention and approval, particularly in political contexts (Kwon and Cho 2017:84). Concisely, in publicly polarised circumstances such as Brexit, incivility may become the expectation such its absence appears dull and demobilising. This is a particularly alarming conclusion because "though some people have high levels of motivation to follow the latest political news, many only pay attention to politics at critical moments, or hardly at all" (Guess et al 2018:3). There are thus important democratic implications of Twitter use at times of high salience for the sizable demographics who do not routinely follow politics. If incivility mobilises engagement but inhibits reception of new information, then it is a recipe for further polarisation and misinformation owing to Twitter's structural affordances (Yardi and Boyd 2010:325). This demonstrates that Twitter interacts with external contexts in important ways, implying that politicians have a duty to be proactive about reducing polarisation at times of high salience and promote engagement through civility. In so doing, they would be aiding Twitter to become a platform which promotes efficacy among the politically unsure, and increases the likelihood of voting based on informed convictions.

Trust, opinion formation, and news consumption

Citizen trust online is a dynamic and evolving entity based on abstract impressions of the collective, of particular importance in spaces like Twitter which facilitate stranger-to stranger interaction (Zubiaga and Ji, 2014). The results showed that tonality of political debate can impact levels of citizen-to-citizen trust irrespective of issue context, civility positively and incivility negatively. Theory suggests that this matters for democratic outcomes. Higher levels of social trust have been shown to be connected to information retention (Carminati et al 2013:33), which feeds political knowledge. Previous work has in turn illuminated the dangerous effects of deepening inequality in political knowledge, potentially precipitating

widespread disengagement and an inability to discern how interests can best be represented (Prior 2007:98). Interest formulation is central to informed citizenship, and if public discourse on Twitter is most often uncivil, exposure to opposing viewpoints is unlikely to be aiding democratic deliberation or attenuating polarisation, as it has done elsewhere (Kim 2015:915).

Citizen trust is an important aspect of social capital which reflects the strength of co-operative infrastructure in a society (Carminati et al 2013:28). The importance of social capital to political stability is long-established within democratic theory (Putnam, 1995), and these results complement the previous section in suggesting that civility online could be a source of this. The positive impact was present in both experiments, and that civility was able to improve trust despite a polarised issue context is further indication of its potential for conflict amelioration. However, as extensively discussed in previous chapters, debate conducted in such a respectful manner is rare on Twitter, and given that incivility may increase as more people with contrasting political allegiances participate in debate (Hopp and Vargo 2017:370), this supports concerns about the impact of Twitter for social cohesion. Subgroup analysis provides further cause for pessimism, given that trust was higher among Leave voters in the EU referendum, right-wingers, and the politically interested. To the extent that left and right-wingers (Conover et al 2011:59), Leavers and Remainers (Polonski 2016b), and political enthusiasts and the uninterested (Leetaru et al, 2013) exist in separate Twitter communities from each other, these differences in public trust can only be expected to widen due to opinion-reinforcing selective exposure (Kelly Garrett: 2009:265; Colleoni et al 2014:280). Moreover, when communicative conflicts do arise between opposing partisans, the relative scarcity of social tools for conflict resolution emanating from echo chambers and position extremity may foster yet further incivility, cyclically eroding public trust (Balliet and Van Lange 2012:8).

Trust in expertise was not as affected by differences in tonality. The lack of significance in these results is surprising, given extensive theory pertaining to the effects of varying civility on trust in experts. In brief, previous studies have shown that political incivility can inhibit rational evaluation of information credibility, rousing a combative form of partisanship that weakens traditional markers of authority (Gervais 2018:1). This increases reliance upon pre-existing viewpoints, which become resistant to challenge and bias information processing (Gerber and Huber 2010:154). Incivility also commonly induces anxiety, fear and anger, emotions which often drive cognitive disengagement and reduce personal efficacy among moderates (see Pingree 2011:26). It would be expected that these impacts might adversely affect trust in expertise, owing to cognitive uncertainty among moderates, increased aggression for partisans, and heavier reliance upon existing views for both groups (Downs 1957:213). Indeed, it has been shown in reference to public assessments of journalism that “incivility in comments had an unconditionally negative effect on the perceived formal quality of an article” (Prochazka et al 2018:62). It is thus somewhat surprising that sample-wide effects of incivility in this study were minor, and when itemised by political views, little evidence is found of motivated reasoning impacting trust. There was evidence that agreement with the expert was strongly conditioned by political priors in all experiments, but no sign that this impacted trust. On the surface, these results perhaps suggest that trust in expertise on Twitter might be more robust than existing literature assumes.

However, averages were comparably low across all groups, perhaps unexpectedly given the prevalence of highly interested liberals in the sample, a group that often advocate for evidence-based policy-making (Pawson 2006:169). This demonstrates the extent of difficulty for experts in an online space overrun by competing information sources, leading to the politicisation of factual reporting (Wood, 2019). Abusive partisan denigration of expert

commentary becomes commonplace in this environment, reinforcing uncertainty about its credibility for observers (Searles et al 2018:4). This speaks to the collapse of traditional information hierarchies and the anticipated effects of network selectivity to encourage motivated reasoning (Sunstein 2018:130). News is just one aspect of a vast content ecosystem on Twitter, and the ease with which user-generated content can be created and disseminated challenges traditional norms of information consumption. Notwithstanding complexities, pre-social media news consumption relied upon print and broadcast media outlets, the best of which "spend significant time and energy deciding what events are newsworthy, turning those events into stories, and assembling a coherent and balanced editorial product. But hierarchy of the news report is considerably lost within the maelstrom of information available on a social media platform" (Boczkowki et al 2018:3533). In combination with fragmented news exposure, these patterns have promoted the growth of informational relativity which directly inhibits the role of experts as an adjudicative heuristic for citizens precisely when it is most needed (Hermida 2012:661).

Relatedly, this study shows that debate about an expert's opinion can weaken credibility assessments, regardless of tonality. Surprisingly, both civil and uncivil debate precipitated a significant decline in trust of expertise for the most politically interested in the poverty-related condition. Furthermore, after debate there was a slight reduction in agreement with the expert across both civil and uncivil treatments. This is supported by similar work, which also "observed a trend showing that the mere presence of comments deteriorates the perceived quality of an article" (Prochazka et al 2018:62). This phenomenon is consistent with the above-referenced loosened informational hierarchies online, while other research alarmingly suggests that Twitter users make more favourable general assessments of trust when lacking information about its source, which can facilitate an inability to distinguish between rumour and fact (Shariff et al:794).

The results indicate that issue context may also matter for reception of expertise, with Brexit-related incivility inducing increased agreement with the expert among the disinterested. This is not found in the poverty-experiment, and could be an encouraging indication that incivility precipitates reliance on authoritative sources as a heuristic for those without the resources or inclination to arbitrate between hostile opposing views on issues of perceived importance. Alternatively, given that the expert's tweet was in favour of leaving the EU and the politically disinterested in the sample were also more likely to be pro-Brexit, it could reflect tendencies for incivility to bias evaluation of information according to ideological priors (Young 2009: 254). Adjudication between these competing mechanisms requires future research, but given the substantial extent of misinformation on social media (Thorson et al 2010:290), research showing that misperceptions are difficult to correct weakens the influence of expertise on Twitter yet further (Nyhan and Riefler, 2010).

In short, while encouraging that trust in expertise was not drastically impacted by incivility or context, the wider picture remains precarious and these results should be viewed with caution. As delineated in the introduction, political discourse on Twitter challenges established norms around information reliability and political socialisation in consequential ways for the evaluation of competing ideas in democracies (Hermida 2012:666). Further work exploring comparative trust of expertise in online and offline populations would be a fruitful endeavour.

A noteworthy gendered effect did emerge from the poverty experiment. Trust in expertise was substantially lower for women than men across all treatment groups, and there was a statistically significant negative treatment effect of incivility for women. This finding supports previous research showing female observers tend to be more critical in political judgements and more cautious believing news stories (Shariff et al 2017:794). If this also applies to

evaluations of expertise however, it serves as further evidence that differentiation between expert and public opinion on Twitter is indeed weakening. Furthermore, younger respondents displayed an indicative move away from agreement with expertise after Brexit-related incivility. Given that younger people are often incidentally exposed to news sources of varying reliability online and especially on Twitter (Fletcher and Nielsen 2018:2462), the above concerns about expert credibility are likely to be especially critical for future generations. Many in this demographic do not actively seek out news, which can foster attitudes that weaken retention of political knowledge (Weeks et al 2017:105). There is evidence that this may translate into a lower propensity for trust in politicians and the media than older generations (Sveningsson 2015:3). This must not be confused with youth disengagement (Teruelle, 2011), but understanding the types of information younger people respond to online and its consequences on their judgements is going to be a crucial determinant of future trends.

Although the position of experts is threatened by the online information environment, levels of trust in politicians to make good judgements were even lower across all groups, irrespective of topic or treatment. Estimates were markedly lower than equivalents for experts or the public. However, this is not a direct consequence of social media and is congruent with previous findings that politicians are the least trusted UK profession (Clemence, 2019). This trend has been worsened by the parliamentary deadlock over Brexit, again showing that social media is interactive with broader contexts. Perhaps reflecting current perceptions of political discourse in Britain, the presence of incivility had no distinguishable effects on trust in politicians. This is suggestive of an entrenched existing association between politics and hostile discussion, such that it exerts no independent effect.

More promisingly, however, alongside its impacts on voting behaviour, the Brexit condition showed that civility can significantly improve these impressions. This supports work positing that argumentatively persuasive information has a greater probability of being influential (Capella et al 2015:411), because civility better allows for the construction of evidence-based arguments. It follows that if the proportion of civil discourse online could be increased, it would possibly have a prominent impact on political trust, matching the earlier findings on citizen-to-citizen trust. Results differed notably according to political orientation and EU allegiance; highest rates of trust were observed for right-wingers and Leave voters, mediated by treatment. Given that The Conservatives have governed for the last decade and Leave won the EU referendum, perhaps being in the political ascendancy increases trust in politicians, especially when arguments are presented in civil terms. This is consistent with well-documented effects of partisanship on performance assessments (Gerber and Huber, 2010).

That civility can lessen mistrust of politicians in the context of Brexit provides support for a slight reformulation of existing work which finds that anger can induce feelings of apathy and aversion (Gervais 2015:167), reducing satisfaction with political discourse (Capella and Jamieson 1997:166). At times of salience when politics is expected to be polarised and uncivil, perhaps apathy and aversion are the norm, such that civility raises trust rather than incivility decreasing it. There is evidence that this is particularly likely on social media, where users now anticipate incivility (Antoci et al, 2019). Reconciling these findings with the previous section suggests that incivility on salient topics may increase engagement while reducing trust. This is a damaging combination for democratic stability, especially given that “those with low levels of trust tend to prefer non-mainstream news sources” which are harder to verify as reliable, and may thus result in voting based on false information (Fletcher and Park 2017:1281). Collectively then, the potential of civility looks somewhat unlikely to be realised, given that network selectivity increases engagement inequality between partisans and moderates, reduce exposure to alternative viewpoints and increases polarisation (Sunstein

2018:99). The net conclusion is that civilised public debate on controversial topics is not promoted by Twitter's structural affordances, irrespective of its hypothetical dividends.

There is also some evidence that the positive potential of civility to improve trust is dependent on the content being debated. Minor negative effects on trust in politicians were found across treatment in the poverty-related debate, indicative of the mediating role that particular issue-frames can play, given that politicians were directly criticised in the stimuli. Civility also displayed a tentative propensity to inculcate the young, women, and those of lowest education into increased relative trust in politicians. These findings were of reasonable magnitude but weakly significant, so caution is required in extrapolation. Nevertheless, they are indicative of civility's propensity to change the attitudes of some of the most disenchanted and underrepresented groups in current discourse. Indeed, proponents of social media hold that it has great potential to widen the scope and quality of public political deliberation (Papacharissi 2004:260; Wojcieszak and Mutz 2009:40), but these results suggest that such hopes are contingent upon the civility of discussion.

Civility, incivility, and the mediating role of issue salience

Against expectations, treatment effects of civility are on balance more prominent than incivility, suggesting again that uncivil political discussion is expected by Twitter users and civility stands out, supported by numerous results. This is complementary to recent research showing that civility on social media can increase trust on a string of measures, while incivility does not elicit significant changes. The authors conclude that "uncivil debate seems to be considered as normal" (Antoci et al 2019:87). Given that activity on social media can direct impact offline participation (Gil de Zuniga et al 2014:612), this thesis argues that efforts to cultivate civility could substantially improve Twitter as a medium for political communication, while the status quo is having actively destabilising impacts.

Collectively, treatment effects are more striking in the poverty-related condition, perhaps evidence of Brexit-related issue fatigue, opinion entrenchment, polarisation and anticipated incivility (Sunstein 2018; Gervais, 2018). These phenomena have been well-documented (Duffy et al 2019:8) in British political discourse since the EU referendum, and are consistent with previous research showing that issue polarisation inhibits the malleability of perceptions to immediate context (Zaller 1992:148). However, as documented by the previous chapter, this did not preclude some interesting results at subgroup level in the Brexit condition.

Furthermore, previous research finds that higher internal and external efficacy are consistent indicators of stronger trust in and engagement with politics. Education, knowledge of political orientation, strength of partisanship and political interest are strong predictors of higher internal efficacy (Kenski and Stroud 2006:175; Morrell, 2003). It is thus reasonable to theorise that participants matching these attributes will have opinions less sensitive to tonality, owing to being better equipped to justify their opinions (Ardevol-Abreu et al 2019:85). A string of results here support this notion, but suggest that differences in the tone of political content can have substantial and contrasting effects on those lowest in efficacy. Civility has the potential to increase trust and participation across a range of measures for political uncertainty. Conversely, incivility maintains or worsens apathetic mistrust for these demographics, and can promote antagonised voting behaviour in polarised contexts (Brexit) unlikely to be based upon reliable information, while making respectful disagreement appear dull and demotivating.

These conclusions carry perhaps alarming democratic implications and make a pressing case for the cultivation of more civilised discourse online . Doing so could combat the ills of network selectivity and reduce the inequalities in political interest driving the current hostile polarisation of political debate on Twitter (Bode 2016:25). At a time when trust is low and politics divided, it

is evident that incivility is being reinforced by the particular characteristics of the political information environment on Twitter, which create “a negative spiral where those who are less engaged participate less” (Kalogeropoulos et al 2017:1).

However, these results clearly show that if civility could emerge it would likely help to inculcate democratic norms that could promote social cohesion (Boyd 2006:863). Indeed, previous work has shown the healthy impact that the presentation of opposing opinions can have on skills of deliberation and critical thinking (Lee et al 2014:707). If this could be conducted with civility, Twitter might begin to realise its democratic potential. As things stand however, the likelihood of this is drastically reduced by the realities of the political atmosphere on the platform.

The political information environment on Twitter: current realities

If political socialisation is occurring in a Twitter information environment that is increasingly consequential owing to its popularity, there is cause for concern. As many authors referenced here have detailed, aspects of discourse on Twitter are marked by declining uptake of political information among general users, declining quality of news, declining diversity of news sources and viewpoint exposure, increasing fragmentation and polarisation, increasing informational relativism, and increasing inequality of political knowledge. The cumulative impact of these trends is affecting offline dynamics and contributing to wider instability in political discourse across established democracies, characterised by partisan simplification of complex phenomena, sharply increasing voter volatility, and the rise of populist politics in various guises (for a detailed discussion see Van Aelst et al, 2017).

This study joins a growing collection of research suggesting that uncivil political discussion is becoming the commonplace expectation of online observers. Results here add some nuance and suggest this is especially pertinent in polarised issue contexts such as Brexit. As shown, current norms of debate are doing little to foster inclusive democratic participation based on the dissemination of reliable information and the deliberative exchange of ideas between citizens (Sunstein 2018:72). Trust and engagement among the politically unsure in the UK are under threat, leading to dangerous inequalities in the articulation of collective interests and the over-representation of extreme viewpoints. Nowhere is this more evident than in the Brexit debate, where compromise has become anathema to the hyper-partisans dominating the discourse on both sides, silencing moderate voices and fuelling polarisation (Murray et al, 2017).

Encouragingly, the central finding of this thesis is that civilised debate could positively influence the dispositions of non-voters, those unsure of their political orientation or partisan identity, and the politically disinterested. These demographics are understandably repelled by current trends, and their political attitudes seem particularly sensitive to differences in civility. In short, it is argued that democratic instability is being exacerbated by engagement inequalities and the continued erosion of trust in politics (Iakhnis et al 2018:3). Online debate is an increasingly influential context for this, and tonality necessarily influences argument quality and accessibility in this forum, which in-turn impacts the likelihood and quality of political exposure for underrepresented groups. Although significant barriers exist, the case for a more civilised discourse online is clear, especially given the widespread role of such platforms in the political socialisation of young people in the UK (Jigsaw Research 2018:2). The concluding chapter brings together these insights to investigate policy solutions to the problematic information environment on Twitter, and suggest avenues for further research into the use of social media for the public good.

Chapter 7: Conclusion – Social Media for the Public Good

As set out at the end of Chapter 1, this thesis makes a threefold contribution to the extant literature. These are: to add to theoretical understanding about the role of social media platforms in highly salient political contexts, to increase consideration of the heterogeneity of information consumers on Twitter, shifting focus beyond direct participants to include effects on observers, and finally to examine pathways for improvement of social media's flaws. Having comprehensively considered the first two by analysing the central problems with the use of Twitter for public political interaction in the UK, this final chapter will outline policy proposals designed to ensure that social media benefits the public interest. This task is complex and multifaceted, leading previous authors to conclude that "given the scale and scope at which these platforms operate, these problems may ultimately be unsolvable" (Napoli 2019:163). The challenge is indeed great, but if policymakers recognise the threat and act decisively, such pessimism is as yet unwarranted.

The numerous measures outlined below are intended to ameliorate the effects of four important interacting determinants of widespread online incivility: misinformation, network selectivity, polarisation, and anonymity. These core issues have been discussed at length throughout this thesis, and lessoning their influence will naturally encourage the emergence of a more civilised discourse, helping facilitate the realisation of its previously hypothesised benefits. Debates about social media reform have generally focussed either on legal and regulatory avenues or civic initiatives designed to equip users to make better use of the platforms. Following Sunstein, this chapter argues that a rigid separation between these approaches is counter-productive, failing to recognise that social change is most often the product of a combination of legal reforms and cultural evolution (2018:250). Policy recommendations in both of these spheres will thus be delineated, before a summary of the key conclusions of this thesis and some reflections on social media's interaction with wider political contexts. The final sections will discuss limitations of this study and avenues for important future research which builds on its findings.

A call for a new and robust regulatory framework

Social media companies are increasingly among the most influential corporations in the world, and they now "operate as key components of the global media economy" (Turner 2018:9). Given that these platforms are highly consequential for public discourse, a strong argument can be advanced that they are currently exercising substantial social and economic power without corresponding accountability. Currently, social network companies are not subject to the same editorial requirements as other media outlets such as TV, radio, and newspapers. In the UK, they are considered 'content intermediaries' with correspondingly lessened responsibilities. Robert Thomson, CEO of US mass media firm News Corp, persuasively argues against this discrepancy: "these companies are in digital denial. Of course they are publishers and being a publisher comes with the responsibility to protect and project the provenance of news" (quoted in Waters et al, 2016). Righting this requires the imposition of obligations on the platforms to deal with false information comparably to the way they are currently mandated to handle illegal content. In a report for the European Parliament, Niklewitz proposes the creation of a 'notice and correct' procedure, under which "social media platforms, like the traditional press, would have to correct (or take down) false information at the request of a genuinely affected party" (2017:41). When a report of misleading information is made, the platforms would have 24 hours to respond, first deciding whether the content should stay up, be forcefully corrected, or deleted by a dedicated fact checking team. Failure to act would be legally and financially punishable for the platform, and if any content is removed its source must have a right of appeal (Niklewitz 2017:42). The coherent operation of

this system would rely on vigilant citizens to report misleading content, recognising that it would be an unrealistic to expect social media companies to proactively monitor all content posted. This would considerably improve existing arrangements which do not compel networks to act decisively, resulting in disputed news stories staying in circulation unless they meet the extreme threshold of illegality. Currently, misleading stories are simply made less prominent or at most flagged as disputed, with worrying evidence that this label can actually precipitate further distribution of the content (Napoli 2019:165). The scale of misinformation online demonstrates the inadequacy of current procedures, and thus platforms must be given editorial responsibility to take more decisive action in-line with obligations on other media.

A very similar system should also be introduced to combat incivility that is corrosive of healthy debate. To achieve this, a two-step implementation process is proposed, beginning with dissemination of the 'Rapoport rules' for political conversation on Twitter. These consist of criteria for respect and decency in disagreement. Before criticism can be advanced, users would be encouraged to seek clarification of an opponent's position, explicitly list points of agreement, and mention something learned from it (Dennett 2014:58). Although neither possible nor desirable to police all political conversation in this way, a baseline code of conduct for political discussion has been shown to improve behaviour in other forums (Polder-Verkiel 2012:140). More concretely, the second component would be a purposeful mechanism to report incivility that does not approach the high threshold of illegality but is clearly unconducive to fruitful debate. To avoid unelected platforms exercising undue political influence, this would be strictly limited to the regulation of tonality and could not apply to the expression of controversial opinions. Akin to the misinformation system above, users could be anonymously reported for uncivil conduct and be forced to issue a private apology if found guilty, or face suspension. If any user committed three breaches in a year, their profile would be automatically removed and IP addresses tracked to prevent continual use of new accounts. For this mechanism to have force and properly promote civility, it is clear that no user could be exempt. This would mean that uncivil behaviour from politicians and other opinion shapers could be properly sanctioned, encouraging elites to lead in diffusing incivility or face public humiliation (Duncombe 2019:419). Politicians acting as role models is a crucial determinant of lasting behavioural change online, discussed later.

However, it is clear from previous chapters that neither incivility nor misinformation are likely to be effectively reduced without challenging network selectivity on Twitter. While impossible to radically enforce this without fundamentally changing the character of the platform, greater exposure to viewpoint-challenging perspectives is easily achievable. Content algorithms could be altered to consider the ideological balance and likely reliability of a user's network, promoting credible and opinion-challenging tweets where necessary. The technological capacity for this already exists, but is currently used for the opposite purpose and recommends that users follow accounts closely matching their existing choices (Majó-Vázquez and González-Bailón 2019:251). It seems increasingly apparent that this is unlikely to serve the public interest. However, any such mechanism must not simply replace hyper-partisan and questionable information with its equivalent of a contrasting ideological impulse. As others have regularly expounded, objectivity is not coterminous with balance, and improvement is contingent upon solutions that only promote the spread of authoritative alternative opinions (Crilley and Gillespie 2019:175).

Twitter must also be obliged to make far greater investment in 'bot' removal and properly verify accounts, to reduce the considerable harms emanating from online anonymity. As previously covered, bots are a growing cause of both disinformation and incivility (Gorodnichenko et al, 2018), and their proliferation demonstrates the failure of existing

attempts to regulate fake content. However, the problem extends far beyond bots. It is increasingly clear that the combination of anonymity and ease of interaction with strangers on Twitter is emboldening the spread of incivility much faster than is found offline (Levmore 2011:54). To disincentivise this, Twitter should introduce stronger mechanisms of user verification by mandating submission of an email address as a prerequisite of an account, removing any profiles that do not provide verification. This should be accompanied by greater tracking of IP addresses to prevent the nefarious duplication of troublesome accounts from the same source, and robust enforcement of laws against knowingly posing as another person online (Scott and Orlikowsi 2014:878). The example of Wikipedia provides a potentially fruitful blueprint in this regard. “Any contribution or revision made to a Wikipedia entry is recorded on a separate, linked history page. This page details the substance of the revision, the date and time it is made, and some information about the contributor. Contributors to Wikipedia essentially agree to be identified on the history page by their IP addresses, and this identification discourages misbehaviour” (Levmore 2011:60). This disciplinarian approach is combined with the cultivation of a shared collaborative ethos, creating a positive climate for civil interaction (Antoci et al 2019:97). By employing a similar mix of methods, Twitter could minimise the direct harms of anonymity without excessively infringing upon the right to privacy.

There are legitimate concerns that attempts to combat the ills of anonymity, network selectivity and misinformation could unintentionally constitute an unacceptable intrusion into democratic freedoms of choice, speech, and privacy (Scaife 2013:130). There is undoubtedly a delicate balance between the promotion of the public interest and the protection of individual autonomy, and reformers should remain cautious of unintended consequences. However, the lack of transparency characterising the current information environment on Twitter brings the notion that existing discourse is the product of democratic freedoms into question (Levmore 2011:67). As discussed earlier with regard to the EU Referendum campaigns, platforms are now awash with targeted advertising, algorithmic behavioural ‘nudges’, and fabricated content. From this perspective, utilising social media’s technological capacities to combat misinformation, encourage diversity in content exposure, and target nefarious misuse of anonymity is the more socially responsible path (Napoli 2019:171). As Niklewitz concludes “it would be rather naïve to guarantee totally unrestricted freedom of speech to those whose long-term aim is to destroy democracy and its freedoms altogether. Unrestricted free speech, devoid of any form of commonly shared rules, would be the equivalent of a snake eating its own tail” (2017:44).

Perhaps the more appropriate apprehension pertains to the concentration of substantial gatekeeping power in the hands of unelected profit-seeking private businesses (Crilley and Gillespie 2019:173). While greater editorial responsibility for social media companies is necessary, there is clear need for a mechanism of comprehensive oversight to monitor such an important social function. Accordingly, the final and perhaps most significant regulatory recommendation is for the establishment of a state funded non-partisan dedicated social media watchdog to oversee the initial implementation and continued operation of the above-suggested measures. State-level law enforcement pertaining to social media is currently weak across advanced democracies, with meaningful remedial action inhibited “as a result of taking a case-by-case approach rather than adopting a framework from which to base decisions” (Scaife 2013:129). Legislation is improving, as shown by Britain’s recent commission of a ‘national security communications unit’ aimed at combatting disinformation, with the potential to expand into other regulatory functions. In 2018, Germany introduced the strongest legal measures on disinformation and vitriol yet seen, a law which “requires social media platforms with more than two million users to remove fake news, hate speech, and other illegal material within twenty-four hours of notification, or receive fines of up to fifty million Euros” (Napoli

2019:187). The proposed independent regulator would reinforce this progress by having the power to compel compliance with directives, unaffected by the biases of party politics. Indeed, the perception of political neutrality will be vital to the success of any such body, and therefore it must be fully removed from political interference but empowered with a clear brief and the means to effectively regulate a fast-moving online space (Levmore 2011:56). The authority of this reform could be well complemented by improved corporate social responsibility among advertisers profiting from social media, who could enhance their brand image by taking a public stand against the excesses of incivility and misinformation online (Turner 2018:11). This combination of measures would ensure that social networks take the expanded responsibilities proposed here seriously, and help these recommendations to have full effect.

Education, active citizenship, and a revitalised civil society

Taken as a package, the above policies highlight the pressing need for social media companies and legislators alike to assume proper responsibility for the current nature of political discourse online. However, formal regulation alone has the potential to be both overbearing, risking undue erosion of democratic freedoms, and counterproductive, by appearing to overly sanitise political discussion (Boyd 2006:863). To avoid these criticisms, better regulation needs to be accompanied by a collective voluntary improvement of the expected standards in political debate. With this, social media could become a powerful vehicle for the public interest, based around "an enhancement of mutual understanding, greater appreciation for differences and the views of others, diminution of racial and gender boundaries, and the building of shared values" (Bimber 1998:145). Achieving this requires new initiatives aimed at revitalising civil society and increasing social capital in growingly divided times. This section will outline some plausible steps to realise this.

Critical to this endeavour is responsible elite leadership and a more respectful discourse among politically influential figures. Modern politics is increasingly personalised, a trend influenced by social media (Enli and Skogerbo 2013:758). A key benefit of platforms like Twitter for politicians is reduced reliance upon journalists to generate publicity; they can now produce and circulate content to large audiences themselves. In this forum, many opinion leaders are failing to behave in a manner which promotes the public interest online; using this advantage to denigrate opponents, encourage hyper-partisanship and fuel polarisation. As Crilley and Gillespie write, "examples of the cyclical reproduction of misinformation and disinformation are abound: from Trump's Tweets spreading misinformation about his inauguration crowd size, refugees and Muslims, to the Republican party's use of social media advertising to suppress voter turnout; from Cambridge Analytica's spread of dark ads, to the Russian state's creation of adverts to influence elections" (2019:175). The general use of Twitter for negative campaigning combined with such instances have been said to "signal the existence of a political establishment that is perhaps more interested in manipulating public opinion than offering fruitful solutions to enduring problems" (Hopp and Vargo 2017:375). Furthermore, a wealth of evidence shows a clear link between elite statements on social media and opinion formation among voters, with additional indications that this may be stronger for demographics that are less politically engaged (Park 2013:1641; Parmelee and Bichard 2011:205-209). It is thus a reasonable assumption that behavioural change in public political debate on social media is unlikely to emerge if public figures do not begin to use their influence to challenge the notion that political discourse is unavoidably polarised and uncivil.

Alongside improving their own conduct, policymakers should consider a major program of society-wide civic education with the aim of creating a more vigilant citizenry and ameliorating the growing inequalities in political knowledge driving social media's negative impacts on democracy (Prior 2007:96). Education has been regularly identified as an influential

antecedent of political participation, as well as likelihood of adherence to civil norms online (Vargo and Hopp 2017:24). This pedagogical initiative would first aim to dramatically improve and update political education in the UK throughout school ages. Given the increasing extent that young people are now politically socialised on online platforms, it is of pressing importance that these demographics are equipped with the necessary tools to be aware of the problematic information environment on Twitter and other sites, and are able to challenge behavioural norms, reshaping online discourse in a more engaging and inclusive manner. This is particularly important in light of recent findings that young people are the most likely demographic to engage in opinion-reinforcing selective curation of their online networks, in order to "navigate the sea of information of uneven quality" (Lee et al 2019:2277). As Niklewitz writes, better education would address the fact that "currently too many users forward content without reading anything but the headline, and too few care about the source of the content. This is a cultural problem, one that is shaped by disconnections in values, relationships and the social fabric" (2017:49). This issue clearly spans generations. Accordingly, a second component of the initiative would expand this sentiment beyond formal education and to all age groups, to include widely publicised campaigns about the dangers of misinformation and network selectivity online, and a large expansion of free citizenship classes through higher education institutions. These would challenge some commonly held negative perceptions about politics, inculcating norms of civility and encouraging greater civic participation.

A key marker of efficacy for this civic education would be its success in encouraging politically moderate voices to feel more confident about contributing to political discourse. A central problem of the online status quo is the overrepresentation of extreme fringes of opinion because moderates are either disinclined to participate in such a frenzied and unwelcoming environment, or they lack the personal efficacy to challenge the excesses of low quality uncivil discourse (Levmore 2011:66). As explained in previous chapters, the net result is too often debate that lacks compromise, civility, or proper consideration of alternative viewpoints. If Twitter is to be successfully remade in the public interest, moderates must be empowered to challenge what they disagree with, call out unnecessary incivility, and build consensus (Khang et al 2014:54; Friedersdorf, 2015). The structural capacities of social media are not an enemy in this regard, greater accessibility of information means that "citizens no longer have to be passive consumers of political party propaganda, government spin or mass media news, but are instead actually enabled to challenge discourses, share alternative perspectives and publish their own opinion" (Loader and Mercea 2011:759). Indeed, the healthy exchange of ideas in a democracy precisely relies upon an active conception of citizenship (Habermas, 2006), with the dangers of passive observation online readily apparent in the fragmentation of recent discourse (Polder-Verkiel 2012:118). Broadened participation would likely be encouraged by a public reassertion of respect for knowledge hierarchies and the important role played by journalists and experts as authoritative sources. Once more, this effort should be particularly targeted towards the young, with evidence that they are more inclined to believe information circulated online by their peers than from a news article (Vromen et al 2015:82). Such attitudes are likely a result of mistrust of information sources on social media, and a cause of aversion towards political discussion (Kruse et al 2018:75). However, if current passive observers were persuaded that their ideas would be listened to and that reference to credible information would be respected, a model of active citizenship becomes easier to promote. Taken together, a successful civic education drive, a more respectful elite discourse, and a public reassertion of the importance of journalists and experts as arbiters of competing knowledge claims, would likely see a vast improvement in the scope and quality of political participation on social media.

Underpinning all of these recommendations is a need for widespread recognition of the inherent tensions between the public interest and unfettered individual consumer freedom online (Sunstein 2018:82). A revitalised civil society is contingent upon an admission that promotion of the public interest requires trade-offs and a conscious decision to prioritise this in the online space (Turner 2018:12). Finding consensus in this area will be a complex task, and to avoid perceptions of political bias and overbearing state intervention, it is very important that the collective reform process has popular legitimacy. Accordingly, the final policy recommendation advanced here is for the establishment of a citizens' assembly dedicated to making social media work in the public interest. The concept has been gaining ground in a variety of political settings in recent years and has proven a useful tool for handling controversial issues requiring democratic consultation, most successfully on constitutional reform in the Republic of Ireland (Farrell et al, 2019). Although currently underused and undertested, there is direct precedent for citizens' assemblies in the UK too; a 2018 group considered action on social care and a forthcoming round will be dedicated to climate change. Following such examples, it would see the creation of a group of 100 deliberators, two thirds of them randomly selected citizens according to a representative sortition procedure among social media users, and the other third a cross-party group of MPs. The body would be given information on the topic before meeting and hear from a number of experts, prior to being charged with a brief to generate policy recommendations on how best to mitigate the problems raised in this thesis. Parliament would then consider them, with a binding commitment to trial and monitor feasible suggestions. If executed properly and given the power to make valued recommendations, an assembly of this kind could be a useful space for an open public dialogue about the compromises and initiatives required to make social media work for the public good.

Summary of conclusions and interaction with broader political trends

Full implementation of the policy package recommended above would go a long way towards this goal. However, this research has shown that political discourse online does not take place in a contextual vacuum, and that the information environment on Twitter interacts with broader political developments in important ways. The positive impacts of Twitter civility on trust and engagement were somewhat reduced when debate centred on a highly salient and divisive issue, in this case Brexit. In support of this finding, Hobolt, Leaper and Tilley assert that "the apparent unwillingness of citizens even to speak across the (Brexit) divide, let alone respect or befriend one another, would seem to undermine the possibility of a deliberative cure". This should serve as further warning to public figures to be cognisant of their future responsibilities in preventing such extreme issue polarisation, for it appears that these contexts inhibit the positive impacts of civil democratic deliberation online, even if it were to emerge. The same authors sagely conclude that "the task may not be to find consensus across the divide, but instead to help citizens to recognise one another not as enemies and out-groups, but as adversaries with a shared collective identity disagreeing over the outcomes of policy debate" (2020:30). The challenges of overcoming divides in the online space are thus neither trivial nor easily solved, but this research suggests that a substantial improvement is not impossible.

The strength of the key results in this thesis are noteworthy, not least because experimentally modelling political attitude change is a challenging endeavour. Dispositions towards trust and engagement are important aspects of an individual's civic identity, and as such they are highly durable and not easy to change. That civility in political conversation repeatedly induced an effect for politically unsure demographics should thus be taken as a finding of serious interest. Previous research has found that interaction with unknown others can provide cognitive resources that individuals do not find in their immediate environment, and citizens who report

larger online and offline conversational circles tend to engage more in civic life (Gil de Zuniga and Valenzuela 2011:413). Social media like Twitter facilitates this naturally, but this study indicates that these benefits are likely conditional on the nature of the engagement. Increased civility in online interactions is argued to be a clear pathway to a more inclusive political discourse that could begin to ameliorate the deepening inequalities in political trust, engagement and knowledge across social media users in Britain. If this starts to arise, there is hope that some of Twitter's long-held potential for community-building, information spreading, and healthy democratic debate could eventually be realised. To the extent that the public sphere can be conceptualised as "a place where private people come together as a public for the purpose of using reason to further critical knowledge which, in turn, leads to political change" (Kruse et al 2018:63), it is clear that Twitter could be transformed for the good of democracy. However, as the recommendations advanced in this chapter demonstrate, there is much work to be done before Twitter becomes a vehicle for the public good in Britain. A robust regulatory framework that has multi-party support and strict enforcement mechanisms is an essential part of such a project, but must be supported by a revitalised civic culture. This is only likely to emerge if political elites lessen their increasing use of uncivil and polarising rhetoric for partisan gain, and enact policies that actively promote the public good over unfettered consumer freedom online. Only this combination of legal and social reform can fully meet the scale of the challenge that the current information environment on Twitter poses for democracy in the UK and beyond.

Limitations and notes of caution

Specific limitations relating to the experimental methodology employed in this thesis were detailed in Chapter 5 and will not be repeated here. However, there are some more general notes of caution which should be attached to these conclusions. Firstly, it should be clearly stated that the direct scope of this study extends only to the use of Twitter in the UK. A string of previous research has shown that different social media platforms have varying characteristics, and effects will inevitably vary by national context (see Blank and Lutz, 2017). While testing the applicability of these findings to different countries and social media platforms would be of great interest, this cannot be known without further research.

Secondly, although the choice to construct a sample using only Twitter users has many more advantages than weaknesses, the data here does have significant skews in favour of young, educated, politically interested, left-wing, and Remain supporting demographics. As mentioned in Chapter 5, a sample size of 1790 split between 6 groups with imbalances in the data meant confidence intervals among less populated subgroups were inevitably larger owing to the shortage of respondents in these categories. This inhibited precise results for many potentially interesting heterogeneous treatment effects, and a repeated experiment with larger sample size would verify these more robustly. Furthermore, the scenarios created for these experiments were 'ideal type' imitations of Twitter debate, and political conversations on Twitter are rarely so clearly structured. In this study, participants were informed about a clear differentiation between 'expert' and 'public', and the tone of debate was either completely civil or completely uncivil. Political discourse on Twitter is in reality not so explicitly defined, and accordingly respondents reactions could have been augmented to a greater or lesser degree than they conventionally would when encountering a political exchange on Twitter.

Finally, the question of how to conceptualise political engagement, particularly among young people, is becoming increasingly important to research such as this. Alongside concerns around voting intention being exaggerated by direct proximity to an election, the measure itself is potentially restrictive. As Loader and Mercea write, "if we move beyond the traditional engagement with mainstream politics, such as voting, party membership, petitioning

representatives and the like, and adopt a more open conception of democratic citizenship, a different focus and set of questions emerge" (2011:761). While there is a strong argument that participation at the ballot box remains the most consequential form of political activity in representative democracy, the conclusions relating to engagement in this study should not be viewed as capturing the full picture, nor should those who expressed an intention not to vote be considered definitively disengaged.

Avenues for further research

It is hoped that the findings and recommendations of this thesis will be of relevance to policymakers and scholars across a variety of disciplines. These include: the study of British politics, political behaviour and psychology, the workings and ethics of computer-mediated interaction, and the challenges of public policy in the digital age. However, the interaction between social media and politics is rapidly evolving, existing conclusions age quickly and new research must keep pace with change. This is not a simple task; social media platforms and national contexts are both highly varied and new forms of social media are arising every year (Blank and Lutz 2017:752). Consequently, many areas remain unexplored and undertheorised, leaving plentiful opportunities for scholars to further the collective understanding of this crucial topic, an endeavour of pressing importance for discerning the determinants of current democratic instability and how to improve it. Such research should embrace the use of original approaches which engage with the dynamic, multiple, and contingent materialities constituting contemporary online realities (Scott and Orlikowski 2014:891). It must also recognise the important role of offline politics in such interactions, namely the "impact of social diversity, inequality and cultural difference as important sources of power influencing democratic innovation" (Loader and Mercea 2011:760).

It is also clear that there is huge variation in the way that Twitter is used by citizens. Further research into the areas covered by this study should investigate how Twitter's impact on trust and engagement is mediated by different usage preferences, for it is likely that entertainment seekers have significantly different expectations of political content than those who actively search for news. Similarly, little is known about the degree that exposure to political information on Twitter facilitates long-term knowledge gains, and under what conditions these are likely to occur. Academic enquiry should also be aware that impressions of social media are necessarily dominated by those who choose to contribute. However, the large proportion of users who are simply observers are being equally exposed to and socialised by such platforms, and continued research into this cohort is vital to understanding this topic. Furthermore, though this study offers a promising starting point, not enough is known about the precise role of issue contexts in mediating the effects of online discourse for democratic attitudes. Understanding the ways in which offline salience and polarisation get projected onto the online space is highly important, and future studies in a similar vein to this one should continue to explore this.

Perhaps most importantly, the central theoretical contribution of this thesis is to show that the predominant focus upon incivility in existing literature is missing an important aspect of the relationship between social media and democratic politics. The possible emotional and behavioural impacts of incivility are very well-documented, but given that it is becoming increasingly normalised online, new research in this area should be probing the direct impacts of civility in more detail. When this understanding is cultivated, it will be possible to make more closely informed judgements on the conditions under which political discourse on Twitter can improve for the good of democracy.

End Matter

Bibliography

Allen, K., Ryan, T., Gray, D., McInerney, D. and Waters, L. (2014). "Social Media Use and Social Connectedness in Adolescents: The Positives and the Potential Pitfalls". *The Australian Educational and Developmental Psychologist*, 31(1): 18-31.

Almond, G. and Verba, S. (1963). *The Civic Culture*. London: SAGE.

Al-Saleh, M., Owaidah, T., Alzahrani, H., Bin Ali, T., Elfaraidi, H., Raish, M., Alkhayal, Z., Al-Momen, A., Mohammed, V. and Al-Zahrani, A. (2017). "Positive Impact of Social Media on Disease Awareness and Management in Patients with Inherited Bleeding Disorders: Saudi Experience from a Tertiary Care Hospital". *Blood,* 1: 23-32.

Altman, D. (2016). "Randomised Trials" in Greenfield, T. and Greener, S. (eds). *Research Methods for Postgraduates*. London: Wiley.

Anderson, A., Brossard, D., Scheufele, D., Xenos, M. and Ladwig, P. (2013). The "'Nasty Effect": Online Incivility and Risk Perceptions of Emerging Technologies'". *Journal of Computer-Mediated Communication* 19(3): 373–387.

Anderson, M. (2015). "Men catch up with women on overall social media use". *Pew Research Centre.* http://www.pewresearch.org/fact-tank/2015/08/28/men-catch-up-with-women-on-overall-social-media-use

Antoci, A., Bonelli, L., Paglieri, F., Reggiani, T. and Sabatini, F. (2019). "Civility and trust in social media". *Journal of Economic Behaviour and Organization*, 160: 83-99.

Ardevol-Abreu, A., Diehl, T. and Gil de Zuniga, H. (2019). "Antecedents of internal political efficacy incidental news exposure online and the mediating role of political discussion". *Politics*, 39(1): 82-100.

Astin, A. and Denson, N. (2009). "Multi-campus studies of college impact: Which statistical method is appropriate?". *Research in Higher Education*, 50: 354-367.

Balliet, D. and Van Lange, P. (2012). "Trust, Conflict, and Cooperation: A Meta-Analysis". *Psychological Bulletin*, December: 1-23.

Bamia, C., White, I. and Kenward, M. (2012). "Some consequences of assuming simple patterns for the treatment effect over time in a linear mixed model". *Statistics in Medicine*, 32: 2585-2594.

Bassilakis, A., Diermeier, M. and Goecke, H. (2018). Converging Media versus Diverging Politics - the Brexit Twitter on Debate". *CESifo Forum*, 19(4): 29-35.

BBC News. (2018). "Vote Leave's targeted Brexit ads released by Facebook". Available at: https://www.bbc.co.uk/news/uk-politics-44966969.

Becker, S., Fetzer, T. and Novy, D. (2017). "Who voted for Brexit? A comprehensive district-level analysis". *Economic Policy*: 601-651.

Benjamin, J. (2011). "Tweets, Blogs, Facebook and the Ethics of 21st-Century Communication Technology" in Noor, A. (eds). *Social Media: Usage and Impact*. Lanham: Lexington Books.

Berinsky, A., Huber, G. and Lenz, G. (2012). “Evaluating Online Labor Markets for Experimental Research: Amazon.com's Mechanical Turk”. *Political Analysis*, 20(3): 351-268.

Bimber, B. (2001). “Information and Political Engagement in America: The Search for Effects of Information Technology at the Individual Level”. *Political Research Quarterly*, 54(1): 53-67.

Bimber, B. (1998). “The Internet and Political Transformation: Populism, Community, and Accelerated Pluralism”. *Polity*, 31(1): 133-160.

Blank, G. (2017). “The Digital Divide Among Twitter Users and Its Implications for Social Research”. *Social Science Computer Review*, 35(6): 679-697.

Blank, G. and Lutz, C. (2017). “Representativeness of Social Media in Great Britain: Investigating Facebook, LinkedIn, Twitter, Pinterest, Google+, and Instagram”. *American Behavioural Scientist*, 61(7): 741-756.

Boczkowski, P., Mitchelstein, E. and Matassi, M. (2018). “’News comes across when I’m in a moment of leisure’: Understanding the practices of incidental news consumption on social media”. *New Media and Society*, 20(10): 3523-3539.

Bode, L. (2016). “Political News in the News Feed: Learning Politics from Social Media, Mass Communication and Society”. *Mass Communication and Society,* 19(1): 24-48.

Bonacchi, C., Altaweel, M. and Krzyzanska, M. (2018). “The heritage of Brexit: Roles of the past in the construction of political identities through social media”. *Journal of Social Archaeology*, 18(2): 174-192.

Borah, P. (2012). “Does it Matter Where you Read the News Story? Interaction of Incivility and News Frames in the Political Blogosphere”. *Communication Research*, 46(6): 809-827.

Boyd, D. and Crawford, K. (2012). “Critical Questions for Big Data: Provocations for a cultural, technological, and scholarly phenomenon”. *Information, Communication and Society*, 15(5): 662-679.

Boyd, R. (2006). “’The Value of Civility’”. *Urban Studies*, 43(5-6): 863-878.

Brandtzaeg, P. (2017). “Facebook is no ‘‘Great equalizer’’: A big data approach to gender differences in civic engagement across countries”. *Social Science Computer Review,* 35(1): 103-125.

Brooks, D. and Geer, J. (2007). “Beyond Negativity: The Effects of Incivility on the Electorate”. *American Journal of Political Science 51*(1), 1–16.

Browning, C. (2019). “Brexit populism and fantasies of fulfilment”. *Cambridge Review of International Affairs*, 32(3): 222-244.

Bruns, A. and Stieglitz, S. (2013). “Towards more systematic Twitter analysis: metrics for tweeting activities”. *International Journal of Social Research Methodology*, 16(2): 91-108.

Cadwalladr, C. (2017). “Cambridge Analytica affair raises questions vital to our democracy”. *The Guardian*. Available at: https://www.theguardian.com/politics/2017/mar/04/cambridge-analytica-democracy-digital-age.

Capella, J. (2017). “Vectors into the Future of Mass and Interpersonal Communication Research: Big Data, Social Media, and Computational Social Science”. *Human Communication Research*, 43: 545–558.

Capella, J. and Jamieson, K. (1997). *Spiral of Cynicism: The Press and the Public Good*. New York: Oxford University Press.

Capella, J., Kim, H. and Albarracin, D. (2015). “Selection and Transmission Processes for Information in the Emerging Media Environment: Psychological Motives and Message Characteristics”. *Media Psychology*, 18(3): 396-424.

Capellini, B., Kravets, O. and Reppel, A. (2019). “Shouting on social media? A borderscapes perspective on a contentious hashtag”. *Technological Forecasting and Social Change*, 145: 428-437.

Capuano, A., Dawson, J., Ramirez, M., Wilson, R., Barnes, L. and Field, R. (2016). “Modelling Likert Scale Outcomes With Trend-Proportional Odds With and Without Cluster Data”. *Methodology*, 12(2): 33-43.

Cardenal, A., Aguilar-Paredes, C., Cristancho, C. and Majó-Vázquez, S. (2019). “Echo-chambers in online news consumption: Evidence from survey and navigation data in Spain”. *European Journal of Communication*, 34(4): 360-376.

Carl, N., Richards, L. and Heath, A. (2019). “Leave and Remain voters’ knowledge of the EU after the referendum of 2016”. *Electoral Studies*, 57: 90-98.

Carminati, B., Ferrari, E. and Viviani, M. (2013). *Security and Trust in Online Social Networks.* California: Morgan and Claypool.

Ceron, A., Curini, L., Iacus, S. and Porro, G. (2013). “Every tweet counts? How sentiment analysis of social media can improve our knowledge of citizens’ political preferences with an application to Italy and France.” *New Media and Society*, 16: 340-358.

Chaffey, D. (2019). “Global social media research summary 2019”. Available at: https://www.smartinsights.com/social-media-marketing/social-media-strategy/new-global-social-media-research/

Clarke, H., Sanders, D., Stewart, M. and Whiteley, P. (2004). *Political choice in Britain.* Oxford: Oxford University Press.

Clemence, M. (2019). “Trust in politicians falls sending them spiralling back to the bottom of the Ipsos MORI Veracity Index”. Available at: https://www.ipsos.com/ipsos-mori/en-uk/trust-politicians-falls-sending-them-spiralling-back-bottom-ipsos-mori-veracity-index.

Clifford, S., Jewell, R. and Waggoner, P. (2015). “Are samples drawn from Mechanical Turk valid for research on political ideology?”. *Research and Politics*, October-December: 1-9.

Colleoni, E., Rozza, A. and Arvidsson, A. (2014). “Echo Chamber or Public Sphere? Predicting Political Orientation and Measuring Political Homophily in Twitter Using Big Data”. *Journal of Communication*, 64: 317-332.

Conover, M., Ratkiewitz, J., Francisco, M., Goncalves, B., Flammini, A. and Menczer, F. (2011). “Political Polarization on Twitter”. *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*: 89-96.

Crilley, R. and Gillespie, M. (2019). “What to do about social media? Politics, populism and journalism”. *Journalism*, 20(1): 173-176.

Curtice, J. (2017). “Why Leave Won the UK’s EU Referendum”. *Journal of Common Market Studies*, 55: 19-37.

Dahl, R. (1998). *On Democracy.* New Haven: Yale University Press.

Daniels, G. (2014). “How far does Twitter deepen democracy through public engagement?: an analysis of journalists’ use of Twitter in the Johannesburg newsroom”. *Journal of African Media Studies*, 6(3): 299-311.

Davenport, T. Gerber, A. and Green, D. (2010). “Field Experiments and the Study of Political Behaviour” in Leighley, J. (ed.). *The Oxford Handbook of American Elections and Political Behaviour*. Oxford: Oxford University Press.

De Vaus, D. (2016). “Survey Research” in Greenfield, T. and Greener, S. (eds). *Research Methods for Postgraduates*. London: Wiley.

Del Vicario, M., Zollo, F., Caldarelli, G., Scala, A. and Quattrociocchi, W. (2017). “Mapping social dynamics on Facebook: The Brexit debate”. *Social Networks*, 50: 6-16.

Dennett, D. (2014). *Intuition Pumps and Other Tools for Thinking*. New York: W.W. Norton.

Downs, A. (1957). *An Economic Theory of Democracy*. New York: Harper Collins.

Druckman, J., Gubitz, S., Levendusky, M. and Lloyd, A. (2019). “How incivility on partisan media (de) polarizes the electorate”. *The Journal of Politics,* 81(1): 291–295.

Duffy, B., Hewlett, K., McCrae, J. and Hall, J. (2019). “Divided Britain? Polarisation and fragmentation trends in the UK”. *Report produced for The Policy Institute, King’s College London*.

Duggan, M. and Smith, A. (2016). “The political environment on social media”. Technical report, *Pew Research Center*: 1-39.

Duncombe, C. (2019). “The Politics of Twitter: Emotions and the Power of Social Media”. *International Political Sociology*, 13: 409-429.

Ecker, A. (2017). “Estimating Policy Positions Using Social Network Data: Cross-Validating Position Estimates of Political Parties and Individual Legislators in the Polish Parliament”. *Social Science Computer Review*, 35(1): 53-67.

Elliot, C. (2017). “All-out war? Brexit, voting and the production of division”. *Renewal*, 25(3): 44-55.

Enli, G. and Skogerbo, E. (2013). “Personalized campaigns in party-centred politics”. *Information, Communication & Society*, 16(5): 757-774.

Evans, G. and Schaffner, F. (2019). “Brexit identities: how Leave versus remain replaced Conservative versus Labour affiliations of British voters”. *The UK in a Changing Europe*. Available at: https://ukandeu.ac.uk/brexit-identities-how-leave-versus-remain-replaced-conservative-versus-labour-affiliations-of-british-voters/

Farrell, D., Suiter, J. and Harris, C. (2019). “‘Systematizing’ constitutional deliberation: the 2016–18 citizens’ assembly in Ireland”. *Irish Political Studies*, 34(1): 113-123.

Fletcher, R. and Neilsen, R. (2018). “Are people incidentally exposed to news on social media? A comparative analysis”. *New Media and Society*, 20(7): 2450-2468.

Fletcher, R. and Neilsen, R. (2017). “Are News Audiences Increasingly Fragmented? A Cross-National Comparative Analysis of Cross-Platform News Audience Fragmentation and Duplication”. *Journal of Communication*, 67: 476-498.

Fletcher, R. and Park, S. (2017). “The Impact of Trust in the News Media on Online News Consumption and Participation”. *Digital Journalism*, 5(10): 1281-1299.

Flynn, P. (2017). “What Brexit should have taught us about voter manipulation”. *The Guardian*. Available at: https://www.theguardian.com/commentisfree/2017/apr/17/brexit-voter-manipulation-eu-referendum-social-media

Friedersdorf, C. (2015). “In Defence of Civility on Twitter”. *The Atlantic,* 4 January. Available at: https://www.theatlantic.com/politics/archive/2015/11/in-defense-of-civility-on-twitter/413860/.

Gayo-Avello, D. (2012): “No, You Cannot Predict Elections with Twitter”. *IEEE Internet Computing*, 16(6): 91-94.

Gerber, A. and Huber, G. (2010). “Partisanship, Political Control, and Economic Assessments”. *American Journal of Political Science*, 54(1): 153-173.

Gervais, B. (2015). “Incivility online: Affective and behavioural reactions to uncivil political posts in a web-based experiment”. *Journal of Information Technology and Politics 12* (2): 1–19.

Gervais, B. (2017). “More than Mimicry? The Role of Anger in Uncivil Reactions to Elite Political Incivility”. *International Journal of Public Opinion Research 29*(3): 384-405.

Gervais, B. (2018). “Rousing the partisan combatant: Elite incivility, anger, and anti-deliberative attitudes”. *Political Psychology*: 1-19.

Gil de Zuniga, H. and Diehl, T. (2015). “Citizenship, Social Media, and Big Data: Current and Future Research in the Social Sciences”. *Social Science Computer Review*, 35(1): 3-9.

Gil de Zuniga, H. and Valenzuela, S. (2011). “The Mediating Path to a Stronger Citizenship: Online and Offline Networks, Weak Ties, and Civic Engagement”. *Communication Research*, 38(3): 397-421.

Gil de Zuniga, H., Molyneux, L. and Zheng, P. (2014). Social Media, Political Expression, and Political Participation: Panel Analysis of Lagged and Concurrent Relationships. *Journal of Communication*, 64: 612-634.

Golder, M. and Ferland, B. (2018). “Electoral Systems and Citizen-Elite Ideological Congruence” in Herron E. (eds.). *The Oxford Handbook of Electoral Systems*. Oxford: Oxford University Press.

Golder, S. and Macy, M. (2012). “Social Science with Social Media”. *American Sociological Association Footnotes*, 40(1): 7.

Gorodnichenko, Y., Pham, T. and Talavera, O. (2018). “Social media, sentiment and public opinions: Evidence from #Brexit and #USElection”. *National Bureau of Economic Research*, Working Paper 24631: 1-39.

Green, J. and Prosser, C. (2016). “Party system fragmentation and single-party government: the British general election of 2015”. *West European Politics*, 39(6): 1299-1310.

Guess, A. (2019). “(Almost) everything in moderation: New evidence on Americans’ online media diets”. (Unpublished manuscript). Retrieved from: https://csdp.princeton.edu/publications/almost-everything-moderation-new-evidence-americans-online-media-diets.

Guess, A., Munger, K., Nagler, J. and Tucker, J. (2019). “How accurate are survey responses on social media and politics?”. *Political Communication,* 36(2): 241-258.

Guess, A., Lyons, B., Nyhan, B. and Reifler, J. (2018). “Avoiding the echo chamber about echo chambers: Why selective exposure to like-minded political news is less prevalent than you think”. *A Technical Report for the Knight Foundation:* 1-25.

Habermas, J. (1989). *The Structural Transformation of the Public Sphere*. Cambridge: MIT University Press.

Hanna, B., Kee, K. and Robertson, B. (2017). “Positive Impacts of Social Media at Work: Job Satisfaction, Job Calling, and Facebook Use among Co-Workers”. *SHS Web of Conferences*, 33: 1-7.

Hänska, M. and Bauchowitz, S. (2017). “Tweeting for Brexit: how social media influenced the referendum*”* in Mair, J. (eds.) *Brexit, Trump and the Media*. Bury St Edmunds: Abramis: 31-35.

Hennig-Thurau, T., Wiertz, C. and Feldhaus, F. (2015). “Does Twitter matter? The impact of microblogging word of mouth on consumer’s adoption of new movies”. *Journal of the Academy of Marketing Science*, 43: 375–394.

Hermida, A. (2012). “Tweets and Truth: Journalism as a discipline of collaborative verification”. *Journalism Practice*, 6(5-6): 659-668.

Hobolt, S., Leeper, T. and Tilley, J. (2020). “Divided by the vote: affective polarisation in the wake of the Brexit referendum”. *British Journal of Political Science* (forthcoming)*:* 1-37.

Hobolt, S., Leeper, T. and Tilley, J. (2018). “Emerging Brexit Identities” in Menon, A. (eds). *Brexit and Public Opinion*. London: KCL.

Hobolt, S. (2016). “The Brexit vote: a divided nation, a divided continent”. *Journal of European Public Policy*, 23(9): 1259-1277.

Hopp, T. and Vargo, C. (2017). “Does negative campaign advertising stimulate uncivil communication on social media? Measuring audience response using big data”. *Computers in Human Behaviour*, 68: 368-377.

Hossain, M., Dwivedi, Y., Chan, C., Standing, C. and Olanrewaju, A. (2018). “Sharing Political Content in Online Social Media: A Planned and Unplanned Behaviour Approach”. *Information Systems Frontiers*, 20: 485-501.

Howard, P. and Kollanyi, B. (2016). “Bots, #StrongerIn, and #Brexit: Computational Propaganda during the UK-EU Referendum”. *COMPROP Research Note,* 1: 1-6.

Huang, F. (2018). “Multilevel Modelling and Ordinary Least Squares Regression: How Comparable Are They?”. *The Journal of Experimental Education*, 86(2): 265-281

Iakhnis, E., Rathbun, B., Reifler, J. and Scotto, T. (2018). “Populist referendum: Was ‘Brexit’ an expression of nativist and anti-elitist sentiment?”. *Research and Politics,* 5(2): 1-7.

Intal, C. and Yasseri, T. (2019). Dissent and Rebellion in the House of Commons: A Social Network Analysis of Brexit-Related Divisions in the 57th Parliament”. *Physics and Society,* 1(2): 1-35.

Iyengar, S. and Westwood, S. (2014). “Fear and Loathing across Party Lines: New Evidence on Group Polarization”. *American Journal of Political Science,* 59(3): 690-707.

Jamieson, K., Volinsky, A., Weitz, I. and Kenski, K. (2017). “The Political Uses and Abuses of Civility and Incivility” in Kenski, K. and Jamieson, K. (eds.). *The Oxford Handbook of Political* Communication. New York: Oxford University Press.

Japec, L., Kreuter, F., Berg, M., Biemer, P., Decker, P., Lampe, C., Lane, J., O’Neil, C. and Usher, A. (2015). “Big data in survey research: AAPOR task force report”. *Public Opinion Quarterly*, 79: 839–880.

Jennings, W. and Stoker, G. (2018). “The Divergent Dynamics of Cities and Towns: Geographical Polarisation and Brexit”. *The Political Quarterly*, 90(2): 155-166.

Jigsaw Research (2018). “News Consumption in the UK: 2018”. *Report produced on behalf of OFCOM,* available at: https://www.ofcom.org.uk/__data/assets/pdf_file/0024/116529/news-consumption-2018.pdf

Kalogeropoulos, A., Negredo, S., Picone, I. and Neilsen, R. (2017). “Who Shares and Comments on News?: A Cross-National Comparative Analysis of Online and Social Media Participation”. *Social Media and Society* (Oct-Dec): 1-12.

Kelly Garrett, R. (2009). “Echo chambers online?: Politically motivated selective exposure among Internet news users”. *Journal of Computer-Mediated Communication*, 14: 265-285.

Kenski, K. and Stroud, N. (2010). “Connections Between Internet Use and Political Efficacy, Knowledge, and Participation”. *Journal of Broadcasting and Electronic Media*, 50:2: 173-192.

Khamis, S., Ang, L. and Welling, R. (2017). “Self-branding, ‘micro-celebrity’ and the rise of Social Media Influencers”. *Celebrity Studies*, 8(2): 191-208.

Khang, H., Han, E. and Ki, E. (2014). “Exploring influential social cognitive determinants of social media use”. *Computers in Human Behaviour*, 36: 48-55.

Khatua, A. and Khatua, A. (2016). “Leave or Remain? Deciphering *Brexit* Deliberations on Twitter”. *IEEE Computer Society*, 16: 428-433.

Kim, Y. (2015). “Does Disagreement Mitigate Polarization? How Selective Exposure and Disagreement Affect Political Polarization”. *Journalism and Mass Communication Quarterly*, 92(4): 915-937.

Kim, Y., Chen, H. and Gil de Zuniga, H. (2013). “Stumbling upon news on the Internet: Effects of incidental news exposure and relative entertainment use on political engagement”. *Computers in Human Behaviour*, 29: 2607-2614.

Kosmidis, S. and Theocharis, Y. (forthcoming). “Can Social Media Incivility Induce Enthusiasm? Evidence from Survey Experiments”. *Political Opinion Quarterly, forthcoming 2020.*

Kruse, L., Norris, D. and Flinchum, J. (2018). “Social Media as a Public Sphere? Politics on Social Media”. *The Sociological Quarterly*, 59(1): 62-84.

Kwak, H., Lee, C., Park, H. and Moon, S. (2010). “What is Twitter, a Social Network or a News Media?”. *World Wide Web*, April: 591-600.

Kwon, K. and Cho, D. (2017). “Swearing Effects on Citizen-to-Citizen Commenting Online: A Large-Scale Exploration of Political Versus Non-political Online News Sites”. *Social Science Computer Review*, 35(1): 84-102.

Lakhiwal, A. and Kar, A.K. (2016). “Insights from Twitter Analytics: Modelling Social Media Personality Dimensions and Impact of Breakthrough Events”. *International Federation for Information Processing*, 9844: 533-544.

LaMarre, H. and Suzuki-Lambrect, Y. (2013). “Tweeting democracy? Examining Twitter as an online public relations strategy for congressional campaigns”. *Public Relations Review*, 39: 360-368.

Lassen, D. and Brown, A. (2011). ‘‘Twitter: The Electoral Connection?’’. *Social Science Computer Review*, 29: 419-36.

Lawrence, E., Sides, J. and Farrell, H. (2010). “Self-Segregation or Deliberation? Blog Readership, Participation, and Polarization in American Politics”. *Perspectives on Politics*, 8(1): 141-157.

Lee, F., Chan, M., Chen, H., Neilsen, R. and Fletcher, R. (2019). “Consumptive News Feed Curation on Social Media as Proactive Personalization: A Study of Six East Asian Markets”. *Journalism Studies*, 20(15): 2277-2292.

Lee, J., Choi, J., Kim, C. and Kim, Y. (2014). “Social Media, Network Heterogeneity, and Opinion Polarization”. *Journal of Communication*, 64: 702-722.

Lee, E. (2007). “Deindividuation Effects on Group Polarization in Computer-Mediated Communication: The Role of Group Identification, Public-Self-Awareness, and Perceived Argument Quality”. *Journal of Communication*, 57: 385–403.

Leetaru, K., Wang, S., Cao, G., Padmanabhan, A. and Shook, E. (2013). “Mapping the global Twitter heartbeat: The geography of Twitter”. *First Monday, 18*(5).

Levmore, S. (2012). “The Internet’s Anonymity Problem” in Nussbaum, M. and Levmore, S. (eds.). *The Offensive Internet: Speech, Privacy, and Reputation*. Cambridge: Harvard University Press.

Llewellyn, C. (2016). “The results are in and the UK will #Brexit: What did social media tell us about the UK’s EU referendum?”. *EU Referendum Analysis 2016.* Available at: https://www.referendumanalysis.eu/eu-referendum-analysis-2016/section-7-social-media/the-results-are-in-and-the-uk-will-brexit-what-did-social-media-tell-us-about-the-uks-eu-referendum/

Loader, B. and Mercea, D. (2011). “Networking Democracy?: Social media innovations and participatory politics”. *Information, Communication & Society*, 14(6): 757-769.

Maher, P., Igou, E., Wijnand, A. and van Tilburg, P. (2018). “Brexit, Trump, and the Polarizing Effect of Disillusionment”. *Social Psychological and Personality Science*, 9(2): 205-213.

Majó-Vázquez, S. and González-Bailón, S. (2019). “Digital News and the Consumption of Political Information” in Graham, M. and Dutton, W. (eds.). *Society and the Internet: How Networks of Information and Communication are Changing Our Lives*. Oxford: Oxford University Press.

Mance, H. (2016). “Britain has had enough of experts, says Gove”. *Financial Times*. Available at https://www.ft.com/content/3be49734-29cb-11e6-83e4-abc22d5d108c

Maor, M. (1997). *Political Parties and Party Systems: Comparative Approaches and the British Experience*. London: Routledge.

Marwick, A. and Boyd, D. (2010). “I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience”. *New Media and Society*, 13(1): 114-133.

Massaro, T. and Stryker, R. (2012). “Freedom of speech, liberal democracy, and emerging evidence on civility and effective democratic engagement”. *Arizona Law Review,* 54: 375-441.

McCormick, T., Lee, H., Cesare, N., Shojaie, A. and Spiro, E. (2017). “Using Twitter for Demographic and Social Science Research: Tools for Data Collection and Processing”. *Sociological Methods and Research*, 46(3): 390-421.

McDonald, M. (2003). “On the Overreport Bias of the National Election Study Turnout Rate”. *Political Analysis*, 11: 180-186.

McGee, L. (2019). “Brexit fatigue is making Britain a weird place to live in”. *CNN*. Available at: https://edition.cnn.com/2019/05/04/uk/brexit-fatigue-is-making-britain-weird-analysis-intl-gbr/index.html

McNealy, J. (2011). “The Realm of the Expected: Redefining the Public and Private Spheres in Social Media” in Noor, A. (eds). *Social Media: Usage and Impact*. Lanham: Lexington Books.

McNeish, D. (2014). “Analyzing clustered data with OLS regression: The effect of a hierarchical data structure”. *Multiple Linear Regression Viewpoints*, 40: 11-16.

Mellon, J. and Prosser, C. (2017). “Twitter and Facebook are not representative of the general population: Political attitudes and demographics of British social media users”. *Research and Politics*, July-September: 1-9.

Mellon, J., Evans, G., Fieldhouse, E., Green, J. and Prosser, C. (2018). “Brexit or Corbyn? Campaign and Inter-Election Vote Switching in the 2017 UK General Election”. *Parliamentary Affairs*, 71: 719-737.

Menon, A. and Salter, J. (2016). “Brexit: initial reflections”. *International Affairs*, 92(6): 1297-1318.

Mirriam-Webster. (2019). “Definition of social media”. Available at: https://www.merriam-webster.com/dictionary/social%20media

Morrell, M. (2003). “Survey and experimental evidence for a reliable and valid measure of internal political efficacy”. *Public Opinion Quarterly, 67:* 589–602.

Moscovici, S. and Zavalloni, M. (1969). “The Group as a Polariser of Attitudes”. *Journal of Personality and Social Psychology*, 12(2): 125-135.

Moulton, B. (1986). “Random group effects and the precision of regression estimates”. *Journal of Econometrics*, 32: 385-397.

Muddiman, A. and Stroud, N. (2017). “News Values, Cognitive Biases, and Partisan Incivility in Comment Sections”. *Journal of Communication,* 67(4): 586–609.

Mullen, A. (2016). “Leave versus Remain: the digital battle”. *EU Referendum Analysis 2016*. Available at: http://www.referendumanalysis.eu/eu-referendum-analysis-2016/section-7-social-media/leave-versus-remain-the-digital-battle/.

Mullinix, K., Leeper, T., Druckman, J. and Freese, J. (2015). “The Generalisability of Survey Experiments”. *Journal of Experimental Political Science*, 2: 109-138.

Munger, K. (2017). “Tweetment Effects on the Tweeted: Experimentally Reducing Racist Harassment”. *Political Behaviour*, 39(3): 629-649.

Murray, I., Plagnol, A. and Corr, P. (2017). “When things go wrong and people are afraid’: An evaluation of group polarisation in the UK post Brexit”. *SSRN:* 1-37.

Mutz, D. (2015). *In-Your-Face Politics: The Consequences of Uncivil Media*. Princeton: Princeton University Press.

Napoli, P. (2019). *Social Media and the Public Interest: Media Regulation in the Disinformation Age*. New York: Columbia University Press.

Niklewitz, K. (2017). “Weeding out Fake News: An Approach to Social Media Regulation”. *A Report for the Wilfried Martens Centre for European Studies*: 1-67.

Nyhan, B. and Reifler, J. (2010). “When corrections fail: The persistence of political misperceptions”. *Political Behavior,* 32(2): 303–330.

Papacharissi, Z. (2004). “Democracy online: civility, politeness, and the democratic potential of online political discussion groups”. *New Media and Society*, 6(2): 259-283.

Park, C. (2013). “Does Twitter motivate involvement in politics? Tweeting, opinion leadership, and political engagement”. *Computers in Human Behaviour*, 29: 1641-1648.

Parmelee, J. and Bichard, S. (2011). *Politics and the Twitter Revolution*: *How Tweets Influence the Relationship between Political Leaders and the Public*. Lanham: Lexington Books.

Pawson, R. (2006). *Evidence-based Policy*. London: SAGE.

Pingree, R. (2011). “Effects of Unresolved Factual Disputes in the News on Epistemic Political Efficacy”. *Journal of Communication*, 61(1): 22–47.

Plattner, M. (2010). “Populism, Pluralism and Liberal Democracy.” *Journal of Democracy*, 21(1): 81-92.

Polder-Verkiel, S. (2012). “Online Responsibility: Bad Samaritanism and the Influence of Internet Mediation”. *Science and Engineering Ethics*, 18: 117-141.

Polonski, V. (2016a). “Impact of social media on the outcome of the EU referendum”. *EU Referendum Analysis 2016.* Available at: https://www.referendumanalysis.eu/eu-referendum-analysis-2016/section-7-social-media/impact-of-social-media-on-the-outcome-of-the-eu-referendum/

Polonski, V. (2016b). “Social Media Voices in the UK’s EU Referendum”. *The Oxford Internet Institute*. Available at: https://www.oii.ox.ac.uk/blog/social-media-voices-in-the-uks-eu-referendum/

Prior, M. (2007). *Post-Broadcast Democracy: How Media Choice Increases Inequality in Political Involvement and Polarizes Elections.* Cambridge: Cambridge University Press.

Prochazka, F., Weber, P. and Schweiger, W. (2018). Effects of civility and reasoning in user comments on perceived journalistic quality. *Journalism Studies,* 19(1): 62-78.

Putnam, R. (1995). “Bowling Alone: America's Declining Social Capital”. *Journal of Democracy*, 6(1): 65-78.

Rauchfleisch, A. and Metag, J. (2016). "The special case of Switzerland: Swiss politicians on Twitter". *New Media and Society*, 18(10): 2413-2431.

Robison, J., Stevenson, R., Druckman, J., Jackman, S., Katz, J. and Vavreck, L. (2018). "An Audit of Political Behaviour Research". *SAGE Open – Research Paper.* July-September: 1-14.

Rubin, D. (1974). "Estimating Causal Effects of Treatments in Randomised and Nonrandomised Studies". *Journal of Educational Psychology*, 66(5): 688-701.

Ruiz-Soler, J. (2017). "Twitter research for social scientists: a brief introduction to the benefits, limitations and tools for analysing Twitter data". *Revista Digitos*, 3: 17-31.

Sabbagh, D. (2019). Election result signifies realignment of UK politics". *The Guardian*. Available at: https://www.theguardian.com/politics/2019/dec/13/election-result-signifies-realignment-of-uk-politics

Scaife, L. (2013). "The interrelationship of platform providers and users in the regulation of Twitter and offensive speech - is there a right to be offensive and offended at content?". *Communications Law*, 18(4): 128-134.

Scott, C., Bay-Cheng, L., Prince, M., Nochajski, T. and Collins, R. (2017). "Time spent online: Latent profile analyses of emerging adults' social media use". *Computers in Human Behaviour*, 75: 311-319.

Scott, S. and Orlikowski, W. (2014). "Entanglements in Practice: Performing Anonymity Through Social Media". *MIS Quarterly,* 38(3): 873-893.

Searles, K., Spencer, S. and Duru, A. (2018). "Don't read the comments: the effects of abusive comments on perceptions of women authors' credibility". *Information, Communication & Society* 0(0): 1–16.

Shah, D., Capella, J. and Russell-Neuman, W. (2015). "Big data, digital media, and computational social science possibilities and perils". *The ANNALS of the American Academy of Political and Social Science*, 659: 6-13.

Shariff, S., Zhang, X. and Sanderson, M. (2017). "On the credibility perception of news on Twitter: Readers, topics and features". *Computers in Human Behaviour*, 75: 785-796.

Sloan, L., Morgan, J., Burnap, P. and Williams, M. (2015). "Who Tweets? Deriving the Demographic Characteristics of Age, Occupation and Social Class from Twitter User Meta-Data". *PLOS One*, March: 1-20.

Snowden, C. (2016). "I'm right, you're wrong, and here's a link to prove it: how social media shapes public debate". *The Conversation*. Available at: https://theconversation.com/im-right-youre-wrong-and-heres-a-link-to-prove-it-how-social-media-shapes-public-debate-65723

Sobieraj, S. and Berry, J. (2011). From Incivility to Outrage: Political Discourse in Blogs, Talk Radio, and Cable News. *Political Communication* 28(1): 19-41.

Spears, R. and Smith, H. (2001). "Experiments as Politics". *Political Psychology*, 22(2): 309-330.

Statista (2019). "Leading countries based on number of Twitter users as of October 2019". Available at: https://www.statista.com/statistics/242606/number-of-active-twitter-users-in-selected-countries/

Sunstein, C. (2018). *#Republic: Divided Democracy in the Age of Social Media.* Princeton: Princeton University Press.

Surridge, P. (2019). "The left-right divide" in Menon, A. (eds). *Brexit and Public Opinion 2019*. London: KCL.

Sveningsson, M. (2015). "'It's Only a Pastime, Really: Young People's Experiences of Social Media as a Source of News about Public Affairs'". *Social Media and Society*, July-December: 1-11.

Sydnor, E. (2018). Platforms for incivility: Examining perceptions across different media formats. *Political Communication*, 35 (1): 97–116.

Taber, C. and Lodge, M. (2006). "Motivated Skepticism in the Evaluation of Political Beliefs". *American Journal of Political Science*, 50, 3: 755–769.

Teruelle, R. (2011). "Social Media and Youth Activism" in Noor, A. (eds). *Social Media: Usage and Impact*. Lanham: Lexington Books.

Theocharis, Y., Barbera, P., Fazekas, Z., Popa, S. and Parnet, O. (2016). "A Bad Workman Blames His Tweets: The Consequences of Citizens' Uncivil Twitter Use When Interacting with Party Candidates". *Journal of Communication 66*(6): 1007–1031.

Thorson, K., Vraga, E. and Ekdale, B. (2010). "Credibility in Context: How Uncivil Online Commentary Affects News Credibility". *Mass Communication and Society*, 13(3): 289-313.

Tumasjan, A., Sprenger, T., Sandner, P. and Welpe, I. (2010). "Predicting elections with Twitter: What 140 characters reveal about political sentiment". *International Conference on Weblogs and Social Media*, 10: 178–185.

Turner, G. (2018). "The media and democracy in the digital era: is this what we had in mind?". *Media International Australia*, 168(1): 3-14.

Usherwood, S. (2016). "Talking past each other: the Twitter campaigns". *EU Referendum Analysis 2016*. Available at: https://www.referendumanalysis.eu/eu-referendum-analysis-2016/section-7-social-media/talking-past-each-other-the-twitter-campaigns/

Vaccari, C., Smets, K. and Heath, O. (2020). "The United Kingdom 2017 election: polarisation in a split issue space". *West European Politics*, 43(3): 587-609.

Valentino, N., Hutchings, V., Banks, A. and Davis, A. (2008). "Is a Worried Citizen a Good Citizen? Emotions, Political Information Seeking and Learning via the Internet". *Political Psychology*, 29(2):247–273.

Valeriani, A. and Vaccari, C. (2016). "Accidental exposure to politics on social media as online participation equalizer in Germany, Italy, and the United Kingdom". *New Media and Society*, 18(9): 1857-1874.

Van Aelst, P., Stromback, J., Aalberg, T., Esser, F., de Vreese, C., Matthes, J., Hopmann, D., Salgado, S., Hube, N., Stepinska, A., Papathanassopoulos, S., Berganza, R., Legnante, G., Reinemann, C., Sheafer, T. and Stanyer, J. (2017). "Political communication in a high-choice media environment: a challenge for democracy?". *Annals of the International Communication Association* 41(1): 3–27.

Vargo, C. and Hopp, T. (2017). “Socioeconomic Status, Social Capital, and Partisan Polarity as Predictors of Political Incivility on Twitter: A Congressional District-Level Analysis”. *Social Science Computer Review*, 35(1): 10-32.

Vromen, A., Xenos, M. and Loader, B. (2015). “Young people, social media and connective action: from organisational maintenance to everyday political talk”. *Journal of Youth Studies*, 18(1): 80-100.

Ward, S. and McLoughlin, L. (2017). “Turds, Traitors and Tossers: the abuse of UK MPs via Twitter”. *USIR*: 1-18.

Waters, R., Garrahan, M. and Bradshaw, T. (2016). “Harsh Truths About Fake News for Facebook, Google, and Twitter”. *Financial Times,* 21 November. Available at: https://www.ft.com/content/2910a7a0-afd7-11e6-a37c-f4a01f1b0fa1.

Webb, P. (2000). *The Modern British Party System*. London: Sage Publications.

Webb, P. (2004). “Party Responses to the Changing Electoral Market in Britain” in Mair, P. et al. (eds.). *Political Parties and Electoral Change: Party Responses to Electoral Markets*. London: Sage Publications.

Weeks, B., Gil de Zuniga, H. and Ardevol-Abreu, A. (2017). “Effects of the News-Finds-Me Perception in Communication: Social Media Use Implications for News Seeking and Learning About Politics”. *Journal of Computer-Mediated Communication*, 22: 105-123.

Williams, R. (2016). “Understanding and interpreting generalised ordered logit models”. *The Journal of Mathematical Sociology*, 40(1): 7-20.

Wojcieszak, M. and Mutz, D. (2009). “Online Groups and Political Discourse: Do Online Discussion Spaces Facilitate Exposure to Political Disagreement?”. *Journal of Communication*, 59: 40-56.

Wood, M. (2019). “The political dilemma of expertise – More than just public trust in experts”. *A blog for the London School of Economics*. Available at: https://blogs.lse.ac.uk/impactofsocialsciences/2019/06/17/the-political-dilemma-of-expertise-more-than-just-public-trust-in-experts/

Wu, T. (2017). “How Donald Trump Wins by Losing”. *The New York Times*, 03/03/17. Available at: https://www.nytimes.com/2017/03/03/opinion/sunday/how-donald-trump-wins-by-losing.html

Yardi, S. and Boyd, D. (2010). “Dynamic Debates: An Analysis of Group Polarization Over Time on Twitter”. *Bulletin of Science, Technology and Society*, 30(5): 316-327.

YouGov. (2016). “How Britain voted at the EU referendum”. Available at: https://yougov.co.uk/topics/politics/articles-reports/2016/06/27/how-britain-voted

Young, A. (2017). “The Constitutional Implications of Brexit”. *European Public Law*, 23(4): 757–786.

Young, M. (2009). “Issue Publics in the New Information Environment: Selectivity, Domain Specificity and Extremity”. *Communication Research,* 36(2): 254-284.

Zaller, J. (1992). *The Nature and Origins of Mass Opinion*. Cambridge: Cambridge University Press.

Zhu, Q. (2017). “Citizen-Driven International Networks and Globalization of Social Movements on Twitter”. *Social Science Computer Review*, 35(1): 68-83.

Zubiaga, A. and Ji, H. (2014). “Tweet, but verify: epistemic study of information verification on Twitter”. *Social Network Analysis*, 4: 163-175.

Methodological Appendix 1: Survey Content and Study Configuration

Survey Content

A full version of the survey, with randomised treatment exactly as presented to participants, is available at: https://oxforddpir.eu.qualtrics.com/jfe/form/SV_7342bQ4yreF6gC1

Treatments

Expert Tweet (Brexit)

T1: Civil debate (Brexit)

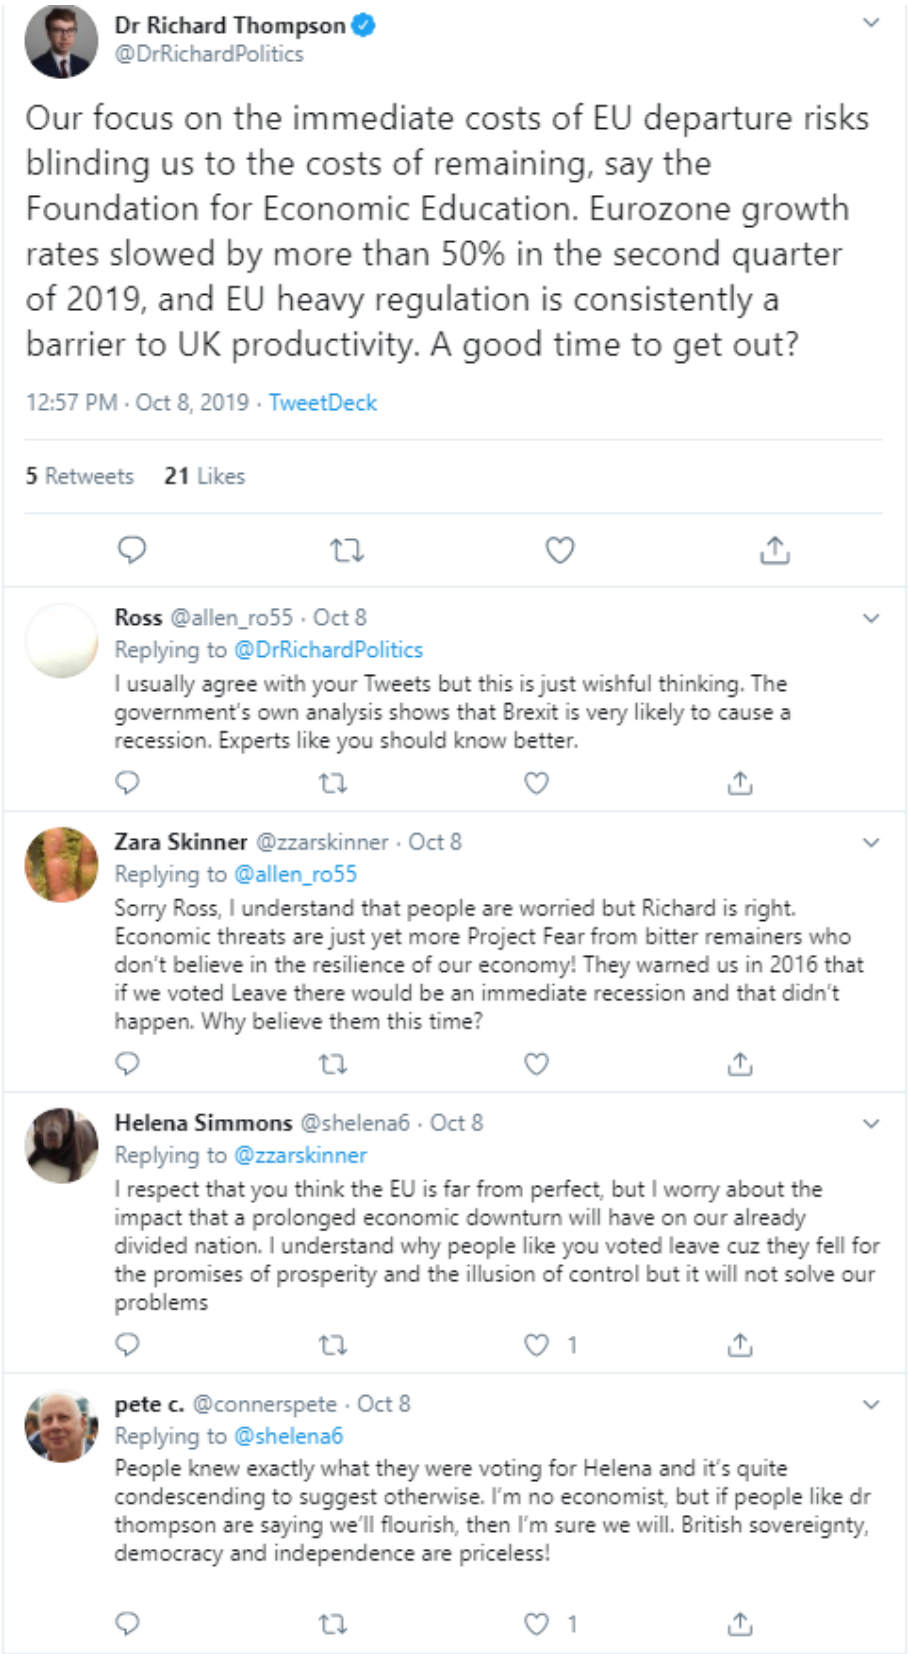

Dr Richard Thompson ✓
@DrRichardPolitics

Our focus on the immediate costs of EU departure risks blinding us to the costs of remaining, say the Foundation for Economic Education. Eurozone growth rates slowed by more than 50% in the second quarter of 2019, and EU heavy regulation is consistently a barrier to UK productivity. A good time to get out?

12:57 PM · Oct 8, 2019 · TweetDeck

5 Retweets **21** Likes

Ross @allen_ro55 · Oct 8
Replying to @DrRichardPolitics
I usually agree with your Tweets but this is just wishful thinking. The government's own analysis shows that Brexit is very likely to cause a recession. Experts like you should know better.

Zara Skinner @zzarskinner · Oct 8
Replying to @allen_ro55
Sorry Ross, I understand that people are worried but Richard is right. Economic threats are just yet more Project Fear from bitter remainers who don't believe in the resilience of our economy! They warned us in 2016 that if we voted Leave there would be an immediate recession and that didn't happen. Why believe them this time?

Helena Simmons @shelena6 · Oct 8
Replying to @zzarskinner
I respect that you think the EU is far from perfect, but I worry about the impact that a prolonged economic downturn will have on our already divided nation. I understand why people like you voted leave cuz they fell for the promises of prosperity and the illusion of control but it will not solve our problems

1

pete c. @connerspete · Oct 8
Replying to @shelena6
People knew exactly what they were voting for Helena and it's quite condescending to suggest otherwise. I'm no economist, but if people like dr thompson are saying we'll flourish, then I'm sure we will. British sovereignty, democracy and independence are priceless!

1

T2: Uncivil debate (Brexit)

Dr Richard Thompson ✓
@DrRichardPolitics

Our focus on the immediate costs of EU departure risks blinding us to the costs of remaining, say the Foundation for Economic Education. Eurozone growth rates slowed by more than 50% in the second quarter of 2019, and EU heavy regulation is consistently a barrier to UK productivity. A good time to get out?

12:57 PM · Oct 8, 2019 · TweetDeck

5 Retweets **21** Likes

Ross @allen_ro55 · Oct 8
Replying to @DrRichardPolitics
This is absolute bollocks, even this corrupt government admit that Brexit will cause a recession! Lying 'experts' like you are overpaid, out of touch and know fuck all about real life

Zara Skinner @zzarskinner · Oct 8
Replying to @allen_ro55
This is just yet more Project Fear from snowflakes like ross who don't believe in the resilience of our economy! The remoaner establishment warned us in 2016 that if we voted out of the EU dictatorship there would be an immediate recession and that didn't happen. Why believe the lying bastards now?

Helena Simmons @shelena6 · Oct 8
Replying to @zzarskinner
Dictatorship? Hahaha the only fascists here are leavers like you that think an economic catastrophe is a price worth paying to get rid of hard working immigrants. Airheads that voted Leave were tricked into it on the promise of prosperity and the illusion of control but it is not the solution! The racist brexiteers should never be forgiven

1

pete c. @connerspete · Oct 8
Replying to @shelena6
Patriotic leavers knew exactly what they were voting for, you condescending bitch. Dr Thompson is a true patriot. Clueless economists and politicians will be wrong again, and you can't put a price on British sovereignty, democracy and independence ! We need to TAKE BACK CONTROL from liberal traitors like you!

1

Expert Tweet (Poverty)

T1: Civil debate (Poverty)

Dr Richard Thompson
@DrRichardPolitics

Food bank use in UK reaches highest rate on record as benefits fail to cover basic costs, new research shows. The Tory welfare 'strategy' of the last decade lacks any human compassion, and the inevitable rise in poverty-related deaths will be blood on their hands. Government must do more, and ministers should be held responsible.

12:57 PM · Oct 8, 2019 · TweetDeck

5 Retweets **21** Likes

Ross @allen_ro55 · Oct 8
Replying to @DrRichardPolitics
Poverty is hard for everyone, but unfortunately the state cannot solve this problem. No amount of government subsidies can substitute for a lack of individual determination and self-reliance.

Zara Skinner @zzarskinner · Oct 8
Replying to @allen_ro55
While I respect your view that people should take responsibility for themselves, there are millions who have a job and are doing all they can to survive but still don't have enough. You cannot ignore the facts from experts, this country should look after it's poorest and the government have abandoned them!

Helena Simmons @shelena6 · Oct 8
Replying to @zzarskinner
Well said Zara. Poverty is not caused by laziness, the laziest people in this country are the bankers who earn millions at boozy lunches while people with jobs and young children are starving. By helping the poor, we are helping everyone who might hit hard times
1

pete c. @connerspete · Oct 8
Replying to @shelena6
I have to disagree on this helena. Anyone who works hard enough can make something of themselves, and benefit payments just increase dependency on the government and prevent these people from realising their full potential
1

T2: Uncivil debate (Poverty)

Dr Richard Thompson ✓
@DrRichardPolitics

Food bank use in UK reaches highest rate on record as benefits fail to cover basic costs, new research shows. The Tory welfare 'strategy' of the last decade lacks any human compassion, and the inevitable rise in poverty-related deaths will be blood on their hands. Government must do more, and ministers should be held responsible.

12:57 PM · Oct 8, 2019 · TweetDeck

5 Retweets **21** Likes

Ross @allen_ro55 · Oct 8
Replying to @DrRichardPolitics
Poverty is not the government's problem. No amount of free handouts can solve laziness. I don't want my taxes bailing out scroungers who are too lazy to get off their arses and work, just like lying 'experts' like you!

Zara Skinner @zzarskinner · Oct 8
Replying to @allen_ro55
Bullshit. There are millions who have a job and are doing all they can to survive but still don't have enough to get by, you heartless pig. You cannot ignore experts just because they point out the injustice that small-minded attitudes like yours are responsible for

Helena Simmons @shelena6 · Oct 8
Replying to @zzarskinner
Too bloody right zara. Poverty is not caused by laziness, the laziest people in this country are the wanker bankers like ross who earn millions at boozy lunches while people with actual jobs and young children are starving. I don't know how cold hearted tossers like you sleep at night!
1

pete c. @connerspete · Oct 8
Replying to @shelena6
Oh shut up you lefty lunatic! Anyone who works hard enough can make something of themselves, and benefit handouts just increase dependency on the government. Free-riders who choose to do fuck all with their lives are not my problem
1

Study Configuration

Three minor amendments were made to the Qualtrics dataset prior to analysis. These were the removal of 14 participants who did not consent to participation in the survey and did not answer any questions. They were thus discounted and replaced in the sample. The second change was the renaming of the column pertaining to gender identification, as 'Q4_1' as opposed to 'Q4', which was an errant duplication of question numbers that would have been problematic to rename elsewhere. Thirdly, rows 2 and 3, referring in the initial dataset to "{"ImportId":"startDate","timeZone":"Europe/London"}" and the written text of the questions respectively, were both removed from the spreadsheet because they were causing variables to be wrongly scaled in R.

Participants were categorised as "conflict averse" if **both** of the following conditions were met:

- They answered either "Strongly disagree" or "Somewhat disagree" to the statement 'I enjoy challenging the opinions of others'
- They answered either "Strongly disagree" or "Somewhat disagree" to the statement 'It doesn't bother me to be in a situation where others are arguing'

Participants were categorised as "conflict agreeable" if **both** of the following conditions were met:

- They answered either "Strongly agree" or "Somewhat agree" to the statement 'I enjoy challenging the opinions of others'
- They answered either "Strongly agree" or "Somewhat agree" to the statement 'It doesn't bother me to be in a situation where others are arguing'

For ease of interpretation and presentational benefit, categories containing very few respondents were not included in all models and graphs relating to heterogeneous treatment effects. These included: people who consider themselves gender non-binary, those aged 65+, those without a formal qualification, and those who were either not eligible to vote or who did not vote in the 2016 EU Referendum.

Methodological Appendix 2: Results

Dataset Composition Statistics and Randomisation Checks

Note: Tables refer to the complete sample where no treatment group is specified

Table 1: Age Composition

X	Frequency
18-24	411
25-35	646
36-45	393
46-55	234
56-64	88
65+	18

Table 2: Gender Composition

X	Frequency
Female	1119
Male	646
Non-binary	25

Table 3: Highest Educational Qualification

X	Frequency
No formal qualification	12
GCSE or equivalent	174
A-level or equivalent	541
University degree or above	1063

Table 4: Self-Reported Political Interest

X	Frequency
Not at all interested	42
Not interested	94
Variable	457
Interested	657
Very interested	540

Table 5: Self-Reported Political Orientation

X	Frequency
Left-wing	621
Centre-left	561
Prefer not to say/I don't know	253
Centre-right	276
Right-wing	79

Table 6: Self-Reported Partisan Identity

X	Frequency
The Labour Party	789
The Conservative Party	240
Liberal Democrats	217
The Green Party	212
The Scottish National Party	75
The Brexit Party	58
Other	40
Prefer not to say/I don't know	159

Table 7: 2017 General Election Party Support

X	Frequency
The Labour Party	819
The Conservative Party	295
Liberal Democrats	133
The Green Party	94
The Scottish National Party	71
UKIP	31
Other	30
Prefer not to say/I don't know	37
Did not vote	280

Table 8: 2016 EU Referendum Vote Allegiance

X	Frequency
Remain	1099
Leave	368
Was not eligible	183
Did not vote	140

Table 9: Most Frequent News Source

X	Frequency
News Websites	489
Social Media	479
Newspapers	391
TV	228
Friends/Family	203

Table 10: Conflict Disposition

X	Frequency
Conflict averse	337
Neutral	1062
Conflict agreeable	391

Table 11: Sample Size by Treatment Group - Experiment 1 (Brexit)

X	Frequency
Control	316
Civil	288
Uncivil	290

Table 12: Sample Size by Treatment Group - Experiment 2 (Poverty)

X	Frequency
Control	316
Civil	296
Uncivil	284

Note: Treatment group sample size discrepancies result from the uneven numbers of respondents that failed the manipulation check in each.

Table 13: Multinomial Randomisation Check - Covariate p-values by Treatment Group (Brexit and Poverty)

	Covariate p-values								
	(Intercept)	**Age**	**Gender**	**Education**	**Political Interest**	**Political Orientation**	**Conflict Disposition**	**Party ID**	**2016 EU Allegiance**
Brexit T1: Civility	0.122	0.382	0.546	0.437	0.548	0.388	0.757	0.865	0.040
Brexit T2: Incivility	0.874	0.442	0.526	0.996	0.078	0.114	0.856	0.170	0.076
Poverty T1: Civility	0.109	0.443	0.188	0.353	0.891	0.697	0.941	0.892	0.090
Poverty T2: Incivility	0.955	0.797	0.749	0.993	0.545	0.794	0.720	0.705	0.650

Note: The randomisation check holds across a strong majority of covariates in the treatment groups. The only covariate that reaches a significance threshold is 2016 EU Referendum allegiance, which is unbalanced in both T1 (civil treatment) and T2 (uncivil treatment) for the Brexit experiment, and in T1 (civil treatment) for the poverty condition.

Alternative Model Specifications – Ordinal Logistic Regression

Tables 14 to 19 below display full model specifications for each of the dependent variables relating to trust in both experiments, using ordinal logistic regression. The results are strongly supportive of the OLS analysis. The coefficients labelled 'T1: Civility' and 'T2: Incivility' in the top sections of the tables show the likelihood (scaled in terms of logs) of members of that treatment group having a positive perception of trust in the cohort of interest relative to control.

More intuitive to interpret are the predicted probabilities of any given member of a treatment group being in each of the five possible answer categories. The sum of probabilities across the categories for each treatment group = 1. They can thus be interpreted as proportions, and provide insight into the distribution of responses. These probabilities are organised to show the direct comparisons for each treatment group by category.

It was not possible to calculate predicted probabilities in the multivariate models without specifying particular conditions for each covariate (e.g. gender = male). Such heterogeneous treatment effects are displayed more efficiently later in this appendix using other methods.

Table 14: Full Sample Trust in Britons Ordered Logit (Brexit)

Trust in Britons		
	Bivariate Ordered Logit	*Multivariate Ordered Logit*
	(1)	(2)
T1: Civility (Brexit)	0.325**	0.370**
	(0.148)	(0.150)
T2: Incivility (Brexit)	-0.168	-0.147
	(0.149)	(0.150)
Predicted Probabilities: Strongly disagree	T1: 0.168 T2: 0.248	-
Predicted Probabilities: Somewhat disagree	T1: 0.389 T2: 0.425	-
Predicted Probabilities: Neither agree nor disagree	T1: 0.187 T2: 0.153	-
Predicted Probabilities: Somewhat agree	T1: 0.220 T2: 0.152	-
Predicted Probabilities: Strongly agree	T1: 0.036 T2: 0.022	-
Observations	894	894
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: The table shows that civility increased trust, while incivility decreased it. Response ratios are higher in T1 relative to T2 in the two categories of agreement and vice versa for disagreement, reflecting the overall pattern, which is itself congruent with the findings reported in-text using OLS.

Table 15: Full Sample Trust in Britons Ordered Logit (Poverty)

Trust in Britons		
	Bivariate Ordered Logit	*Multivariate Ordered Logit*
	(1)	(2)
T1: Civility (Poverty)	0.472***	0.459***
	(0.149)	(0.151)
T2: Incivility (Poverty)	-0.367**	-0.368**
	(0.149)	(0.152)
Predicted Probabilities: Strongly disagree	T1: 0.112 T2: 0.225	-
Predicted Probabilities: Somewhat disagree	T1: 0.381 T2: 0.466	-
Predicted Probabilities: Neither agree nor disagree	T1: 0.232 T2: 0.167	-
Predicted Probabilities: Somewhat agree	T1: 0.254 T2: 0.132	-
Predicted Probabilities: Strongly agree	T1: 0.021 T2: 0.009	-
Observations	896	896
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: Civility increased significantly increased trust in Britons while the opposite held for incivility. Response ratios are again higher in T1 relative to T2 in the two categories of agreement and vice versa for disagreement, more markedly than findings in the Brexit condition. The directionality and magnitude of both the coefficients and predicted probabilities are further support for the in-text analysis.

Table 16: Full Sample Trust in Experts Ordered Logit (Brexit)

Trust in Experts		
	Bivariate Ordered Logit	*Multivariate Ordered Logit*
	(1)	(2)
T1: Civility (Brexit)	-0.013	0.017
	(0.148)	(0.150)
T2: Incivility (Brexit)	-0.146	-0.189
	(0.150)	(0.151)
Predicted Probabilities: Strongly disagree	T1: 0.067 T2: 0.076	-
Predicted Probabilities: Somewhat disagree	T1: 0.206 T2: 0.225	-
Predicted Probabilities: Neither agree nor disagree	T1: 0.191 T2: 0.197	-
Predicted Probabilities: Somewhat agree	T1: 0.441 T2: 0.419	-
Predicted Probabilities: Strongly agree	T1: 0.095 T2: 0.084	-
Observations	894	894
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 17: Full Sample Trust in Experts Ordered Logit (Poverty)

Trust in Experts		
	Bivariate Ordered Logit	*Multivariate Ordered Logit*
	(1)	(2)
T1: Civility (Poverty)	-0.163	-0.148
	(0.149)	(0.151)
T2: Incivility (Poverty)	-0.222	-0.240
	(0.151)	(0.152)
Predicted Probabilities: Strongly disagree	T1: 0.066 T2: 0.069	-
Predicted Probabilities: Somewhat disagree	T1: 0.206 T2: 0.214	-
Predicted Probabilities: Neither agree nor disagree	T1: 0.201 T2: 0.204	-
Predicted Probabilities: Somewhat agree	T1: 0.454 T2: 0.443	-
Predicted Probabilities: Strongly agree	T1: 0.073 T2: 0.069	-
Observations	896	896
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 18: Full Sample Trust in Politicians Ordered Logit (Brexit)

Trust in Politicians		
	Bivariate Ordered Logit	*Multivariate Ordered Logit*
	(1)	(2)
T1: Civility (Brexit)	0.276*	0.298*
	(0.150)	(0.152)
T2: Incivility (Brexit)	-0.060	-0.089
	(0.150)	(0.152)
Predicted Probabilities: Strongly disagree	T1: 0.331 T2: 0.409	-
Predicted Probabilities: Somewhat disagree	T1: 0.395 T2: 0.378	-
Predicted Probabilities: Neither agree nor disagree	T1: 0.183 T2: 0.146	-
Predicted Probabilities: Somewhat agree	T1: 0.080 T2: 0.059	-
Predicted Probabilities: Strongly agree	T1: 0.011 T2: 0.008	-
Observations	894	894
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Points of interest: The coefficients here show the propensity for civility to reduce mistrust of politicians, although overall responses are still indicative of low levels of trust, as shown by the probabilities. This provides strong support for the relationships found using OLS.

Table 19: Full Sample Trust in Politicians Ordered Logit (Poverty)

Trust in Politicians		
	Bivariate Ordered Logit	*Multivariate Ordered Logit*
	(1)	(2)
T1: Civility (Poverty)	-0.096	-0.112
	(0.148)	(0.149)
T2: Incivility (Poverty)	-0.157	-0.161
	(0.151)	(0.153)
Predicted Probabilities: Strongly disagree	T1: 0.361 T2: 0.376	-
Predicted Probabilities: Somewhat disagree	T1: 0.383 T2: 0.380	-
Predicted Probabilities: Neither agree nor disagree	T1: 0.182 T2: 0.175	-
Predicted Probabilities: Somewhat agree	T1: 0.068 T2: 0.064	-
Predicted Probabilities: Strongly agree	T1: 0.005 T2: 0.005	-
Observations	896	896
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Tables and Figures Referenced In-Text

Dependent Variable 1: Trust in Britons

Figure 3: Trust in Britons by Age (Brexit)

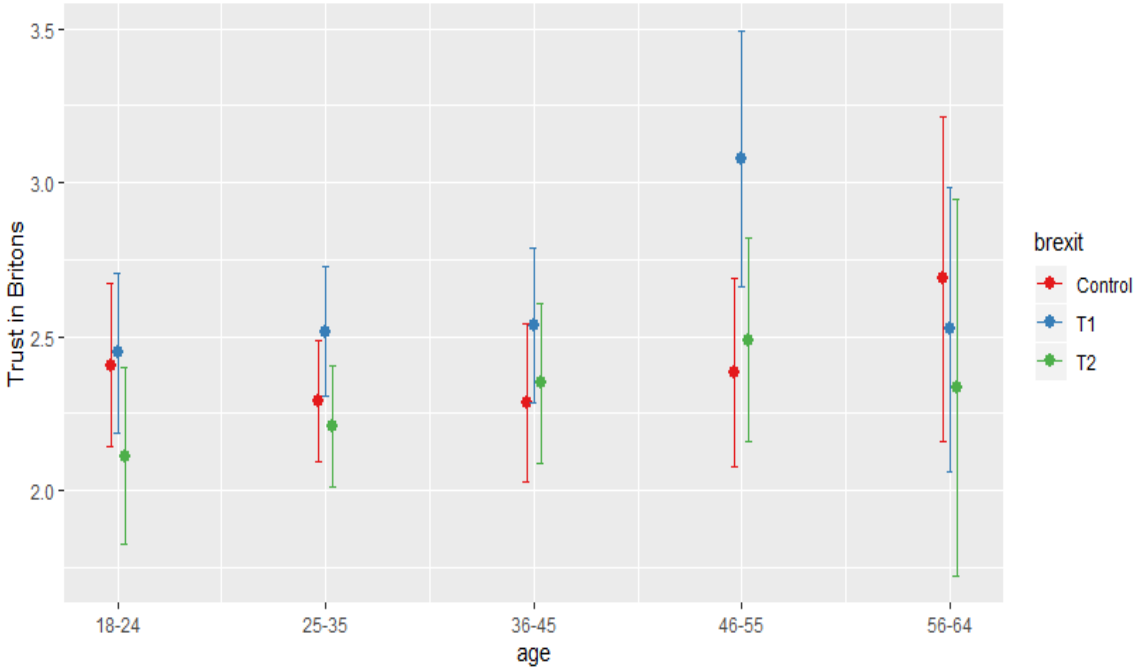

***Interpretation guide for all graphs**: Red = Control, Blue = T1 (Civil treatment), Green = T2 (Uncivil treatment). The word "brexit" (or "poverty" where appropriate) above the legend denotes the topic of the experiment. Centre points represent the average estimate for the group, while the stems show 95% confidence intervals. Method: Bivariate OLS.*

Note: The relationships shown in Table 20 hold across age groups, with treatment effects strongest for the youngest.

Figure 4: Trust in Britons by Age (Poverty)

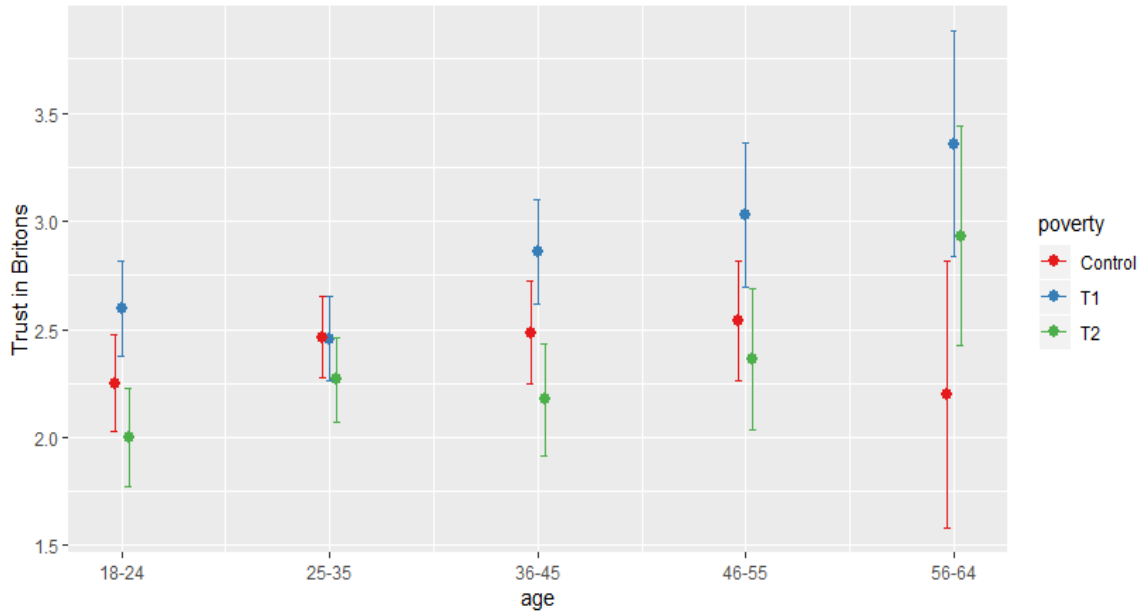

Note: The statistically significant difference between the averages in T1 (civility) and T2 (incivility) for 18-24 year olds = 0.59.

Figure 5: Trust in Britons by Highest Educational Qualification (Brexit)

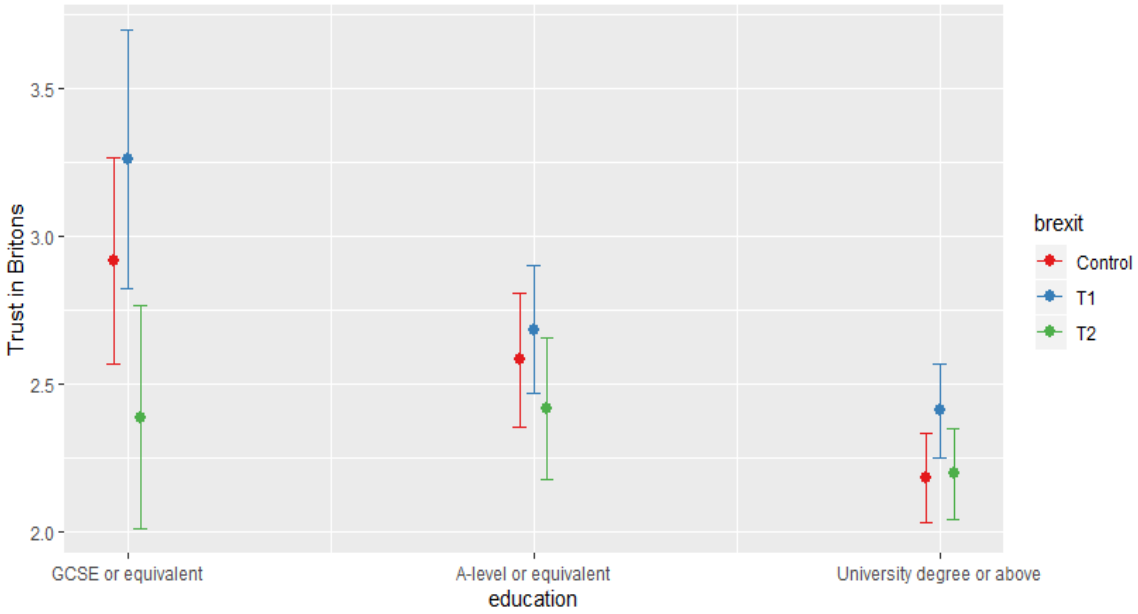

Note: The significant difference between the averages in T1 (civility) and T2 (incivility) for the least educationally qualified = 0.87.

Table 21: Trust in Britons by 2016 EU Allegiance (Brexit)

Trust in Britons		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (Control Group: Leave)	3.333***	3.516***
	(0.135)	(0.529)
T1: Civility (Leave)	0.029	-0.021
	(0.180)	(0.187)
T2: Incivility (Leave)	-0.368*	-0.388**
	(0.188)	(0.193)
T1: Civility (Remain treatment effect differential)	0.162	0.230
	(0.208)	(0.214)
T2: Incivility (Remain treatment effect differential)	0.354*	0.374*
	(0.214)	(0.219)
Control Group: Remain	-1.255***	-1.204***
	(0.152)	(0.159)
Observations	735	710
R^2	0.195	0.229
Adjusted R^2	0.190	0.207
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:	*p<0.1; **p<0.05; ***p<0.01	

Interpretation guide for all tables detailing heterogeneous treatment effects: The constant denotes the control group estimate for the first subgroup listed in the table (in this case Leave voters). Following regular tables, the T1 and T2 coefficients directly below this denote the difference between this figure and the average for this subgroup (Leave) in the respective treatment group.

The coefficient(s) for the control condition in other subgroups (in this case Remain) found in the bottom half of the table shows the difference between the group represented by the constant (Control: Leave) and the average for the named group (Remain) in the control. The T1 and T2 coefficients for the subgroups listed above this in the middle of the table (Remain) should be interpreted as the difference between the treatment effects for the same condition (T1 or T2) in the group represented by the constant (Leave) and the treatment effects for the listed group, relative to the control average for that group.

From here, averages and treatment effects (β) for all subgroups can be calculated. Using the above table as an example, estimates (calculations in parentheses) are as follows (shown as a graph in Figure 6):

Control: Leave (constant): 3.333

T1: Leave: 3.362, β = 0.029

T2: Leave: 2.965, β = -0.368

Control: Remain: 2.078 (3.333 – 1.255)

T1: Remain: 2.269, β = 0.191 (To calculate β: 0.029 + 0.162. To calculate the estimate: 2.078 + 0.191)

T2: Remain: 2.064, β = -0.014 (a. -0.368 + 0.354 b. 2.078 – 0.014)

Note: Following the method set out above, all future tables of this type (both in-text and in Appendix 2) will be labelled with estimates for the control group and treatment effects for T1 and T2 for any subgroup not represented by the constant which is relevant to the analysis, but no further explanation will be provided as that would simply repeat the procedure set out here.

Table 22: Trust in Britons by 2016 EU Allegiance (Poverty)

Trust in Britons		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (Control Group: Leave)	2.897***	2.697***
	(0.130)	(0.354)
T1: Civility (Leave)	0.237	0.231
	(0.174)	(0.176)
T2: Incivility (Leave)	-0.206	-0.185
	(0.187)	(0.190)
T1: Civility (Remain treatment effect differential)	-0.030	-0.042
	(0.203)	(0.206)
T2: Incivility (Remain treatment effect differential)	0.043	0.045
	(0.213)	(0.217)
Control Group: Remain	-0.588***	-0.495***
	(0.148)	(0.153)
Observations	732	709
R^2	0.090	0.125
Adjusted R^2	0.083	0.100
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: Estimates for the Remain subgroup in the table above are as follows (shown as a graph in Figure 7):

Control: 2.309 (2.897 – 0.588)

T1: β = 0.207 (0.237 – 0.030)

T2: β = -0.163 (-0.206 + 0.043)

Figure 8: Ordered Logit of Trust in Britons by Political Orientation (Brexit)

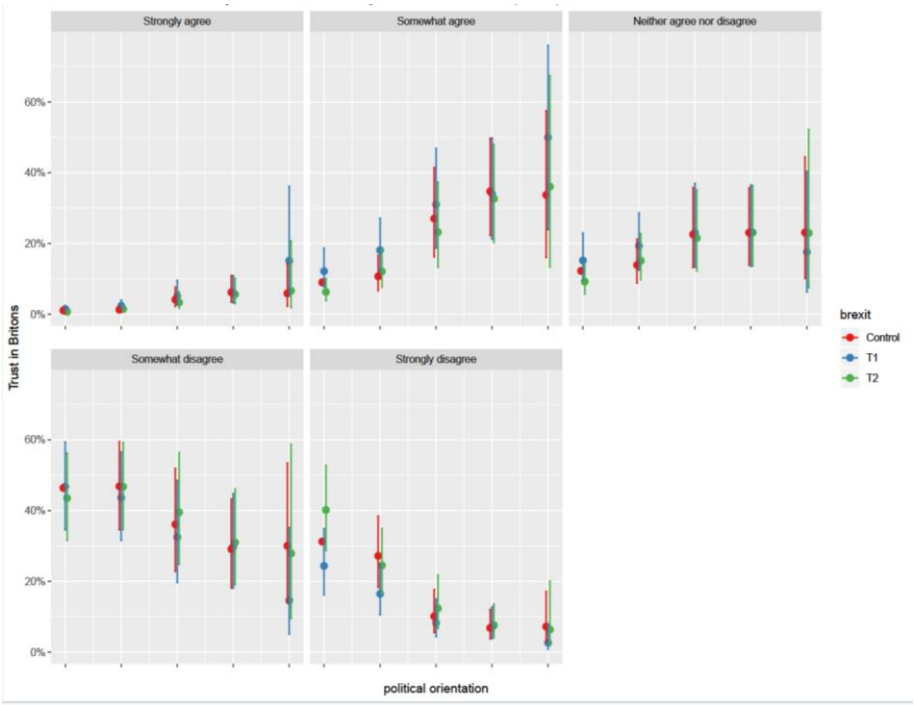

X-axis = Political Orientation (L to R: Left, Centre-left, Do not know, Centre-right, Right)

Graphs (Top L to Bottom R): Strongly agree, Somewhat agree, Neither agree nor disagree, Somewhat disagree, Strongly Disagree)

Note: This shows the distribution of responses in each of the five answer categories, organised by political orientation and treatment group. Across treatment groups, there is a clear pattern of higher disagreement among the left-wing categories, and higher agreement in the right-wing categories.

Figure 9: Trust in Britons by Political Interest (Brexit)

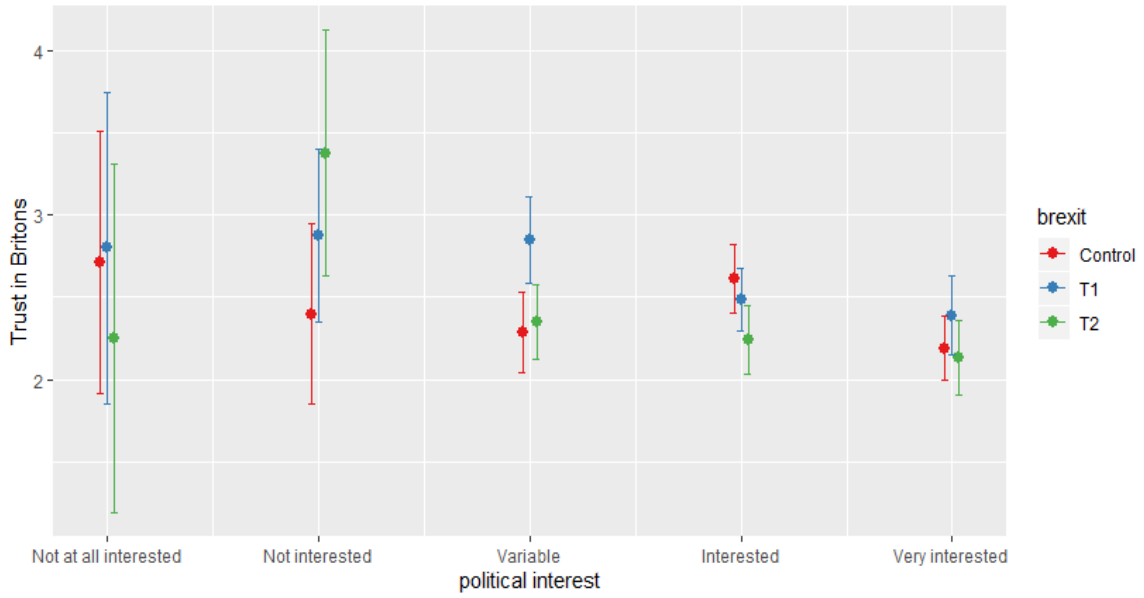

Note: β = -0.366 (SE = 0.201) for Interested participants in the uncivil treatment (T2).

Figure 10: Trust in Britons by Political Interest (Poverty)

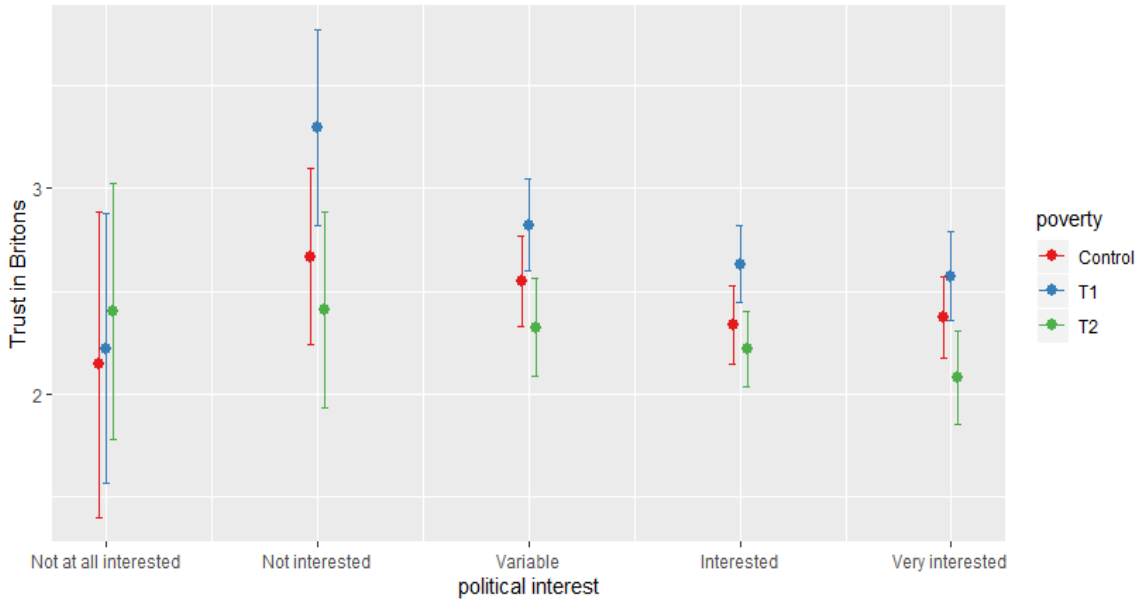

Note: The estimate for Very Interested respondents in T2 (incivility) is 0.493 lower that that found for T1 (civility), and the clear distance between confidence intervals is indicative of statistical significance at the 95% level. This pattern is also found for Interested respondents.

Dependent Variable 2: Trust in Experts

Table 23: Full Sample Trust in Experts (Brexit and Poverty)

Trust in Experts						
	Bivariate OLS (Brexit)	*Bivariate GLM (Brexit)*	*Multivariate OLS (Brexit)*	*Bivariate OLS (Poverty)*	*Bivariate GLM (Poverty)*	*Multivariate OLS (Poverty)*
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	3.297***	3.297***	3.313***	3.358***	3.358***	3.021***
	(0.062)	(0.062)	(0.558)	(0.060)	(0.060)	(0.093)
Brexit T1: Civility	0.001	0.001	0.020			
	(0.090)	(0.090)	(0.088)			
Brexit T2: Incivility	-0.091	-0.091	-0.108			
	(0.090)	(0.090)	(0.088)			
Poverty T1: Civility				-0.101	-0.101	-0.083
				(0.086)	(0.086)	(0.085)
Poverty T2: Incivility				-0.125	-0.125	-0.129
				(0.087)	(0.087)	(0.086)
Observations	894	894	894	896	896	896
R^2	0.001		0.064	0.003		0.049
Adjusted R^2	-0.001		0.049	0.0004		0.037
Multivariate Controls	Age, Education, Gender, Conflict Disposition					
Note:						*p<0.1; **p<0.05; ***p<0.01

Note: In the Brexit condition, civil debate has indistinguishable impact while incivility induces a minor and insignificant negative effect. Both treatments exert a minor negative impact in the poverty-related condition.

Figure 12: Trust in Experts by Highest Educational Qualification (Poverty)

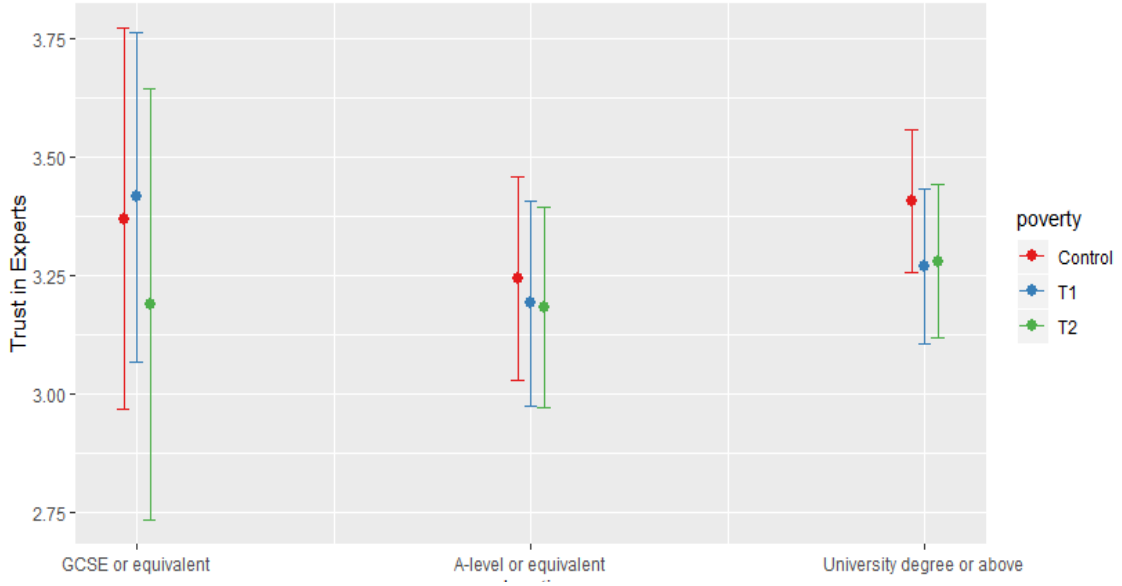

Note: β = -0.18 (SE = 0.31) for the least educated participants in the uncivil treatment (T2).

Figure 13: Trust in Experts by 2016 EU Allegiance (Brexit)

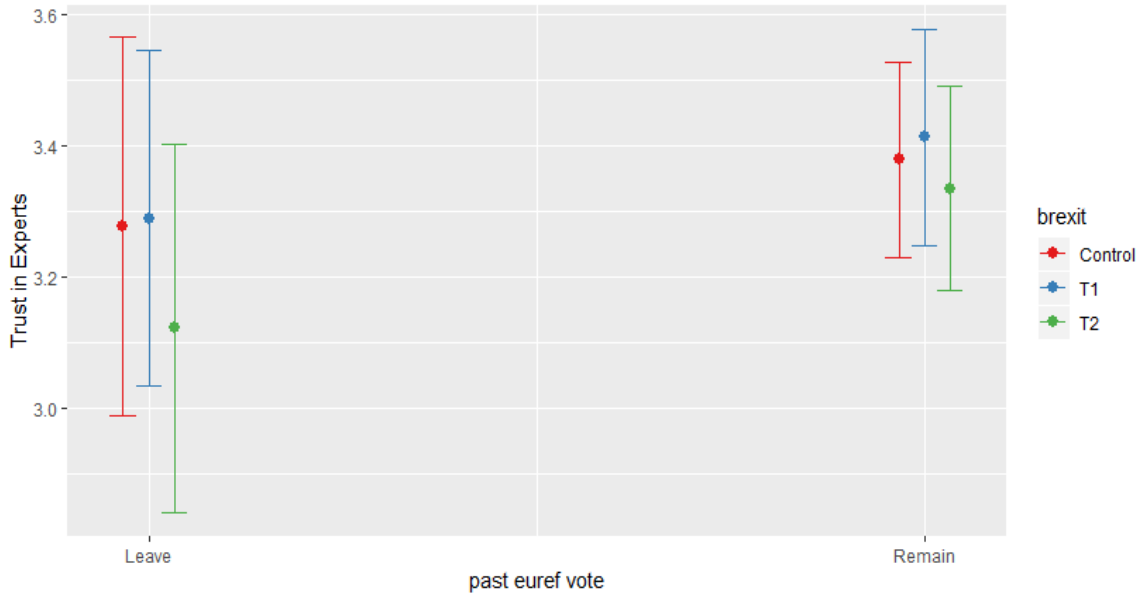

Note: Average estimates are marginally larger across treatment groups among Remain voters in spite of their likely disagreement with the expert.

Figure 14: Trust in Experts by Political Interest (Poverty)

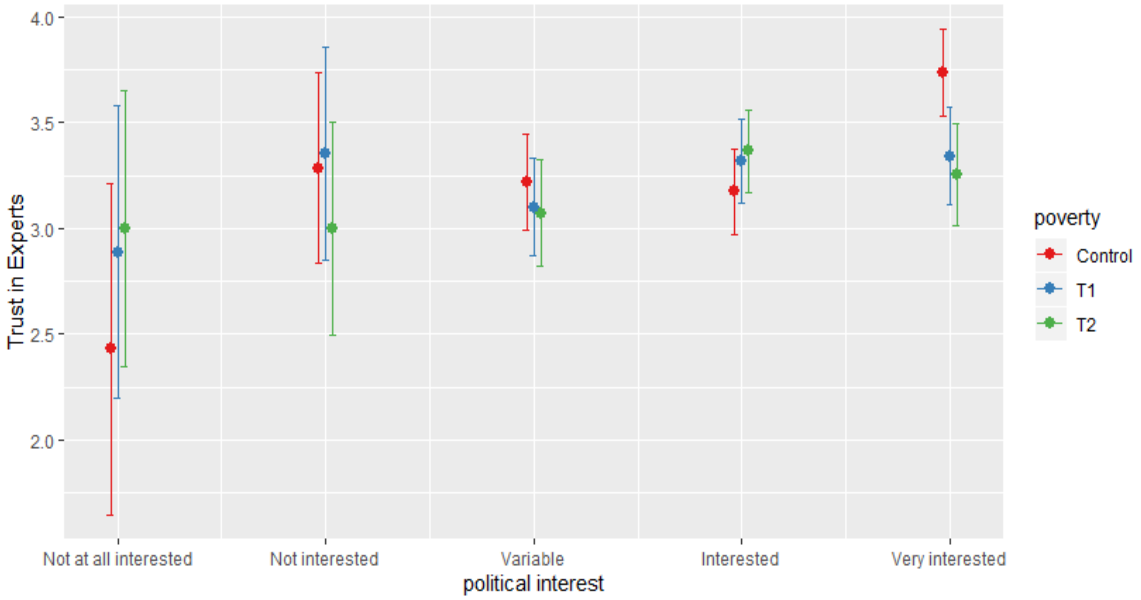

Note: β = -0.479 (significant) for the most interested participants in the uncivil treatment (T2). Full estimates in Table 24.

Figure 15: Trust in Experts by Age (Poverty)

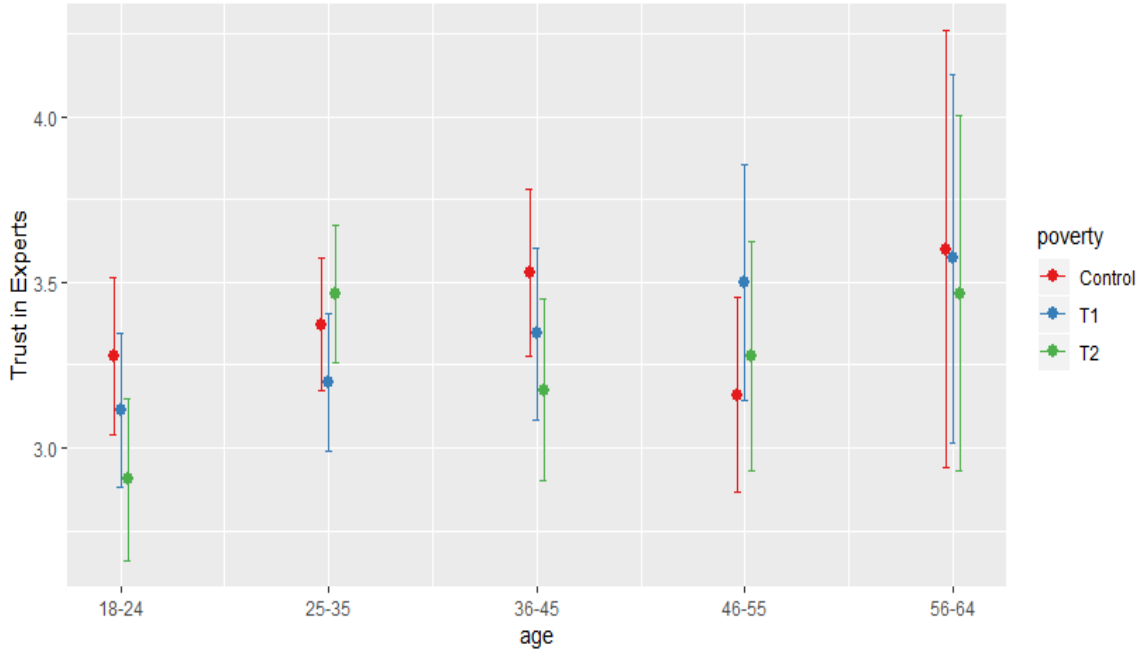

Note: β = -0.372 (SE = 0.174, p = < 0.05) for the 18-24 subgroup in the uncivil treatment (T2).

Figure 16: Trust in Experts by Gender (Poverty)

Note: Women were both less likely to trust experts than men on average across all groups, and more impacted by incivility. T2: β = -0.202, SE = 0.111 (significant at p = < 0.10).

Dependent Variable 3: Expert Tweet Opinion Change

Table 25: Full Sample Expert Tweet Opinion Change (Brexit and Poverty)

	Expert Tweet Opinion Change					
	Bivariate OLS (Brexit)	*Bivariate GLM (Brexit)*	*Multivariate OLS (Brexit)*	*Bivariate OLS (Poverty)*	*Bivariate GLM (Poverty)*	*Multivariate OLS (Poverty)*
	(1)	(2)	(3)	(4)	(5)	(6)
Constant (T1: Civility)	-0.024	-0.024	-0.113	-0.071**	-0.071**	-0.119
	(0.035)	(0.035)	(0.306)	(0.029)	(0.029)	(0.260)
Brexit T2: Incivility	-0.072	-0.072	-0.076			
	(0.049)	(0.049)	(0.049)			
Poverty T2: Incivility				0.018	0.018	0.014
				(0.041)	(0.041)	(0.042)
Observations	578	578	578	580	580	580
R^2	0.004		0.026	0.0003		0.021
Adjusted R^2	0.002		0.004	-0.001		-0.002
Multivariate Controls	Age, Education, Gender, Conflict Disposition					
Note:						*p<0.1; **p<0.05; ***p<0.01

Note: In the Brexit condition, civil debate has indistinguishable impact on opinion change while incivility induces a minor negative effect that is significant from zero. Both treatments exert a minor and negative impact in the poverty-related condition, civility significantly.

Figure 18: Expert Tweet Agreement by Political Interest (Brexit)

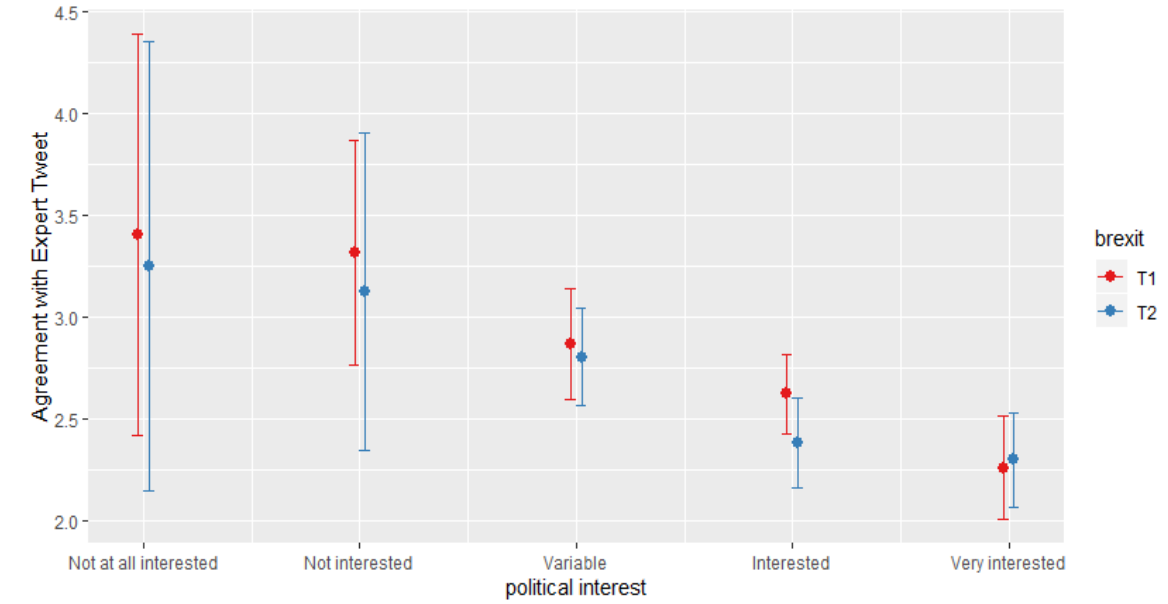

For graphs of this dependent variable only, Red now represents T1 (civility) while Blue shows estimates for T2 (incivility).

Note: Average levels of agreement with the expert opinion decline as levels of political interest increase.

Figure 19: Expert Tweet Opinion Change by Political Interest (Brexit)

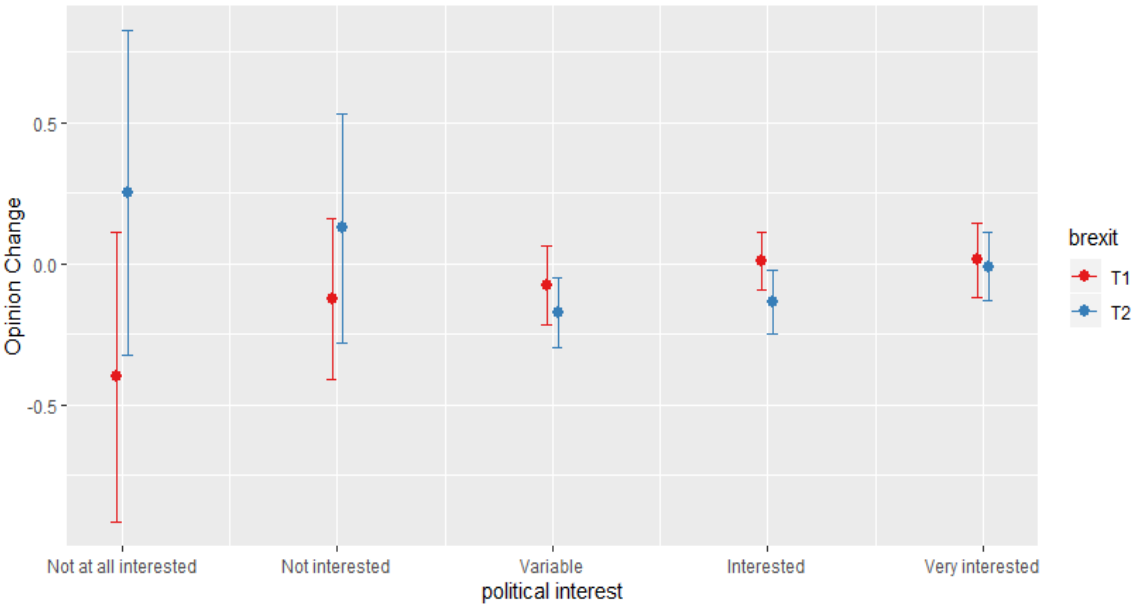

Note: There is a vast effect treatment effect differential between T1 and T2 for the Not at all interested subgroup (0.650), as civility induces opinion change away from agreement with the expert (-0.40), while incivility has the opposite effect (0.25). Full estimates in Table 26.

Figure 20: Expert Tweet Opinion Change by Political Interest (Poverty)

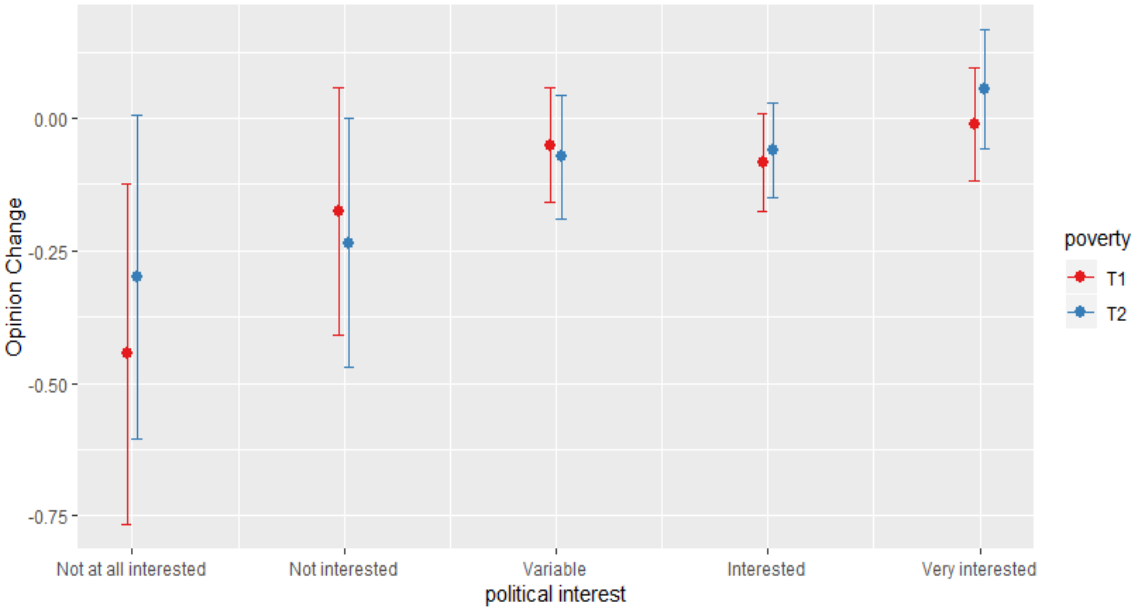

Note: β = -0.444 for the least interested participants in the civil treatment (T1), and β = -0.300 in the uncivil treatment (T2). Full estimates in Table 27.

Table 27: Expert Tweet Opinion Change by Political Interest (Poverty)

Expert Tweet Opinion Change		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (T1: Civility - Not at all Interested)	-0.444***	-0.560***
	(0.164)	(0.194)
T2: Incivility (Not at all Interested)	0.144	0.182
	(0.226)	(0.237)
T1: Civility (Very Interested)	0.432**	0.493***
	(0.173)	(0.187)
T2: Incivility (Very Interested treatment effect differential)	-0.078	-0.121
	(0.239)	(0.251)
Observations	580	565
R^2	0.027	0.050
Adjusted R^2	0.012	0.019
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: Estimates for the 'Very interested' subgroup are as follows:

T1: -0.012 (-0.444 + 0.432)

*T2: 0.054. (**a.** 0.144 – 0.078 = 0.066. **b.** -0.012 + 0.066)*

Figure 23: Expert Tweet Agreement by Party ID (Poverty)

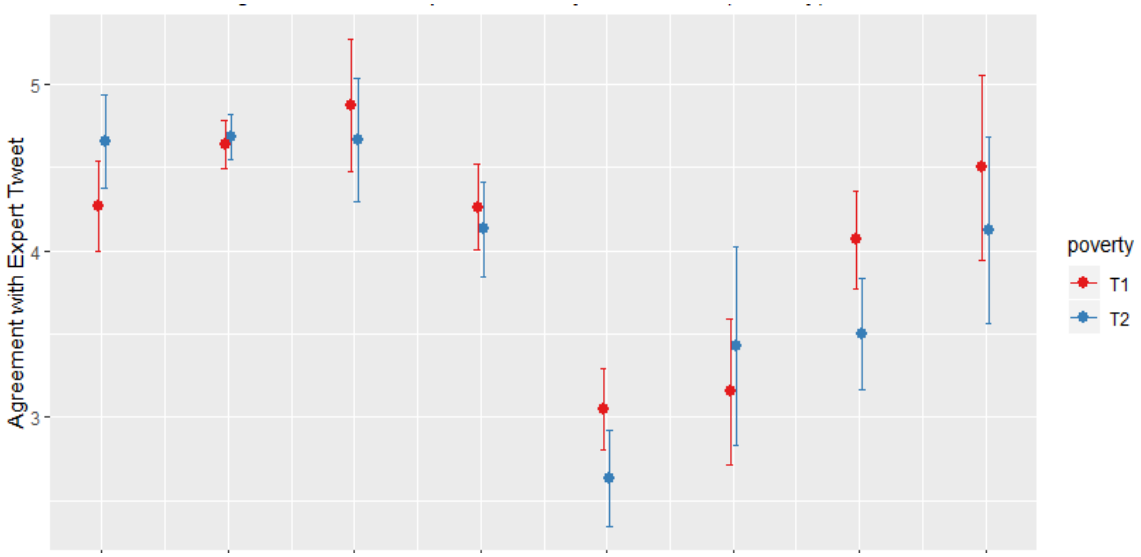

X-axis = Party ID (L to R): Green Party, Labour Party, Scottish National Party, Liberal Democrats, Conservative Party, Brexit Party, Do not know, Other

Note: Levels of agreement with the expert tweet are notably higher among supporters of left-wing parties (Green, Labour, SNP) than right-wing parties (Conservative, Brexit Party).

Table 28: Expert Tweet Opinion Change by Age (Brexit)

Expert Tweet Opinion Change		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (T1: Civility – 18-24)	-0.077	-0.136
	(0.073)	(0.219)
T2: Incivility (18-24)	-0.164	-0.138
	(0.108)	(0.110)
T1: Civility (25-35)	0.018	0.054
	(0.093)	(0.099)
T2: Incivility (25-35 treatment effect differential)	0.197	0.151
	(0.134)	(0.139)
Observations	573	562
R^2	0.019	0.036
Adjusted R^2	0.003	0.004
Multivariate Controls	Education, Gender, Conflict Disposition, Political Interest	
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: Treatment effects of note apply to the 18-24 subgroup (T2: β = -0.241). These findings do not apply to other age groups, as seen by the estimate for the 25-35 category below:

T1: -0.059 (-0.077 + 0.018)

*T2: -0.026. (**a.** -0.164 + 0.197 = 0.033. **b.** -0.059 + 0.033)*

Table 29: Expert Tweet Opinion Change by Conflict Disposition (Poverty)

Expert Tweet Opinion Change		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (T1: Civility - Neutral)	-0.065*	-0.061
	(0.038)	(0.090)
T2: Incivility (Neutral)	0.014	0.007
	(0.053)	(0.054)
T1: Civility (Conflict Averse treatment effect differential)	0.115	0.128
	(0.074)	(0.078)
T2: Incivility (Conflict Averse treatment effect differential)	-0.197*	-0.198*
	(0.111)	(0.114)
T1: Civility (Conflict Agreeable treatment effect differential)	-0.132*	-0.140*
	(0.071)	(0.073)
T2: Incivility (Conflict Agreeable treatment effect differential)	0.183*	0.171
	(0.102)	(0.104)
Observations	580	565
R^2	0.017	0.054
Adjusted R^2	0.009	0.026
Multivariate Controls	Age, Education, Gender, Political Orientation	
Note:	*p<0.1; **p<0.05; ***p<0.01	

Estimates for the Conflict Averse and Conflict Agreeable subgroups in the table above are as follows (shown as a graph in Figure 24):

Conflict Averse:

T1: 0.050 (-0.065 + 0.115)

*T2: -0.133 (**a.** 0.014 -0.197 = -0.183. **b**. 0.050 - 0.183)*

Conflict Agreeable:

T1: -0.197 (-0.065 – 0.132)

*T2: 0.000 (**a**. 0.014 + 0.183 = 0.197. **b.** -0.197 + 0.197)*

Dependent Variable 4: Trust in Politicians

Figure 25: Full Sample Trust in Politicians by Treatment Group (Brexit)

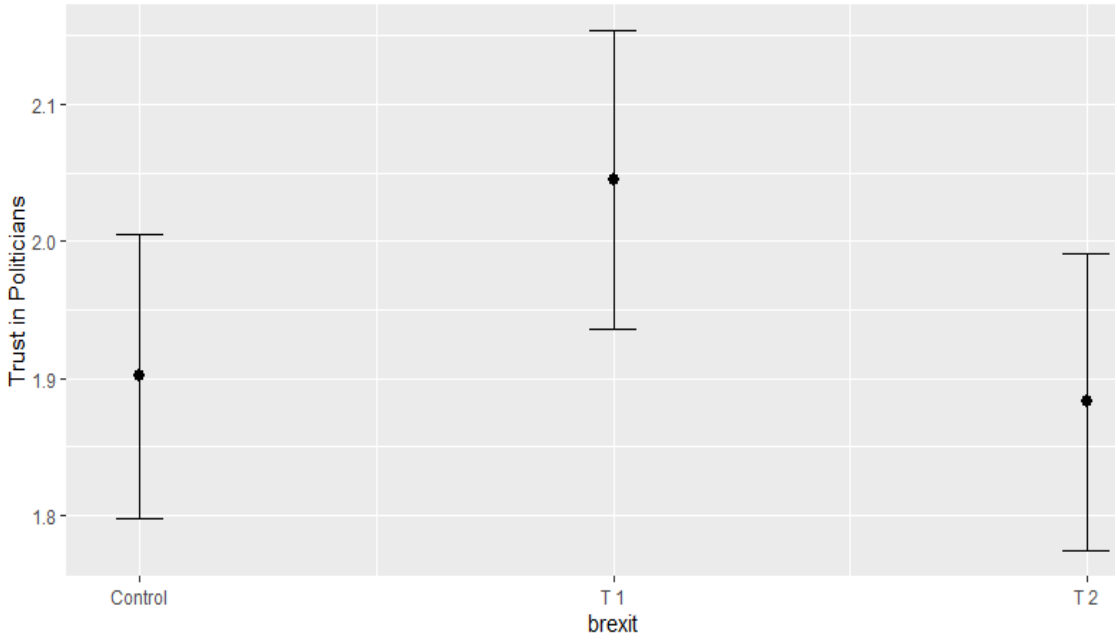

Note: T1 (civility): β = 0.143, SE = 0.077. T2 (incivility) shows a null effect. Full estimates in Table 30.

Figure 26: Trust in Politicians by Political Interest (Brexit)

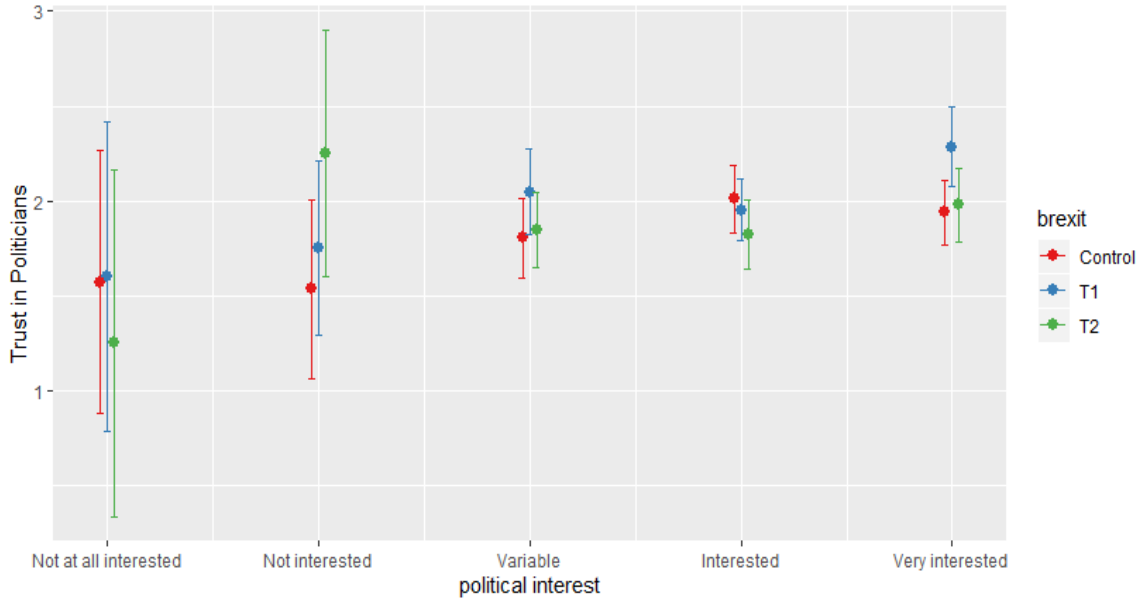

Note: Trust in politicians increased on average as levels of political interest increased, although in a minor and statistically insignificant fashion.

Figure 27: Trust in Politicians by Political Interest (Poverty)

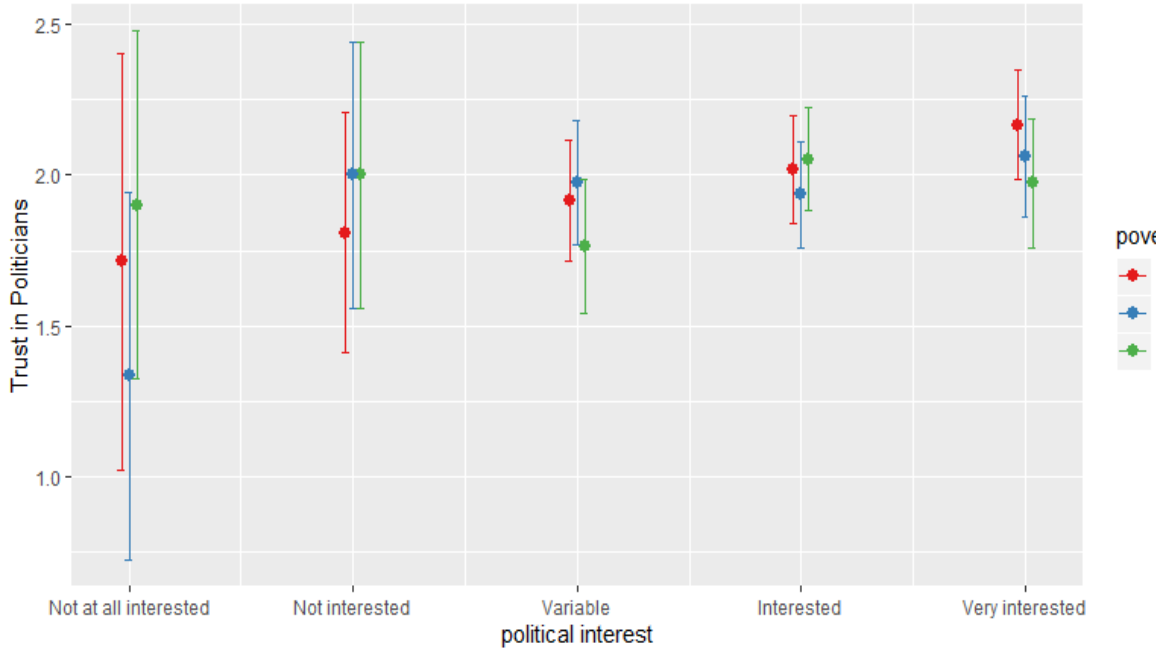

Note: Trust in politicians increased on average in the control condition as levels of political interest increased, although in a minor and statistically insignificant fashion that is not found in the other treatment groups.

Figure 28: Trust in Politicians by Political Orientation (Brexit)

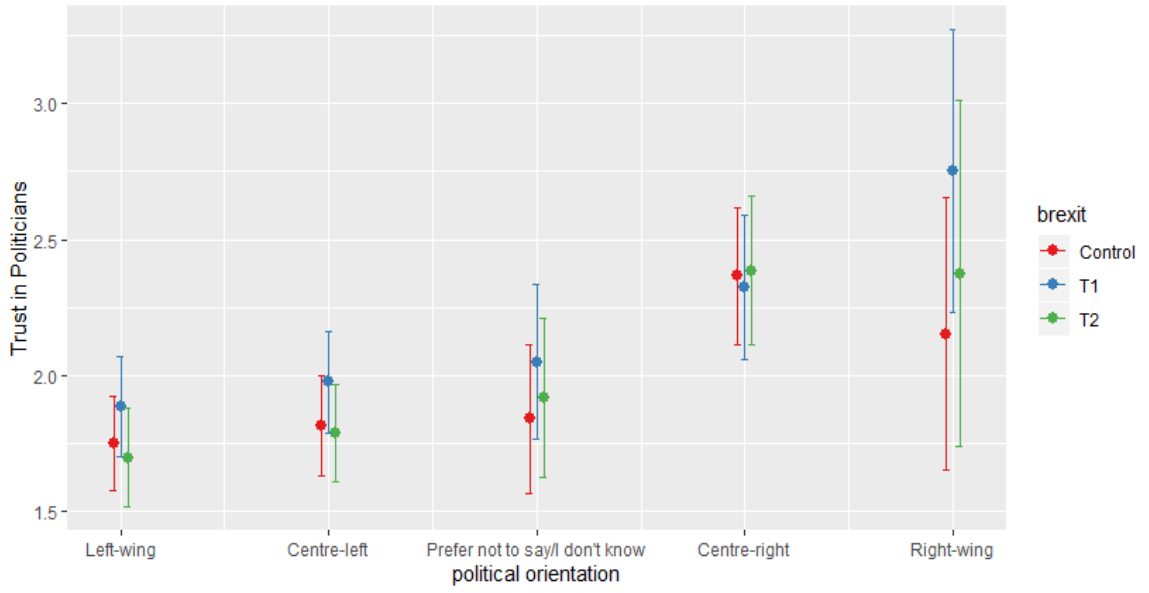

Note: On average, trust in politicians was higher in the two right-wing categories than their equivalents on the left across treatment groups.

Figure 29: Trust in Politicians by 2016 EU Allegiance (Brexit)

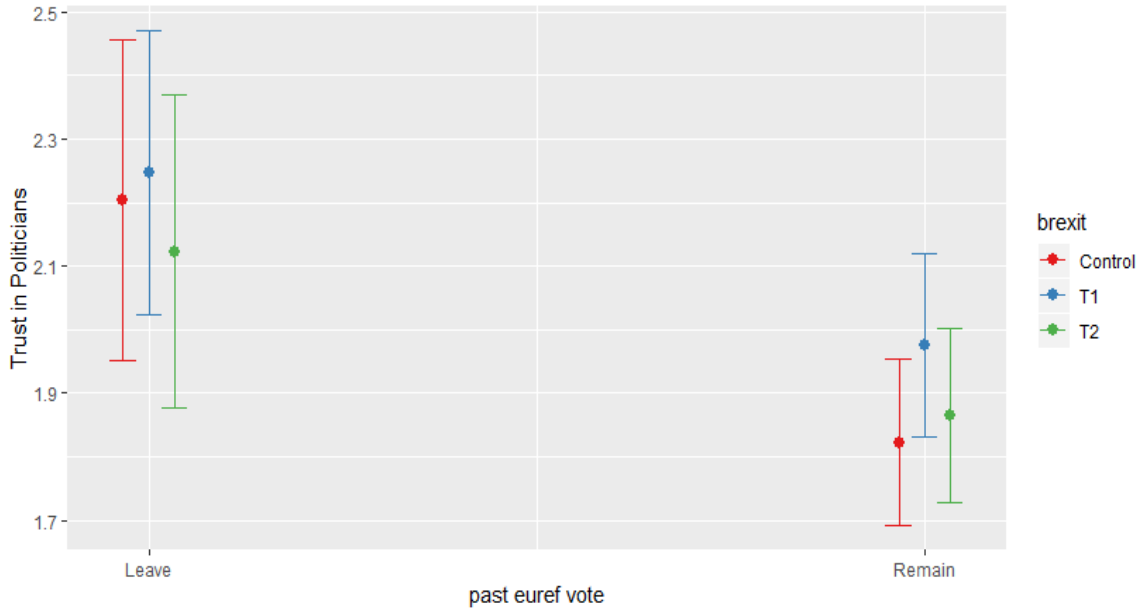

Note: Values are higher in the Leave subgroup across treatment group, significantly in the Control condition (β = 0.384).

Figure 30: Trust in Politicians by 2016 EU Allegiance (Poverty)

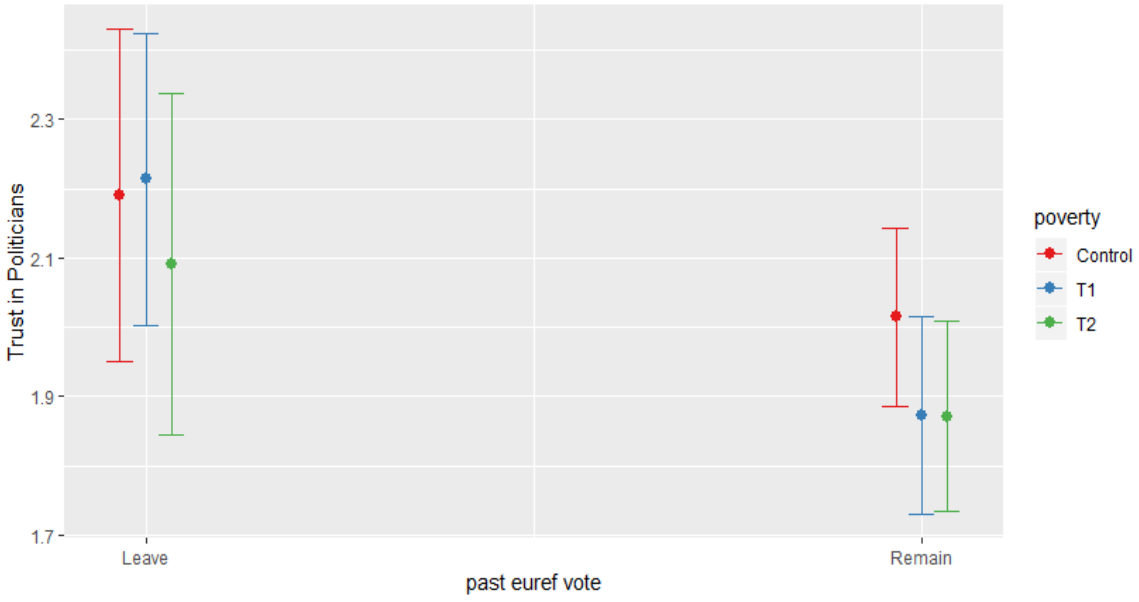

Note: Values are higher in the Leave subgroup across treatment group, significantly in T1 (civil debate).

Table 31: Trust in Politicians by Gender (Brexit)

Trust in Politicians		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (Control Group: Female)	1.794***	1.714***
	(0.066)	(0.142)
T1: Civility (Female)	0.254***	0.251***
	(0.095)	(0.093)
T2: Incivility (Female)	0.032	0.018
	(0.096)	(0.094)
T1: Civility (Male treatment effect differential)	-0.331**	-0.309*
	(0.162)	(0.161)
T2: Incivility (Male treatment effect differential)	-0.181	-0.179
	(0.160)	(0.158)
Control Group: Male	0.346***	0.251**
	(0.112)	(0.114)
Observations	882	868
R^2	0.019	0.085
Adjusted R^2	0.013	0.067
Multivariate Controls	Age, Education, Conflict Disposition, Political Orientation	
Note:	*p<0.1; **p<0.05; ***p<0.01	

Note: Estimates for the Male subgroup are as follows (shown as a graph in Figure 31):

Control: 2.140 (1.794 + 0.346)

T1: β = -0.077 (0.254 – 0.331)

T2: β = -0.149 (0.032 – 0.181)

Figure 32: Trust in Politicians by Highest Educational Qualification (Brexit)

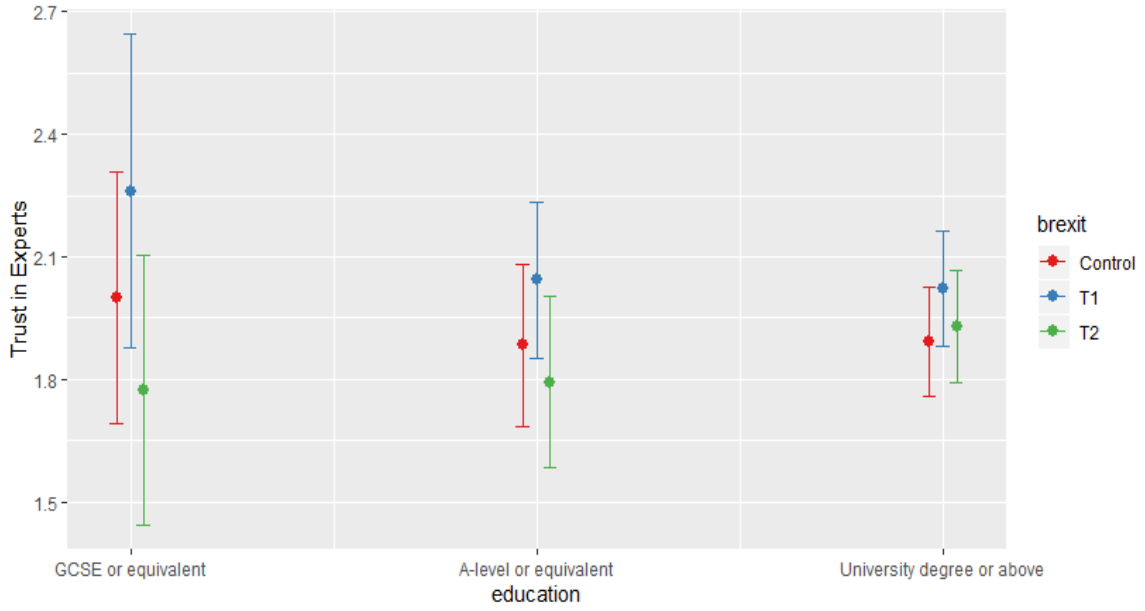

Note: For the lowest educationally qualified subgroup, in civil treatment (T1) β = 0.261, SE = 0.251, while in uncivil treatment (T2) β = -0.226, SE = 0.230. This is an overall difference between T1 and T2 of 0.487.

Figure 33: Trust in Politicians by Age (Brexit)

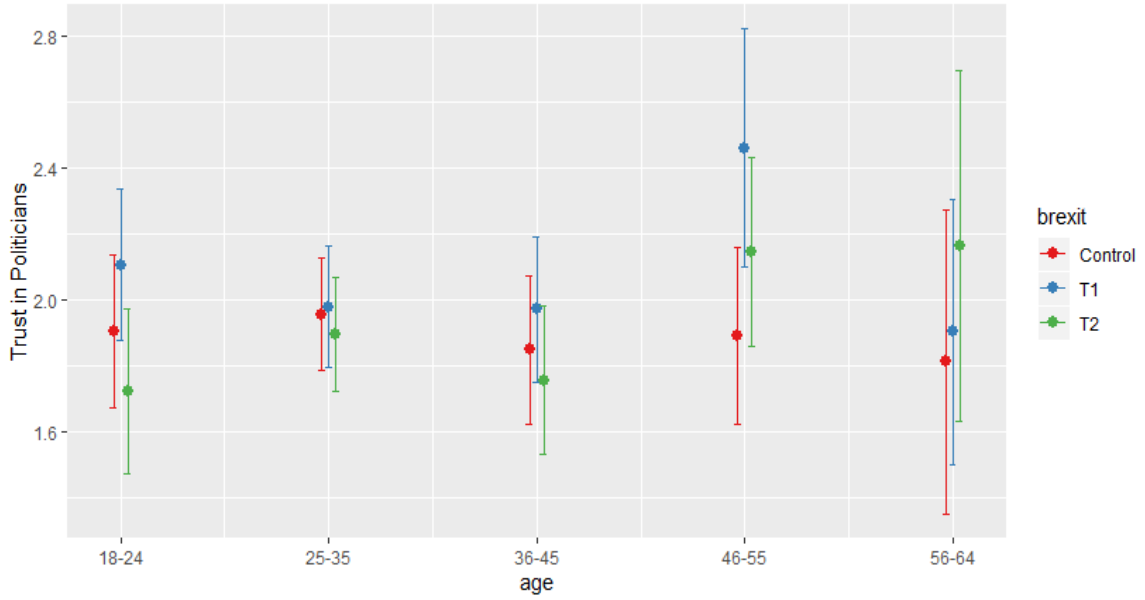

Note: For the 18-24 subgroup, in civil treatment (T1) β = 0.201, SE = 0.165, while in uncivil treatment (T2) β = -0.184, SE = 0.174. This is an overall difference between T1 and T2 of 0.385.

Dependent Variable 5: Voting in the Next General Election

Table 32: Full Sample Likelihood to Vote in Next General Election (Brexit and Poverty)

Likelihood to Vote in Next General Election						
	Bivariate OLS (Brexit)	*Bivariate GLM (Brexit)*	*Multivariate OLS (Brexit)*	*Bivariate OLS (Poverty)*	*Bivariate GLM (Poverty)*	*Multivariate OLS (Poverty)*
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	0.959***	0.959***	0.989***	0.940***	0.940***	0.974***
	(0.012)	(0.012)	(0.020)	(0.013)	(0.013)	(0.022)
Brexit T1: Civility	-0.007	-0.007	-0.007			
	(0.017)	(0.017)	(0.016)			
Brexit T2: Incivility	-0.0002	-0.0002	-0.001			
	(0.017)	(0.017)	(0.016)			
Poverty T1: Civility				0.006	0.006	0.010
				(0.019)	(0.019)	(0.019)
Poverty T2: Incivility				0.0003	0.0003	-0.001
				(0.019)	(0.019)	(0.019)
Observations	894	894	894	896	896	896
R^2	0.0003		0.092	0.0001		0.055
Adjusted R^2	-0.002		0.076	-0.002		0.039
Multivariate Controls	Age, Education, Gender, Conflict Disposition					
Note:						*p<0.1; **p<0.05; ***p<0.01

Note: This question was answered using a binary yes/no choice, and subsequently estimates can be interpreted as proportions. Proportions were very high, there was little variation across experiments, and no treatment effects of note.

Table 33: Likelihood to Vote in Next General Election by Political Interest (Poverty)

Likelihood to Vote in Next General Election		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (Control Group: Not at all Interested)	0.286***	0.305***
	(0.082)	(0.082)
T1: Civility (Not at all Interested)	0.492***	0.477***
	(0.109)	(0.108)
T2: Incivility (Not at all Interested)	0.214**	0.203*
	(0.106)	(0.105)
T1: Civility (Not Interested treatment effect differential)	-0.478***	-0.459***
	(0.130)	(0.129)
T2: Incivility (Not Interested treatment effect differential)	-0.318**	-0.304**
	(0.128)	(0.127)
T1: Civility (Very Interested treatment effect differential)	-0.565***	-0.542***
	(0.113)	(0.113)
T2: Incivility (Very Interested treatment effect differential)	-0.228**	-0.219**
	(0.111)	(0.110)
Control Group: Not Interested	0.524***	0.513***
	(0.094)	(0.093)
Control Group: Very Interested	0.714***	0.704***
	(0.084)	(0.084)
Observations	896	875
R^2	0.161	0.175
Adjusted R^2	0.147	0.155
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Estimates for the Not Interested and Very Interested subgroups are as follows (depicted in Figure 34):

Not Interested:

Control: 0.810 (0.286 + 0.524)

T1: β = 0.014 (0.492 – 0.478)

T2: β = -0.104 (0.214 – 0.318)

Very Interested:

Control: 1.000 (0.286 + 0.714)

T1: β = -0.073 (0.492 – 0.565)

T2: β = -0.014 (0.214 – 0.228)

Table 34: Likelihood to Vote in Next General Election by Political Interest (Brexit)

Likelihood to Vote in Next General Election		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (Control Group: Not at all Interested)	0.429***	0.431***
	(0.070)	(0.071)
T1: Civility (Not at all Interested)	-0.029	-0.048
	(0.109)	(0.108)
T2: Incivility (Not at all Interested)	0.071	0.058
	(0.116)	(0.116)
T1: Civility (Not Interested treatment effect differential)	-0.084	-0.058
	(0.128)	(0.128)
T2: Incivility (Not Interested treatment effect differential)	0.004	0.037
	(0.142)	(0.142)
T1: Civility (Very Interested treatment effect differential)	0.011	0.045
	(0.112)	(0.112)
T2: Incivility (Very Interested treatment effect differential)	-0.063	-0.050
	(0.119)	(0.119)
Control Group: Not Interested	0.371***	0.354***
	(0.085)	(0.085)
Control Group: Very Interested	0.563***	0.569***
	(0.072)	(0.072)
Observations	894	872
R^2	0.187	0.207
Adjusted R^2	0.174	0.187
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Estimates for the Not Interested and Very Interested subgroups are as follows (depicted in Figure 35):

Not Interested:

Control: 0.800 (0.429 + 0.371)

T1: β = -0.113 (-0.029 – 0.084)

T2: β = 0.075 (0.071 + 0.004)

Very Interested:

Control: 0.992 (0.429 + 0.563)

T1: β = -0.018 (-0.029 + 0.011)

T2: β = 0.008 (0.071 – 0.063)

Table 35: Likelihood to Vote in Next General Election by Past Voting Turnout (Poverty)

Likelihood to Vote in Next General Election		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (Control Group: Past GE Vote: No)	0.500***	0.541***
	(0.034)	(0.037)
T1: Civility (Past GE Vote: No)	0.260***	0.232***
	(0.051)	(0.051)
T2: Incivility (Past GE Vote: No)	0.100*	0.080
	(0.051)	(0.052)
T1: Civility (Past GE Vote: Yes treatment effect differential)	-0.289***	-0.259***
	(0.054)	(0.054)
T2: Incivility (Past GE Vote: Yes treatment effect differential)	-0.109**	-0.089
	(0.054)	(0.054)
Control Group: Past GE Vote: Yes	0.496***	0.477***
	(0.036)	(0.037)
Observations	821	801
R^2	0.273	0.280
Adjusted R^2	0.269	0.269
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:	*p<0.1; **p<0.05; ***p<0.01	

Note: Estimates for voters in the 2017 General Election are as follows (depicted in Figure 36):

Control: 0.996 (0.500 + 0.496)

T1: β = -0.029 (0.260 – 0.289)

T2: β = -0.009 (0.100 – 0.109)

Table 36: Likelihood to Vote in Next General Election by Past Voting Turnout (Brexit)

Likelihood to Vote in Next General Election		
	Bivariate OLS	*Multivariate OLS*
	(1)	(2)
Constant (Control Group: Past GE Vote: No)	0.652***	0.593***
	(0.036)	(0.043)
T1: Civility (Past GE Vote: No)	-0.176***	-0.160***
	(0.052)	(0.053)
T2: Incivility (Past GE Vote: No)	-0.061	-0.032
	(0.051)	(0.053)
T1: Civility (Past GE Vote: Yes treatment effect differential)	0.182***	0.165***
	(0.054)	(0.055)
T2: Incivility (Past GE Vote: Yes treatment effect differential)	0.064	0.034
	(0.054)	(0.055)
Control Group: Past GE Vote: Yes	0.333***	0.339***
	(0.037)	(0.039)
Observations	837	813
R^2	0.303	0.320
Adjusted R^2	0.298	0.308
Multivariate Controls	Age, Education, Gender, Conflict Disposition	
Note:		*p<0.1; **p<0.05; ***p<0.01

Note: Estimates for voters in the 2017 General Election are as follows (depicted in Figure 37):

Control: 0.985 (0.652 + 0.333)

T1: β = 0.006 (-0.176 + 0.182)

T2: β = 0.003 (-0.061 + 0.064)

Figure 39: Likelihood to Vote in Next General Election by Political Orientation (Poverty)

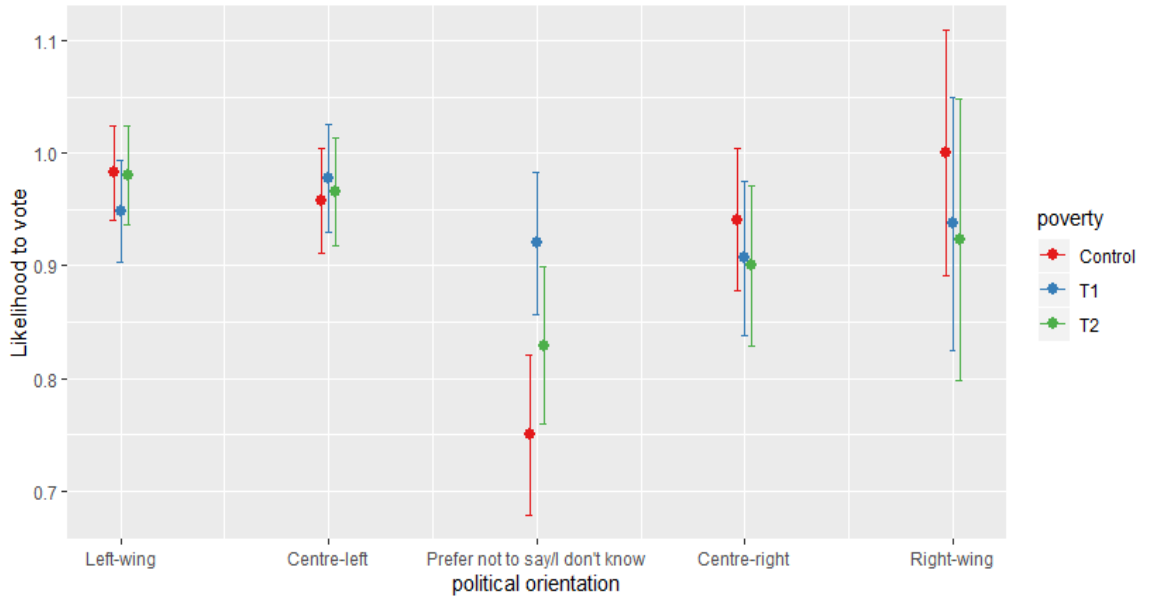

Note: There are clear positive treatment effects of civility (T1) on those unsure of political orientation vis-à-vis both Control and incivility (T2) in the poverty experiment.

Figure 40: Likelihood to Vote in Next General Election by Party ID (Brexit)

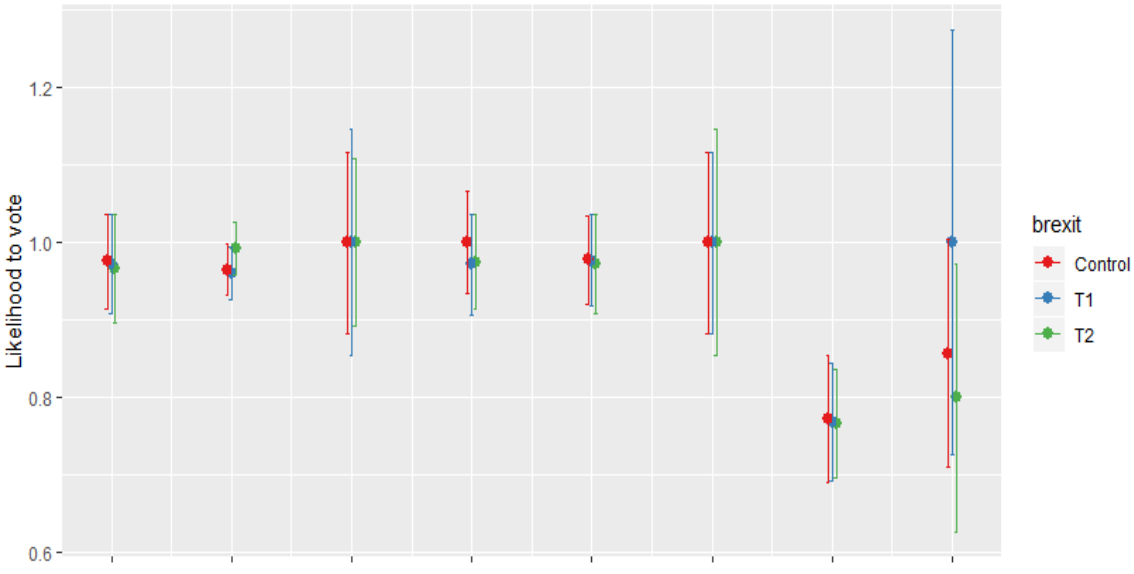

X-axis = Party ID (L to R): *Green Party, Labour Party, Scottish National Party, Liberal Democrats, Conservative Party, Brexit Party, **Do not know**, Other*.

Note: The second category from the right shows that those without a partisan identity are significantly less likely to vote than supporters of various political parties. Unlike the poverty experiment, there is no positive treatment effect after civility (T1).