DEPARTMENT OF GOVERNMENT

UNIVERSITY OF ESSEX

Doctoral Thesis

How motivated reasoning leads to tolerance of false claims:

Three experimental tests of mechanisms

Author:

Christine STEDTNITZ

Supervisory board:

Rob Johns

Patrick LOWN

Dominik DUELL

A thesis submitted for the degree of Doctor of Philosophy

Submitted 30 Sep 2019

Abstract

Christine STEDTNITZ

How motivated reasoning leads to tolerance of false claims: Three experimental tests of mechanisms

Many recent electoral events have been characterised by false claims which, despite abundant fact-checking, were often widely believed. This led to much talk about a 'post-truth' politics. Meanwhile, an extensive literature confirms that false beliefs voters have about political issues are often highly resistant to correction. The drivers of that resistance are well-known: people engage in motivated reasoning. They are prone to accept as fact those claims that confirm their opinions, and dismiss claims that challenge them.

This dissertation investigates reactions to two types of challenging information: false information that affirms individuals' political opinions, and factual information that challenges them. More specifically, it explores how the situation individuals are in when exposed to challenging information affects their susceptibility to engage in motivated reasoning. In a series of experiments, I explore the effect of three types of situations: stress, low group status, and exposure to post-truth comments. In chapter 1, I explore the effect of stress on tolerance of false claims. I describe a survey experiment (n=380) fielded shortly after the 2016 Brexit referendum in which I attempted to induce stress to investigate the effect of stress on belief in false campaign claims among Leave voters. I found no effect.

In chapter 2, I explore the effect of low group status on tolerance of false claims. To test the pure causal mechanism in a non-political environment I designed a laboratory experiment (n=277). The set-up was a pub quiz in which one team was disadvantaged, received more difficult questions and, therefore, a lower payoff than the other team. After the quiz players were shown two sets of feedback – one person called the quiz 'fair play'; the other called it 'unfair' and demanded a top-up for the disadvantaged team. Crucially, one of the two feedback givers made false claims. As hypothesized, the disadvantaged team overlooked false claims coming from the person who sided with their team (but not from the other team). I discuss implications for 21st century political campaigns in which false facts about and in which some groups of voters are structually disadvantaged.

In chapter 3, I shift to reactions to factually *accurate* information that challenges political beliefs. This chapter is joint work with John Bartle and Rob Johns, Professors of Politics in the Government Department at the University of Essex. We conducted a nationally representative survey experiment in the UK (n=2936) in which we corrected false beliefs about immigration. We identified false beliefs on both sides of the immigration debate (pro and anti-immigration) and showed respondents a fact check challenging a false claim they had rated as true. We went beyond classic misperception-correction studies in two key ways: First, we used a more nuanced, seven point true to false scale, which allowed us to trace smaller changes in the perceived accuracy of false claims. Second, we mimicked 'post-truth' surroundings in which expert information is rarely the 'final' word: Some of our respondents saw not only the fact check but also a post-truth comment encouraging them to retain their old beliefs. We found that fact checks significantly reduced belief in false claims. Exposure to post-truth comments cancelled out some but not all of this effect. We discuss what these results mean for 21st century election campaigns and the public broadcasters who accompany them.

Acknowledgements

I wish to thank the faculty and the PhD students in the Government Department at the University of Essex for helpful comments on this dissertation: My supervisors, Rob Johns and Patrick Lown for invaluable feedback, support, and guidance throughout all stages of this dissertation, Dominik Duell, John Bartle, and Adam Peresman, for reading some or all draft papers, providing helpful comments and pointing me to books worth reading, Maja Erikkson Stern for real-world exposure, Masoud Farokhi, Baris Ari, Michele Fenzl, Larissa Kersten, and Somaia El Sayed for discussion groups and many conversations that helped me articulate ideas, see weak points, and improve designs, Andrea Gonzales, Andres Dockendorff, and Andres Ballesteros, and the Latinos for the flair that makes research fun, Nelson Fernandez for bringing them all together, my PhD cohort, and the ones who came before and after, including Aziz Alghashian, Sayan Banerjee, Gwang-Eun Choi, Joe Greenwood, Kana Inata, David Liao, Wakako Maekawa, Katariina Mustasilta, Phillip Nelson, Marina Petrova, Paula Szewach, Chi-Lin Tsai, Linda Urselmans, Farzad Vaziry, and many others for helpful comments on various HEROES and conference presentations, Iulia Popescu, Roxi Dumitrescu and many other GV207 students for pre-testing experiments, Raynee Gutting for helpful comments on the experimental paper, Federica Genovese and

the quantitative social sciences crowd for helpful comments at various Eastern Academic Research Council conferences, Elliot Brearey, Emma McClelland, Nicola Rowley, Jamie Seakens, Kimberley Stockley, Sallyann West, and everyone else in the admin team who work tirelessly behind the scenes and without whom the Department would collapse in a matter of days. I acknowledge that the third paper of this dissertation is co-authored work with with Rob Johns and John Bartle; all co-authors contributed equal parts. This dissertation was fully funded by the Eastern Academic Research Council. I gratefully acknowledge their support. Additional research funding for the laboratory experiments came from an ESSEXLab Seedcorn Grant, which I gratefully acknowledge, too.

Contents

Abstract					
Acknowledgements					
1	The	Effect	fect of Stress on Belief in False Facts in the Brexit Debate		
	1.1	Abstra	1ct	23	
	1.2	Introd	uction	24	
	1.3	Metho	d	29	
	1.4	Result	S	34	
	1.5	Discus	ssion and Conclusion	39	
2	The	e effect of boosting in-group identities on tolerance of false			
	facts	ts			
	2.1	Abstra	nct	43	
	2.2	Introd	oduction		
	2.3	Theoretical Background		46	
		2.3.1	Global Status Threat	49	
		2.3.2	Ethnic status threat	51	
		2.3.3	Status Threat and Vote Choice	52	
		2.3.4	Proposed mechanism	54	
		2.3.5	Method	58	

		2.3.6 Hypotheses	61			
		2.3.7 Results	64			
	2.4	Discussion				
	2.5	Conclusion				
	2.6	Acknowledgements	79			
3	A Post-Truth Public? Investigating the Mechanisms of Resis- tance to Factual Correction					
	3.1	Abstract	81			
	3.2	Introduction				
	3.3	Reactions to uncongenial expert information				
	3.4	Method				
	3.5	Hypotheses				
	3.6	Results	98			
		Effect of the fact-check on perceptions about the				
		fact-checked fact	100			
		3.6.1 Evidence of post-truth reasoning	106			
		Limitations	109			
	3.7	Discussion and Conclusion	111			
A	Cha	pter 1	123			
B	Chapter 2		137			
	B.1	Screenshots of the laboratory experiment	137			
C	C Chapter 3 – Graphs					
Bi	Bibliography 15					

x

Introduction

In the aftermath of the Brexit and the Trump election, two long-standing liberal democracies find themselves deeply divided, with half the electorate puzzled at how the other half voted. As in any election, remain voters had their reasons to vote remain; leave voters had their reasons to vote leave; Donald Trump's voters had their reasons to vote for Donald Trump and Hillary Clinton's voters had their reasons to vote for Hillary Clinton. What made these elections stand apart, however, was the level of factual inaccuracy in many of the campaign claims. On both sides of the Atlantic, campaigners on all sides – in particular the Leave and the Trump campaign – were widely accused of lying.

False information in politics is not new. Even the amount of false information in modern campaigns is not unprecedented. What is new is the way in which many politicians respond when confronted with accusations that a piece of information they used is false: To many, it does not seem to matter. That is at the heart of a phenomenon often described as 'post-truth'. Matthew d'Ancona describes post-truth as a "crash in the value of truth, comparable to the collapse of a currency or a stock ... Honesty and accuracy are no longer assigned the highest priority in political exchange" (D'Ancona 2017b, p.8).

The dawn of the post-truth age is difficult to date. Some journalists

have pointed to Vladimir Putin as a precursor of modern post-truthism (D'Ancona 2017a). Two years before Britain's EU referendum and Trump's ascent to the United States Presidency, Kremlin-affiliated media outlets waged what NATO's top commander General Philip Breedlove called 'the most amazing information warfare blitzkrieg we have ever seen in the history of information warfare' (Pomerantsev 2014). In February and March of 2014, an army of 'little green men' – soldiers in military gear without insignia, carrying Russian weapons - swarmed over the peninsula of Crimea, seizing its Supreme Court, its Council of Ministers, and its parliament building. Putin referred to them as 'local self-defense forces', or as 'volunteers'. In a televised question and answer session on 16 April 2014, the Russian President said, "I will say this clearly. There are no Russian troops in Ukraine" (Oliphant 2015; Engel 2015). A day later, he said, "Of course our troops stood behind Crimea's self-defence forces" (Anischchuk 2014). In a documentary aired on state television a year after the annexation, titled "The Way Home", Putin went a step further, explaining that specialists were needed to protect Russians in Crimea from violence and repression by Ukrainian nationalists. "That's why I gave orders to the Defense Ministry – why hide it? – to deploy special forces of the GRU (military intelligence) as well as marines and commandos there ..." (Schreck 2019). The following year, at a televised press conference in December 2015, Putin denied again that 'regular forces' were involved in the conflict. When a Ukrainian journalist asked about two Russian officers who had been captured in Crimea and were being held in Ukraine, Putin replied, "We never said that there weren't people there dealing with certain tasks, including in the military sphere. (...) But that doesn't mean

there are regular Russian forces there. Feel the difference" (Oliphant 2015). The Soviet-born British journalist Peter Pomerantsev pointed out how Putin's propaganda is different from Soviet propaganda: "For the Soviets, the idea of truth was important– even when they were lying. Soviet propaganda went to great lengths to 'prove' that the Kremlin's theories or bits of disinformation were fact." Putin barely tried to cover up his own lies.

Comparing campaigns in liberal democracies like the United Kingdom and in the United States to campaigns in autocratic regimes like Russia is, in many ways, comparing apples to oranges. But the importance that top government officials in all three countries attach to the factual accuracy of their own statements bears resemblance. The way the Leave campaign dealt with the revelation that one of their key campaign promises was based on a false number is a case in point. On his blog, Dominik Cummings, campaign director of Vote Leave in 2015-16 laid out the five foundations of his campaign. First, "'Let's take back control'. The overall theme." Second, "The official bill of EU membership is £350 million per week – let's spend our money on our priorities like the NHS instead." The £350 million claim was debunked on numerous occasions. Even the UK Statistics Authority wrote to Vote Leave during the referendum campaign, saying the claim was misleading and undermined trust in official statistics. It excluded the UK's rebate negotiated by Margaret Thatcher in 1984 (the UK pays significantly less than the 1 per cent of national GDP that member states are normally expected to pay), and it failed to account for payments received by the UK from the EU. The figure also failed to take into account the contribution to the treasury of trade and business

that would not have happened without the EU single market (Henley 2016). And yet, many Leave campaigner continued to repeat this false claim. The statement 'We send the EU £350 million a week. Let's fund the NHS instead' adorned the Leave campaign bus until the day of the referendum.

Why would campaigners keep repeating a number that is known to be false? According to Cummings' own account, the Leave campaign used the false number because 'it worked'. He wrote, "The Treasury gross figure is slightly more than £350m of which we get back roughly half. (...) Sometimes we said 'we send the EU £350m' to provoke people into argument. This worked much better than I thought it would." On 23 June 2016 51.9 per cent of UK voters voted to leave the European Union. Cummings commented, "Would we have won without immigration? No. Would we have won without £350m/NHS? All our research and the close result strongly suggests No." In January 2018 – a year and a half after the EU referendum – Boris Johnson, then Foreign Secretary, reopened the row about the £350m/NHS claim. In an interview with the Guardian he claimed the figure was an underestimation: "There was an error on the side of the bus. We grossly underestimated the sum over which we would be able to take back control" (Stewart and Asthana 2018).

Five months after the UK's EU referendum, Donald Trump was voted into office. If his take on the relevance of factual accuracy was not clear by then it became clear on day one of his presidency. Trump's presidency began with a fight over whose inauguration crowd was bigger: His or that of former U.S. President Barack Obama. Photographic evidence suggested the latter. When photos of the two crowds were published side by side (showing without any doubt that Obama had drawn the larger crowd) the newly inaugurated President resorted to creative evidencemaking: He asked a government photographer to edit the official pictures of the inauguration (Swaine 2018).

White House Secretary Sean Spicer called a special press conference on the matter. He insisted that Trump's inauguration had drawn "the largest audience to ever witness an inauguration, period, both in person and around the globe". The ungrateful task of explaining why Sean Spicer had lied at his very first press conference fell to Kellyanne Conway, senior aid to the president. The following day, on NBC's Meet the Press, Conway coined what would become the catchphrase of the year: "Don't be so overly dramatic about it, Chuck. You're saying it's a falsehood. (...) Sean Spicer, our press secretary, gave alternative facts to that" (NBC News 2017a). Just like Putin, Conway suggested that the number – the facts – did not matter. Her point was clear: You have your version of what happened, I have mine, and that is fine.

Trump's own interactions with journalists paint a similar picture. One of the earlier cases in which the United States President had to explain one of his false statements occurred at a White House news conference in February 2017. NBC's Peter Alexander enquired after Trump's statement that he had achieved the biggest electoral college win since Ronald Reagan. Alexander began to quote the number of votes Obama had received in 2008 when Trump interrupted him, saying he had been talking about Republicans. In reply, Alexander started reading out statistics about the number of votes George H.W. Bush had won in 1988. The President interrupted him again, saying: "I was given that information. I was just given

it. We had a very, very big margin." Alexander proceeded to ask: "... why should the American people trust you when you accuse the information they receive as being fake when you're providing information that's not accurate?" Trump repeated: "I was given that information. I actually, I've seen that information around. But it was a very substantial victory, do you agree with that?" (NBC News 2017b). This widely cited interaction reveals two things: First, the United States President used false numbers to exaggerate his victory. (More harshly, one might say he lied.) Second, and more importantly, the United States President did not seem to care whether his claim was correct. He had compared the two crowds to make a point. What mattered was the point, not the facts.

In these three cases the politicians' false claims were widely corrected. The identity of the 'little green men' who annexed Crimea has been an open secret from the day (or night) they appeared. The £350 million a week claim was the most fiercely debunked claim of the Brexit referendum. Countless broadcasters and newspapers reported that the United Kingdom's cheque to Brussels was far lower than £350 million a week. And the side-by-side photos of Trump's and Obama's inauguration crowds that went viral on social media left no doubt as to whose crowd was bigger. In all three cases, journalists and fact-checking charities went to great lengths to verify information. Fact checking became a profession. Fact checkers checked politicians' statements on a daily basis; major debates were sometimes even fact checked in real time. In all three countries, the majority of voters have access to the internet. The facts were only just a google search away. And yet, election results in all three countries seemed to suggest that voters did not punish politicians for lying. In his book, 'Post-Truth' Matthew d'Ancona quotes the Serbian-American writer Steve Tesich who coined the word 'post-truth' in a 1992 article. Tesich wrote that Watergate and the Iran-Contra scandal had traumatized the American public to the extent that they had started to discredit the value of truth. He wrote: "We are rapidly becoming prototypes of a people that totalitarian monsters could only drool about in their dreams. All the dictators up to now have had to work hard at suppressing the truth. We, by our actions, are saying that this is no longer necessary, that we have acquired a spiritual mechanism that can denude truth of any significance. In a very fundamental way we, as a free people, have freely decided that we want to live in some post-truth world."

This dissertation is about the citizen consumers of false facts. I ask: When do voters tolerate false information and when do notice, or punish it? How does the situation they are in when they are confronted with challenging information affect the way in which they process false information? What makes voters overlook false claims? And how do voters react to information that is challenging in a different way: not because it is false but because it challenges what they thought was true? What do voters make of statistics that challenge information they were convinced was true? What if their political opinions are built on that information, which, as they now find out, the experts say is false? How much trust to modern voters place in experts, anyway? And what if the experts do not have the 'final word'? What if people around them doubt the experts? Do post-truth surroundings undermine whatever credibility is left in a post-truth age?

This work builds on a large body of research on motivated reasoning

(Kunda 1990; Kraft, Lodge, and Taber 2015; Lodge and Taber 2007; Lodge and Taber 2013). As its name suggests, the theory holds that "all reasoning is motivated". We pursue not only accuracy goals but also directional goals: "people are more likely to arrive at conclusions that they want to arrive at" (Kunda 1987; Kunda 1990). Kahan (2016b) defines motivated reasoning as "the tendency of individuals to unconsciously conform their assessment of information to some goal *collateral* to determining its *truth*." The goal, he notes, "can be myriad": individuals may wish to maintain a positive self-conception, to rationalize self-serving behaviour, to avoid anticipated stress or anxiety of unwelcome news, or to perceive coherence rather than complexity in information that is relevant to important decisions (see Kahan 2016b, p.2). In 'politically motivated reasoning' the goal is 'the formation of beliefs that maintain a person's connection to and status within an identity-defining affinity group whose members are united by shared values' [ibid]. In other words, the decision whether to credit or to discredit a new piece of information does not only depend on my perception of how true or false it is. It also depends on how I believe that crediting or discrediting that piece of information will affect my standing in a group of people who share my values, for instance, the political party I support (see also Klein 2015; Harari 2018). Naturally, this affects the way I process new information: A long tradition of research has found evidence of 'confirmation bias' - selectively crediting information that confirms prior beliefs and discrediting information that disconfirms it (Nickerson 1998).

The classic demonstration of motivated reasoning and one of the first studies in the field is Ross, Lepper and Hubbard's dealth penalty study

(Ross, Lepper, and Hubbard 1975). The authors sought out individuals with strong views in favour or against capital punishment and asked them to read two reports about the effectiveness of death penalty as a deterrent against crime. One report found it was effective; the other found it was ineffective. Both opponents and proponents of death penalty argued that the study that was consistent with their own opinion was most persuasive, and found fault with the respective other study. Since then, a large number of other studies have consistently shown that voters tend to see information through a partisan lens (Wlezien, Franklin, and Twiggs 1997; Bartels 2002; G. Evans and Andersen 2006; G. Evans and Pickup 2010; Lavine, Johnston, and Steenbergen 2013; De Vries, Hobolt, and Tilley 2015) and that false beliefs, once established, tend to resist correction (Lewandowsky, Stritzke, et al. 2005; Gaines et al. 2007; Sides and Citrin 2007; Jacobson 2010; Nyhan 2010; Thorson 2015; Schaffner and Roche 2017; Jerit and Barabas 2012; Nyhan, Reifler, and Ubel 2013; Lodge and Taber 2013; Nyhan and Reifler 2010; Eric C. Nisbet, Cooper, and Garrett 2015; Nyhan and Reifler 2015a; De Vries, Hobolt, and Tilley 2015; Nyhan and Reifler 2017; Barrera et al. 2018; Flynn and Krupnikov 2019).

Westen et al. (2006) provided the first neuroimaging evidence for motivated reasoning. During the 2004 U.S. presidential elections, the authors conducted experiments, in which they asked supporters of U.S. Presidential Candidates George Bush and John Kerry to consider consonant and dissonant information about both candidates. Magnetic resonance imaging showed that the reasoning areas of partisans' brains shut down when they were considering information that was threatening to their own candidate. In contrast, when they saw positive information about their preferred candidate, the emotional areas of the brain lit up. Hence, it seems that a certain proclivity to motivated reasoning is hard-wired in our DNA (see also Van Veen et al. 2009).

This dissertation contributes to a nascent body of literature investigating factors that mitigate (or worsen) motivated reasoning biases. To date, most of this research evolves around the question of how to present written information in a non- (or less) threatening way. A number of guide-lines for journalists and public broadcasters have been established: If possible, corrective information ought to include graphs, come from a credible source, and close causal gaps left by retracted information. False information ought not to be repeated; instead, correct information ought to be confirmed (see Nyhan and Reifler (2013) for a summary of research findings).

Another strand of research investigates the effect of accuracy incentives on accurate answers. Scholars in this field have questioned an underlying assumption of much of the motivated reasoning literature: that respondents who give wrong, partisan answers actually believe what they say. Bolsen, Druckman, and F. L. Cook (2014) found that a written accuracy incentive ('Please try to view the policy in an evenhanded way and from various perspectives') reduced partisan motivated reasoning. Similarly, Prior, Sood, and Khanna (2015) and Bullock et al. (2015) found that monetary accuracy incentives can reduce the gap between Democrats and Republicans in responses to 'partisan' factual questions. Schaffner and Luks argue that at least some of the false beliefs partisans report in survey result not from genuinely held misperceptions but from partisan cheerleading (Schaffner and Luks 2018). In contrast, in a series of experiments using a slightly different design Berinsky (2018) found little evidence of expressive responding.

A smaller strand of research has explored the effect of personality differences and, in particular ideology on motivated reasoning biases – here, some early studies have found that conservative respondents were more likely to engage in motivated reasonings (McCright and Dunlap 2011). More recent studies found no effect of partisanship (Berinsky 2018; Kahan 2013). (Berinsky 2018) found that *liberals* were slightly more likely to engage in motivated reasoning (Berinsky 2018).

This dissertation seeks to contribute to a a nascent strand of research investigating how situational factors affect motivated reasoning biases. A few factors have been tested already. Nyhan and Zeitzoff (2017) tested if feelings of control affected individuals' susceptibility to denying wrong-doing that their group has inflicted on others. They found no effect. Lee, Kim, and Schwarz (2015a) explored the effect of a 'fishy' smell on motivated reasoning biases. They found that it increased the likelihood that respondents detected a semantic distortion (the 'Moses illusion') and the likelihood that they engaged in negative hypothesis testing (falsifying their first hunch). Others investigated the effect of self-affirmation exercises: G. L. Cohen et al. (2007) and Koningsbruggen, Das, and Roskos-Ewoldsen (2009) found a positive effect on accepting counterattitudinal messages.

I explore the effect of three situational factors on motivated reasoning biases: stress, low status, and post-truth surroundings. All three have received little attention as potential predictors of motivated reasoning biases. To date, there is one study on the effect of stress (or time pressure) on the likelihood of detecting false facts (Bago, Rand, and Pennycook (2020), in press at the Journal of experimental psychology). A few studies, in particular Mutz (2018) and Major, Blodorn, and Major Blascovich (2016) have investigated the effect of status threat on voting behaviour in election campaigns with abundant false information. I am not aware of any studies testing the effect of post-truth messages on tolerance of false claims or on reactions to expert advice.

Examining external factors that might affect voters' confirmation bias or the likelihood of detecting dodgy information in political campaigns I make no assumptions as to who is more or less susceptible to these biases. I do not assume that voters on either one side of the Brexit or Trump electorate were any more gullible than voters on the other side. Hence, I challenge an implicit assumption of much of the ensuing post-truth talk: that Brexit voters believed in the false claims spread by the Leave campaign's lies and that Trump voters believed in his false claims. Instead, I seek to investigate mechanisms that lead *any* individual, or any group of individuals to overlook false claims in political campaigns.

Because I test causal mechanisms I rely on randomized experiments. Randomly assigning respondents to a treatment or a control groups allows researchers to control for bias from unknown confounders. Therefore, randomized controlled experiments are often seen as the 'gold standard' of causal inference (Brass, Nunez-Neto, and Williams 2006; Bickman and Reich 2009). Two of the three studies described in this dissertation are survey experiments (chapter 1 and chapter 3); one is a laboratory experiment (chapter 2).

Chapter 1 is about the effect of stress on motivated reasoning biases. Building on dual-process models of information processing I argue that stress, or time pressure might shift respondents to an intuitive (system 1) rather than analytical (system 2) thinking style and, thereby, increase confirmation bias. To test this hypothesis I needed a population of voters with strong political views and objectively false information that challenged their political views. I found both in the Brexit referendum. I hypothesized that stress caused Leave voters to overlook factual inaccuracies in Leave campaign claims. I used a popular platform for political discussions to recruit participants: Twitter. Recruiting participants through twitter is a low (or, in this case no)-cost way to obtain a convenience sample (see Mullinix et al. (2015) on the generalizability of survey experiments using convenience samples). What is more, using twitter allowed me to selectively contact voters who held strong views in the Brexit debate: I tweeted at twitter users who had previously used an EU referendum-related hashtag. The survey experiment I report on in this chapter was embedded in a larger survey by Rob Johns, Professor of Politics at the University of Essex, Marcel van Egmond, Senior Lecturer at the University of Amsterdam, and Heinz Brandenburg, Senior Lecturer in Politics at the University of Strathclyde. About a third of the twitter users who clicked on the link we sent them were directed to my short survey experiment; the rest participated in Johns, van Egmond's, and Brandenburg's survey. We fielded the survey in June 2016, shortly after

the Brexit referendum (n=277 including 99 Leave voters). My study focused on respondents who had just voted to leave the European Union. I hypothesized that exposure to a subliminal stressor would increase a gut feeling they had just expressed at the polls: that the Leave campaign claims were true. Creating a subliminal stressor in a survey experiment proved difficult. I relied on time pressure, a common stress manipulation: Respondents in the stress treatment group were given 35 seconds to answer a few questions about the European Union. (I chose EUrelated questions to remain in line with the announced topic of the survey). While they were answering those questions a timer counted down to 0; if they reached 0 without answering a question they were automatically forwarded to the next page. Control group respondents answered the same three questions, but without the timer. Next, respondents were shown a list of statements including a few false claims. Results showed no effect of the stress treatment on belief in these false claims. I argue that these null findings are not conclusive evidence: the time respondents in both groups spent answering the quiz indicates that the 35 second time limit was too long. I conclude that the time pressure probably failed to induce stress and that more research is needed to clarify the role of stress in motivated reasoning biases.

Chapter 2 is about the effect of status differences on motivated reasoning biases. The theoretical starting point is social identity theory. A number of studies in the social identity theory tradition have found that individuals are willing to make sacrifices for in-groups, even in minimal group settings (such as groups that are assigned based on over or undestimating the number of dots in a frame, or preference for an abstract painting), and even at their own expense (that is, if out-group punishment entails a cost, see Tajfel and Turner (1986) for a summary of social identity theory and early findings). My broader, general hypothesis is that the sacrifices individuals are willing to take to express in-group preference extend to overlooking false information coming from in-group members or any favourable individual. In the experiment I report on I do not test this general hypothesis but a more specific version, adding status differences to the equation.

According to social identity theory, groups are particularly likely to rise up if they perceive status differences to be unstable and illegitimate. I speculate that this higher willingness to engage in conflict goes hand in hand with cognitive biases in information processing. If a group is low in status, if they perceive status differences to be illegitimate, if they have reason to believe that those status differences can be overturned, and if they are then confronted with an individual who acknowledges their hardship and calls for justice then group members will overlook things that might otherwise set off alarm bells - such as, for instance, false claims. The same ought to apply to groups who see themselves not at the bottom of the social ladder, but toward the middle, or even the higher ranks and who fear for their status. A few studies have found support for the hypothesis that status loss (e.g. a declining financial position, or the perception that one's country was loosing status) was associated with voting for Trump (Mutz 2018), or for Brexit (Antonucci, Horvath, and Krouwel 2017). I speculate that the missing link is overlooking 'fishy' information: Status threat may have been one of the factors that led Brexit and Trump voters to overlook, ignore, or pardon factual inaccuracies in these campaigns – and, consequently, to vote for them despite those factual inaccuracies.

The hypothesis I seek to test in this study is simple: Low status or status threat causes group members to overlook factual inaccuracies in political campaigns that promise to raise their group's status. It is generally acknowledged that the best way to test causal claims in an observational setting is using panel data. In this case, one would need to set up a multi-year panel study, in which some groups experience a decline in group status, and track reactions to false claims in political campaigns. (I would expect threatened groups to become more willing to overlook false claims, and I would not expect to see any difference among groups whose status is not under threat).

Post hoc it is obviously impossible to set up a panel study to explore reactions to false claims in the Brexit or Trump campaing. Diana Mutz, however, came close to the ideal research design: In her groundbreaking study on the effect of status loss she used a panel of U.S. respondents who were interviewed in Oct 2012 and again in Oct 2016, just before Donald Trump was elected President. The interviews included questions that allowed her to estimate perceived status threat. (Her outcome variable was not tolerance of false facts but support for Trump. Mutz found that those who lost jobs or experienced stagnant wages due to the loss of manufacturing jobs punished the incumbent (Democratic) party.)

Because a panel study was not feasible I opted for an experiment, manipulating status and observing reactions to false information rather than

Contents

waiting for group status to change and observe reactions to false information. There are two approaches to manipulate status in a lab: Working with real group identities or creating group identities. The first approach relies on existing group identities, such as a person's partisan identity, their religious, ethnic, or cultural identity, their family, gender, profession, sports team, club, or any other group of people they identify with. The researcher would recruit a sample of individuals belonging to the same group and randomly assign some of them to a treatment group and others to a control. The difficulty is in designing a treatment that effectively threatens group members perception of their status. For instance, to threaten the status of natives scholars have often resorted to statistics about rising numbers of immigrants. This is the approach Major, Blodorn, and Major Blascovich (2016) took. The authors reminded their sample of white Americans that whites in the United States would soon be outnumbered by non-whites. (Their results showed that among Americans high in ethnic identification this racial status threat increased support for Donald Trump.)

The second approach, creating group identities, has a long tradition in social psychology and behavioral economics. Scholars have found that behaviour in minimal groups closely resembles behaviour in real social groups. Even if group assignment is arbitrary (e.g. based on preference for a painting) group members show in-group favoritism, inter-group competition, and out-group hostility (Tajfel 1970; Tajfel and Turner 1986; Eckel and Grossman 2005). Creating identities in a laboratory setting rather than working with real identities has several advantages: On a

practical side, it does not require finding a large-enough sample of respondents who share a common group identity. Given the ease with which group identity is created in a lab *any* individual, including undergraduate students from across the world, can participate in the study (Tajfel 1970; Tajfel and Turner 1986; Eckel and Grossman 2005; Chen and Li 2009).

What is more, a laboratory study allows the researcher a degree of control that minimizes noise (or the share of subjects who respond to the treatment) Experiments that manipulate real group status typically suffer from a fairly high level of noise. For instance, in the study mentioned above, (Major, Blodorn, and Major Blascovich 2016) not *every* white reader will have felt threatened after reading that whites will soon be outnumbered. The fewer respondents are expected to respond to a treatment the higher the required sample size. In contrast, a laboratory setting allows the reseracher to manipulate status for all group members and to control to what extent status is manipulated. Finally, testing a causal mechanism in laboratory experiment allows the researcher to test the barebones version of the mechanism, disregarding any known or unknown group-specific peculiarities that might affect the outcome variables.

For the purposes of the study I report on in this dissertation I took the second approach, testing the mechanism in a controlled laboratory experiment fielded in Colchester in the summer of 2018 (n=277). Respondents were randomly assigned to one of two groups: Team A and Team B. I used a pub quiz-style geography quiz to induce status differences. The set-up was simple: Respondents were randomly assigned to Team A or Team B. What they did not know was that the questions were drawn

from two different sets, one easy and one difficult. For team A, the odds of drawing a difficult question were 1:4 for Team B, they were 4:1. Because pay-off hinged on individual and group performance Team B had a much lower average payoff.

To mimic a politician who would promise to raise a group's status (but also lie) I exposed respondents to two types of feedback, allegedly coming from fellow players. One feedback giver sided with the advantaged team, calling the game 'fair play' and congratulating the winning team. The other feedback giver sided with the disadvantaged team, calling the quiz 'unfair' and demanding Team B be paid more to offset the disadvantage. Crucially, one of the two lied: I randomly assigned a false claim and a grossly exaggerated claim to one of the two opinions. As expected, results showed that low status affected tolerance of false facts: Members of the disadvantaged team overlooked false claims if those false claims came from the person who had sided with their team and demanded they be paid more. If, however, they came from the person who had sided with the other team they were more likely to be noticed.

Chapter 3, the final chapter is about the effect of post-truth surroundings on motivated reasoning biases. I report results from a study exploring the effect of post-truth surroundings on reactions to unsolicited, counterattitudinal expert advice. Here, my interest was not a subset of a population, but the entire population of voters in England, Scotland, and Wales. This is a co-authored project with Rob Johns and John Bartle, Professors of Politics in the Government Department at the University of Essex. To combine the causal power of experiments the generalizability of population-based samples we used a population-based experiment

(Mutz 2011).

We operated on the assumption that fact checks are useful but rarely have the last word: In a post-truth world, news are heavily commented, and not all comments acknowledge the authority of expert opinion. Hence, we ask: Does exposure to comments telling you not to take an expert too seriously make you more likely to reject unwanted statistics? To find out, we ran a large-n study (n=2936), fielded by Deltapoll to a representative sample of residents in England, Scotland, and Wales in June 2018. We used a topic on which misperceptions abound among liberals and conservatives: immigration.

At the beginning of the survey, we identified false beliefs and showed participants a fact-check correcting one of their false beliefs. Next, we showed some (but not all) a post-truth comment, saying that 'I would take these statistics with a pinch of salt'. In a 2*3 design we varied the authoritativeness of the post-truth commenter (LSE Professor / Blogger) and the reason why they said they would take the statistics with a grain of salt. We found that fact-checks worked: All four expert statements significantly reduced belief in the corrected false claims. However, posttruth comments worked, too: They eliminated some (but not all) of the effect of the fact check. Whether the post-truth commenter was an LSE professor or a a blogger had no effect. The reason they gave to ignore the expert did not matter either: Saying the fact checker was probably biased was just as effective as saying their information did not match people's everyday experience or saying that 'when in doubt, it's best to trust your instinct.' An electronic version of this dissertation, as well as supporting documents (all questionnaires, all datasets, and all R code) is available at https://github.com/julchenhagenstroem/.

Chapter 1

The Effect of Stress on Belief in False Facts in the Brexit Debate

1.1 Abstract

This study investigated the effect of stress on belief in false facts in the leadup to Britain's EU referendum. An online survey was distributed via twitter to United Kingdom residents in the immediate aftermath of the EU referendum. Participants (n=380, including 99 leave voters) were randomly assigned to a stress treatment (a timed quiz about the EU) or a control condition (the same quiz, not timed). Contrary to expectations, this study did not find any evidence that stress increased belief in false campaign claims. The null results may be due to an ill-designed stress treatment: The 35 second timer proved too long to induce stress. Hence, a need for further research on the effect of stress on political information processing persists.

1.2 Introduction

The United Kingdom's EU referendum was characterized by a high level of uncertainty, and high stakes. Leave campaigners portrayed a 'remain' vote as surrendering to Brussels, and as sending the nation's wealth to Brussels at the expense of the national health system The Sun (2016). Leaving the EU, they said, would free up £350 million a year for the country's cash-strapped National Health Service (Reuben 2016). Remain campaigners warned of an unprecedented economic downturn. Leaving the EU, they said, could cost families £4,300 a year. Leave campaigners, in turn, called those predictions 'fearmongering' (Reuben 2016). Voters were in unchartered territory. Many felt uninformed, and overwhelmed. In early June, an LSE survey found that up to 30 per cent of voters would wait for the last week to decide how to vote, with half of those only deciding finally on the day of the referendum (Helm 2016). The referendum was also characterized by false information. The two sides accused each other of lying (Giles 2016). Independent fact-checking charities and public broadcasters debunked some of the key promises of the Leave campaign as factually inaccurate. In particular, the claim that the UK sent the EU GBP 350 a week was, simply, false. And yet, Leave campaign continued to repeat the false number; and even kept it on their campaign bus. Leave won.

This study explores a causal link between the type of stress or mental fatigue voters experience as a result of being overwhelmed by difficultto-digest information and their susceptibility to overlooking factual inaccuracies in campaign claims. It builds on a large body of literature

suggesting that individuals use one of two routes to process information: thinking 'fast' or 'slow' (Gilovich, Griffin, and Kahneman 2002; Kahneman 2011; J. S. B. T. Evans 2008; Petty 1994). The slow route ('system 2', 'peripheral route', or 'heuristic processing') is quick and dirty: Individuals rely on heuristics, or mental shortcuts to evaluate information. Decisions are taken instinctively. The slow route ('System 2' (Kahneman 2011), 'central route' (Petty 1994), or 'systematic processing' (Chaiken 1980) entails careful, effortful, and logical deliberation. In an ideal world, voters would use the slow route to process campaign claims. They would seek out information about the arguments on both sides of a debate, they would think about it, fact check information as needed, accept what is true, reject (and forget) what is false and come to a well-informed conclusion as to which argument is most convincing. Then, they would vote. In the real world, voters rarely have the time, energy, or motivation to think carefully about every snippet of political information they are consciously or subconsciously exposed to.

A large body of political science literature testifies to the fact that voters use heuristics to take decisions, effectively outsourcing the cognitive work of seeking out and evaluating information to sources they trust, be it political parties, television channels, newspapers, blogs, or other individuals or groups of individuals (Robert Johns 2009). Generally, this works surprisingly well: "People can be knowledgeable in their reasoning about political choices without necessarily possessing a large body of knowledge about politics" (Sniderman, Brody, and Tetlock 1991, p.19, see also Gigerenzer and Brighton 2009). However, the quality of a system 1 voting decision depends on the person or the institution whose judgment a voter trusts. If a voter trusts a particular party to pursue policies that are in their best interest and if that party actually does pursue these policies then the voter can save themselves the time and effort it takes to evaluate the policy proposals. The problem is that this is not always the case.

If British voters who put their trust in tabloids such as The Sun or The Daily Mail or in populist politicians such as Nigel Farage, or Boris Johnson took a quick system 1 route when faced with the decision how to vote they may have overlooked a crucial detail: that many of their promises were based on false figures. For instance, if a voter saw the £350m claim and took the fast route to process it they may have thought that this was an excellent reason to leave, because that is a lot of money and the NHS needs it. (This gut reaction is especially likely to be true for the many British citizens waitlisted to see a specialist or have an operation.) It takes a system 2 route to see potential problems. In this case, thinking slowly would not help detect the wrong figure. (Research on innumeracy has found that the human brain is unable to imagine large numbers (Paulos 2001).) However, it might lead voters to entertain the idea that Britain would not be able to divert the *entire* budget that is now transferred to the EU over to the NHS. For instance, thinking slowly might have led them to consider that there is a 'divorce bill' which makes large payments to the NHS unlikely in the short term. Thinking slowly about the longterm consequences might have helped voters see that at least some funds would be needed to offset other payments from the EU, or that, or that leaving might have a negative effect on trade, which would reduce the
per-capita income, and, thereby, tax revenues. Thinking slowly does (and should) not lead everybody to the same conclusion. In this case, Thinking slowly about this argument and the many other arguments by forward by the two campaigns will have led some to vote leave and others to vote remain. In both cases, it would have been a well-informed decision. Many voters took a well-informed decision – but not all. Hence, I hypothesize that feeling stressed, or rushed, nudges voters to use the fast route and overlook factual inaccuracies.

A growing body of scholars in psychology and behavioural economics study the effect of stress on decision-making. Experimental treatments vary considerably, as do scholars' definitions of stress. Seymour Levine noted that "Attempts at definitions of stress have bewildered many an illustrious scholar" (Levine 1985). Some scholars define (and measure) stress as an increase in cortisol or other physiological measures, others point to behavioral features (see Buckert, Oechsler, and Schwieren 2014), still others do not define it at all. In particular, the difference between 'stress', 'cognitive load' and 'mental fatigue' are unclear; the terms are often used interchangeably, experimenal treatments of 'mental fatigue' closely resemble those others use to manipulate 'stress'. For the purposes of this study I follow Marcora, Staiano, and Manning's definition of mental fatigue: 'a psychobiological state caused by prolonged periods of demanding cognitive activity' (Marcora, Staiano, and Manning 2009).

The overall finding is that stress impairs deliberate reasoning. Instead,

decisions are taken instinctively. Individuals under stress scan fewer alternatives (rather than considering all options), show an impaired memory performance, think less strategically and often take suboptimal decisions (Janis 1982; Keinan, Friedland, and Ben-Porath 1987; Dougherty and Hunter 2003; Galván and Rahdar 2013; Staal 2004; Starcke and Brand 2012; Leder, Häusser, and Mojzisch 2013; Gok and Atsan 2016). Yu (2016) proposed that stress elicits a switch from analytic system 2 thinking to intuitive system 1 thinking. (He predicted that this switch led to less activity in the prefrontal executive control regions and more activity in subcortical reactive emotion brain areas and called for research examining the underlying neural mechanism proposed.) The effect of stress on political behaviour is less studied. In particular, we know very little about how stress affects political information processing. To date, I am only aware of one study that explored the effect of stress on detecting false information in political statements: In a 2019 study in the U.S. Bago, Rand, and Pennycook (2020) tested the effect of giving respondents more thinking time. The authors presented respondents with a series of headlines and asked them to give an initial, intuitive response under time pressure and concurrent working memory load. Later on, they allowed respondents to re-think their answers. As hypothesized, rethinking corrected intuitive mistakes: Respondents believed false headlines (but not true headlines) more when they had to give a quick answer than when they were allowed to take time to think about them.

I tested the effect of stress on tolerance of false claims in the United Kingdom during the Brexit debate. The Brexit referendum lent itself to an

analysis as a number of false claims had been spread and had been corrected. I focused on false facts spread by the leave side – not because all Remain claims were factually accurate but because the controversial claims concerned predictions about the economic impact of Brexit (which were, obviously, not falsifiable). I used an online survey experiment to test to what extent a subliminal stressor made twitter-recruited Leave voters believe in the false facts of the Leave campaign. Results showed that among the whole sample the stressor – a short timed quiz about the EU – had no affect on belief in false facts. However, it found heterogeneous treatment effects: subjective social status moderated the effect of the stress treatment on belief in false facts. Subjects who saw themselves at the bottom of British society performed worse under stress. This was true for leave voters; and, surprisingly, also for remain voters. Similarly, citizens who expected their personal finances to get worse as a result of Brexit and citizens who were low in self-esteem showed significantly higher levels of belief in false facts than their peers in the control group. Implications of the difficulty of these subgroups (or of the left behind) to evaluate facts under stress are discussed.

1.3 Method

Twitter was used to recruit an opinionated sample of leave voters. To induce a stress, I used time pressure. Experimental researchers have used a number of manipulations to induce stress; most were lab-based and, therefore, not applicable for this study. ¹ To date there is no standard

¹ Researchers in psychology often use a combination of public speaking and/or arithmetic tasks in front of a panel to induce psychosocial stress (for instance, the Trier

approach for survey experiments. A common denominator in many of these manipulations is a cognitive task and time pressure. For the purposes of this online study I therefore designed a cognitively demanding taks to be completed under time pressures. For a cognitively challenging task that would not deviate from the announced topic of the survey I chose the three standard questions the Eurobarometer uses to measure knowledge of the EU. In 2015, only 27 per cent of British respondents answered all three questions correctly (Hix 2015).

Data were collected through a qualtrics survey fielded to twitter users shortly after the 23 June 2016 EU referendum. Twitter was chosen as a recruitment tool in order to obtain a sample of British residents who held strong opinions about Brexit. This project was embedded in a larger project about public opinion dynamics during the EU referendum campaign by Rob Johns, Professors of Politics in the Government Department at the University of Essex, Heinz Brandenburg, Senior Lecturer in Politics at the University of Strathclyde, and Marcel van Egmond, Senior Lecturer in Research Methods, Department of Communication Science, University of Amsterdam (Brandenburg, Egmond, and Rob Johns 2017).

We created five matching twitter accounts with handles that reflected

Social Stress requires participants to make an interview-style presentation, followed by a surprise mental arithmetic test, in front of an interview panel who do not provide feedback or encouragement, see Cingl 2018; Leder, Häusser, and Mojzisch 2013). Physical treatments are also common; one method is to induce active stress is the Cold Pressor Test (CPT), in which respondents are asked to submerge their forearm in ice water for 3 minutes (room-temperature water for the control group) (FeldmanHall et al. 2015, e.g.), others involve running on a treadmill. (Hepler and Kovacs 2017, e.g.). Cognitive stress manipulations are also common, for instance, studies have used arithmetic tasks (Hepler and Kovacs 2017, e.g.), such as repeating numbers (S. Duffy and J. Smith 2014), serial substraction tasks (Hepler and Kovacs 2017, e.g.); or colour-naming Stroop (CNS) tasks (M. R. Smith et al. 2016; Buckert, Oechsler, and Schwieren 2014).

the intent of the survey (BrexitStudy, BrexitSurvey, eurefstudy, EURef-Surv2016, and brexit_survey). In lieu of a bio we wrote a short description, 'A research project from the Universities of Essex, Strathclyde & Amsterdam. We don't take sides, but we're keen to hear which side you're on.' A link to the Government Department at the University of Essex was provided underneath this description. A picture of an EU and a British flag was used as a profile picture, 10 Downing Street was used as a background picture to capture the attention of subjects interested in politics (see A.1a). Potential respondents were identified based on their tweeting behaviour: We sent out direct messages to 6651 individuals who had used Brexit-related hashtags in the weeks before the referendum. Our messages included a short invitation text ². and an individualized link (so that it could only be taken once). About two third of respondents (n=652) were directed to Brandenburg, Egmond, and John's study, the remaining third was directed to this study.

Data collection started on 25 June and ended on 2 August 2016; most subjects participated in the first two weeks after the referendum (see figure fig:timing). Participants were not paid. Most respondents used the individualized link, some used an open link which we sent to respondents who were unable to open the individualized link in their browser. Data were pre-screened for completeness and residence in the United Kingdom. The final sample consisted of 357 subjects, aged 16 to 75, 63 per cent male. 319 lived in England, 26 in Scotland, nine in Wales, and three

²The wording of the invitations sent to our subjects varied slightly; one of five versions read "@TwitterName: Would you participate in a survey on the EU referendum? Please click here: [link to qualtrics survey]"

in Northern Ireland.³

The survey began with questions on demographics, interest in the EU referendum, expectations about what would happen after the referendum, vote choice, and attitude strength ('How sure were you about your referendum vote choice?'). We then asked about respondents' feelings about the outcome of the referendum, whether they thought the decision was final or whether they thought there might be a second referendum, and what their most important sources of information about the referendum were. Next, we asked about a few demographics (gender, age, marital status, income) and assessed a few personality measures: the big 5 personality traits (we used the classic Ten-Item Personality Inventory, Gosling, Rentfrow, and Swann 2003), trust, self-esteem, subjective social status (Adler and Steward 2007), ⁴ life satisfaction, need to evaluate, and need for closure.

Next, subjects were randomly assigned to a treatment group (n=183) or

³278 UK residents completed the survey using links connected to their twitter accounts. We received a few messages from individuals who had problems accessing their individual links (mainly Safari users) or who wanted to re-tweet the link to the survey. These individuals were sent an open link to the survey. To guarantee that no twitter user took the survey twice we included a question asking for subjects' twitter handles in the open link version of the survey. 103 participants completed the survey using this open link. Their twitter names were scanned for personal acquaintances. One personal acquaintance and nine subjects not residing in the UK were excluded. The samples accessing the survey via individual links and the sample accessing it via the open link were similar across the main political and demographic measures and were pooled.

⁴ The MacArthur Scale of Subjective Social Status was used to capture subjects' perception of their own social status across SES indivators. It shows a ten-step "social ladder"; participants place themselves on of the ten steps. The description read: "Think of this ladder as representing where people stand in the United Kingdom. At the top of the ladder are the people who are the best off – those who have the most money, the most education, and the most respected jobs. At the bottom are the people who are the worst off – who have the least money, least education, and the least respected jobs or no job. The higher you are on this ladder, the closer you are to the people at the very top; the lower you are, the closer you are to the people at the very bottom. Where would you place yourself on this ladder? (Please write below a number between 1 and 10.)"

a control group (n=174). All respondents answered the Eurobarometer's three EU knowledge questions ("The EU currently consists of 28 member states."; "The members of the European parliament are directly elected by the citizens of each member state." and "Switzerland is a member state of the EU."). The control group saw these questions in the same format as all other questions; introduced as "True or False?". For the treatment group, the same questions were introduced as a timed quiz. A clock started ticking as soon as they reached the page. They had 35 seconds to answer the questions. Because this treatment was designed as a subliminal stressor I did not add any manipulation checks. I did, however, time the time spent on the (timed or not timed) quiz as well as the time spent on the subsequent questions measuring belief in false campaign claims. The timed version and the not timed version of the quiz is shown in figure A.2 in the appendix.

Immediately after the stress manipulation participants were presented with a battery of seven statements about the EU, and were asked to rate how true or false they were. Four of these statements (the ones of interest to this study) were factually incorrect campaign statements propagated by the Leave campaign, and corrected by remain campaigners, fact checkers and national and international public broadcasters: "Leaving the EU frees up GBP 350m a week for the NHS."; "If Britain had remained in the EU it would have had to accept Turkish. membership."; "The EU could have forced British soldiers to join a European army."; "The EU could have made Britain join the Euro.". The remaining statements were distractor items. Two were correct: "The pound plunged to its lowest level since 1985 after Britain voted to leave the EU", and "All council areas in Scotland voted to remain". The final statement had been subject to speculation ever since it was reported in a tabloid: "The Queen backs Brexit".

The seven statements appeared in a random order. Respondents rated each statement as "definitely true", "probably true", "probably false", "definitely false", or "don't know". A hidden timer measured how long subjects spent evaluating them. The survey ended with questions on how much they thought they had in common with poeple who wanted to leave and with people who wanted to remain (apart from what they thought about Europe) and, finally, measures of curiosity, need for cognition, and need for evaluation.

1.4 Results

As expected, the overwhelming majority of respondents were highly interested in the referendum. 93 per cent of both leave and remain voters were "very interested" in the EU referendum (see figure A.4a in the appendix). Contrary to expectations, the distribution of remain to leave voters did not mirror the results of the referendum. 225 subjects, i.e. two thirds of the sample reported they had voted to remain. 99 subjects had voted to leave, 5 had not voted. Participants' age ranged from 16 to 75; most were in the 20s to 50s. As in the referendum results, the ratio of leave voters to remain voters increased with age. Among participants in their teens and twenties, less than a third voted to leave. In contrast, a third of those in their thirties, a forth of those in their forties, and half of those in their sixties and seventies reported to have voted to leave. Attitude strength was high: 93 per cent of leave voters and 95 per cent of remain voters scored 5 or 6 on a 6-point slider question where 1 meant 'not sure at all' and 6 meant 'absolutely sure'. Two thirds of leave voters saw no reasons at all or not very many reasons to remain, only two per cent saw 'very many' reasons to remain. Remain voters were even more convinced of their cause: 90 per cent saw 'no reasons at all' or 'not very many reasons' to leave; only one per cent saw 'very many' reasons to leave. Among leave voters the forecast for a post-Brexit Britain was positive and testified to a high level of belief in the campaign's promises: 59% of leave voters [versus 16% of remain voters] believed immigration would be curbed after Brexit, 56% [1%] thought that the NHS would be better off; 61% [2%] expected the general economic situation to improve, and 25 % [1%] believed that that their personal financial situation would improve as a result of Britain leaving the European Union.

The main independent variable was subjects' assessment of five false claims that had been spread by Leave campaigners. A simple additive index was created to measure belief in these false facts. Each false claim was converted to a numeric variable where 'Definitely true' was assigned a 5, "'Probably true''' a 4, "'Do not know''' a 3, "'Probably false''' a 2, and "'Probably false''' a 1. The resulting additive index ranged from 5 (all five statements rated as "definitely false") to 20 (all five statements rated as "definitely true"). The index was reliable (Cronbach's α = .81(95 per cent confidence interval: .77- 84) As expected, leave voters reported greater levels of belief in false facts about the EU than did remain voters (cf. distribution in figure A.11b).

Contrary to expectations, the stress treatment had no effect on how Leave



FIGURE 1.1: Belief in false facts by vote choice

voters evaluated these false claims. Figure 1.1 shows no difference between leave voters who had been exposed to time pressure and leave voters who had not been exposed to time pressure. (Similarly, Remain voters in the treatment group showed just about the same levels of belief as remain voters in the control group.) Figure 1.2 shows the effect of stress by attitude strength, where 'absolutely sure' remain or leave voters refer to those who ticked the maximum of 6 on a 0 (not at all) to 6 (absolutely sure) scale). It shows that even the most convinced Leave and Remain voters were unaffected by the stress treatment: the average level of belief in false facts was almost the same in the treatment and remain groups.

Next, I looked at the effect among a subgroups of leave voters who might have been particularly drawn to the Leave campaign's promises: subjects



FIGURE 1.2: Belief in false facts about the EU by vote choice and certainty about vote choice

who saw themselves in the lower echelons of society. To these individuals, the status-boosting tone of the Leave campaign might have been particularly appealing, which means that for them, in particular, the gut reaction when confronted with these claims ought to be that they are true. I test if these individuals were any more responsive to the stressors than the general sample. To extimate the effect of those who were low in subjective social status I created a dummy variable distinguishing between subjects of low subjective social status (steps 1 to 5) and subjects of high subjective social status (steps 6 to 10). Figure 1.4 shows slightly higher levels of belief among low-status leave voters than among high-status leave voters. (This can be traced back to differences in education). More importantly, though, it shows small differences in how high and lowstatus individuals respond to the stress treatment: Among those who see themselves in the middle or at the bottom of Britain's society, those who



FIGURE 1.3: Belief in false facts about the EU by vote choice and subjective social status

were exposed to time pressure show slightly higher levels of belief in the false facts than their peers who were not exposed to any time pressure. I looked at a few other personality measures and demographics to find potential heterogeneous treatment effects. I found no significant effect of gender, age, neuroticism, life satisfaction, generalised social trust, need for cognition or need to evaluate on responses to stress.



FIGURE 1.4: Belief in false facts about the EU by vote choice and subjective social status

1.5 Discussion and Conclusion

This study set out to investigate the effect of a stress on motivated reasoning biases among leave voters. Results show no effect. I found no evidence that the treatment that was designed to induce stress – time pressure on a short quiz on the EU – impaired respondents' ability to accurately assess factually incorrect statements about the EU. For two reasons, these null findings ought to be taken with a pinch of salt: First, the sample was small, unrepresentative and unusually opinionated – among leave voters levels of belief in these false facts were so high that at least some of the non-effect may be due to ceiling effects. (Belief was too high for the stressor to have an additional effect). More importantly, it is highly likely that the stress manipulation was too weak. The sample was highly educated and had no difficulty answering the EU knowledge questions correctly. Moreover, the time limit was too short: On average, treated respondents spent 22 seconds answering these questions (sd=7); control group respondents only spent marginally longer: Including outliers, control group respondents spent 32 seconds (sd=31), exluding those who spent longer than 1.5 minutes the average was only 26 seconds (sd=17). Hence, it is very probable that this treatment did not make respondents feel rushed. It may even have made respondents more alert, acting as a wake-up call in a long study. If so then the effect of this alertness may have offset the effect of stress. Therefore, this study ought not to be seen as a evidence that stress has no effect on belief in debunked facts. Future research with a more stressful stress manipulation is needed.

Nonetheless, even this relatively un-stressful stressor did affect evaluations among a few subpopulations. It never improved judgment. If it had an effect, it impaired judgment. There were three groups of people who showed higher levels of belief in the false facts under stress: individuals who were low in subjective social status, individuals who anticipated their personal finances to get worse after Brexit, and individuals low in self-esteem. Notably, this was true for Leave *and* Remain voters. (Among the subset of Remain voters, generalised social trust and life satisfaction also moderated the effect of the stress treatment on belief in false facts: stressed Remain voters who were low in trust or low in life satisfaction did significantly worse on average than unstressed remain voters who were also low in trust or life satisfaction.) The factors that moderated the effect of stress – feeling at the bottom of society, being less educated, having worse jobs and less money, and expecting it all to get worse are signs of feeling left behind. Hence, this study points to a need for further research on how subpopulations who feel left behind, or disadvantaged, process political information.

Chapter 2

The effect of boosting in-group identities on tolerance of false facts

2.1 Abstract

Many recent electoral events have been characterised by false claims which, despite abundant fact-checking, were often widely believed. This paper asks why voters do not punish false facts in political campaigns. Building on findings that suggests a link between status threat and Trump/Brexit voting it investigates a simple hypothesis: That low group status or threatened high group status makes voters susceptible to overlooking false claims coming from politicians who boosts their group identities. A laboratory experiment (n=277) was conducted to test this hypothesis. The set-up was a pub-quiz style geography quiz. Unbeknownst to the participants the likelihood of drawing easy or hard questions was varied so that one team would do worse and receive a lower payoff than the other team. After the quiz, both teams were shown two pieces of feedback: One person called the quiz 'fair play'; the other person called it 'unfair' and demanded a pay rise for the disadvantaged team. One of the two feedback givers were randomly assigned to make false claims. Results showed that the disadvantaged overlooked false claims coming from the person who sided with their team: There was no significant difference in how disadvantaged team members rated the accurate and the inaccurate version of the 'unfair' feedback. If, however, the false facts came from the person who sided with the advantaged team the disadvantaged were more likely to notice the false facts: Those on the disadvantaged team who saw the version of the 'fair play' feedback that included false claims were less agreed (p=0.02), thought it was less accurate (p=0.07), and were slightly (but not significantly) less likely to rate its author as a good team representative (p=0.22) than their peers who saw the factually accurate version. If generalisable to electoral campaigns these results are sobering: If candidates recognize the disadvantage of those who feel disadvantaged then lying will not cost them any votes among that group.

2.2 Introduction

Politics is not known to be the most ethical of professions. For the longest time, a lie (or two) in a political speech did not raise an eyebrow. But in 2016, the amount of lies in political campaigns seemed to have crossed a threshold: The factual inaccuracies in the UK's Brexit campaign and, subsequently, in Donald Trump's election campaign led to a public outcry over politicians spreading false information. What is startling is that on both sides of the Atlantic the campaigns that were most accused of making false or misleading claims carried away the elections: in June, the Leave campaign won the EU referendum, having toured the United Kingdom in a bus with a false figure of the the weekly cost of EU membership on it. In November, Donald Trump, whose claims about attendance at his inauguration would lead to the coining of the infamous word "alternative facts", won the White House. This project is about the citizen voters confronted with false facts in political debates. It asks: Why would voters vote for a candidate who makes claims that are, simply, false? This paper suggests a link between low group status or threatened group status and tolerance of false facts in political campaigns that side with one's in-group.

Both the Trump campaign and the Leave campaign argued that 'ordinary citizens', and the British or, respectively, the American people had had the rough end of the stick for too long. Both campaigns appealed to segments of their societies who were concerned about natives losing jobs to immigrants, and promised to improve the status of natives relative to immigrants. Both campaigns also appealed to segments of the British and American society who mourned the loss of Britain's or America's status in the world, and promised to recover a once-held and now lost superior status of their nation. It seems that this campaign strategy paid off: Voters who were concerned about the status of natives or the loss of America's or Britain's status were more likely to vote for Trump, or, respectively, Brexit. What did these voters make of the false claims in these two campaigns? The motivated reasoning-informed hypothesis advanced in this paper is: nothing. They overlooked them. The hypothesis is simple: The mere feeling that the status of a group individuals identify with is lower than it ought to be or lower than it used to be makes voters susceptible to overlooking false facts coming from politicians who offer positive social identity, for instance by recognizing a group's hardship or promising to improve their status.

A laboratory experiment is conducted to test the causal mechanism on a non-political issue in a controlled laboratory setting at a UK university. The focus is on low status groups: How do groups that are lower in status react to false facts in statements from people who promise to raise their group status?

2.3 Theoretical Background

This study borrows from two theoretical traditions: social identity theory, and motivated reasoning theory. Social identity theory argues that people make comparisons between not only themselves and others, but between their groups and other groups. It is based on a number of group experiments, in which subjects favoured in-groups and discriminated against out-groups even if group assignment was as flimsy and unimportant as preference for a painter (Tajfel 1974; Tajfel 1982; Tajfel and Turner 1986). Tajfel and Turner's social identity theory assumes that part of peoples' self-image is based on group membership (this is their social identity) and that groups are associated with (socially consensual) positive or negative value connotations. Comparing favourably with relevant out-groups leads to high prestige and positive social identity; comparing unfavourably leads to low prestige and negative social identity (Tajfel and Turner 1986, p. 40). The authors (Tajfel and Turner 1986) theorize that individuals seek to achieve or maintain positive social identity: "The aim of the differentiation is to maintain or achieve superiority over an out-group on some dimensions." (Tajfel and Turner 1986, p. 41).

Unfavourable comparisons (i.e. negative social identity) are said to lead to three types of behaviour: (1) individual mobility (leaving a group to join a more positively distinct group), (2) social creativity (e.g. comparing themselves on a different measure or with a different out-group, or reinterpreting characteristics formerly seen as inferior), or (3) social competition (fighting to improve the group's negative image or position). Much of social identity theory is concerned with the conditions under which negative social identity generates conflict over scarce resources (cf. Tajfel and Turner 1986, p. 44f and Tajfel 1974, p. 76f).

According to Tajfel and Turner's theory, two types of groups are particularly prone to conflict: low-status groups who perceive the status differences as unstable and/or illegitimate and high-status groups whose higher status is threatened. Low status groups are competetive toward the dominant group to the degree that "(a) subjective identification with the subordinate group is maintained; and (b) the dominant group continues or begins to be perceived as a relevant comparison group" (Tajfel and Turner 1986, p. 45). They are most likely to rise up if they are unable to join other groups (i.e. individual mobility is unavailable) and if they perceive status difference as both unstable and illegitimate. (Tajfel and Turner (1986, p.45) argued that this was probably the set of conditions that underlay the development of ethnocentrism among Black Americans, French Canadians, and New Zealand Maoris.) High-status groups whose high status is under threat are most likely to 'react in an intensely discriminatory fashion' to attempts by the subordinate group to reverse their group status when they perceive their superiority as legitimate (Tajfel and Turner 1986, p.45-46). This paper does not seek to explain conflict. It takes a step back and considers the cognitive consequences of low status and threatened high status in election campaigns.

On this front, Motivated Reasoning Theory offers helpful insights. It argues, as its name suggests, that "all reasoning is motivated": "people are more likely to arrive at conclusions that they want to arrive at" (Kunda 1987; Kunda 1990). Kahan (2016b, p.2) defines motivated reasoning as "the tendency of individuals to unconsciously conform their assessment of information to some goal collateral to determining its truth." Abundant research confirms that individuals see information through an ideological lens. (cf. also Fischle 2000; Bolsen, Druckman, and F. L. Cook 2015; Kull, Ramsay, and Lewis 2003; Nyhan, Reifler, and Ubel 2013; Bartels 2002; Kahan 2016b; Westen et al. 2006). Lodge and Taber 2013 present a number of experiments that provide consistent evidence that people do not only defend their prior attitudes but that this happens outside of conscious awareness: we know how we *feel* about a piece of information before we know what we *think* of it. They theorize that conscious deliberation is but a way of rationalizing these unconscious processes to find reasons why we feel the way we feel about something (Kraft, Lodge, and Taber 2015, p.129).

Drawing on social identity theory and motivated reasoning theory this paper assumes that (a) voters identify with groups, (b) voters have a need to see their in-groups compare well with out-groups, and that (c) voters are unable to assess information that threatens their in-group in an evenhanded way. This paper explores the cognitive side-effects of coping with negative social identity – the feeling that one's in-group does *not* compare well with relative out-grous – in a political environment in which lying politicians offer positive social identity. It suggests that the low or threatened group status blurs the lens through which one sees politicians who side with one's in-groups.

The Trump campaign and the Brexit campaign seemed to appeal to people who felt economically disadvantaged, or left behind. It also seemed to appeal to people who felt that the status of their people was threatened. This paper focuses on two types of status threat: Global status threat (their country's status in the world was under attack) and racial / ethnic status threat (the status of ordinary (implied, white) citizens was under attack).

2.3.1 Global Status Threat

Both campaigns addressed their electorate in terms of their national (British or American) identity. On its own, this is unremarkable. What makes it remarkable is that both campaigns presented their national identity as being under threat. For instance, Nigel Farage, the key figurehead for the Leave campaign declared that "We are British; we are not going to be bullied by anybody" – implying that Britain had been bullied before (ITV 2016). Similarly, the slogan 'Take *Back* Control' implied that Britain had lost control. The same way, Donald Trump promised to 'Make America

Great *Again'*, implying that America had lost at least some of its greatness. Trump's campaign launch speech in New York City in June 2015 was a vivid example: 112 words into the speech, he declared "We don't have victories anymore. We used to have victories, but we don't have them" (Phillips 2017). In the same speech, Trump exclaimed that 'Sadly, the American Dream is dead.' – shattering an idea, the American Dream, that is at the very heart of America's national identity.

Both Trump and the Leave campaign offered a way out of the dilemma, promising to free this besieged American or British identity. The Trump campaign conveyed the impression that other nations, including China, had trampled on America's status in the world. The Brexit campaign conveyed the impression that the EU had trampled on Britain's status in the world. Both campaigns promised to un-do this injustice and to return their nations to their rightful position as great nations (implied, better than other nations). Given the history of the American and the British nation it is unsurprising that this rhetoric fell on fertile ground: The Leave campaign operated in a former (one might argue a fallen) empire, a former colonizer that has ceded at least some of its power to the European Union; Trump spoke to a current superpower, the world's largest economy and the mightiest military – but also a country that has tarnished its image in the wars of the past few decades and is losing ground, politically and economically, to other rising powers – above all, China.

Populist campaigns beyond Britain and the US reveal similar patterns. For instance, when Tayyip Erdoğan and Vladimir Putin address their diasporas in Europe they consistently depict their countries as belittled by the West. The message they have been sending (be it in person or through their various state-controlled media channels) has been clear: their diasporas belonged to a great nation and ought to be proud of it. (And, to express their Turkish or Russian identity, they ought to vote for them.) Both Erdoğan and Putin have been surprisingly successful in gathering support amongs their expat communities across continental Europe, in particular in Germany, where the social and economic status of citizens of Turkish and Russian decent tends to be below average.

2.3.2 Ethnic status threat

Trump and Brexit campaigners also appealed to people in terms of their identity as 'ordinary' citizens – which, at least in the US context, is often understood to mean 'white' (Tesler 2016b). Both campaigns bemoaned, subtly or not so subtly, the declining status of white men.

For instance, the Leave campaign sided with 'working people'. In a 3 June 2016 interview, Michael Gove warned that "you can say, 'their concerns don't matter' [...]. You can dismiss the concerns of working people. [...] You are dismissing the concerns of working people." In good populist fashion he continued, 'You're on the side of the elites, I am on the side of the people" Sky News 2016. Two weeks later (15 June), he pointed directly at European immigrants: "At the moment all our public services – the NHS and education are under strain, as a result of unlimited free movement from the EU" BBC 2016. Meanwhile, Nigel Farage complained that "it is wrong, wrong, that for average, decent families in this country, their living standards have fallen by 10 per cent over the course of the last few years. And it is about time [...] we started thinking about [...] not just about the rich getting richer, but about ordinary, decent Britons who have got a rotten time. And they really have, too. " ITV 2016. Nigel Farage's 'Breaking Point' campaign poster was less subtle: It showed a photograph of migrants crossing the Croatia-Slovenia border in 2015. The only white person in it was obscured by a box of text, saying 'We must break free of the EU and take back control of our borders' (see Stewart and Mason 2016). Hence, leading Leave campaigners appealed to voters in terms of their identity as 'working people'; they appreciated their hardship, and they vowed to fight to improve their living standards.

Donald Trump's rhetoric was similarly to the point. He promised to bring back 'our' jobs, 'our' manufacturing (namely, Ford), and 'our' military. In his campaign launch speech Trump shouted, "We need money. (...) And we need the right people." In the same speech he described his (mainly White) audience as 'the best and the finest', and contrasted them with 'the people Mexico sends. (...) they're not sending their best. (...) They're bringing drugs. They're bringing crime. They're rapists' (Phillips 2017).

2.3.3 Status Threat and Vote Choice

Post-election research showed that perceived discrimination against whites was indeed correlated with support for Brexit and Trump. Sides, Tesler, and Vavreck (2018) examined British Election Study data. They found that white voters who thought there was a lot of discrimination *against* white people were over 60 percentage points more likely to support Brexit than white voters who thought there was a lot of discrimination in *favour* of white people in the UK (Sides, Tesler, and Vavreck 2018, p.216-217).

Tesler (2016a) analysed the 2016 American National Election Pilot Study. They found that in the US, whites who thought it was extremely likely that "many whites are unable to find a job because employers are hiring minorities instead" were over 50 points more likely to support Trump than their country men and women who disagreed (Tesler 2016a). Fowler, Medenica, and C. J. Cohen (2017) coined the term 'white vulnerability': "the perception that whites, through no fault of their own, are losing ground to other groups". To measure it, the authors constructed a threeitem scale, asking whether whites were "economically losing ground through no fault of their own"; whether discrimination against whites was "as big a problem as that against Blacks and other minorities"; and whether minorities overtaking whites as the majority of the U.S. population by 2050 would "strengthen or weaken the country." Controlling for the usual predictors of Trump support (e.g. partisanship, racial resentment, living in the South, gender, and employment status) white vulnerability strongly and significantly predicted support for Donald Trump. In the US, the perception that whites are being discriminated against is growing; and it is increasingly tied to voting Republican (Sides, Tesler, and Vavreck 2018, p.170f). Americans who voted for Trump saw whites as more discriminated against than Muslims of black Americans (Edwards-levy 2016).

Naturally, none of these correlational studies can establish if white vulnerability, or status threat *caused* citizens to vote for Trump. However, there are a few experimental and panel-data studies – and they do point to a causal connection: First, and before the elections, Major, Blodorn, and Major Blascovich 2016 found experimental evidence that reminding whites that they would soon be outnumbered by non-whites made white Americans high in ethnic identification more supportive of Donald Trump. (It also increased support for anti-immigrant policies and opposition opposition to political correctness among this group.)

Second, and after the elections, Diana Mutz found panel data evidence that rising white and global status threat made Americans shift toward Trump. Mutz argued that white Americans were facing two types of status threat: they were losing in numbers, and they were losing in global status (the US is gradually losing its status as the number one world power). As hypothesized, both types of perceived status threat led people to vote for Trump: People who scored higher on social dominance orientation in 2016 than they did in 2012 (a proxy for racial status threat) were more likely to shift toward Trump. Similarly, people who saw China as more threatening in 2016 than they did in 2012 and who grew less supportive of free trade agreements in 2016 were more likely to shift toward Trump.

2.3.4 Proposed mechanism

This paper suggests a causal connection between the kind of group status threat many UK and US voters felt in 2016 (be it the status of 'ordinary' people, 'white' people, or the status of the 'American' or 'British' nation) and the fact that they overlooked dubious information in campaigns that promised to raise or recover their group status.

False claims coming from politicians who have boosted one's in-group's status poses a dilemma: On the one hand, individuals have a need for factual accuracy (Kunda 1990; Taber and Lodge 2006). On the other

hand, according to social identity theory, they have a need for positive group distinctiveness. Acknowledging that a politician is wrong about one thing (e.g. that they don't get their facts right) would mitigate the positive effect of their status-boosting speech on in-group members' social identity. If, as motivated reasoning research has shown, in-group members are willing to sacrifice money to see their in-group fare better than out-groups (Tajfel 1970) (and if, as history has shown, people sacrifice their lives in wars protecting their in-groups) it seems plausible that they will – at least to some extent – sacrifice their need for factual accuracy as well.

Politicians can nudge group categorization by addressing individuals as members of an in-group. They can nudge negative social identity by pointing out to a low-status group ways in which they do not compare well with relevant outgroups or by pointing out to high-status groups that its higher status is threatened. (Both happened in 2016.) And then they can alleviate these feelings of negative social identity. If, for instance, politicians praise an in-group's hard work, recognize their hardship, or promise to improve their socio-economic standing relative to other groups then they effictively provide a straw for in-group members to clutch at. In-group members will *want* them to be right because they have made them feel better about the groups they belong to. But what if these political elites then say or do something that *ought to* set off alarm bells? What if they use a piece of information that isn't entirely true, or if they advocate a political idea that you do not quite subscribe to, or if they make an over-the-top statement about other groups? What if, ultimately, they propagate political extremism or violence? The very beginning of this chain is politicians making benign but false claims. This paper suggests, simply, that in-group members have a motivation to overlook such false claims – not necessarly to believe them, but to overlook them.

This process ought to work at different levels of conscious awareness. At a pre-conscious level individuals will simply not see or feel any signs of fishiness. At a somewhat higher level of conscious awareness citizens may sense something suspicious but may chose not to spend too much time or too many thoughts on it. At an even higher level of awareness people may sense the fishiness, look into the facts, acknowledge them to be false, but decide to vote for the candidate nonetheless. Put differently, they will over-value the positive impact of the candidate's identityjustifying information and under-value the importance of his or her factual accuracy for their vote choice.

This does not negate the well-established research finding that people do have accuracy goals (Kunda 1990; Taber and Lodge 2006). It does, however, posit that people's need for positive group distinctiveness (or the need to recover lost status) exceeds their need for factual accuracy. Citizens may not believe in everything their leaders say. But as long as these leaders promise to boost their in-group's low or lost status then the nitty-gritty of their talk becomes irrelevant.

A laboratory experiment was designed to test the causal mechanism in a non-political context focussing on economic differences. ¹

¹The oTree code is available here; the pre-registration is available here.



FIGURE 2.1: Groups of interest



FIGURE 2.2: Proposed causal process – Low-status groups



FIGURE 2.3: Proposed causal process – Threatened highstatus groups

2.3.5 Method

In a 2*2 design, this study varied a) advantage or disadvantage at a moneypaying real effort task (a pub quiz-style geography quiz) and b) exposure to false facts by a person siding with the advantaged or the disadvantaged team. The experiment began with a pub quiz-style geography quiz, desigend to establish status differences. Subjects were randomly allocated to an advantaged team (presented as 'Team A') or a disadvantaged team (presented as 'Team B'). Both teams answered 12 multiple-choice geography questions (3 rounds * 4 questions). The questions were randomly drawn from two sets of thirty questions: an easy set containing questions such as 'What is the capital city of Germany?' and a difficult set containing questions such as 'What is the capital city of Liechtenstein?'. Unbeknownst to the participants, the odds of drawing easy or difficult questions varied by team. For each question, the advantaged team had an eighty per cent chance of drawing an easy question, and a twenty per cent chance of drawing a difficult question. The disadvantaged team had an eighty per cent chance of drawing a difficult question, and a twenty per cent chance of drawing an easy question. (Therefore, on average, the advantaged team would answer 9 easy questions and 3 difficult questions, whereas the disadvantaged team would answer 9 difficult questions and 3 easy questions.) The quiz was timed: Participants had thirty seconds to tick one of the four answers and hit a 'submit' button. (If they failed to submit their answer within thirty seconds they were automatically forwarded to the next page.) After submitting each question and before moving on to the new question respondents saw the respective question for the other team (see page 137 in the appendix for example questions).

The quiz was divided into three rounds with four questions each. After each round players were shown summary statistics for that round: the correct answers to both teams' questions and, for each question, the percentage of respondents who submitted a correct answer. After the last round, i.e. after 12 questions players were shown their own payoffs and the mean payoffs for both teams (see page 138). Payoffs were designed to depend on participants' own performance as well as their teams' performance: Each correct answer was worth GBP 1.00. Participants kept half of their earnings and contributed the other half to their team's pot, which was evenly divided amongst all group members. To mimic a pub quiz group sizes were set at 4-8 team members.

Following the quiz, respondents were asked for feedback. (They received three questions, asking if they thought any questions were too easy (i), too difficult (ii), and what they thought about the thirty second time limit (iii). Next, they were asked to consider two other people's feedback. One person's feedback called the game 'fair play', saying Team A (the advantaged team) had had the 'luck of the draw' (see page 142); the other person called the game 'unfair', suggesting that Team B (the disadvantaged team) be paid an additional GBP 5.00 on top of their payments (see page 145). Crucially, one of the two feedback pages were randomly assigned to contain one false claim and one gross exaggeration. The false claim stated that "Some of the places in that quiz don't even exist." The exaggerated claim varied depending on team membership: If assigned to the 'unfair' treatment, the last feedback point said the thirty seconds were "barely enough time to read the questions." If assigned to the 'fair play' treatment, it said they were "more than enough", and that "10 seconds

would have been plenty."²

This study attempted to capture tolerance of false facts on a spectrum ranging from *overlooking* false claims to *believing* them. First, subjects were asked if they 'generally agreed' with both people's feedback (see page 145). Second, they rated to what extent both authors were 'factually accurate', and 'a good representative of my team' (see page 145). (To disguise the intent of the question participants were also asked to rate how 'educated' the author was.) Third, respondents rated each point, including the false claim, on a four-point scale from 'definitely true' to 'definitely false' (see page 145). The experiment ended with a short question-naire (stage 4) including demographics and various control questions.

A difficulty that is inherent in studying respondents' susceptibility to overlooking false claims is that it is imperative that they see the false claims (and do not just skim-read the treatment without noticing the false claim). At the same time, one cannot make the false claim stand out too much and one cannot simply ask by way of a manipulation test if they read it. (Because printing a false claim in red letters or even just asking if they read this claim would draw their attention to it. That of course would defeat the purpose of measuring if they noticed it.) To maximise the likelihood that respondents read the feedback it was kept short. The four sentences that included a false claims were layouted as bullet points. Hence, even skim-readers ought to at least have skim-read the false claim.

²NB: Informal feedback from participants on the advantaged team showed that the exaggerated claim in the 'fair play' feedback was slightly less exaggerated than intended: Some of their questions were so easy that they did not need the entire 30 seconds. (Nonetheless, 10 seconds would not have been enough to answer any of the directions questions or the questions from the difficult set.) Generally, answers to the open-ended feedback question about the time limit revealed that almost all respondents thought 30 seconds was a fair limit.

	False 'fair play'	False 'unfair'
Advantaged team	1	2
(Team A)	n=69	n=64
Disadvantaged team	3	4
(Team B)	n=83	n= 61

TABLE 2.1: Research Design - Study 1

A second difficulty that arises when respondents rate two different people's opinions is that they need to remember which person they are rating. Therefore, when respondents answered questions about either feedback they were able to scroll down to see the respective feedback.

2.3.6 Hypotheses

This experiment sought to create two status groups: an economically disadvantaged group and an economically advantaged group. The general hypothesis depends on status: For the low-status group (the disadvantaged team), it was hypothesized that exposure to a message recognizing the team's disadvantage and suggesting to raise their payoffs makes team members overlook false claims by the source of the message. For the threatened high-status group (the advantaged team), it was hypothesized that exposure to a message legitimising their higher status makes team members overlook false claims by the source of the message. The following hypotheses will be tested across three dependent variables: general agreement, rating the feedback as 'factually accurate', and rating its author as a 'good team representative'.

Hypothesis 1 considers the hypothetical scenario in which a person who

sides with your team either makes false claims or doesn't: It is hypothesized that factual accuracy has no effect on perceived suitability as a team representative:

Hypothesis 1a: Disadvantaged team members who see a factually inaccurate 'unfair' feedback rate the author of the 'unfair' feedback more favourably than their peers who see a factually accurate 'unfair' feedback.

Hypothesis 1b: Advantaged team members who see a factually inaccurate 'fair play' feedback rate the author of the 'fair play' feedback more favourably than their peers who see a factually accurate 'fair play' feedback.

Hypothesis 2 considers the hypothetical scenario in which a person who sides with the *other* team either makes false claims or does not. In this case, it is hypothesized that in-group members will notice (and punish) the false claims in the other camp:

Hypothesis 2a: Disadvantaged team members who see a factually inaccurate 'fair play' feedback rate the author of the 'fair play' feedback more favourably than their peers who see a factually accurate 'fair play' feedback.

Hypothesis 2b: Advantaged team members who see a factually inaccurate 'unfair' feedback rate the author of the 'unfair' feedback more favourably than their peers who see a factually accurate 'unfair' feedback.

Hypothesis 3 zooms in on the groups that were randomly assigned to see false facts in the feedback that sided with their group (and no false facts
in the other person's feedback).

Hypothesis 3a: Those on the disadvantaged team who are exposed to a false 'unfair' feedback and an accurate 'fair play' feedback will rate the author of the (false) 'unfair' feedback more favourably than the author of the (accurate) fair play feedback.

Hypothesis 3b: Those on the advantaged team who are exposed to a false 'fair play' feedback and an accurate 'unfair' feedback will rate the (false) 'unfair' feedback more favourably than the author of the (accurate) fair play feedback.

2.3.7 Results

The lab experiment was fielded to a convenience sample sample of 277 subjects at ESSEXLab, Colchester, UK in the summer of 2018 (June-Sep). ³ Subjects were recruited from ESSEXLab's participants' pool. 48 per cent were undergraduate students; 62 per cent female, 54 per cent British. ⁴

³Results presented here exclude a first round with 264 participants, conducted June 5th-7th, 2018. The first round showed the disadvantaged group was slightly more tolerant of facts by the person who recognized their in-groups' hardship than by the person who called the game 'fair play'. In contrast, the advantaged game was slightly more likely to tolerate false facts in the 'unfair' feedback than in the 'fair play' treatment. However, group differences were far from statistical significance. Post-experimental feedback from individual participants suggested that the null results stemmed from cognitive overload: the treatment, that is, the two feedback pages were fairly long (around 11 lines), and many participants only skim-read them. What is more, there was no 'back' button so that participants who did not read the feedback pages were unable to refer back to them as they were answering questions about the two feedback givers. Two measures had been taken to ensure participants knew which feedback they were rating: First, respondents were shown each feedback before they were asked questions about ("As you will remember, this was feedback 1/2. (...)". Second, each feedback was assigned a colour, which was not associated with any particular political party: The 'fair play' feedback and all questions about it were shown on a teal-coloured background, whereas the 'unfair' feedback and all questions about it were shown on a beige background. Because these measures proved to be insufficient two further measures were taken before the study was re-run: First, to encourage respondents to read the feedback in the first place both feedback pages were shortened; important sentences were printed in bold, and the feedback points that contained false or exaggerated claims were formatted as bullet points. Second, to ensure respondents knew which feedback they were rating and could refer back to it if they had not read it they were shown the feedback on two of the three pages on which they were asked questions about it. Hence, when answering if they generally agreed with a feedback and when assessing how true or false each individual claim was respondents were able to scroll down to view the respective feedback, introduced as: "For your reference, this was feedback 1:". The feedback was not however reprinted on the second page, asking if 'The points this person makes are factually accurate' and if 'This person is a good representative of my team.' These questions were intended to measure to what extent subjects pre-contiously overlooked the false claims. In this case, showing the feedback again would have allowed respondents to check how factually accurate the feedback was and, thereby, would have defeated the purpose of the question.

⁴34% postgraduate students, 10% administrative staff, 3% faculty, 84% affilated with the University of Essex, mean age=31 years (min age=19 years, max age=80 years, two thirds in their twenties), 54% British (15% Eastern European, 6% Southern European, 4% Western European).







(B) Rating the 'fair play' feedback

FIGURE 2.4: Disadvantaged team



FIGURE 2.5: Disadvantaged players who saw false claims in the 'unfair' feedback (group 4) rating the (false) 'unfair' feedback and the (correct) 'fair play' feedback

Table 2.1 shows the number of participants in each group: 133 participants were randomly assigned to the advantaged team. 69 of them saw false claims as part of the feedback that called the game 'fair play'; 64 saw false claims as part of the feedback that called it 'unfair'. 144 participants were assigned to the disadvantaged team. 83 of them saw false claims in the 'fair play' feedback; 61 saw false facts in the 'unfair' feedback. ⁵

In all sessions, the average payoffs for advantaged players exceeded that of the disadvantaged. Advantaged team players earned GBP 8.90 (min=GBP 5.75, max=GBP 11.80) on average; disadvantaged players earned GBP

⁵(Due to a mistake one session was larger with 9 participants on the advantaged team and 10 on the disadvantaged team.)



(A) Rating the 'fair play' feedback





FIGURE 2.6: Advantaged team



FIGURE 2.7: Advantaged players who saw false claims in the 'unfair' feedback (group 1) rating the (false) 'unfair' feedback and the (correct) 'fair play' feedback

4.80 (min=GBP 2.25, max=8.40)). 67 A large majority (260/277) of players correctly remembered that on average, Team A had received a higher payoff than Team B.

Did players notice the false claim? Did they punish it? Three questions were asked to assess to what extent players noticed, or overlooked this false claim: On a first page, respondents were asked if they agreed: 'Generally speaking, do you agree with the author of this feedback?' (1='Strongly Disagree', 2='Disagree', 3='Slightly Disagree', 4='Slightly Agree', 5='Agree', 6='Strongly Agree'). Next, respondents were to rate

⁶Calculated payoffs differed from paid payoffs in two instances: First, because of a minimum payoff policy respondents who earned less than GBP 5.00 (n=78) were paid GBP 5.00. Second, due to maintenance problems two studies started 5 and 10 minutes late, respectively. Subjects in those studies received a GBP 1.00 or GBP 2.00 top-up to compensate for their extra time.

⁷The difference in payoffs varied quite substantially: In two sessions, it was less than 1 GBP (0.21 and 0.80, respectively). This was mainly due to the fact that the game manipulated the *chances* of drawing easier or more difficult questions, not the numer of questions.

the author of each feedback along three dimensions: education (a distractor), factual accuracy ('The points this person makes are factually accurate'), and suitability as a team representative ('This person is a good representative of my team.'). All three questions were measured on a scale from 0 (not at all) to 100 (very). Due to space contraints the results reported here below focus mainly on suitability as a team representative; the full results are shown in the appendix.

Figure 2.4 shows results for the hypotheses for the disadvantaged team (1a, 2a, and 3a). Figure 2.6 show results for the advantaged team (1b, 2b, and 3b). Beige bars represent ratings of the 'unfair' feedback; teal-coloured bars represent ratings of the 'fair play' feedback.

The main question this project sought to answer was: Do players overlook false facts that appear in the feedback that sides with their team? As hypothesized, they do. Figure 2.4 shows evidence for the disadvantaged team: It compares how the disadvantaged team rated the accurate version of the 'unfair' feedback ("no false facts") and the inaccurate version of it ("false facts"). For each evaluation the graph shows two bars, representing those who saw (and rated) the accurate version ("no false facts") and the inaccurate version of it ("false facts"). As expected, there is absolutely no difference between the two groups. Whether a disadvantaged member saw an accurate version of the 'unfair' feedback or a version with false claims in it had absolutely no effect on how they evaluated it. Those on the disadvantaged team who rated an 'unfair' feedback that contained false claims gave it an average rating of 66/100. Their peers who rated the factually accurate version of the same feedback gave it a mere two points more (68/100). The difference in means was far from statistical significance (t(131)=-0.53, p = 0.59). (The results for the other two dependent variables are shown in the appendix.)

The same was true for the advantaged team: As shown in figure 2.6 those on the advantaged team who rated a 'fair play' feedback that contained false claims evaluated it just as favourably as their peers who had seen the factually accurate version of it. Those on the advantaged team who rated a 'fair play' feedback that contained false claims gave it an average rating of 45. Their peers who rated the factually accurate version of the same feedback gave it an average representation rating of 47. Again, this two point-effect of factual accuracy was far from significance (t(123)=-0.24, p = 0.80).

Hypotheses 1a and 1b are therefore confirmed. Factual accuracy has no effect on how players rate the feedback that favours their team. This finding implies that the effect of lying – or *not* lying – on individuals' chances as being seen as a good representative are negligeable.

What about false claims in the *other* camp? Do players notice and punish false facts that appear in the feedback that favours the *other* team? The disadvantaged do (at least a little); the advantaged do not. The middle graph in figure **??** shows evidence for the disadvantaged team, comparing ratings of the 'no false facts' and the 'false facts' version of the feedback that favoured the other team. It shows that the disadvantaged were generally less agreed with the author of the inaccurate 'fair play' feedback; they rated them as less accurate and an (even) worse representative than the author of the accurate version of it. However, only the difference

in general agreement reached statistical significance at a 95 per cent confidence interval; the difference in accuracy ratings reached significance at a 90 per cent confidence interval (see appendix). The difference in representation ratings (42 v. 48/100, on average) did not reach significance (t(135)=-1.23, p = 0.22). This is suggestive, but not conclusive evidence – Hypotheses 2a cannot be confirmed.

When it comes to the advantaged team the data is clear : Contrary to hypotheses there is no evidence that the advantaged punished the person who favoured the other team for lying. Advantaged team membes gave the 'unfair' feedback accuracy ratings in the lower 60s and representation ratings in the mid-40s – regardless of whether or not that feedback contained false claims . Hence, hypothesis 2b is rejected.

Hypotheses 1 and 2 examined counterfactual situations in which a potential representative either makes false claims or does not make false claims. Hypothesis 3 examines a more common situation: one in which one person lies and another person tells the truth. If the lying person is the one who favours an out-group then in-group members have every incentive to notice the lies. But what if the lying person is the one who favours one's in-group? Half the players in this experiment were nudged to be motivated reasoners: Those who saw false claims as part of the feedback that favoured their team (groups 1 and 4 in table 2.1). If motivated reasoning extends to overlooking false claims then these ought to turn a blind eye to the false claims and rate the feedback that favours their own team more favourably than the feedback on the other side. That is exactly what they did. As shown in graph of figure 2.5, disadvantaged players who saw a lying feedback favouring them and an honest feedback favouring the other team showed significantly higher levels of agreement with the former than the latter. They rated it as slightly (but not significantly) more accurate, and deemed its author a much better team representative than the author of the feedback that favoured the opposing team ($M_{correct 'fair play'}=48/100$, $M_{false 'unfair'}=66/100$, t(87)=-3.66, p = 0.00.).

The same was true for the advantaged team: As shown in figure 2.7 those who were nudged to be motivated reasoners, that is, those who were confronted with a 'fair play' feedback that contained false claims and an 'unfair' feedback that did not contain any false claims evaluated the former a better team representative than the latter: $M_{correct 'unfair'}=45/100$, $M_{false 'fair play'}=64/100$, t(132)=4.58, p = 0.00. This confirms hypotheses 3a and 3b.

2.4 Discussion

This laboratory study was designed to test the effect of boosting in-group status on tolerance of false facts among the in-group. It asked two main research questions: First, do groups that are economically disadvantaged relative to other groups overlook false facts by leaders who recognize their disadvantage and who promise to improve their group's relative status? Second, do advantaged groups that fear losing their advantage overlook false facts by leaders who justify their group's advantage and promise to maintain it? The core of the experiment was a pub quiz-like geography quiz in which the chances of receiving easy or difficult questions were manipulated: One (advantaged) team had a higher chance of receiving easier questions, and therefore a higher average payoff than the other (disadvantaged) team. After the quiz, respondents were exposed to two people's feedback. One person called the quiz 'unfair'; the other called it 'fair play'. One of the two feedback pages was randomly assigned to include a false claim, complaining that "some of the places don't even exist", and an exaggerated claim, complaining that the thirty second time limit was too short ('barely enough time to read the question') or, if assigned to the 'fair play' feedback, too long ('10 seconds would have been plenty').

Hypothesis 1 examined the (counterfactual) scenario in which a person who favours one's in-group either makes false claims or does not make any false claims. As hypothesized, it didn't affect their standing: Both the disadvantaged and the advanted agreed with the feedback that favoured their team, evaluated their author as a good team representative and stated that 'the points this person makes are factually accurate' – regardless of whether they were accurate or not. This finding confirms recent findings that loyalty trumps honesty Hildreth and Anderson (2018) Implications for political campaigns are sobering: If grossly exaggerating or making false claims does not affect perceived leadership qualities then politicians have no incentive *not* to lie. As long as they are loyal to their target in-group and as long as their target in-group is large enough to win them the election they will get elected.

Hypothesis 2 examined the (counterfactual) scenario in which a person who favours the *other* team either makes false claims or does not make any false claims. In this case, it was hypothesized that players would notice and punish those false claims. Results are inconclusive. The disadvantaged team showed some signs of sanctioning lies in the opposing team; the advantaged team showed no signs at all. On average, the disadvantaged team 'slightly agreed' with the accurate 'fair play' feedback. They 'slightly disagreed' with the inaccurate 'fair play' feedback. Here, the difference in means was statistically significant (see results in the appendix). When asked to assess how accurate the feedback was they gave the accurate feedback a score of 60, and the inaccurate feedback a score of 50. Here, the difference was significant at a 90 per cent confidence interval (p=0.07, see appendix). Yet when it came to evaluating how apt the person was to act as a team leader respondents seemed to care a little less about how honest they were: The disadvantaged rated the author of the incorrect 'fair play' feedback as slightly, but not significantly less accurate and as a slightly, but not significantly worse team representative than the author of the accurate version of the 'fair play' feedback.

Hence, hypothesis 2 could not be confirmed. For two reasons, these (partially) null findings ought to be taken with a grain of salt: First, there were clear signs that the disadvantaged rated the feedback that favoured the advantaged less favourably when it contained false facts. The fact that group differences did not reach significance across all indicators ought to be seen as a reflection of the small sample size. A larger sample size would most probably have pushed the them toward significance. Second, and more importantly, there was a flaw in the design. Somewhat unfortunately, the false claims were not 'neutral'. The statement that 'some of these places don't even exist' was clearly false. However, it implied that the game was rigged – and members of the disadvantaged team must have felt cheated. This claim recognized their hardship – even if it appeared as part of the 'fair play' feedback. Future research ought to investigate reactions to more 'neutral' false claims.

Similarly, the fact that the advantaged team let false facts in the 'fair play' feedback slide should not be interpreted as final evidence that status threat does not affect the way individuals evaluate a person who favours the out-group. It is likely to the failure of this study to credibly threaten team A's status.

Hypothesis 3 zoomed in on those who were nudged to be motivated reasoners. The disadvantaged behaved as expected: They rated the lying 'unfair' person as a far better team representative and as slightly more accurate than the honest 'fair play' person. If generaliseable to election campaigns this is consequential. In our experiment voters did not have to choose between two candidates – in real elections they do. This makes it all the more likely that they will overlook false facts. Therefore, these results suggest that if voters who identify with a low-status group have the choice between a candidate who wants to raise their group status but lies and a candidate who wants to maintain the status quo and is honest then they will overlook the false claims and vote for the former.

The high-status group behaved in the same way: Advantaged players who saw a 'fair play' feedback that contained false claims rated the author of that 'fair play' feedback as a better team representative (albeit not as more accurate) than the author of the (correct) 'unfair' feedback. This, too, is alarming. When voters who identify with a high-status group have the choice between a candidate who promises to maintain their relatively higher status and a candidate who wants to level out status differences they, too, will overlook the false claims and elect the candidate on their side.

Two limitations ought to be noted. The main caveat concerns the advantaged group. This study was designed to create a low status group and a threatened high-status group. It did create status differences. The low status group's low status was built up during a twelve-question quiz that culminated in a very low payoff – on average, four pounds lower than the other team's payoff and, in many cases lower than the minimum payoff. (Which made them eligible for the lab equivalent of welfare benefits; i.e. the un-earned minimum payoff of GBP 5.00). The advantaged group's higher status was similarly consolidated: Their success was consolidated: They beat the other team across most of the twelve questions, and made more than they may have expected to make in a 30-minute study. However, the attempt to threaten the high status group's relatively higher status seems to have been less successful. The status threat depended entirely on a few sentences in the 'unfair' feedback: The demand to top up Team B's payoffs was assumed to act as a status threat to Team A. This is unlikely to have worked. The advantaged group had nothing to lose: their own payoffs were safe and even in the unlikely event that the experimenter had topped up team B's payments they were still part of the winning team. It is very likely that the threat to their higher status did not felt very real. Future studies ought to create a more realistic sense of status threat, ideally coming from above, not from a fellow participant.

A second caveat is the small sample size. Larger studies are needed to

detect small effects and heterogenous treatment effects. In particular, the sample size was too small to examine the effect of status-boosting speech among players who according to social identity theory were most likely to engage in violence: those who strongly identified with their team, who thought status differences were illegitimate and who believed it was possible that the status differences that resulted from the quiz might be overturned. (I would expect these players to be particularly immune to false claims in the feedback on their side.)

2.5 Conclusion

This study contributes to research in an important way: It showed that low status makes individuals vulnerable to overlooking factual inaccuracies in statements from people who sympathize with their group. In light of the rising levels of inequality in Western Democracies and the rise of populist leaders appealing to voters who believe to be socially or economically disadvantaged this is important to know. If feeling disadvantaged makes voters vulnerable to overlooking false facts in campaigns that promise to raise their status then it is unlikely that the present tide of populist politicians with agendas that are based on facts dubious quality will turn.

This laboratory experiment was designed to test a causal mechanism. As with any laboratory study the flipside of isolating the mechanism in a controlled lab environment is that the setting deviates from a real-world setting. For external validity, the next step in this research avenue is to test this mechanism in a setting that is closer to the everyday life of disadvantaged voters. For instance, online survey experiments could manipulate low group status and then expose respondents to a hypothetical election in which one campaign recognizes their hardship and seeks to level out inequality. It could vary the factual accuracy of their campaign promises.

In addition, more research is needed on the effect of status threat. Both laboratory and survey experiments could be used to test if status threat makes voters vulnerable to overlooking false facts: Do voters who believe to have lost in status or who fear losing status overlook false facts from politicians who promise to move their group up again? Here, too survey experiments around hypothetical elections are a promising route for future research. Future studies could try to nudge feelings of ethnic or global status loss and investigate the effect of that feeling on tolerance of false facts in political campaigns.

Future studies should also tease out the effect of status threat and status boost: Is threatening status sufficient to make voters overlook false claims? Is boosting status sufficient? How do the two interact? What are the minimal conditions under which voters overlook false facts? Could in-group membership or siding with the right team be a sufficient condition? In this case, If none of the teams had been disadvantaged and if both teams had had the same payoff would team members still have overlooked false claims from a person who sided with their team? If so, how far does this willingness stretch: Would minimal (e.g. Klee/Kandinsky) groups do the trick, too? In addition, future studies should investigate the effect of *perceived* status differences: Do subject who *believe* to be disadvantaged but are not actually disadvantaged also overlook false claims by leaders who promise to raise their group status or fight off threats to it? Exploratory analyses from this study suggest that the driving factor is not actual disadvantage but perceived disadvantage: Because this study varied the likelihood of receiving difficult or easy questions (but not the number of difficult or easy questions) the level of advantage and disadvantage differed across the games. Both the *somewhat* disadvantaged and the *very* disadvantaged overlooked false claims in the feedback that favoured their team. (This question could easily be answered in a variant of these experiments in which both teams receive the same number of easy and difficult questions; and a flat payoff.)

Finally, scholars ought to investigate the extent to which voters tolerate false claims coming from a person who recognizes their struggles: If they are willing to support someone who makes false but benign claims are they also willing to support them if the claims become less benign? For instance, will they continue to support their candidate if the candidate makes derogatory comments about out-group members, if they condone violence against out-group members, or, ultimately, if they incite violence?

2.6 Acknowledgements

This research was supported by an ESSEXLab Seedcorn grant as well as research funding by the Eastern Academic Research Council. I gratefully acknowledge this support.

Chapter 3

A Post-Truth Public? Investigating the Mechanisms of Resistance to Factual Correction

3.1 Abstract

Many recent electoral events have been characterised by false claims which, despite abundant fact-checking by experts were apparently widely believed. This led to much talk about 'post-truth' politics. An extensive literature confirms that political misperceptions often resist correction. But how strong is this tendency? We conducted a representative survey experiment in Britain (N=2900) testing belief in common misperceptions about immigration. We followed the classic setup of misperception-correction studies but added a twist: First, we identified false beliefs and provided expert information countering one of those false beliefs. Second, we approximateed 'real world' conditions, where expert information is rarely the final word. Respondents in our treatment group saw a 'comment' saying 'I would take these statistics with a big pinch of salt'. We varied two things: who made the comment (an authoritative or a non authoritative source) and the reason they gave to take the statistics with a pinch of salt. Our results show that the fact checks worked: Whenever the expert had the final word average perceptions shifted from just below 'definitely true' to just above 'don't know'. However, the post-truth comments worked, too: Whenever the commentator had the final word average perceptions remained well within the 'true' side of our 'true' to 'false' scale, albeit closer to false than before they had seen the fact check. Our findings illustrate the relevance of fact-checking false claims in a posttruth era.

3.2 Introduction

'Facts don't work.' This was the advice Arron Banks, co-founder of the Leave.EU campaign received from the political strategists Goddard Gunste. They heeded the advice. In an interview with the Guardian, Banks explained how the Leave campaign differed from the Remain campaign: 'The Remain campaign featured fact, fact, fact, fact. It just doesn't work. You've got to connect with people emotionally. It's the Trump success.' (Booth, Travis, and Gentleman 2016). Throughout the campaign, newspapers, public broadcasters, and independent charities fact checked the arguments, the statistics, and the projections made on both sides. In particular, fact checkers went to great lengths to correct one false claim that adorned the bus in which the Leave campaign toured the country: 'We send the EU £350m a week. Let's fund the NHS instead.' Numerous

journalists explained Margaret Thatcher's rebate, and laid out why even the amount Britain actually transferred to the EU could not simply be re-directed to the NHS. Even the UK Statistics Authority intervened, saying the figure was misleading (Henley 2016; Reuben 2016; Arnold 2016; Full Fact 2017). Nonetheless, the false figure remained on the bus, leading leave campaigners continued to repeat the false claim (Stewart and Asthana 2018), and on 23 June 2016, Britain voted to leave the European Union. Two years after the referendum, a representative survey found that 42 per cent of those who had heard the 350 million claim believed it was true (Stone 2018). The unwillingness of the leave campaign to refrain from using an incorrect statistic and the share of the public who think this number is correct points to a puzzle: If politicians at the highest levels of government disregard expert advice and encourage voters to do the same (Watts 2016) then who will voters trust? Is fact checking worth the effort?

Francis Fukuyama described the 'post-fact world' as a world in which 'virtually all authoritative information sources are called into question and challenged by contrary facts of dubious quality and provenance' (Fukuyama 2017). This paper is about the citizen consumers of authoritative information called into question by contrary facts of dubious quality and provenance. It contributes to a nascent body of literature investigating the effect of a post-truth context ¹ on political behaviour. We ask: How does exposure to post-truth comments shape the way that we process challenging factual information? If expert information suggests we

¹@Rob & John: I am still not sure how to call this. Post-truth surroundings? Context? Arguments? Era?

are wrong but others put forward essentially fact-free arguments countering the expert information whom do we believe? The expert whose information is accurate but challenges our beliefs or the non-expert whose information is of dubious quality but affirms our beliefs? And what if the non-expert is an authoritative person?

A large body of literature in the motivated reasoning tradition has shown that humans are biased information processors, pursueing not only accuracy goals but also directional goals: "People are more likely to arrive at conclusions that they want to arrive at" (Kunda 1990). We speculate that where accuracy goals compete with directional goals post-truth surroundings might tilt the balance toward the latter. In other words, we expect post-truth comments to nudge individuals to ignore the expert and retain their old beliefs. We explored these questions using a representative-sample survey experiment in Britain (N=2,936). We chose a topic British voters on both sides of the Brexit debate feel strongly about: immigration. We aimed high, investigating the effect of not just any fact check but a fact check that respondents thought was true and which affirmed their stance on immigration. We began by asking respondents to rate four false claims about immigration on a scale from 0 (Definitely false) to 6 (Definitely true). Two false statements were catered to respondents who held pro-immigration opinions and two were catered to respondents who held anti-immigration opinions. We then directed respondents to a fact check correcting the false claim they had rated closest to 'true'.

We had three main goals: First, we sought to find if the fact checks affected respondents' perceptions of how true or false the false claim were. Hence, we asked respondents to re-rate the false claims after they had seen the fact check. To capture smaller changes in perceptions we departed from the classic dichotomous 'true' or 'false' question, and asked respondents to rate the claims on a scale from 0 (Definitely false) to 6 (Definitely true). Second, we explored how robust the effect of the fact check (if any) was to post-truth comments. To that end, we showed four out of five respondents an extra piece of information before we asked them to re-rate the false claims: a 'comment', saying "I would take these statistics with a pinch of salt'. In a 2*3 design we varied the authoritativeness of the post-truth commenter ('Professor of Economics' / 'Blogger') and the reason why they said they would take the statistics with a grain of salt. In one variant, the commenter argued that the expert was probably biased, in another variant they pointed out that the expert's statistics did not match people's everyday experience. In the third variant the comment did not doubt the expert information at all, saying that it was best to trust one's feelings even if it looks like the facts are different. Our final goal was to explore evidence of a post-truth mindset. At the end of the survey, we asked a few questions explicitly testing whether respondents agreed with classic post-truth arguments.

Our results showed that fact checks successfully reduced misperceptions. On average, respondents placed the false facts closer to 'false' after they saw the fact-check than before they saw it. This was the case for respondents on both sides of the immigration debate (those wanting to reduce immigration and those wanting to maintain current levels of immigration or increase them). The fact check did not cause a complete u-turn: respondents did not shift from rating the false claim as 'true' to rating it as 'false'. But the effect was substantial: It reduced misperceptions by almost two (1.75) points on the seven-point scale. Before respondents saw the fact check, the average veracity rating was just a point below '6 - Definitely true' (5.1). After they saw the fact check, it was just above the midpoint between '0 - Definitely false' and '6 - Definitely true' (3.4). Hence, the fact check shifted respondents from being almost absolutely sure the false fact was true to almost admitting they did not know if it was true or false.

As hypothesized, exposure to post-truth comments mitigated the effect of the expert statement on respondents' perceptions of how true or false the false claims were. The post-truth comment canceled out a small but significant part of the fact check: about half a point, or a third or the 1.75 point effect. Contrary to our expectations, authoritativeness did *not* moderate the effect of the post-truth comment: Our post-truth professor was not any more convincing than our post-truth blogger. The reason why they said they would 'take those statistics with a pinch of salt' did not seem to matter either. Arguing that an expert was biased was just as effective as arguing that their statistics did not match people's everyday experience or that their statistics were probably right but when in doubt, it is best to trust one's feelings. Overall, these findings are reassuring: Expert opinion mitigates misperceptions *even if* the expert does not have the 'final word'.

Nonetheless, we found a level of agreement with post-truth arguments that supports the notion a post-truth public. We asked respondents to rate the statement they had rated as true and read a fact check about on a scale from '0 - Purely a matter of fact' to '6 - Purely a matter of opinion'.

Among the control group (that is, among respondents who skipped the post-truth comment) only 17 per cent evaluated the fact-checked statement as 'purely a matter of fact'. 35 per cent deemed it closer to a matter of opinion than a matter of fact (4, 5, or 6); 13 per cent saw it as 'purely a matter of opinion'. Next, we asked if the statistics were in line with what they believed. If the answer was 'no' we followed up, asking respondents to choose one of three statements to describe where they stood. About a third (32%) said 'The statistics made me change my mind.' Another third (36%) said 'The statistics are probably wrong.' A third third (32%) said 'The statistics are probably wrong.' A third third (32%) said 'The statistics are probably defined up, asked respondents that 'It's OK to disagree with the facts if that's what you believe.' 46 per cent agreed. 18 per cent *strongly* agreed.

3.3 Reactions to uncongenial expert information

A large body of social science literature testifies to the fact that misperceptions tend to resist correction. It has long been known that individuals often continue to rely on information even after it has been retracted or proven wrong (Ross, Lepper, and Hubbard 1975; Carretta and Moreland 1983; Wyer and Budesheim 1987). Individuals are particularly unlikely to accept information that challenges factual beliefs they feel strongly about or that are tied to their identity (Kull, Ramsay, and Lewis 2003; Thorson 2016; Jacobson 2010; Hart and Erik C. Nisbet 2012; Nyhan, Reifler, and Ubel 2013; Nyhan, Reifler, Richey, et al. 2014; Kraft, Lodge, and Taber 2015; Thorson 2015; Rampell 2016). The drivers of this resistence to counter-attitudinal facts are well-known: Human beings are motivated reasoners. We engage with new information with a motivation to reach a particular conclusion: We pursue not only accuracy goals but also partisans, or directional goals (Kunda 1990; Kahan 2016b; Kahan 2016a; Lodge and Taber 2013; Kraft, Lodge, and Taber 2015). Brain imaging studies suggest that a proclivity to motivated reasoning is hard-wired in our brains. For instance, Westen et al. (2006) found that the areas of partisans' brains responsible for reasoning shut down when they looked at information that was threatening to their own candidate. In contrast, when they were considering consonant information, the emotional areas of their brains lit up. The results is 'fact polarization': Voters find themselves divided over questions of fact – an issue Kahan (2016b) called 'one of the signature features of contemporary democratic political life'.

Our study investigates if exposure to post-truth comments exacerbates fact polarization. The idea is simple: If information that challenges political opinions sets off an internal battle between accuracy and directional goals then seeing someone else doubt the statistics ought to tilt the balance in favour of the directional goals. As long as individuals have a strong enough motivation to reject the evidence then even weak arguments, no arguments, or fact-free arguments ought to work as a nudge to maintain one's prior beliefs.

Our study contributes to a growing body of literature on poltical misperceptions. Much of this research has focused on finding ways to communicate political information more effectively (see Seifert 2002; J. Cook and Lewandowsky 2011; Nyhan and Reifler 2012; Nyhan 2012; Lewandowsky,

Ecker, et al. 2012; Nyhan and Reifler 2013 for summaries of research findings). Less is known about how individuals react to conflicting information about factual matters. In particular, few studies have examined reactions to factual, but difficult-to-digest information that is followed up by non-factual, but easier-to-digest information. Our study builds on a few studies in the field of science communication, which examined perceptions about vaccines and about anthropogenic climate change (ACC). A major finding of the past few decades concerns the negative side effects of journalistic 'false balance': If the weight of evidence points to one conclusion then reporting on views that are not supported by overwhelming evidence (or, as Dearing (1995) put it, giving a platform to 'maverick scientists') can lend credibility to the minority view and reduce confidence in expert consensus (Dearing 1995; M. T. Boykoff and J. M. Boykoff 2004; C. E. Clarke 2008; C. E. Clarke et al. 2015; Dixon and C. E. Clarke 2013; Dixon, Mckeever, et al. 2015; Kortenkamp and Basten 2015; McCright, Charters, et al. 2016; Koehler 2016; Linden et al. 2017).

McCright, Charters, et al. (2016) provided the first experiment directly testing to what extent reporting on the views of ACC denial activists affected belief in human-caused climate change and views on climate policies (McCright, Charters, et al. 2016, p.82). The authors showed respondents a fabricated article about a report from the Intergovernmental Panel on Climate Change, which, according to the article provided 'unequivo-cal evidence that climate change is happening now, is caused by humans, and is producing harmful societal impacts.' Among other things, the authors varied whether the article gave a voice to those denying ACC. For some respondents (but not all) the article ended with an additional

paragraph, saying that 'most conservative leaders and Republican politicians believe that so-called climate change is vastly exaggerated' and that 'some scientists ... are quick to point out that the Earth hasn't actually warmed in the last decade'. The authors found clear evidence of motivated reasoning. Right-leaning respondents reported weaker belief in the reality of ACC. They were also more affected by exposure to climate change denying views. (Among extremely conservative respondents seeing the extra paragraph about ACC denyers tipped average perceptions of the veracity of a number of factual statements on ACC from just above 'I don't know', tending toward 'agree' to just below 'I don't know', tending toward 'disagree'. Left-leaning respondents showed much higher levels of belief in ACC, regardless of whether or not the article included critical voices.

In a similar study, the effect of exposure to science-denying voices was even stronger: Linden et al. (2017) investigated whether information about scientific consensus on climate change survived counter-messages. They did not. The authors showed respondents a pie chart stating that '97% of climate scientists have concluded that human-caused climate change is happening'. Next, they showed some respondents (but not all) a (real) petition urging the U.S. government to reject the Kyoto protocol because 'there is no convincing scientific evidence that human release of carbon dioxide ... will ... cause catastrophic heating of the Earth's atmosphere ...'. Here, the main dependent variable was not belief in human-cause climate change but perceptions of scientific consensus. Among those who had only seen the initial pie chart, perceived consensus shifted from 71%, on average, to 90%, on average. That effect was *completely* canceled out by

	Authoritative	Not authoritative
	(Professor)	(Blogger)
Biased source		
Personal experience		
Alternative facts		

TABLE 3.1: Experimental Design. (Treatment groups only.Control group receives no second opinion.)

the countermessage: On average, those who had seen the pie chart, followed up by the petition stated that 73% scientists agreed that humans caused climate change, both before and after they read the article.

Our research goes beyond these studies in a few important ways: First, we conduct research in a country other than the United States: the United Kingdom. Second, we assess reactions to a topic that both right and left-leaning individuals tend to feel strongly about and hold misperceptions about: immigration. Third, and most importantly, we assess reactions to a post-truth, rather than a denial counter-frame. We are unaware of any studies testing counter-messages that do not deny facts but provide other, essentially fact-free reasons to disregard them. Our study seeks to contribute to filling this gap.

3.4 Method

To test the effect of post-truth comments on citizen reactions to inconvenient expert information we looked for a topic British people were misinformed about and felt strongly about so that they would have a motivation to want to disregard factual information. We chose immigration. The issue was salient and polarizing, as the United Kingdom was (and is) still recovering from the Brexit referendum, in which immigration played a major role (Goodwin and Milazzo 2017; H. D. Clarke, Goodwin, and Whiteley 2018; Hobolt 2016). We knew that significant parts of the population believed in false claims exaggerating the negative side-effects of immigration or denying the positive (B. Duffy 2019). We suspected that voters on the other side of the immigration divide would be equally likely to believe in false claims denying the negative side-effects (or exaggerating the positive). We pre-tested a number of false claims on both sides of the debate and found fairly high levels of belief across the political spectrum. For the main study we chose to focus on two false claims that were congenial to those with anti-immigration attitudes ('There has been a sharp rise in the number of people applying for asylum in the UK in the past ten years.'; 'European immigrants receive more in benefits and services than they pay in taxes.') and two that were congenial to those with pro-immigration attitudes ('There has been a sharp rise in the number of people applying for asylum in the UK in the past ten years."; "European immigrants receive more in benefits and services than they pay in taxes.") In the following, we refer to these claims as 'pro' and 'anti-immigration claims'. We designed an experiment in which we tested belief in each of these four claims; then provided a fact check countering one of them, provided a post-truth counter-message for some, and re-tested belief in the false claim.

The experiment consisted of seven stages. In **stage 1**, we asked a few background questions, including questions about immigration: Should Britain allow more immigrants, fewer immigrants, or should the number

of immigrants stay the same? We also asked how important respondents thought this issue was to the country and to them. In **stage 2**, we tested belief in false claims about immigration. Respondents were presented with a list of eight statements, which they were asked to rate on a scale from 0 (Definitely false) to 6 (Definitely true). Four of these statements were the above-mentioned misperceptions about immigration. The remaining four were unrelated distractor items, including 'Fracking causes earthquakes' and 'England's plastic bag usage dropped 85% since the 5p charge was introduced.' In **stage 3**, respondents were directed to a fact check challenging the false claim they had rated closest to '6 - Definitely true'. (If veracity scores on two or more items were tied respondents were randomly assigned to see one of the respective fact checks.)

The four fact checks were fabricated but the content was correct. They were designed to be as effective and as realistic as possible, following best practices and expert advice on countering misinformation (Nyhan and Reifler 2013). Copying the format of popular fact checking sites in the UK such as the BBC's Reality Check and Full Fact the fact checks began with a short summary of the finding ('This statement is false.') and then provided a more detailed description, quoting offical statistics, and including a graph. The false claims were not repeated; instead, correct claims were affirmed. Wherever the corrective information may have left a causal gap in the narrative an alternative causal explanation was provided (Schwarz et al. 2007; Lee, Kim, and Schwarz 2015b, see). For instance, to counter the claim that the UK had experienced a high influx of asylum seekers in the past ten years the fact check conceded that that there had indeed been a sharp rise of asylum seekers crossing into

Europe. However, most of them never reached the UK but stayed in Germany, Sweden, or Hungary. All fact checks were attributed to the same authoritative source: an Oxford Professor of Economics who holds a PhD from Harvard University and who previously worked as a consultant for the Office of National Statistics (ONS). The expert had a generic English name (Professor Richard Clarke). The fact checks included a small picture of a middle-aged professor in front of a black board. After reading the fact check respondents answered a short battery of demographic questions designed as distractors.

In stage 4, we tested how resistent factual corrections were to fact-free dispute. At this point, the sample was divided: The control group (25% of the sample) skipped stage 4 and moved directly to stage 5. The treatment group (75% of the sample) saw a comment about the fact check they had just read. Across all variations of the comment the commenter was introduced as 'David Williams' with a small picture of a white, greyhaired man sitting on an outside chair wearing a beret and holding a cigar in his right hand. Across all variations, Mr. Williams' comment started with the words, 'I would take these statistics with a big pinch of salt'. In a 2*3 design, we varied the authoritativeness of the source and the content of the message. David William was either described as a 'Professor at the London School of Economics' (authoritative) or as a 'Blogger' (not authoritative). He gave one of three reasons to take these statistics with a pinch of salt: In the 'Biased Source' condition he wrote, 'The fact that someone is a professor doesn't mean that they don't have an agenda. And we all know that there is a lot of scope to choose and present statistics so that they end up saying just what you want them to say.' In the 'Personal Experience' condition he held that, 'A graph might say one thing but the experience of people's everyday lives could be quite different. And I think that a lot of people reading those statistics will say: that doesn't sound like the world I live in.' In the 'Alternative Facts' condition Mr. Williams did not deny the statistics but, in classic post-truth fashion recommended to ignore them nonetheless: 'There's so much information and so many statistics out there that it can be hard to know what to believe. In that case, I think it's best to trust your instincts even if looks as if the facts are different.'

In stage 5, respondents re-evaluated the corrected claim on a scale from 0 (Definitely false) to 6 (Definitely true). In stage 6, we looked for signs of post-truth reasoning. First, we presented respondents with a second seven-point scale. This time, we asked them to rate the false statements on a scale from 0 (Purely a matter of Fact) to 6 (Purely a matter of Opinion). Next, we asked if the statistics they had seen were consistent with what they believed. If the answer was 'No' we followed up, asking respondents to choose which of the following three statements best described where they stood: a) 'The statistics are probably right but I believe something different.', b) 'I think that the statistics are wrong', or c) 'The statistics made me change my mind.' Finally, we asked if respondents agreed or disagreed that 'it is OK to disagree with the facts if that's what you believe'. In stage 7, we asked two questions about the source of the fact check (the Oxford professor) and the source of the post-truth comment (the LSE professor or the blogger): How accurate did respendents think the information was and how much they would generally trust what each of those people said about the issue of immigration?

The full questionnaire is provided on page ?? of the appendix.

3.5 Hypotheses

We test three main hypotheses. First, we expect our fact check to reduce misperceptions: Fact Check Hypothesis (H1). Exposure to the fact check reduces belief in false claims. In most misperception-correction studies the percentage of respondents who *radically* change their perceptions after seeing corrective information is low (e.g. Nyhan and Reifler 2015a; Nyhan and Reifler 2015b; Wood and Porter 2016; Barrera et al. 2018). Given that our topic is particularly controversial we do not expect our fact checks to shift respondents from rating the false facts as 'true' to rating them as 'false'. We do, however, expect them to reduce the certainty with which respondents believe in false claims. To capture those smaller changes in perceptions we refrained from using the classic dichotomous 'true' or 'false' scale, relying instead on a more fine-tuned scale allowing respondents to place the false statements on a scale from 0 (Definitely false) to 6 (Definitely true). To measure the effect of the fact check we compare the pre and post fact check veracity scores among control group respondents (who saw a fact check only, i.e. no post-truth follow-up). On average, we expect them to rate the false claim somewhat closer to 'false' after they see the fact check than before they saw it.

Our main hypothesis concerns the effect of exposure to post-truth comments: **Post-truth Comment Hypothesis (H2). Exposure to post-truth comments mitigates the effect of the fact check.** To measure if the posttruth comment made a dent in the effect of the fact check we compare the

effect of the fact check (that is, the difference in pre and post-fact check veracity scores) among those who saw the post-truth comment and those who skipped it. We expect it to be larger among the latter, i.e. we expect the fact check to be more effective among those who skipped the posttruth comment. Just how effective the post-truth comment is ought to depend on who makes it. Reputable sources have long been found to be more persuasive, and more effective correcting false perceptions (Guillory and Geraci 2013; Berinsky 2015; Nyhan and Reifler 2015a; Swire et al. 2017; Schaffner and Luks 2018). Hence, we expect a comment coming from an authoritative person (in our case, a professor) to carry more weight than a comment coming from a non-authoritative person (a blogger): Authoritativeness Hypothesis (H2a). Authoritativeness moderates the effect of the post-truth comment. We also vary the content of the message. Two of the messages we test are arguments (albeit weak arguments) to disregard the evidence: one points out that the expert might be biased; another one is based on anecdotal evidence disconfirming the expert's statistics. The third message is essentially fact-free; it does not doubt the statistics but suggests to trust one's instinct regardless of the statistics. If our hypothesis is correct, that is, if people would clutch at any straw to maintain their beliefs then even this message ought to nudge them to disregard the evidence.

Finally, we attempted to gauge the state of post-truth reasoning in the UK: How common is it to 'disagree' with facts, or to denigrate them as matters of opinion? This question is purely exploratory. We rely on our control group (individuals who were not exposed to a post-truth comment) for population estimates of agreement with post-truth statements.

We also assess the effect of exposure to post-truth comments on agreement with post-truth statements. To that end, we compare average agreement with the post-truth arguments among those who saw the post-truth comment and those who did not. We expect to see higher levels of agreement among the former than among the latter. **Post-truth Mindset Hypothesis (H3). Exposure to post-truth comments nudges respondents to a) categorise factual questions as matters of opinion, b) disagree with the facts, and c) agree that it is okay to disagree with the facts.**

3.6 Results

We cooperated with Deltapoll who fielded the survey to a representative sample British citizens across England, Scotland, and Wales in June 2019 (n=2,936). ² 54% of respondents were female, 51% university-educated, age ranged from 18 to 99 (mean age: 45 years, SD=17). Among those who reported to have voted in the 2016 Brexit referendum, 52% voted to leave the European Union. In the 2017 general election, 26% had voted Conservative, 30% Labour, 6% Liberal Democratic, 6% UKIP, and 3% had voted for the Green Party. Half of the sample (47 per cent) thought Britain should take in *fewer* immigrants (values between '-3' and '-1'), a quarter (25%) thought Britain should take in *'many* fewer' ('-3'). About a fifth of the sample (21%) thought Britain should take in *more* immigrants (values between '+1' to '+3', see histogram in figure 1 in the appendix). We

²The original sample size was 2,938; two participants were excluded because they selected 'don't know' or the top option on almost all questions, resulting in a sample of 2,936 respondents.
consider respondents who thought Britain should take in fewer immigrants to have 'anti-immigration' attitudes and respondents who thought Britain should take in more immigrants or who wanted no change to have 'pro-immigration' attitudes. In total, we categorized 1388 respondents (47%) as 'anti-immigration', and 1548 (52%) as 'pro-immigration'.

Both pro- and anti-immigration respondents showed clear signs of motivated reasoning. Generally, anti-immigration respondents thought the two anti-immigration claims were true; pro-immigration respondents thought the two pro-immigration claims were true. Respondents had been shown the fact check corresponding to the claim they had rated closest to '6 -Definitely true'. Almost ninety per cent of respondents saw a fact check about a false fact they thought was somewhere on the 'true' side of the 'false to true' scale, i.e. a claim they had rated as 4, 5, or 6 - Definitely true. 40 per cent saw a fact check about a statement they were absolutely convinced was true ('6 - Definitely true'); about a fourth (23%) saw a fact check about a a statement they had rated as 5, just a point below definitely true, and another fourth (24%) saw a fact check about a statement they had rated as '4', i.e. just a point above the midpoint of the scale. 12 per cent saw a fact check challenging a statement they were not sure about ('3' on the scale from 0 to 6). Less than one per cent received statistics about a statement they had rated as closer to 'false' than to true. Only one respondent correctly rated all four false facts as '0 - Definitely false'.

One aspect of our research proved problematic: We directed respondents to see the fact-check that corresponded to the false fact they rated closest to true. We *assumed* that this would be one of the false facts that confirmed their immigration opinions. This was not always the case: 18 per cent of our *anti*-immigration respondents saw a fact-check that either said that most crime in London was committed by ethnic minorities, not white people, or that immigration had a negative effect on low-paid wages. Even more concerningly, 45 per cent of our *pro*-immigration respondents saw a fact-check that said the number of asylum seekers had not risen at all in last ten years or that immigrants pay more in taxes than they receive in benefits and services. Leaving these respondents in the sample would have inflated the effect of the fact-check because these people were motivated to believe the expert. Similarly, leaving them in the sample would have deflated the effect of the post-truth comment because they had no interest in retaining their old beliefs. Therefore, we restricted the main analyses to the n=1982 respondents who saw statistics that challenged their opinions. This final sample includes 1156 anti-immigration and 826 pro-immigration respondents.

Effect of the fact-check on perceptions about the fact-checked fact

All four fact checks significantly reduced belief in false claims. We looked at respondents in the control group (who were not exposed to a post-truth comment) to assess the effect of the fact check. The most successful fact check was the one countering the (anti-immigration) claim that immigration had risen sharply over the past ten years. It shifted respondents from an average rating of 5.34 on a scale from 0 (Definitely false) to 6 (Definitely true) to an average of 3.16. The least successful fact check was the one countering the (pro-immigration) claim that immigration had no effect on low-paid wages. Yet even that fact check moved respondents from an initial rating of 4.91 to a post- fact check rating of 3.55 (see figure 5 in the appendix). The difference between pre- and post- fact check ratings were significant in all four cases, and the data was pooled for subsequent analyses.

On average, the fact checks made respondents shift 1.77 points closer to 'false' (paired t-test: t(491)=18.705, p=0.00, 95% confidence interval: 1.59-1.96). Before they saw the fact check respondents rated the false claims just about a point under 'definitely true' ($M_{T1}=5.11$, SD=1.02). After they saw the fact check they rated it just above the mid-point of the scale ($M_{T2}=3.37$, SD=1.74, see histogram in figure **??**). In general, it is difficult to change peoples' perceptions. Immigration is an issue that engages strong attitudes. We take the effect of our expert fact check as a success.

Two important factors predicted responsiveness to challenging statistics: gender and attitude strenght. The gender gap appeared post-fact check: Women were much more more responsive to the expert telling them the statement was wrong. We observed no difference in how men and women rated the false claims the first time they saw them. At the beginning of the survey, both placed the false claims they would later see corrected at around 5, that is, a point below 'Definitely true'. After seeing the fact check, the men in our sample moved 1.5 points closer to 'false' (M_{T2, Male}=3.59, SD=1.93), remaining well within the 'true' side of the spectrum. Meanwhile, the women in our sample moved a whole 2 points closer to false, reaching the midpoint of the scale (M_{T2, Female} =3.2, SD=2.17). This gap was largely driven by a gender gap in reactions to the two anti-immigration fact-checks: Women were more convinced by the expert saying that immigrants paid more in taxes than they received in benefits and services and by the expert saying that the number of asylum seekers had remained constant of the past ten years (see 6). ³

Consistent with motivated reasoning research, attitude extremity also predicted people's responsiveness to challenging statistics. The stronger the anti-immigration attitudes and the stronger the pro-immigration attitudes, the more they resisted the fact-check. Respondents who thought Britain should take in *many fewer* immigrants only moved a point and a half closer to the false end of the scale. In contrast, respondents who only wanted to reduce immigration a little bit, answering '-1' on a scale from -3 (many fewer) to +3 (many more) moved a full 2.5 points closer to false (see figure 7). ⁴ We observed a similar, albeit less pronounced pattern among pro-immigration respondents: Those who thought Britain should take in *many more* immigrants were half a point less convinced than those who thought Britain should keep current levels of immigration. ⁵

	belief T1	sd	belief T2	sd	diff	trust	sd	accurate	sd	n
no comment	5.12	1.02	3.37	1.02	1.75	3.26	1.50	2.80	0.68	490
comment	5.04	1.05	3.87	1.05	1.16	2.98	1.48	2.67	0.69	1492

TABLE 3.2: Effect of the post-truth comment

³ The gender gap persists when controlling for attitude extremity: Even among the most staunchly anti-immigration respondents women shifted significantly further toward true than men: $M_{diff, men, '-3 many fewer'} = 1.24$, $M_{diff, women, '-3 many fewer'} = 2.09$, t(149)=-2.51, p=0.01, see table **??**.

⁴ Those who wanted to reduce immigration the most moved from definitely true $(M_{T1, '-3'}=5.68)$ to probably true, on average $(M_{T2, '-3'}=3.95, diff=1.73)$. Those who wanted to reduce immigration a little bit were fairly certain the false claims were true before they saw it $(M_{T1, '-1'}=4.93)$ and shifted to the false side of the scale, rating it as probably false after they saw the fact check $(M_{T2, '-1'}=2.42.)$

⁵ Those who thought Britain should take in *many more* immigrants rated the false facts as almost definitely true at the beginning of the survey ($M_{T1, \text{ pro, }+3}$ =5.45), and moved almost two points closer to the midpoint ($M_{T2, \text{ pro, }+3}$ =3.23). Those who wanted no change in immigration numbers rated the false facts as most probably false ($M_{T1, \text{ pro, no change}}$ =4.67), and moved to 'don't know' after they saw the fact check ($M_{T2, \text{ pro, no change}}$ =2.95).



FIGURE 3.1: Perceived veracity of false claims, after exposure to a fact check

Eighty pecent of our sample saw not only an expert statement correcting a false fact but also a post-truth comment casting doubt on the expert fact check. To replicate common post-truth comments we varied the reason why the commenter said they would 'take these statistics with a grain of salt': They argued that a) the expert was biased, b) the expert's statistics did not match their personal experience, or c) suggested that when it doubt it was best to trust one's instincts.

The post-truth comment did significant damage to the effectiveness of the fact-check. It did not cancel out out the *entire* effect of the fact-check – but about half a point (or 45%) of it (see figure 3.1 and table 3.2). The fact check alone made respondents to shift 1.75 points closer to 'false'. If it was followed up by a post-truth comment it made respondents shift 1.16 points closer to 'false'. This half point difference is crucial: it kept

respondents from moving to 'don't know'. Those who were exposed to the post-truth comment moved from thinking the false claim was almost certainly true ($M_{T1, no comment}=5.04$) to thinking it was probably true ($M_{T1, no comment}=3.87$). Their peers who had *not* seen the post-truth comment shifted much closer to the midpoint of the scale, almost admitting that they did not know whether the claim was true or false ($M_{T2, no comment}=3.37$). ⁶ Hence, we accept hypothesis 2a: Exposure to a post-truth comment reduced the effect of the correction on belief in false facts.

Exposure to a post-truth comment also decreased trust in the source of the fact check and the perceived accuracy of the expert's information: On a scale from 0 (would not trust at all) to 6 (would trust a great deal) those who only saw the fact check rated the Oxford professor closer to trustworthy than to not trustworthy, $M_{no comment}=3.26$. Those who also saw a post-truth comment rated him right between the two ends of the scale, signalling they were not sure if this expert was trustworthy or not, $M_{comment}=2.98$ (see table 3.2). ⁷ To measure the perceived accuracy of the expert's information we used a four-point scale, where 4=very accurate, 3=fairly accurate, 2=not very accurate, and 1=not at all accurate. On average, respondents who only saw the fact check rated the information just below 'fairly accurate', $M_{no comment}=2.80$; respondents who also saw a post-truth comment rated it as $M_{comment}=2.67$. ⁸.

What the post-truth commenter said, or who they were, did not seem

⁶ The difference in how far those who saw the post-truth comment and those who did not see it shifted toward 'false' was statistically significant (t(688)=5.64, p=0.00).

⁷ The difference in means was statistically significant t(818)=3.52, p=0.00.

⁸ t(834)=3.52, p=0.00.

to matter. A post-truth comment coming from 'David Williams, Professor of Economics at the London School of Economics' did not have a greater effect than a comment coming from 'David Williams, Blogger'. Both pushed people back to think that the statement they just seen factchecked was, after all, probably true (average scores in all three conditions were at about 4, see figure **??**. Hence, we reject hypothesis 2b: authoritativeness did not moderate the effect of the post-truth comment in our experiment.

The content of their message was equally irrelevant. None of the three arguments – about bias, personal experience, or that it was okay to disagree with the facts – was significantly more or less effective than the other two (see figure **??**. ⁹

Responsiveness to the post-truth comment depended on some of the same factors that predicted responsiveness to the fact check. Women were far more likely to adapt their perceptions after seeing the fact check, but they were also more swayed by the post-truth comment encouraging them to retain their old beliefs. Figure **??** shows pre- and post-fact check veracity scores depending on exposure to a post-truth comment. Overall, the post-truth comment levelled out all gender differences in post-fact check perceptions: after seeing the post-truth comment both men and women thought the false claim they had seen fact-checked was probably true (M_{T2, men/comment}=3.90, M_{T2, women/comment}=3.85, see table **??**). The gender gap was largely driven by women with anti-immigration attitudes:

⁹ We checked if gender affected responsiveness to the different sources and messages. It did not. Both male and female respondents were just as willing to listen to a professor as they were to listen to the blogger, and both men and women were equally responsive to any of the three arguments.



FIGURE 3.2: Average post-fact check veracity scores depending on gender and immigration opinions

Among this group, the post-truth comment canceled out about half of the effect of the fact-check (diff_{no comment}=2.26 - diff_{comment}=1.17). Tables ?? and ?? show OLS regression models indicating the effect of exposure to a post-truth comment as well as a number of demographic variables on the difference in veracity scores before and after exposure to a fact check.

3.6.1 Evidence of post-truth reasoning

Our final goal was to capture the extent of post-truth thinking in the UK and to see if exposure to a post-truth comment might increase agreement with classic post-truth arguments.

We rely on our control group (i.e. those who skipped the post-truth comment) to explore the extent to which British citizens agree with classic post-truth arguments. After they saw the fact check we gave them two seven-point scales: one ranging from 'true' to 'false', and one ranging from 'fact' to 'opinion'. Overall, 40 per cent thought of the corrected claim as belonging in the realm of a 'matter of fact', 25 per cent were undecided, placing it in between a fact and an opinion, and 35 per cent deemed it a 'matter of opinion'. A full 12 per cent ticked the extreme end of the scale, saying the corrected claim as '6 - Purely a matter of opinion'.

Next, we asked respondents to think back about the statistics they had seen: 'Would you say that the statistics here were consistent with what you believed?' For 9 out of 10 the correct answer would have been 'no': 89 per cent had rated the false claim on the 'true' side of the scale. Exactly half of those for whom the correct answer would have been 'no' said 'no'. We followed up and gave the no sayers three statements, asking which described best where they stood. 20 per cent chose 'The statistics are probably right, but I believe something different. 51 per cent chose 'I think that the statistics are wrong'. 29 per cent chose the post-truth themed statement 'The statistics are probably right, but I believe something different'.

Finally, we asked, outright, if they agreed or disagreed that "It's OK to disagree with the facts if that's what you believe". Two thirds of our sample (68 per cent) agreed or strongly agreed. (51 per cent agreed; 17 per cent who*strongly* agreed.)

The distribution was very similar across respondents evaluating the four statements that challenged false claims on both sides of the spectrum of immigration attitudes: The percentage of respondents who thought the statement was a 'matter of opinion' only varied between 32 (whiteCrime) and 35 per cent (costImmig). It varied slightly more depending on opinions about immigration: Among those who wanted to raise immigration, 35 per cent said the statement they saw was a matter of opinion; among those who wanted to curb immigration, 40 per cent said the statement they saw was a matter of opinion.

These numbers refer to respondents who saw a fact check that challenged their beliefs (n=490). Two things are noteworthy: First, the population estimates, that is, the statistics for the whole control group sample including those who saw statistics that did not challenge their beliefs (n=737) are only marginally different. ¹⁰ Second, the distribution among those who saw a post-truth comment was almost identical, too. Exposure to post-truth comments did not make people any more likely to rate the corrected false claim as as a matter of opinion rather than a matter of fact or to agree that it is 'okay to disagree with the facts if that's what you believe: There was no significant difference in fact/opinion ratings and in agreement that it is okay to disagree with facts. These null findings may be due to ceiling effects: Given the extraordinarily high levels of agreement with these post-truth statements in the control group there simply wasn't any scope for the post-truth comment to increase agreement any for the post-truth comment to increase agreement and for the post-truth comment to increase agreement any for the post-truth comment agreement agreement and for the post-truth comment to increase agreement any for the post-truth comment agreement agreement agreement agreement agreement agreement agreement agreement agreement agreement

further.

¹⁰Among the entire sample, 35 per cent thought the corrected statement was a matter of opinion (same as the motivated to reject). 47 per cent of those who had rated the false statement as 'true' noticed that the statistics were inconsistent with what they had thought, and of those, 32 per cent thought the post-truth statement best described where they stood ('The statistics are right but I believe something different.') 64 per cent agree that it's okay to disagree with the facts.

Limitations

A few limitations ought to be noted. One concerns our sample: We had to exclude a significant part of our sample (n=489/2936, i.e. 16 per cent) from the analyses: respondents who believed in false facts on the 'wrong' side of the immigration debate and, therefore, had no motivation to disbelieve the expert that proved them wrong. The main culprit was our asylum claim: Respondents on both sides of the spectrum evaluated the claim that the number of asylum seekers had risen sharply in the past ten years as true. Obviously, this affected the representative nature of our sample. (Given the large size of our sample this is not dramatic: Our purpose was not primarily to estimate the precise extent of fact-avoidance in the electorate, but to explore how individuals react to facts when they have the motivation to believe the post-truth comment rather than the fact.) Nonetheless, future researchers ought to be careful to use false facts that discriminate between opposing camps in a political debate.

Furthermore, the post-truth mindset we encountered in this experiment ought to be interpreted with caution. Our aim was to explore if exposure to post-truth surroundings created something of a post-truth citizen. The people who answered our post-truth questions had just come across information that disconfirmed their factual beliefs and their opinions. Due to financial contraints we were unable to include a control group that would only have answered the post-truth questions.

Another limitation lies in the difficulty of studying the effect of post-truth surroundings in a 2019 environment where all respondents will have been exposed to some sort of a post-truth surroundings. We are well aware that it is impossible to isolate the effect of post-truth surroundings. Our control group is not a pure control group; and our post-truth comments did not create post-truth surroundings. They made them more salient. Our study provides a first snapshot of post-truth thinking among motivated reasoners in Brexit-era post-truth Britain. To fully understand the extent of post-truth thinking and reactions to post-truth comments we need panel data covering times of high and low post-truth popularity.

What is more, our study was confined to the effect of post-truth comments on belief in debunked false facts. We did not measure how posttruth surroundings affect political beliefs that are based on false facts. It is well known that accepting factual corrections does not usually make people change their political beliefs Gaines et al. 2007. Future studies ought to examine the effect of post-truth surroundings on self-exposure to expert information, on the way people process expert information, on perceptions about beliefs on the other side of the political divide, on willingness to engage with people on the other side of the debate and, ultimately, on political opinions. Creative experimental designs experiments are needed to understand how exposure exposure to post-truth surroundings affects political behaviour in as close to real-world settings as possible. Finally, future research ought to investigate ways in which the effect of post-truth surroundings on reactions to expert information can be restrained.

3.7 Discussion and Conclusion

This project set out to explore the effect of post-truth surroundings on reactions to unsolicited expert information that contradicts prior beliefs. We chose immigration as a salient, and important issue on which voters across party lines hold misperceptions. In a nationally representative survey we tested belief in false facts and exposed respondents to an authoritative expert statement correcting the one false fact people rated as closest to 'definitely true'. To mimick real-world settings in which news rarely go un-commented and, in particular, to mimick the post-truth era we then followed up with a post-truth comment, which we attributed to either a highly authoritative source (a Professor of Economics at the LSE) or a not authoritative source (a Blogger). We varied the reason the gave to 'take these statistics with a pinch of salt': The blogger/professor either suggested that 'just because someone is a professor doesn't mean they don't have an agenda', or that 'that doesn't sound like the world I live in' or, simply that 'there is so much information out there that it can be hard to know what to believe. In that case, I think it's best to trust your instincts even if it looks as if the facts are different.'

Our findings have three important implications: First, not all hope is lost: contrary to many correction-misperception studies, we find that citizens *do* listen to expert information. They may not shift from rating a false fact as 'Definitely true' to rating it as 'Definitely false'. However, our nuanced scale shows that fact-checks do move people away from the 'Definitely true' end of the scale and closer to 'don't know': On average, respondents shifted from 5 – just a point below 'Definitely true' – to 3, the midpoint

of the scale. We are not at all alarmed that they did not shift any further toward the 'Definitely false' side of the scale. Given the relevance of priors in information processing we are perfectly content to see citizens move from thinking they know a false fact is true to knowing they do not know if it is true or false. (We would of course hope to see citizens reach out for more information and, eventually, after having seen more reliable information move toward rating a false facts as 'false'.

Second, people are not immune to post-truth surroundings. Our respondents did not *only* listen to experts providing expert information; they also listened to non-experts providing non-expert information. Wherever the expert information was followed up by a post-truth comment respondents were less willing to adapt their factual beliefs. On average, the post-truth comment cancelled out about half a point of the effect of the expert statement (.33 points among respondents on the favourable side of the immigration debate and about .73 points among respondents on the sceptical side of the debate). This effect is not dramatic: Contrary to prior research (Linden et al. 2017) our second opinion did not cancel out the *entire* effect of the correction. It was, nonetheless, consequential: The post-truth comment caused respondents to remain on the 'true' side of the scale where the expert information on its own would have led them to shift to evaluating a false claim as right between 'true' and 'false'. Notably, people seem to respond to *any* comment telling them to disregard the expert – regardless of how authoritative the source is or what they say. When we attributed the comment that "I would take these statistics with a pinch of salt" to a professor it did just as much damage as when we attributed it to a blogger. And when they continued saying that the expert was biased they had the same effect as when they continued saying that the statistics did not match their personal experience or that 'I think it's best to trust your instincts even if it looks as if the facts are different'.

Third, we find a level of agreement with post-truth arguments that supports the notion a post-truth public. We asked our respondents to rate the false fact they had just read about on a scale from '0 - purely a matter of fact' to '6 - purely a matter of opinion'. 35 per cent deemed it closer to a matter of opinion than a matter of fact (i.e. rating it as 4, 5, or 6), with 13 per cent evaluating it as '6 - purely a matter of opinion'. Next, we asked if the statistics were in line with what they believed. If they were not we followed up, asking respondents to choose one of three statements to describe where they stood. 32%) said "The statistics are probably right but I believe something different." Finally, we presented respondents with the statement that "It's OK to disagree with the facts if that's what you believe." 46 per cent agreed. 18 per cent strongly agreed.

Democracies depend on voters' ability to aquire accurate information. If politicians use false facts and continue to use them after they have been debunked then this information environment places a great deal of weight on voter's shoulders. If online newspapers, blogs, and social media are slowly replacing state-regulated media and if, as we found, virtually any post-truth comment erodes trust in experts, then it is not surprising that it has become increasingly difficult for expert information to reach the masses. And if, as we found, the line between fact and opinon has become blurry, if voters believe it is okay to disagree with facts, if that's what you believe, then this does not bode well for the future of democratic decision making.

Our findings allow for one conclusion: Fact-checking is as important as ever. We find that even under the most post-truth circumstances, experts are still heard. We have exposed citizens to an expert statement that disconfirmed their factual beliefs, challenged their political beliefs, and followed up with a post-truth comment saying it was okay to ignore the expert. And yet, the expert statement made them shift slightly closer to 'don't know'. We conclude that the battle is not lost – but it requires some serious fighting. Fact-checking charities and public broadcasters such as the BBC Reality Check are fighting Goliath – they deserve our gratitude and they deserve our public funding.

Conclusion

Humans are not known to be unbiased information processors. Decades of research in psychology and political science confirm the confirmation bias we see in others and others see in us. This dissertation set out to explore how the sitation we are in when exposed to challenging information affects the way we react to it. I examined the effect of three potentially detrimental situations, stress, status threat, and post-truth surroundings, on two motivated-reasoning-related outcomes: the likelihood of detecting (and punishing) false facts from friendly sources and the likelihood of accepting expert information that challenges one's political views.

In chapter 1, I attemted to examine the effect of stress on belief in false claims propagated by the Leave campaign in the United Kingdom's June 2016 EU referendum. The stress treatment, exposure to a 35 second time limit on three questions about the EU, had no effect on belief in false facts among the 99 twitter-recruited leave voters. I argued that the null findings were due to an unstressful stress treatment. Hence, the effect of stress on confirmation bias in poltical campaigns remains to be studied in future research. Nonetheless, exploratory analyses revealed heterogeneous treatment: Among low status individuals the relatively unstressfull stressr did have an effect: Examining a subset of individuals who had placed themselves on the lower echelons of a social status latter those who had been timed were slightly (but not significantly) more likely to believe in false claims. This pointed to an avenue for future research: Could the perception that one's group is structurally disadvantaged impair people's ability (or willingness) to detect false facts?

In chapter 2, I explored the effect of status differences on information processing. The proposed mechanism combined insights from motivated reasoning and social identity theory: if you feel that your group is disadvantaged relative to relevant other groups and if a politician recognizes this disadvantage and promises to raise your group's status (ideally above and beyond the status of other groups) then you will want to believe them. If that leader then says or does something that ought to set off alarm bells – for instance, if they say something that is, simply, false - then you will put on blinders and ignore it. The same should hold for individuals who are (or believe to be) part of a high-status groups whose higher status is threatened: If you feel that your group is about to lose its superior (e.g. cultural or socio-economical) standing or if feel that your group has already lost their legitimately higher status and if you are then exposed to a politician who promises to prevent or undo this status loss then you will want to believe them. If these leaders say anything that could compromise their credibility then you have every reason to ignore those cues. To test this hypothesis, I designed a laboratory experiment. To manipulate status differences I set up a rigged pub quiz. To test tolerance of false claims I exposed respondents to false claims coming from either a person who congratulated the winning team, or a person who complained that this quiz was unfair and that the disadvantaged team should receive an extra payment to compensate for their disadvantage. The design proved (very) successful at creating low status: The disadvantaged team performed much worse than the advantaged team and received a much lower payoff, on average. It was also successful a creating high status: The advantaged team performed much better and made much more. It did not however succeed at creating a feeling of status threat among the advantaged group. (Hence, the effect of status threat on information processing remains to be studied.) The low-status team behaved as expected: They overlooked false facts in the feedback that favoured their own team but noticed (and punished) the same false facts when they appeared in the feedback that favoured the other team. These results do not bode well for the future of factual accuracy in political debates in an increasingly populist world.

In chapter 3, I investigated the effect of post-truth surroundings. This time, the outcome variable was not tolerance of false facts but tolerance of correct facts – more specifically, the willingness to give up factual beliefs after seeing official statistics that prove them wrong. In a large-scale survey experiment respondents were presented with expert opinions refuting a false belief they held about immigration. Next, some (but not all) respondents were exposed to a post-truth comment nudging them to disregard the expert advice and return to their original beliefs. As expected, the post-truth cancelled out some – but not all – of the effect of the expert statement.

Overall, this dissertation adds to a growing body of research that testifies to the difficulty of accepting challenging political information. Much of this research focuses on individual characteristics of those who are most likely to overlook false information, and on ways to design a bullet-proof fact check (Nyhan 2012; Nyhan and Reifler 2013; Lewandowsky, Ecker, et al. 2012; Nyhan and Reifler 2012; Graves, Nyhan, and Reifler 2016). The findings are invaluable, and have informed fact checking websites. (Although some findings, such as the recommendation not to repeat false information, are taking longer to find their way into public broadcasting.) The study I reported on in chapter 3 testifies to the importance of fact checks: People *do* have accuracy goals and reading these fact checks *does* reduce misperceptions, even among those who are least inclined to believe them. However, the problem in the Trump election, or the Brexit referendum was not that voters read a fact-check and did not believe it. The problem was that most voters never read those fact checks, however well designed.

This dissertation contributes to research on politically motivated reasoning by pointing to a so-far understudied field of potential predictors: the situation a person is in when they are exposed to challenging information. The three studies presented here above provide but a glimpse into the role of stress, group dynamics, and other people's opinions. Nonetheless, this first glimpse is enough to suggest that scholarly attention ought to shift toward exploring situational factors. Enough research has attempted to find *who* turns a blind eye to false facts – even if some may be more or less open to challenging information it is now clear that we all are vulnerable to these biases. Acknowledging this fact makes it all the more important to turn to the situations that worsen these biases. Ultimately, the goal of research on politically motivated reasoning is not to blame individuals for being guilty of motivated reasoning biases. It is to find ways to reach those who happen to be in a situation that makes them susceptibile to overlooking dodgy information, or to accepting false information. The goal is to devise policy recommendations that make it possible to communicate potentially threatening information in a nonthreatening way. To get there, we need to identify the situations in which individuals are particularly vulnerable to misinformation, and we need to understand how they react to it in those situations (for instance tracing information search and evaluating real data), and we need to create situations that alleviate these biases.

A particularly promising avenue for future research is to explore the role of group dynamics and inequality on politically motivated reasoning. If, as the findings of chapter 2 suggest structural group disadvantage makes individuals more susceptible to overlooking false facts then this points to an urgent need for research. A number of questions ought to be answered: First, is the mere perception that one's group is disadvantaged sufficient to trigger (or worsen) politically motivated reasoning? Second, is status threat, that is, the fear that one's group may soon be sliding down the social ladder sufficient to trigger (or worsen) politically motivated reasoning? Third, and most importantly, is there anything policy makers can do to facilitate reception of challenging political information in difficult group settings? For instance, might only just acknowledging a disadvantaged group's disadvantage help group members accept difficult-to-digest information? Given what we already know about how group dynamics affect cognition and given the current trend for state or other actors to channel group-targed misinformation to specific groups and to induce among those groups a feeling of structural disadvantage, unfair treatment, or lack of due respect it is of vital importance for future research to find ways to counter these trends. We need to find ways to reach those who are in a situation that leaves them vulnerable to actors exploiting their status for political gain.

A promising way forward is for political scientists to team up with researchers and practicioners to devise and test strategies. Research findings ought to inform public broadcasting on how to report information that may be threatening to some groups. We ought to team up with journalists to test ways to communicate challenging information effectively (and, on a side note, to test the effectiveness of existing strategies and the many new ways in which journalists are already attempted to counter false information). In addition, we ought to team up with organizations at the forefront of misinformation campaigns, evaluate ways in which they try to counter misinformation, and test (and evaluate) new ways to do so. Political scientists should also cooperate with educators. Longterm, research findings need to find their way into public education and into textbooks. Here, the goal is that facts, such as historical events that may be threatening to some are not only taught to all but accepted by all and, even more importantly, that children and young people are taught to be literate information processors, that is, that they learn the tools they need to evaluate political information and to be en guard against false information.

Nonetheless, if the recent outcy over post-truth politics has shown one thing, it is that voters – not all, but many – notice deviations from the norm of truthfulness in political speech. Enough of the public has taken offence for this phenomenon to be named: *post-truth*. If voters take offense at political campaigns printing false figures on their buses, if they

protest when state-regulated media channels spread false stories, or when politicians tweet falsehoods into the world then this shows that our era is not *entirely* post-truth: Even if the value of truth is under attack we do, still, defend it.

Appendix A

Chapter 1



(A) Twitter account used to launch the survey



(B) Timing of participation

FIGURE A.1: Survey Experiment

True or false?

We've introduced a time limit to make this a bit more exciting! You have 35 seconds to answer these questions.

	True	False	Don't know
Switzerland is a member state of the EU.	0	0	0
The members of the European parliament are directly elected by the citizens of each member state.	0	0	0
The EU currently consists of 28 member states.	0	0	0

Timing

These page timer metrics will not be displayed to the recipient.

0 seconds

0 seconds

0 clicks

First Click 0 seconds Last Click Page Submit Click Count



(A) Treatment group

True or false?					
	True	False	Don't know		
The EU currently consists of 28 member states.	0	0	0		
The members of the European parliament are directly elected by the citizens of each member state.	0	0	0		
Switzerland is a member state of the EU.	0	0	0		

(B) Control group

FIGURE A.2: Stress treatment

A lot has been said about Britain and the EU. What would you say -- are the following statements true or false?

	Definitely true	Probably true	Probably false	Definitely false	× Don't know
If Britain had remained in the EU it would have had to accept Turkish membership.	0	0	0	0	0
The EU could have forced British soldiers to join a European army.	0	0	0	0	0
Leaving the EU frees up £350m a week for the NHS.	0	0	0	0	$^{\circ}$
The EU could have made Britain join the Euro.	0	0	0	0	0
The Queen backs Brexit.	0	0	0	0	0
After Britain voted to leave the EU the British pound hit its lowest value against the US dollar since 1985.	0	0	0	0	0
All council areas in Scotland voted to remain in the EU.	0	0	0	0	0

FIGURE A.3: False (and correct) claims



(A) How interested were you in the EU referendum that was held in the UK on June 23rd?



(B) How sure were you about your vote choice?

FIGURE A.4: Interest and Vote Choice



(A) How many good reasons were there to leave?







Time treatment group (35 second timer) spent answering three quiz questions



(B) no timer

FIGURE A.6: Time spent answering quiz questions







FIGURE A.7: Time spent answering quiz questions (low status respondents)



Time treatment group (35 second timer) spent answering three quiz questions



(B) no timer

FIGURE A.8: Time spent evaluating false campaign claims (and distractors)



(A) If Britain had remained in the EU it would have had to accept Turkish membership.



(B) The EU could have forced British soldiers to join a European army. FIGURE A.9: Belief in false campaign claims



(A) Leaving the EU frees up GBP 350m a week for the NHS.



(B) The EU could have made Britain join the Euro.FIGURE A.10: Belief in false campaign claims (continued)



(B) Misperceptions index

FIGURE A.11: Belief in false campaign claims (continued)


(A) All council areas in Scotland voted to remain in the EU.



(B) The pound plunged to its lowest level since 1985 after Britain voted to leave the EU.

FIGURE A.12: Belief in correct statements (continued)



FIGURE A.13: Time spent evaluating false campaign claims (and distractors) and belief in false claims

Appendix **B**

Chapter 2

B.1 Screenshots of the laboratory experiment

Question 1 of 12: Team B Time left to complete this page: 0:25 Concord is the capital of which U.S. state? Vermont New Hampshire Maine New York Next

Question 1 of 12: Team A

And here is the other team's question.
Note that this is just for your information not for you to answer!
What is the capital city of Germany?
 Frankfurt/Main Berlin Munich Bonn Please click 'Next' to see your team's next question.
Next

FIGURE B.1: Example quiz questions (as shown to a Team B player)

Predictions

Before we move on to the last question: If you had to predict how easy or how difficult the last question will be would you say that:

- O Both teams' questions will be about the same difficulty.
- Team A will get an easier question than Team B.
- \bigcirc Team B will get an easier question than Team A..
- One team will get an easier question than the other team. Whether that is Team A or Team B is a 50-50 chance.
- O Don't know. / Impossible to say.

Next

FIGURE B.2: Perceived bias (shown before the last question)

Payoffs

Payoffs are calculated as follows: Each correct answer is worth £1.00. Half of what you earn is yours. The other half is put into your team's pot, which is evenly divided among all team members.

Your Correct Answers

In total, you answered 5 out of 12 questions correctly.

That means that you have earned £5.00. Half of this, i.e. £2.50, is yours. The other half has been put into your team's pot to be evenly divided among all team members.

Check the following table (the last column) to see how much you got out of your team's pot:

Your Team's Pot

Team	# Team Members	Joint # of Correct Answers	Joint Earnings	£ in Team's Pot (Joint Earnings / 2)	£ for each Team Member (Pot / # Team Members)
A	1 team members	10 correct answers	£10.00	£5.00	£5.00
В	1 team members	5 correct answers	£5.00	£2.50	£2.50

Your Payoffs

As a member of Team B you have received $\pounds 2.50$ out of your team's pot.

Therefore, your total payoff for this game is $\pounds 2.50 + \pounds 2.50 = \pounds 5.00$.

Note that we have a £5.00 minimum payoff policy for all experiments at ESSEXLab. That means that if your payoff is less than £5.00 you will still be paid £5.00.

Average Payoffs

Team	Average Payoffs		
A	£10.00		
В	£5.00		
Disco click Worth to continue			

Next

FIGURE B.3: Payoffs



(A) Number of difficult questions by team



FIGURE B.4: Number of difficult questions and payoffs for both teams

Your Feedback





<form><form><form><form><form><form>

FIGURE B.6: Perceived legitimacy of difference in payoffs

Luck of the Draw

Now sometimes it happens in a pub quiz that one team has the luck of the draw, and others do not. Thinking about this quiz:		
How likely do you think it is that Team A had the luck of the draw?		
In other words:		
How likely do think it is that this game was unbalanced and Team A was favoured?		
Please use the following scale from 0 ('Definitely by chance: Team A had luck of the draw') to 100 ('Definitely unbalanced: No way this happened by chance; Team A was favoured') to indicate your answer.		
0 = Definitely by chance. 100 = No way this happened by chance.		
50		

FIGURE B.7: Perceived Bias

Different Payoffs



FIGURE B.8: Emotional Reactions to Different Payoffs

Does Feedback Make a Difference?
At university (and elsewhere!) you are often asked for feedback. For instance, restaurants ask for feedback on their menu, and professors ask for feedback on their classes. We asked you for feedback about a geography quiz.
Thinking in particular about the person who thought this quiz was unfair: How likely do think it is that this person's voice is heard?
Put differently: How likely do you think it is this quiz will be changed according to this person's suggestions?
○ Very Likely
Somewhat Likely
Somewhat Unlikely
O Very Unlikely
Next



Other People's Feedback





(B) Including false claims

FIGURE B.10: 'Fair Play' Feedback

This game was unfair.

The questions Team A got were far easier than the questions Team B got. Team B had no chance of getting anywhere near Team A's payoffs, let alone the maximum payoff. Considering the difficulty of the questions Team B got anything more than zero correct answers is impressive. Anyone can see from the payments that Team B has done extraordinarily well.

Team B should have their payments topped up. The least you can do is give each member of Team B an additional flat payment of GBP 5.00 on top of their payoffs.

To answer the points you wanted feedback on:

- The questions about European capitals were quite easy.
- Some of the questions about state capitals or provincial capitals were too difficult.
- For example, 'Concord is the capital of which U.S. state?' People are very unlikely to know that unless they are
- American.

 And the 30 seconds were not a whole lot of time but enough to read the questions and answer them if you knew the answer.

(A) Not including any false claims

This game was unfair.

The questions Team A got were far easier than the questions Team B got. Team B had no chance of getting anywhere near Team A's payoffs, let alone the maximum payoff. Considering the difficulty of the questions Team B got anything more than zero correct answers is impressive.

Anyone can see from the payments that Team B has done extraordinarily well. Team B should have their payments topped up. The least you can do is give each member of Team B an additional flat payment of GBP 5.00 on top of their payoffs.

To answer the points you wanted feedback on:

- One of the first questions was 'What is the capital city of Germany'. That was pretty easy.
- Too hard: How to get from one place you have never heard of to another place you have never heard of.
 Some of the places in that quiz don't even exist.
- And the 30 seconds we had was barely enough time to read the questions

(B) Including false claims

Fair or Unfair?			
The two people's feedback we showed you differed in how fair they thought this quiz was.			
Thinking about the person whose feedback you just evaluated: How fair or unfair do you think THIS PERSON thought the quiz was?			
We would like you to place this person's thoughts about the quiz on a scale from 'definitely unfair' (0) to 'definitely fair' (100). Note that this is not a true-false question: we are just interested in how this feedback came across to you.			
0 = This person thought the quiz was definitely unfair. 100 = This person thought the quiz was definitely fair.			
50			
Next			

(C) Attention Check

FIGURE B.11: 'Unfair' Feedback

Next

Your Views on Feedback 1

Generally speaking, do you agree with the author of this feedback?

- Strongly agreeAgree
- Slightly agree
- Slightly disagree
- Disagree
- Strongly disagree

Your Views on Feedback 1

Here is a scale from 0 (very low) to 100 (very high). Thinking again about the person who wrote this feedback: How would you rate this person on the following:

0 = Very Low. | 100 = Very High. Education: This person is well educated.

O	50	
Accuracy: The points this person makes are	factually a	ccurate.
0	50	

Representation: This person is a good representative of my team.

Your Views on Feedback 1

Now we are interested in what you think about the last four points in this person's feedback. How true or false do you think they are?

The questions about European capitals were quite easy.
O Definitely True
O Probably True
O Probably False
O Definitely False
Some of the questions about state capitals or provincial capitals were too difficult.
O Definitely True
O Probably True
O Probably False
O Definitely False
For example, 'Concord is the capital of which U.S. state?' People are very unlikely to know that unless they are American.
O Definitely True
O Probably True
O Probably False
O Definitely False
And the 30 seconds were not a whole lot of time but enough to read the questions and answer them if you knew the answer.
O Definitely True
O Probably True
O Probably False
O Definitely False

FIGURE B.12: DVs – 'Fair Play' Feedback

Your Views on Feedback 2

Here is a scale from 0 (very low) to 100 (very high). Thinking again about the person who wrote this feedback: How would you rate this person on the following:

0 = Very Low. 10	00 = Very	High
--------------------	-----------	------

Education: This person is well educated.				
	50			
Accuracy: The points this person makes are factually accurate.				
0	50			
Representation: This person is a good representative of my team.				
O	50			

Your Views on Feedback 2

Now we are interested in what you think about the last four points in this person's feedback. How true or false do you think they are?

One of the first questions was 'What is the capital city of Germany'. That was pretty easy.

- O Definitely True
- O Probably True
- O Probably False
- O Definitely False

Too hard: How to get from one place you have never heard of to another place you have never heard of.

- O Definitely True
- O Probably True
- O Probably False
- O Definitely False

Some of the places in that quiz don't even exist.

- O Definitely True
- O Probably True
- O Probably False
- O Definitely False

And the 30 seconds we had was barely enough time to read the questions.

- Definitely True
- O Probably True
- O Probably False
- O Definitely False



Appendix C

Chapter 3 – Graphs



FIGURE 1: Immigration opinions



FIGURE 4: Fact checked statements by immigration opinions



(A) 'There has been a sharp rise in the number of people applying for asylum in the UK in the past ten years.'



(B) 'Immigrants receive ore in benefits and services than they pay in taxes.'

FIGURE 2: False claims that affirm anti-immigration opinions



(A) 'The majority of crimes in London are committed by white people, not ethnic minorities.'



(B) 'Immigration to the UK does not affect the wages of the low-paid.'

FIGURE 3: False claims that affirm pro-immigration opinions



FIGURE 5: Effect of the fact-check on belief in the four false claims (Control group only)



FIGURE 6: Effect of the fact check depending on gender and immigration opinions









FIGURE 7: Effect of the fact check depending on attitude strength

Bibliography

- Adler, Nancy and Judith Steward (2007). *The MacArthur Scale of Subjective Social Status*. URL: http://www.macses.ucsf.edu/research/ psychosocial/subjective.php%7B%5C#%7Drelationhealth (visited on 01/25/2017).
- Anischchuk, Alexei (2014). "Putin admits Russian forces were deployed to Crimea". In: *News agency company Reuters* 09, pp. 1–6. URL: http:// uk.reuters.com/article/russia-putin-crimea-idUKL6N0N921H20140417.
- Antonucci, Lorenza, Laszlo Horvath, and André P. M. Krouwel (2017). Brexit was not the voice of the working class nor of the uneducated – it was of the squeezed middle.
- Arnold, Phoebe (2016). UK Statistics Authority calls GBP 350m EU membership fee claim 'potentially misleading'. URL: https://fullfact.org/blog/ 2016/apr/uk-statistics-authority-350-million-eu-membershipfee-potentially-misleading/ (visited on 05/24/2020).
- Bago, Bence, David G. Rand, and Gordon Pennycook (2020). "Fake news, fast and slow: Deliberation reduces belief in false (but not true) news headlines". In: *Journal of experimental psychology. General*, pp. 1–18. DOI: https://doi.org/10.1037/xge0000729.

- Barrera, Oscar et al. (2018). "Facts, Alternative Facts, and Fact Checking in Times of Post-Truth Politics". In: *SSRN Electronic Journal* 646662. DOI: 10.2139/ssrn.3004631.
- Bartels, M (2002). "Beyond the Running Tally. Partisan Bias in Political Perceptions". In: *Political Behavior* 24.2, pp. 117–150. ISSN: 1572-0241.
 DOI: 10.1038/ajg.2011.482.
- BBC (2016). BBC's Question Time EU Special, Nottingham, 15 June 2016. URL: https://www.youtube.com/watch?v=%7B%5C_%7D9-jDRQuToo.
- Berinsky, Adam J. (2015). "Rumors and Health Care Reform: Experiments in Political Misinformation". In: *British Journal of Political Science* June, pp. 1–22. ISSN: 0007-1234. DOI: 10.1017/S0007123415000186. URL: http://www.journals.cambridge.org/abstract%78%5C_%7DS0007123415000186.
- (2018). "Telling the truth about believing the lies? Evidence for the limited prevalence of expressive survey responding". In: *Journal of Politics* 80.1, pp. 211–224. ISSN: 14682508. DOI: 10.1086/694258.
- Bickman, Leonard and Stephanie M. Reich (2009). "Randomized controlled trials: A Gold Standard With Feet of Clay?" In: *What counts as credible evidence in applied research and evaluation practice*? Ed. by Stewart I. Donaldson, Christina Christie, and Melvin M. Mark. Thousand Oaks, CA: Sage. ISBN: 9789811052514. DOI: 10.1007/978-981-10-5251-4_94.
- Bolsen, Toby, James N. Druckman, and Fay Lomax Cook (2014). "The Influence of Partisan Motivated Reasoning on Public Opinion". In: *Political Behavior* 36.2, pp. 235–262. ISSN: 01909320. DOI: 10.1007/s11109-013-9238-0.

- (2015). "Citizens', Scientists', and Policy Advisors' Beliefs about Global Warming". In: *The ANNALS of the American Academy of Political and Social Science* 658.1, pp. 271–295. ISSN: 0002-7162. DOI: 10.1177/0002716214558393.
 URL: http://ann.sagepub.com/cgi/doi/10.1177/0002716214558393.
- Booth, Robert, Alan Travis, and Amelia Gentleman (2016). *Leave donor* plans new party to replace Ukip 4 possibly without Farage in charge. London. URL: https://www.theguardian.com/politics/2016/jun/29/leavedonor-plans-new-party-to-replace-ukip-without-farage.
- Boykoff, Maxwell T. and Jules M. Boykoff (2004). "Balance as bias: Global warming and the US prestige press". In: *Global Environmental Change* 14.2, pp. 125–136. ISSN: 09593780. DOI: 10.1016/j.gloenvcha.2003. 10.001.
- Brandenburg, Heinz, Marcel van Egmond, and Rob Johns (2017). "Brexit: Analyzing tweets, surveying tweeters to understand public opinion dynamics during the EU referendum campaign." Berlin.
- Brass, C. T., B. Nunez-Neto, and E. D. Williams (2006). "Congress and program evaluation: An overview of randomized controlled trials (RCTs) and related issues". Washington, DC.
- Buckert, Magdalena, Jörg Oechsler, and Christiane Schwieren (2014). "Imitation under stress". In: WZB Discussion Papers 309.December.
- Bullock, John G et al. (2015). "Partisan Bias in Factual Beliefs about Politics". In: *Quarterly Journal of Political Science* 10, pp. 519–578.
- Carretta, Thomas R and Richard L. Moreland (1983). "The Direct and Indirect Effects of Inadmissible Evidence1". In: *Journal of Applied Social Psychology* 13.4, pp. 291–309. ISSN: 0021-9029. DOI: 10.1111/j.1559-

1816.1983.tb01741.x.URL:http://doi.wiley.com/10.1111/j.1559-1816.1983.tb01741.x.

- Chaiken, Shelly (1980). "Heuristic versus systematic information processing and the use of source versus message cues in persuasion". In: *Journal of Personality and Social Psychology*. ISSN: 00223514. DOI: 10.1037/ 0022-3514.39.5.752.
- Chen, Yan and Sherry Xin Li (2009). "Group identity and social preferences". In: *American Economic Review*. ISSN: 00028282. DOI: 10.1257/ aer.99.1.431.
- Cingl, Lubomír (2018). "Social learning under acute stress". In: *PLoS ONE* 13.8, pp. 1–26. ISSN: 19326203. DOI: 10.1371/journal.pone.0202335.
- Clarke, Christopher E. (2008). "A question of balance: The autism-vaccine controversy in the British and American elite press". In: *Science Communication*. ISSN: 10755470. DOI: 10.1177/1075547008320262.
- Clarke, Christopher E. et al. (2015). "Including "Evidentiary Balance" in News Media Coverage of Vaccine Risk". In: *Health Communication* 30.5, pp. 461–472. ISSN: 15327027. DOI: 10.1080/10410236.2013.867006.
- Clarke, Harold D, Matthew Goodwin, and Paul Whiteley (2018). "Why Britain Voted for Brexit : An Individual- Level Analysis of the 2016 Referendum Vote". In: April 2017, pp. 439–464. DOI: 10.1093/pa/gsx005.
- Cohen, Geoffrey L. et al. (2007). "Bridging the partisan divide: Self-affirmation reduces ideological closed-mindedness and inflexibility in negotiation."
 In: Journal of Personality and Social Psychology 93.3, pp. 415–430. ISSN: 0022-3514. DOI: 10.1037/0022-3514.93.3.415. URL: http://doi.apa.org/getdoi.cfm?doi=10.1037/0022-3514.93.3.415.

- Cook, John and Stephan Lewandowsky (2011). The Debunking Handbook. Tech. rep. St. Lucia, Australia: University of Queensland. URL: https: //www.skepticalscience.com/docs/Debunking%7B%5C_%7DHandbook. pdf.
- D'Ancona, Matthew (2017a). *Post-Truth. The new war on truth and how to fight back*. London: Ebury Press.
- (2017b). The New Jerusalem of Brexit is revealing itself to to be a mirage.
- De Vries, Catherine E., Sara B. Hobolt, and James Tilley (2015). "Facing up to the facts. What causes economic perceptions?" In: pp. 1–32. ISSN: 0261-3794. DOI: 10.1016/j.electstud.2017.09.006.
- Dearing, James W. (1995). "Newspaper coverage of maverick science: Creating controversy through balancing". In: *Public Understanding of Science* 4.4, pp. 341–361. ISSN: 09636625. DOI: 10.1088/0963-6625/4/ 4/002.
- Dixon, Graham N. and Christopher E. Clarke (2013). "Heightening Uncertainty Around Certain Science: Media Coverage, False Balance, and the Autism-Vaccine Controversy". In: *Science Communication* 35.3, pp. 358–382. ISSN: 10755470. DOI: 10.1177/1075547012458290.
- Dixon, Graham N., Brooke Weberling Mckeever, et al. (2015). "The Power of a Picture: Overcoming Scientific Misinformation by Communicating Weight-of-Evidence Information with Visual Exemplars". In: *Journal of Communication* 65.4, pp. 639–659. ISSN: 14602466. DOI: 10.1111/jcom. 12159.

- Dougherty, Michael R P and Jennifer Hunter (2003). "Probability judgment and subadditivity: the role of working memory capacity and constraining retrieval." In: *Memory & cognition* 31.6, pp. 968–982. ISSN: 0090-502X. DOI: 10.3758/BF03196449.
- Duffy, Bob (2019). Brexit Misperceptions. Tech. rep. London: The UK in a Changing Europe. The Policy Institute. King's College London, pp. 9– 11. DOI: 10.1017/CB09781107415324.004. arXiv: arXiv:1011.1669v3. URL: https://www.kcl.ac.uk/policy-institute/research-analysis/ the-publics-brexit-misperceptions.
- Duffy, Sean and John Smith (2014). "Cognitive load in the multi-player prisoner's dilemma game: Are there brains in games?" In: *Journal of Behavioral and Experimental Economics*. ISSN: 22148051. DOI: 10.1016/j.socec.2014.01.006.
- Eckel, Catherine C. and Philip J. Grossman (2005). "Managing diversity by creating team identity". In: *Journal of Economic Behavior and Organization*. ISSN: 01672681. DOI: 10.1016/j.jebo.2004.01.003.
- Edwards-levy, By Ariel (2016). Nearly Half Of Trump Voters Think Whites Face A Lot Of Discrimination. New York. URL: https://www.huffingtonpost. com/entry/discrimination-race-religion%78%5C_%7Dus%78%5C_ %7D5833761ee4b099512f845bba.
- Engel, Pamela (2015). Putin: 'I will say this clearly: There are no Russian troops in Ukraine'. URL: http://www.businessinsider.com/putini-will-say-this-clearly-there-are-no-russian-troops-inukraine-2015-4.

- Evans, Geoffrey and Robert Andersen (2006). "The political conditioning of economic perceptions". In: *Journal of Politics*. ISSN: 00223816. DOI: 10.1111/j.1468-2508.2006.00380.x.
- Evans, Geoffrey and Mark Pickup (2010). "Reversing the causal arrow: The political conditioning of economic perceptions in the 2000-2004 U.S. presidential election cycle". In: *Journal of Politics* 72.4, pp. 1236– 1251. ISSN: 00223816. DOI: 10.1017/S0022381610000654.
- Evans, Jonathan St. B. T. (2008). "Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition". In: *Annual Review of Psychology*. ISSN: 0066-4308. DOI: 10.1146/annurev.psych.59.103006.093629.
- FeldmanHall, Oriel et al. (2015). "The Effects of Social Context and Acute Stress on Decision Making Under Uncertainty". In: *Psychological Science* 26.12, pp. 1918–1926. ISSN: 14679280. DOI: 10.1177/0956797615605807.
- Fischle, Mark (2000). "Mass Response to the Lewinsky Scandal : Motivated Reasoning or Bayesian Updating ?" In: *Political Psychology* 21.1, pp. 135–159. ISSN: 1467-9221. DOI: 10.1111/0162-895X.00181.
- Flynn, D.J. and Yanna Krupnikov (2019). "Misinformation and the Justification of Socially Undesirable Preferences". In: *Journal of Experimental Political Science* 6.1, pp. 5–16. ISSN: 2052-2630. DOI: 10.1017/xps.2018.
 12.
- Fowler, Matthew, Vladimir E Medenica, and Cathy J Cohen (2017). *Why* 41 percent of white millennials voted for Trump.
- Fukuyama, Francis (2017). *The Emergence of a Post-Fact World*. (Visited on 08/29/2019).

- Full Fact (2017). GBP 350 million EU claim 'a clear misuse of official statistics'. URL: https://fullfact.org/europe/350-million-week-borisjohnson-statistics-authority-misuse/.
- Gaines, Brian J. et al. (2007). "Same facts, different interpretations: Partisan motivation and opinion on Iraq". In: *Journal of Politics* 69, pp. 957– 974. ISSN: 00223816. DOI: 10.1111/j.1468-2508.2007.00601.x.
- Galván, A. and A. Rahdar (2013). *The neurobiological effects of stress on adolescent decision making*. DOI: 10.1016/j.neuroscience.2012.09.074.
- Gigerenzer, Gerd and Henry Brighton (2009). "Homo Heuristicus: Why Biased Minds Make Better Inferences". In: *Topics in Cognitive Science* 1.1, pp. 107–143. ISSN: 17568757. DOI: 10.1111/j.1756-8765.2008. 01006.x. arXiv: arXiv:1011.1669v3.
- Giles, Chris (2016). *Reality check: the FT dissects pro and anti-Brexit claims*. London. URL: https://www.ft.com/content/133da9c0-f7f7-11e5-803c-d27c7117d132.
- Gilovich, Thomas, Dale W. Griffin, and Daniel Kahneman (2002). *Heuristics and Biases The Psychology of Intuitive Judgment*. Ed. by Thomas Gilovich, Dale W. Griffin, and Daniel Kahneman. Cambridge: Cambridge University Press. ISBN: 9780511808098. DOI: https://doi.org/10.1017/ CB09780511808098.
- Gok, Kubilay and Nuray Atsan (2016). "Decision-Making under Stress and Its Implications for Managerial Decision-Making: A Review of Literature". In: *International Journal of Business and Social Research* 6.3, p. 38. ISSN: 2164-2540. DOI: 10.18533/ijbsr.v6i3.936.
- Goodwin, Matthew and Caitlin Milazzo (2017). "Taking back control? Investigating the role of immigration in the 2016 vote for Brexit". In:

British Journal of Politics and International Relations 19.3, pp. 450–464. ISSN: 1467856X. DOI: 10.1177/1369148117710799.

- Gosling, Samuel D., Peter J. Rentfrow, and William B. Swann (2003). "A very brief measure of the Big-Five personality domains". In: *Journal of Research in Personality*. ISSN: 00926566. DOI: 10.1016/S0092-6566(03) 00046-1.
- Graves, Lucas, Brendan Nyhan, and Jason Reifler (2016). "Understanding Innovations in Journalistic Practice: A Field Experiment Examining Motivations for Fact-Checking". In: *Journal of Communication* 66.1, pp. 102–138. ISSN: 14602466. DOI: 10.1111/jcom.12198.
- Guillory, Jimmeka J. and Lisa Geraci (2013). "Correcting erroneous inferences in memory: The role of source credibility". In: *Journal of Applied Research in Memory and Cognition* 2.4, pp. 201–209. ISSN: 22113681. DOI: 10.1016/j.jarmac.2013.10.001. URL: http://dx.doi.org/10.1016/ j.jarmac.2013.10.001.
- Harari, Yuval Noah (2018). "Yuval Noah Harari extract: 'Humans are a post-truth species'". In: *The Guardian*, pp. 1–5. URL: https://www.theguardian.com/culture/2018/aug/05/yuval-noah-harari-extract-fake-news-sapiens-homo-deus.
- Hart, P. S. and Erik C. Nisbet (2012). "Boomerang Effects in Science Communication: How Motivated Reasoning and Identity Cues Amplify Opinion Polarization About Climate Mitigation Policies". In: *Communication Research* 39.6, pp. 701–723. ISSN: 0093-6502. DOI: 10.1177/0093650211416646.
- Helm, Toby (2016). Third of EU referendum voters won't make up their minds until week before poll. London. URL: https://www.theguardian.com/

politics/2016/jun/11/brexit-eu-referendum-vote-last-minutedecision-decide-lse.

- Henley, Jon (2016). Why Vote Leave's GBP 350m weekly EU cost claim is wrong. London. URL: https://www.theguardian.com/politics/ reality-check/2016/may/23/does-the-eu-really-cost-theuk-350m-a-week.
- Hepler, Teri J. and Attila J. Kovacs (2017). "Influence of acute stress on decision outcomes and heuristics". In: *Journal of Sports Medicine and Physical Fitness*. ISSN: 18271928. DOI: 10.23736/S0022-4707.16.06554-3.
- Hildreth, John Angus D. and Cameron Anderson (2018). "Does loyalty trump honesty? Moral judgments of loyalty-driven deceit". In: *Journal* of Experimental Social Psychology 79.July 2017, pp. 87–94. ISSN: 10960465. DOI: 10.1016/j.jesp.2018.06.001. URL: https://doi.org/10.1016/ j.jesp.2018.06.001.
- Hix, Simon (2015). Britons among least knowledgeable about European Union. URL: https://www.theguardian.com/news/datablog/2015/nov/ 27/brits-least-knowledgeable-european-union-basic-questions (visited on 09/28/2019).
- Hobolt, Sara B. (2016). "The Brexit vote: a divided nation, a divided continent". In: *Journal of European Public Policy* 23.9, pp. 1259–1277. ISSN: 14664429. DOI: 10.1080/13501763.2016.1225785. URL: http://dx. doi.org/10.1080/13501763.2016.1225785.
- ITV (2016). The ITV #Brexit "Debate" David Cameron & Nigel Farage on the EU Referendum 7 June 2016. URL: https://www.youtube.com/watch?v=baCwRZ7Fbvc.

- Jacobson, Gary C. (2010). "Perception, Memory, and Partisan Polarization on the Iraq War". In: *Political Science Quarterly* 125.1, pp. 31–56. ISSN: 00323195. DOI: 10.1002/j.1538-165X.2010.tb00667.x.
- Janis, Irving L. (1982). Groupthink : psychological studies of policy decisions and fiascoes. 2nd ed. Boston : Houghton Mifflin, ISBN: 0395317045 (pbk.) URL: https://searchworks.stanford.edu/view/2781636.
- Jerit, Jennifer and Jason Barabas (2012). "Partisan Perceptual Bias and the Information Environment". In: *The Journal of Politics* 74.3, pp. 672–684. ISSN: 0022-3816. DOI: 10.1017/S0022381612000187.
- Johns, Robert (2009). "Tracing foreign policy decisions: A study of citizens' use of heuristics". In: *British Journal of Politics and International Relations* 11.4, pp. 574–592. ISSN: 13691481. DOI: 10.1111/j.1467-856X.2009.00388.x.
- Kahan, Dan M. (2013). "Ideology, motivated reasoning, and cognitive reflection". In: *Judgment and Decision Making* 8.4, pp. 407–424. ISSN: 1556-5068. DOI: 10.2139/ssrn.2182588. arXiv: arXiv:1011.1669v3.
- (2016a). "The Expressive Rationality of Inaccurate Perceptions". In: *Behavioral and Brain Sciences*. ISSN: 0140-525X. DOI: 10.1017/S0140525X15002332.
- (2016b). "The Politically Motivated Reasoning Paradigm". In: *Emerging Trends in Social & Behavioral Sciences*, pp. 1–16. DOI: 10.1002/9781118900772.
 etrds0417.
- Kahneman, Daniel (2011). *Thinking, fast and slow*. New York: Farrar, Straus and Giroux. ISBN: 9780374275631; 0374275637.
- Keinan, Giora, Nehemia Friedland, and Yossef Ben-Porath (1987). "Decision making under stress: Scanning of alternatives under physical

threat". In: *Acta Psychologica* 64.3, pp. 219–228. ISSN: 00016918. DOI: 10.1016/0001-6918(87)90008-4.

- Klein, Rebekka A. (2015). "Joshua Greene. Moral Tribes: Emotion, Reason, and the Gap between Us and Them". In: *Philosophy, Theology and the Sciences*. ISSN: 2195-9773. DOI: 10.1628/219597715x14268452693103.
- Koehler, Derek J. (2016). "Can Journalistic "False Balance" Distort Public Perception of Consensus in Expert Opinion?" In: *Journal of Experimental Psychology. Applied* 22.1, pp. 24–38. ISSN: 1076898X. URL: http://10.0. 4.13/xap0000073%7B%5C%%7D0Ahttps://unco.idm.oclc.org/login? url=http://search.ebscohost.com/login.aspx?direct=true%7B% 5C&%7Ddb=ssf%7B%5C&%7DAN=113732753%7B%5C&%7Dsite=ehost-live.
- Koningsbruggen, Guido M. van, Enny Das, and David R. Roskos-Ewoldsen (2009). "How Self-Affirmation Reduces Defensive Processing of Threatening Health Information: Evidence at the Implicit Level". In: *Health Psychology*. ISSN: 02786133. DOI: 10.1037/a0015610.
- Kortenkamp, Katherine V. and Bianca Basten (2015). "Environmental Science in the Media: Effects of Opposing Viewpoints on Risk and Uncertainty Perceptions". In: *Science Communication*. ISSN: 15528545. DOI: 10.1177/1075547015574016.
- Kraft, P. W., Milton Lodge, and Charles S. Taber (2015). "Why People "Don't Trust the Evidence": Motivated Reasoning and Scientific Beliefs". In: *The ANNALS of the American Academy of Political and Social Science* 658.1, pp. 121–133. ISSN: 0002-7162. DOI: 10.1177/0002716214554758.
- Kull, Steven, Clay Ramsay, and Evan Lewis (2003). "Misperceptions, the Media, and the Iraq War". In: *Political Science Quarterly* 118.4, pp. 569–598. ISSN: 00323195. DOI: 10.2307/30035697.

- Kunda, Ziva (1987). "Motivated inference: Self-serving generation and evaluation of causal theories." In: *Journal of Personality and Social Psychology* 53.4, pp. 636–647. ISSN: 0022-3514. DOI: 10.1037/0022-3514. 53.4.636.
- (1990). "The case for motivated reasoning." In: *Psychological bulletin* 108.3, pp. 480–498. ISSN: 0033-2909. DOI: 10.1037/0033-2909.108.
 3.480.
- Lavine, Howard G., Christopher D. Johnston, and Marco R. Steenbergen (2013). *The Ambivalent Partisan: How Critical Loyalty Promotes Democracy*. ISBN: 9780199979622. DOI: 10.1093/acprof:oso/9780199772759.001. 0001.
- Leder, Johannes, Jan Alexander Häusser, and Andreas Mojzisch (2013). "Stress and strategic decision-making in the beauty contest game". In: *Psychoneuroendocrinology* 38.9, pp. 1503–1511. ISSN: 03064530. DOI: 10. 1016/j.psyneuen.2012.12.016.
- Lee, David S., Eunjung Kim, and Norbert Schwarz (2015a). "Something smells fishy: Olfactory suspicion cues improve performance on the Moses illusion and Wason rule discovery task". In: *Journal of Experimental Social Psychology* 59, pp. 47–50. ISSN: 00221031. DOI: 10.1016/j.jesp. 2015.03.006.URL:http://www.sciencedirect.com/science/article/ pii/S0022103115000281.
- (2015b). "Something smells fishy: Olfactory suspicion cues improve performance on the Moses illusion and Wason rule discovery task". In: *Journal of Experimental Social Psychology* 59, pp. 47–50. ISSN: 00221031. DOI: 10.1016/j.jesp.2015.03.006.

- Levine, Seymour (1985). "A Definition of Stress?" In: *Animal Stress*. DOI: 10.1007/978-1-4614-7544-6_4.
- Lewandowsky, Stephan, Ullrich K. H. Ecker, et al. (2012). "Misinformation and Its Correction". In: *Psychological Science in the Public Interest* 13.3, pp. 106–131. ISSN: 1529-1006. DOI: 10.1177/1529100612451018.
- Lewandowsky, Stephan, Werner G K Stritzke, et al. (2005). "Memory for fact, fiction, and misinformation: The Iraq War 2003". In: *Psychological Science* 16.3, pp. 190–195. ISSN: 09567976. DOI: 10.1111/j.0956-7976. 2005.00802.x.
- Linden, Sander van der et al. (2017). "Inoculating the Public against Misinformation about Climate Change". In: *Global Challenges* 1.2, p. 1600008. DOI: 10.1002/gch2.201600008.
- Lodge, Milton and Charles S. Taber (2007). "The Rationalizing Voter: Unconscious Thought in Political Information Processing". In: *SSRN Electronic Journal*, pp. 1–46. ISSN: 1556-5068. DOI: 10.2139/ssrn.1077972.
- (2013). The Rationalizing Voter. 1st ed. Cambridge: Cambridge University Press. ISBN: 0521763509. URL: https://books.google.de/books/ about/The%7B%5C_%7DRationalizing%7B%5C_%7DVoter.html?id= tq15GxWmIU4C%7B%5C&%7Dpgis=1.
- Major, B., A. Blodorn, and G. Major Blascovich (2016). "The threat of increasing diversity: Why many White Americans support Trump in the 2016 presidential election". In: *Group Processes & Intergroup Relations*. ISSN: 1368-4302. DOI: 10.1177/1368430216677304.

- Marcora, Samuele M., Walter Staiano, and Victoria Manning (2009). "Mental fatigue impairs physical performance in humans". In: *Journal of Applied Physiology*. ISSN: 87507587. DOI: 10.1152/japplphysiol.91324. 2008.
- McCright, Aaron M., Meghan Charters, et al. (2016). "Examining the Effectiveness of Climate Change Frames in the Face of a Climate Change Denial Counter-Frame". In: *Topics in Cognitive Science* 8.1, pp. 76–97. ISSN: 17568765. DOI: 10.1111/tops.12171.
- McCright, Aaron M. and Riley E. Dunlap (2011). "Cool dudes: The denial of climate change among conservative white males in the United States". In: *Global Environmental Change* 21.4, pp. 1163–1172. ISSN: 09593780. DOI: 10.1016/j.gloenvcha.2011.06.003. URL: http://dx.doi.org/ 10.1016/j.gloenvcha.2011.06.003.
- Mullinix, Kevin J. et al. (2015). "The Generalizability of Survey Experiments". In: *Journal of Experimental Political Science*. ISSN: 20522649. DOI: 10.1017/XPS.2015.19.
- Mutz, Diana C. (2011). *Population-based survey experiments*. ISBN: 9780691144511. DOI: 10.1177/0094306112462561s.
- (2018). "Status threat, not economic hardship, explains the 2016 presidential vote". In: *Proceedings of the National Academy of Sciences* 115.19, E4330–E4339. ISSN: 0027-8424. DOI: 10.1073/pnas.1718155115. URL: http://www.pnas.org/lookup/doi/10.1073/pnas.1718155115.
- NBC News (2017a). Conway: Press Secretary Gave 'Alternative Facts'. URL: https://www.nbcnews.com/meet-the-press/video/conway-presssecretary-gave-alternative-facts-860142147643 (visited on 09/01/2019).

- NBC News (2017b). Peter Alexander Fact Checks President Donald Trump On Incorrect Electoral College Margin. URL: https://www.youtube.com/ watch?v=AhZhxgz5yJU.
- Nickerson, Raymond S. (1998). "Confirmation bias: A ubiquitous phenomenon in many guises". In: *Review of General Psychology*. ISSN: 10892680. DOI: 10.1037/1089-2680.2.2.175.
- Nisbet, Eric C., Kathryn E. Cooper, and R. Kelly Garrett (2015). "The partisan brain: How dissonant science messages lead conservatives and liberals to (dis)trust science". In: *The ANNALS of the American Academy of Political and Social Science* 658.1, pp. 36–66. ISSN: 0002-7162. DOI: 10. 1177/0002716214555474.
- Nyhan, Brendan (2010). Why the 'Death Panel' Myth Wouldn't Die: Misinformation in the Health Care Reform Debate. DOI: 10.2202/1540-8884.1354. URL: http://dx.doi.org/10.2202/1540-8884.1354.
- (2012). Journalists, do no harm! If you must cover the birthers, here's an annotated how-to. URL: https://archives.cjr.org/united%78%5C_ %7Dstates%78%5C_%7Dproject/annotated%78%5C_%7Dguide%78%5C_ %7Dcovering%78%5C_%7Dbirthers.php (visited on 08/15/2019).
- Nyhan, Brendan and Jason Reifler (2010). "When corrections fail: The persistence of political misperceptions". In: *Political Behavior* 32.2, pp. 303– 330. ISSN: 01909320. DOI: 10.1007/s11109-010-9112-2.
- (2012). "Misinformation and Fact-checking: Research Findings from Social Science". In: p. 28. URL: http://www.dartmouth.edu/%78~%7Dnyhan/.
- (2013). Which Corrections Work? Research Results and Practice Recommendations for Journalists.
- (2015a). Does correcting myths about the flu vaccine work? An experimental evaluation of the effects of corrective information.
- (2015b). "Estimating Fact-checking's Effects Evidence from a long-term experiment during campaign 2014". In: *American Press Institute*.
- (2017). "The Roles of Information Deficits and Identity Threat in the Prevalence of Misperceptions".
- Nyhan, Brendan, Jason Reifler, Sean Richey, et al. (2014). *Effective messages in vaccine promotion: a randomized trial*. DOI: 10.1542/peds.2013-2365.
- Nyhan, Brendan, Jason Reifler, and Peter A. Ubel (2013). "The hazards of correcting myths about health care reform." In: *Medical care* 51.2, pp. 127–32. ISSN: 1537-1948. DOI: 10.1097/MLR.0b013e318279486b.
- Nyhan, Brendan and Thomas Zeitzoff (2017). "Fighting the Past: Perceptions of Control, Historical Misperceptions, and Corrective Information in the Israeli-Palestinian Conflict". In: *Political Psychology* xx.xx. ISSN: 0717-6163. DOI: 10.1007/s13398-014-0173-7.2. arXiv: arXiv: 1011.1669v3.
- Oliphant, Roland (2015). "Vladimir Putin admits: Russian troops 'were in Ukraine'". In: The Telegraph, pp. 1–29. URL: http://www.telegraph.co. uk/news/worldnews/vladimir-putin/12054164/Vladimir-Putinsannual-press-conference-2015-live.html.
- Paulos, John Allen (2001). Innumeracy: Mathematical Illiteracy and Its Consequences. Hill & Wang. URL: http://www.amazon.co.uk/Innumeracy-Mathematical-Illiteracy-Its-Consequences/dp/0606265953.
- Petty, RE (1994). "Two routes to persuasion: State of the art". In: International perspectives on psychological ... 2, pp. 1–15. URL: http://ebook.lib.sjtu.edu.cn/iupsys/Proc/Bruss2/bpv2ch11.htm.

- Phillips, Amber (2017). 'They're rapists.' President Trump's campaign launch speech two years later, annotated - The Washington Post. URL: https:// www.washingtonpost.com/news/the-fix/wp/2017/06/16/theyrerapists-presidents-trump-campaign-launch-speech-two-yearslater-annotated/?utm%7B%5C_%7Dterm=.7a425dc9faf3 (visited on 12/11/2017).
- Pomerantsev, Peter (2014). "Russia and the Menace of Unreality". In: *The Atlantic*, pp. 1–23. ISSN: 1072-7825.
- Prior, Markus, Gaurav Sood, and Kabir Khanna (2015). "You Cannot be Serious. The Impact of Accuracy Incentives on Partisan Bias in Reports of Economic Perceptions". In: *Quarterly Journal of Political Science* 10.July, pp. 489–518. ISSN: 15540634. DOI: 10.1561/100.00014127. URL: http://www.vanderbilt.edu/csdi/events/MPrior.pdf.
- Rampell, Catherine (2016). Americans especially but not exclusively Trump voters — believe crazy, wrong things. URL: https://www.washingtonpost. com/news/rampage/wp/2016/12/28/americans - especially - but not - exclusively - trump - voters - believe - crazy - wrong - things/ ?utm%7B%5C_%7Dterm=.d22618c0fe92.
- Reuben, Anthony (2016). Reality Check : Are we giving £ 350m a week to Brussels ? London. URL: https://www.bbc.com/news/uk-politicseu-referendum-36110822.
- Ross, Lee, Mark R. Lepper, and M Hubbard (1975). "Perseverance in selfperception and social perception: biased attributional processes in the debriefing paradigm." In: *Journal of personality and social psychology* 32.5, pp. 880–892. ISSN: 0022-3514. DOI: 10.1037/0022-3514.32.5.880.

- Schaffner, Brian F. and Samantha Luks (2018). "Misinformation or Expressive Responding?" In: *Public Opinion Quarterly* 82.1, pp. 135–147. ISSN: 15375331. DOI: 10.1093/poq/nfx042.
- Schaffner, Brian F. and Cameron Roche (2017). "Misinformation and motivated reasoning: Responses to economic news in a politicized environment". In: *Public Opinion Quarterly* 81.1, pp. 86–110. ISSN: 15375331. DOI: 10.1093/poq/nfw043.
- Schreck, Carl (2019). "From 'Not Us' To 'Why Hide It?': How Russia Denied Its Crimea Invasion, Then Admitted It". In: *RadioFreeEurope/RadioLiberty* 09, pp. 1–9.
- Schwarz, Norbert et al. (2007). "Metacognitive Experiences and the Intricacies of Setting People Straight: Implications for Debiasing and Public Information Campaigns". In: Advances in Experimental Social Psychology 39.06, pp. 127–161. ISSN: 00652601. DOI: 10.1016/S0065-2601(06) 39003-X. URL: http://konstanz.summon.serialssolutions.com/2.0. 0/link/0/eLvHCXMwrV07T8MwELZQkRDqwLs8pZsQDClxnDgJY1uqVoCEBEjdIjt2pC7h0TLwA% 7B%5C_%7Djf3DlpUxgRU-Qojh0ffXc-3%7B%5C_%7DeZMRF0fe-XTohsEKtUojXGQZMT70% 7B%5C_%7D0gxxdfa7iWDkWxtXki95KkmVtCSoN73R3feeq7tur1-mUEL8yCtymoaDg3ISQ1bjyI-LHQW8ZdSH6zM.
- Seifert, Colleen M. (2002). "The Continued Influence of Misinformation in Memory: What Makes A Correction Effective?" In: *The Psychology of Learning and Motivation* 41, pp. 265–292. ISSN: 00797421.
- Sides, John and Jack Citrin (2007). "How Large the Huddled Masses?: The Causes and Consequences of Public Misperceptions about Immigrant Populations". In: *Midwestern Political Science Conference*. URL: http:// home.gwu.edu/%20jsides/huddled.pdf.

- Sides, John, Michael Tesler, and Lynn Vavreck (2018). *Identity Crisis. The* 2016 Presidential Campaign and the Battle for the Meaning of America. Princeton and Oxford.
- Sky News (2016). *Michael Gove 'EU: In Or Out?' 3 June 2016*. URL: https://www.youtube.com/watch?v=t8D8AoC-5i8 (visited on 12/20/2017).
- Smith, Mitchell R. et al. (2016). "Mental fatigue impairs soccer-specific decision-making skill". In: *Journal of Sports Sciences*. ISSN: 1466447X. DOI: 10.1080/02640414.2016.1156241.
- Sniderman, Paul M, Richard A Brody, and Philip E Tetlock (1991). Reasoning and Choice: Explorations in Political Psychology. ISBN: 0521407702. DOI: Doi10.2307/1964392.
- Staal, M a (2004). "Stress, cognition, and human performance: A literature review and conceptual framework". In: NASA Technical Memorandum August, p. 168. URL: http://human-factors.arc.nasa.gov/ flightcognition/Publications/IH%7B%5C_%7D054%7B%5C_%7DStaal. pdf%7B%5C%%7D5Cnpapers3://publication/uuid/E92DA994-B825-40CA-9E74-A94101F33496.
- Starcke, Katrin and Matthias Brand (2012). "Decision making under stress: A selective review". In: *Neuroscience and Biobehavioral Reviews* 36.4, pp. 1228– 1248. ISSN: 01497634. DOI: 10.1016/j.neubiorev.2012.02.003.
- Stewart, Heather and Anushka Asthana (2018). Interview. Boris Johnson: '
 Let us have a grown-up conversation with our American friends '. London.
 URL: https://www.theguardian.com/politics/2018/jan/15/borisjohnson let us have a grown up conversation with our american-friends-interview.

- Stewart, Heather and Rowena Mason (2016). Nigel Farage's anti-migrant
 poster reported to police. URL: https://www.theguardian.com/politics/
 2016 / jun / 16 / nigel farage defends ukip breaking point poster-queue-of-migrants.
- Stone, Jon (2018). British public still believe Vote Leave '£350million a week to
 EU' myth from Brexit referendum. URL: https://www.independent.co.
 uk/news/uk/politics/vote-leave-brexit-lies-eu-pay-moneyremain-poll-boris-johnson-a8603646.html.
- Swaine, Jon (2018). "Trump inauguration crowd photos were edited after he intervened | World news | The Guardian". In: The Guardian 10, pp. 4–7. URL: https://www.theguardian.com/world/2018/sep/06/ donald-trump-inauguration-crowd-size-photos-edited.
- Swire, Briony et al. (2017). "Processing political misinformation: comprehending the Trump phenomenon". In: *Royal Society Open Science* 4.3, p. 160802. ISSN: 2054-5703. DOI: 10.1098/rsos.160802. URL: http://rsos.royalsocietypublishing.org/lookup/doi/10.1098/rsos.160802.
- Taber, Charles S. and Milton Lodge (2006). "Motivated skepticism in the evaluation of political beliefs". In: *American Journal of Political Science* 50.3, pp. 755–769. ISSN: 00925853. DOI: 10.1111/j.1540-5907.2006.00214.x.
- Tajfel, Henri (1970). "Experiments in intergroup discrimination." In: Scientific American 223.5, pp. 96–102. ISSN: 0036-8733. DOI: 10.1038/scientificamerican1170-96. arXiv: arXiv:1011.1669v3.
- (1974). "Social identity and intergroup behaviour". In: *Social Science Information* 13.2, pp. 65–93.

- Tajfel, Henri (1982). "Social Psychology of Intergroup Relations". In: Annual Review of Psychology 33.1, pp. 1–39. ISSN: 0066-4308. DOI: 10.1146/ annurev.ps.33.020182.000245. arXiv: arXiv:1011.1669v3.
- Tajfel, Henri and John C. Turner (1986). The social identity theory of intergroup behavior. DOI: 10 . 1111 / j . 1751 - 9004 . 2007 . 00066 . x. arXiv: arXiv:1011.1669v3. URL: http://psycnet.apa.org/psycinfo/2004-13697-016.
- Tesler, Michael (2016a). *Economic anxiety isn't driving racial resentment*. *Racial resentment is driving economic anxiety*.
- (2016b). Trump Voters Think African Americans Are Much Less Deserving Than 'Average Americans'. URL: https://www.huffingtonpost.com/ michael-tesler/trump-voters-think-africa%78%5C_%7Db%78%5C_ %7D13732500.html.
- The Sun (2016). We urge our readers to beLEAVE in Britain and vote to quit the EU on June 23. London. URL: https://www.thesun.co.uk/news/ 1277920/we-urge-our-readers-to-believe-in-britain-and-voteto-leave-the-eu-in-referendum-on-june-23/.
- Thorson, Emily (2015). "Belief Echoes: The Persistent Effects of Corrected Misinformation". In: *Political Communication* 4609.November, pp. 1–21. ISSN: 1058-4609. DOI: 10.1080/10584609.2015.1102187. URL: http://dx.doi.org/10.1080/10584609.2015.1102187.
- (2016). "Belief Echoes: The Persistent Effects of Corrected Misinformation". In: *Political Communication* 33.3, pp. 460–480. ISSN: 10917675. DOI: 10.1080/10584609.2015.1102187.

- Van Veen, Vincent et al. (2009). "Neural activity predicts attitude change in cognitive dissonance". In: *Nature Neuroscience*. ISSN: 10976256. DOI: 10.1038/nn.2413.
- Watts, Joe (2016). Brexit: Michael Gove clings on to Vote Leave's discredited £350m NHS claim. London. URL: https://www.independent.co.uk/ news/uk/politics/brexit-latest-news-michael-gove-nhs-claim-350m-twitter-vote-leave-eu-uk-a7498651.html.
- Westen, D et al. (2006). "The neural basis of motivated reasoning: An fMRI study of emotional constraints on political judgment during the U.S. Presidential election of 2004." In: J. Cognitive Neuroscience. 18, pp. 1947– 1958.
- Wlezien, Christopher, Mark Franklin, and Daniel Twiggs (1997). "Economic perceptions and vote choice: Disentangling the endogeneity".
 In: *Political Behavior*. ISSN: 01909320. DOI: 10.1023/A:1024841605168.
- Wood, Thomas and Ethan Porter (2016). "The Elusive Backfire Effect: Mass Attitudes' Steadfast Factual Adherence". In: SSRN Electronic Journal. ISSN: 1556-5068. DOI: 10.2139/ssrn.2819073. URL: https://www. ssrn.com/abstract=2819073.
- Wyer, R. S. and T. L. Budesheim (1987). "Person memory and judgments: The impact of information that one is told to disregard". In: *Journal* of Personality and Social Psychology 53, pp. 14–29. ISSN: 0022-3514. DOI: 10.1037/0022-3514.53.1.14. URL: http://www.ncbi.nlm.nih.gov/ pubmed/3612487.
- Yu, Rongjun (2016). "Stress potentiates decision biases: A stress induced deliberation-to-intuition (SIDI) model". In: *Neurobiology of Stress* 3, pp. 83–

95. ISSN: 23522895. DOI: 10.1016/j.ynstr.2015.12.006. URL: http: //dx.doi.org/10.1016/j.ynstr.2015.12.006.