Essays on Information and Fairness

Stefan Meissner

DEPARTMENT OF ECONOMICS

NHH Norwegian School of Economics

A thesis submitted for the degree of

Philosophiae Doctor (Ph.D.)

BERGEN 2020

Acknowledgements

First and foremost, I want to thank my main supervisor, Alexander W. Cappelen, for his continued support, countless hours of insightful discussions and most importantly an incredible amount of energy and motivation. Without your guidance and your relentless optimism and encouragement, my Ph.D. experience would have been immeasurably less inspiring and, frankly, fun.

I am also deeply grateful to Erik Sørensen for his patience and his deep insight. The work presented in this dissertation and my approach to scientific problems more generally have benefited immensely from your advice. Your critical but constructive comments have stopped some terrible ideas in their tracks while greatly improving the better ones.

Furthermore, I want to thank Bertil Tungodden for his support, in particular his role in the second chapter of this dissertation, and his work in creating the wonderful research environment in FAIR. Your suggestions greatly improved this chapter and offered me an amazing chance to learn and grow. The same can be said about Alexander's role in this project. Moreover, I want to extend a special thank you to my co-authors in Chapter 3 of the dissertation, Hanna Krasnova and Katharina Baum. Katharina is not only a co-author but also a great friend whose sharp insight generated great improvements.

I also want to thank my external supervisor, Björn Bartling, who was very generous with his advice and support and whose insightful comments helped me tremendously. My colleagues at the NHH have been very supportive

and helpful in their advice as well. In particular, I want to thank Hallgeir, Henning and Thomas who have been very supportive over the course of my Ph.D.. I am also very grateful for the support I received from the administrative teams of the department and FAIR. In particular I want to thank Janina, Sebastian, Anna and Adriana for the great work they did facilitating my experiments.

I also want to extend a big thank you to the fellow Ph.D. students in my cohort: Ceren and Sara. Without the two of you the Ph.D. would have been so much more difficult and dull. Furthermore, I want to thank all of my fellow Ph.D. students for contributing to a great social environment. In particular, I am thankful to Oda, Mads, Ingvild, Felix, Ingar, Xu, Alessandro, Daniel, Rene, Shrey, Charlotte, Ranveig, Kjetil and Mirjam for being good colleagues and great friends.

During my Ph.D., I was very fortunate to get the opportunity to go to different places and learn from inspiring researchers. In particular, I want to thank Georg Weizsäcker and Jana Friedrichsen for giving me the opportunity to spend time at the Humboldt-University Berlin and to learn from the great work that is done there. Furthermore, I want to thank Alex Imas, George Loewenstein and Linda Babcock for the time at Carnegie Mellon University and Hanna Krasnova, again, for the time at Weizenbaum Institute. I made some great memories during that time and grew a lot as researcher thanks to these opportunities.

I would also extend big thanks to my great friends and my family. Ich möchte Axel und Rita dafür danken, mich stets in meiner Neugier bestärkt, mir unglaublichen Rückhalt geboten und mir immer das Gefühl gegeben zu haben, erreichen zu können, was ich will. My most special thanks goes to

Nadja. I cannot even begin to list the things I want to thank you for and I cannot and do not want to imagine doing this without you.

Contents

Introduction

> Man the food-gatherer reappears
> incongruously as
> information-gatherer
>
> Marshall McLuhan

We live in the information age (Fuchs, 2007). New technological, scientific and societal developments have reconstructed main pillars of modern society like business and government (Gleick, 2011; Lyon, 2013). The production, provision and access to information has become a key economic activity that drives a substantial part of wealth creation (Floridi, 2010). Governments make use of mandatory information provision policies and increasingly utilize information interventions to reach a diverse set of policy goals, ranging from environmental protection to unemployment reduction (Bar-Gill et al., 2019; Ben-Shahar and Schneider, 2011; Sunstein, 2016). Furthermore, the rise of social media and the internet has made it clear that the use and abuse of information can shape outcomes of elections and referendums, incite violence and even lead to large geopolitical crises (Allcott and Gentzkow, 2017; Brooking and Singer, 2016; Lazer et al., 2018).

People play several important roles with regards to information. They are, at the same time, recipients as well as providers of information (Floridi, 2010, 2014). In their role as recipients, people acquire information for various

purposes, for example to help them make decisions or update their beliefs (Stigler, 1961). As information providers, people must decide which and how much of their information to provide to others, shaping their decisions and beliefs. To which information people have access can profoundly influence all sorts of important outcomes for them (Thaler and Sunstein, 2009). It can affect their professional lives, for example by contributing to their choice of profession (Coffman et al., 2017), their health, for example by informing people about calorie counts of food (Elbel et al., 2009) or their financial well-being, for example by promoting retirement saving (Fernandes et al., 2014). This illustrates that the impact that the informational environment that surrounds people and to which they contribute is immense (Webster, 2014). Especially with the rise of government-mandated information provision, people increasingly play a third role as regulators of informational environments. They have to decide whether they support or oppose policies that provide information, for example on retirement plans, energy use or product safety (Reisch and Sunstein, 2016) or how social media platforms, an increasingly important source of information, should be governed (Kruse et al., 2018).

This thesis contributes to our understanding of the three roles that people play in the information society. Traditionally, economists have viewed information as an instrumental means to reach better decisions and nothing else (Golman et al., 2017; Stigler, 1961). Information acquisition was described as a process in which people rationally updated their beliefs based on Bayes' rule if they received relevant information while irrelevant information was discarded (Benjamin, 2019). Therefore, receiving additional information was always either positive or neutral for the recipient as it either

improved their decisions or did not affect them. Information provision was seen as a way to achieve self-interested goals and maximize utility of the provider without regards to the recipient (Crawford and Sobel, 1982; Jin et al., 2015; Oehler and Wendt, 2017). These views are increasingly challenged by a large and growing literature in behavioral and experimental economics (DellaVigna, 2009; Loewenstein et al., 2014). Through the use of data that is collected in laboratory, field and online experiments, economists have vastly extended the understanding of how people act in their roles of recipients, providers and regulators of information. Behavioral economics has contributed in two main ways to the understanding of people in their different roles. The first strand of literature has documented that people are not perfectly rational when updating their beliefs based on new information (Tversky and Kahneman, 1983). Rather, they make systematic errors in that process, undervaluing older signals (Benjamin, 2019), ignoring the absence of information (Enke, 2020) or focusing too strongly on information that confirms pre-existing beliefs (Bénabou and Tirole, 2016). Due to these mistakes, information provision can negatively impact recipients, leading to worse decisions or biased beliefs. The second main contribution of behavioral economics is the study of people's social preferences. This literature has shown that people's actions deviate from pure self-interest. Rather, they take the outcomes of others into account when they make decisions (Charness and Rabin, 2002). Given that information provision strongly impacts recipients' outcomes, people might take these consequences into account when deciding which and how much information to provide. A large body of literature has empirically documented that actions are often motivated by factors beyond narrow self-interest and that people consider the outcomes of others when making important choices (Almås et al., 2010; Charness and Rabin, 2002; Fehr

and Schmidt, 1999; Fehr and Fischbacher, 2003). Other moral considerations also influence their decisions on which and how much information to provide. Previous research has documented that people are reluctant to lie and provide false information even when that would be in their own financial interest (Abeler et al., 2019; Gneezy, 2005; Gneezy et al., 2018). These findings illustrate the need to account for the absence of perfectly rational updating and social preferences when studying people's behavior in their role as recipients, providers and regulator of information. This thesis combines insights from both strands of the behavioral economics literature. It explores human behavior that is motivated by social preferences or biased beliefs. Thereby, this dissertation provides new insight into how people decide which policies to support, how they make decisions that influence outcomes for themselves and others and what they consider fair or unfair. More precisely, the thesis consists of three chapters that describe and understand behavior in settings in which people make information-related choices. It shows how people's actions as recipients, providers and regulators of information are profoundly shaped by their concern for others, their biases in dealing with information and the need and desire for transparent and truthful information.

To investigate how people act as recipients, providers and regulators of information, the present thesis uses experimental methods. The use of these methods was essentially nonexistent in economics for a long time. In the 1985 version of their popular undergraduate textbook, Samuelson and Nordhaus wrote that "[Economists] cannot perform the controlled experiments of chemists or biologists because they cannot easily control other important factors" (Samuelson and Nordhaus, 1985). The paradigm that economics cannot rely on data that is collected in experimental settings shifted in the 1970s and

1980s due to economists' increasing attention to systematic deviations from standard economic theory in people's behavior (Barberis, 2018). In a vast array of fields, economists documented that people are motivated by factors that lie outside the standard assumptions of rational models. Researchers increasingly made use of laboratory, online and field experiments to test theories and gain novel insight into the ways people form attitudes and beliefs, make decisions and to test the influence of certain factors in complex settings (List, 2011). In particular, the opportunity to tightly control the situation in which decisions are made has empowered economists to develop new models that better reflect human behavior, to understand market interactions and to develop new, efficient policy measures to achieve important societal goals (List, 2008). Furthermore, due to random assignment of participants, experiments allow researchers to draw causal conclusions from their results (Thaler, 2015). This makes it possible to directly infer the impact of different factors and conditions on people's behavior, something that is often very difficult or even impossible to do outside the realm of experimental settings. To ensure robustness of results and engagement of participants, economists often use incentives, for example financial compensation or the promise of real-world impact, when they run experiments (Camerer and Hogarth, 1999). These incentives can take different forms. Participants' decisions can either directly impact outcomes for themselves (stakeholders) or they can act as impartial third-parties whose actions influence other people's outcomes (spectators) (Cappelen et al., 2007; Konow, 2009). Both methods allow economists to answer important questions about human behavior. Stakeholder designs enable the study of people's actions in situations in which their outcomes are directly affected, for example in their roles as market participants (Bartling et al., 2015; Falk and Szech, 2013; Thaler, 1980) or voters (Bond et al., 2012;

DellaVigna et al., 2016). Spectator designs however allow economists to measure people's attitudes, for example towards fairness, because their actions are unaffected by considerations of self-interest (Almås et al., 2020). The three recent Nobel Prizes to Daniel Kahneman and Vernon Smith (2002), Richard Thaler (2015) and Abhijit Banerjee, Esther Duflo and Michael Kremer (2019) reflect that the use of experiments has become one key method in the toolkit of economics and has already provided new and crucial findings. The present thesis follows the experimental tradition. The three chapters of this thesis report results from four tightly-controlled, incentivized experiments. Both methods, spectator and stakeholder designs were used to shed light on people's behavior in their roles as recipients, providers, and regulators of information and to causally investigate how this behavior is affected by social preferences and biased beliefs.

Economists' increasing focus on factors beyond self-interest and the growing use of experimental methods has broadened the topics that economists study (Angrist and Pischke, 2010; Angrist et al., 2017). While in the past, economics was largely insulated from other social and behavioral sciences (Fourcade et al., 2015; Pieters and Baumgartner, 2002), resulting in important gaps in economists' understanding of key social phenomena (Akerlof, 2020), it has become increasingly common for economists to rely on work done in other fields, especially psychology and political science (Angrist et al., 2020). The more comprehensive scope of economic research and the inclusion of non-economic work enables economists to contribute to the solution of pressing societal issues, including income and wealth inequality (Alesina et al., 2018), discrimination (Acquisti and Fong, 2020), and the new challenges that social media poses to well-being (Allcott et al., 2020) and democracy (Enikolopov et

al., 2020). The three chapters in this dissertation are motivated and informed by research that is conducted in fields beyond economics. They build on work that has been done by economists, political scientists, psychologists and other scientists alike. All chapters benefited from helpful comments from researchers from a variety of fields. This wide range of insights makes it possible to present findings that directly contribute to important current political debates.

The following questions are addressed in the three chapters below:

Chapter 1: *Do people perceive information provision as a moral obligation?*

Chapter 2: *Are voluntary transactions in which one side misunderstands information and is taken advantage of by the other perceived as acceptable?*

Chapter 3: *What drives people's opposition to political microtargeting, a way to address audiences in a highly granular way with tailor-made information and messages?*

Although all chapters address very different research questions, they are connected by important commonalities. The most striking of course is the focus on information. However, the common threads run deeper than that. All the chapters present results from large incentivized online experiments with general population samples of the United States. The use of experimental methodology offers the opportunity to investigate the influence of attitudes and beliefs on people's behavior. Due to random assignment of participants to treatment conditions and a high level of control over the situation in which

decisions are made, it is possible to draw causal conclusions about the impact that certain factors have on people's behavior with regards to their roles as recipients, providers and regulators of information. While the first two chapters rely on spectator designs as they study people's attitudes towards fairness, the third chapter focuses more directly on people's support for rules that could impact their own lives. The use of general population samples allows for gaining a broader understanding of public attitudes that go beyond the relatively narrow confines of student populations. Participants in all of the studies make decisions that will have a real-world impact. To contribute to transparency in empirical research, experimental designs and main hypotheses of all studies were pre-registered at the AEA RCT Registry (Chapters 1 and 3) or the Open Science Foundation (Chapter 2) before data analysis (Munafò et al., 2017; Nosek et al., 2015).

The first chapter of the dissertation is titled *"There is something you should know - the moral obligation to provide information"*. This chapter studies whether people are willing to provide information even when that leads to outcomes that they dislike. It is motivated by the mounting evidence that information provision can have negative consequences for some recipients while helping others (Persson, 2018; Roberts, 2017). Therefore, information provision can contribute to or create inequality between recipients. A large body of literature has shown that people in general are averse to inequality (Cappelen et al., 2007, 2013b; Charness and Rabin, 2002; Engelmann and Strobel, 2004; Fehr and Schmidt, 1999). As a consequence, they could be willing to withhold part of their information to avoid undesirable distributive outcomes. However, a vast literature in philosophy and medical ethics posits that people could feel morally obliged to provide information for reasons that are unrelated

to the material outcomes information provision produces, for example, as a means to provide other people with autonomy (Cox and Fritz, 2016; Kymlicka, 2002).

The chapter reports the results of two large incentivized experiments that investigate whether people are willing to provide information that has distributive consequences to which they are averse. Participants in the experiment make decisions that have a real impact on the outcomes of two other people. The results of both studies strongly support the hypothesis that people have the attitude that there is a moral obligation to provide information. As a consequence of this perception, a significant share of people are willing to accept outcomes they dislike when they are the consequence of information provision. This finding is robust across settings and demographic and political groups. Further exploratory data analysis reveals a close link between people's attitudes towards autonomy and the belief that there is a moral obligation to provide information. In line with the discussion in philosophy, participants that prefer autonomy over outcomes also seem more likely to accept outcomes that they dislike if they are the consequence of information provision (Fallis, 2018; Harris and Keywood, 2001). This suggests that people see the provision of information as a way to provide other people with autonomy and that the respect for autonomy is an important source of the perception that there is a moral obligation to provide information.

The second chapter of this thesis has the name *"Cancel the deal? An experimental study on the exploitation of irrational consumers"* and is co-authored with Alexander W. Cappelen and Bertil Tungodden. Economists pay increasing attention to situations in which companies take advantage of consumers who misunderstand information about products or mispredicting their own

future behavior, thereby increasing profits (Akerlof and Shiller, 2015; Heidhues and Kőszegi, 2018). In these settings, companies and consumers agree to one-sided transactions which only benefit the sophisticated seller while harming the naive buyers (Korobkin, 2003). Yet, often these transactions are uncoerced and voluntary, without either side lying to the other. Therefore, different conceptions of fairness could drive people's attitudes towards these transactions. This study focuses on the question on whether people perceive these voluntary but exploitative transactions as acceptable or not.

The study introduces an experimental framework that allows to measure people's acceptance of transactions in which a seller exploits the naivety of a buyer and to disentangle their reasoning for that judgment. Participants can decide whether they want to cancel a transaction between a pair of other people, a buyer and a seller. In the experimental setting, the seller takes advantage of the buyer who misunderstood information about the value of the product. The findings reveal that a large majority of Americans perceives such transactions as unacceptable and are willing to cancel them. This attitude seems to be driven by aversion to the one-sided outcomes of such transactions. However, about a third of Americans considers these transactions as fair, even when the seller behaved unethically towards the buyer by obfuscating important information about the value of the product and proposing a one-sided deal. This shows that for a substantial share of the population, fairness concerns are closely connected with considerations of autonomy.

In summary, the paper shows that a large share of participants perceives it as unfair if people who do not properly understand information are exploited by others. This strongly suggests that full, transparent information provision is seen as an important foundation for a fair transaction.

The third chapter of this dissertation is co-authored with Hanna Krasnova and Katharina Baum and is called *"Partisan self-interest is an important driver for people's support for the regulation of targeted political advertising"*. This chapter examines the role of partisan self-interest, the drive to benefit one's own political party, in people's attitude towards microtargeted political advertising online. Microtargeted political advertising has become an increasingly important way for politicians to reach out to potential and likely voters (Aral and Eckles, 2019). It requires the collection and use of vast amount of personal information about voters, raising concerns about personal privacy (Acquisti et al., 2016; Boerman et al., 2017). Furthermore, because messages and information are designed for potentially small groups of voters and are inaccessible for others, this creates an intransparent public discourse (Sunstein, 2018). As a consequence of this lack of transparency, people have to form beliefs about how others are influenced by political advertising and how that will affect political outcomes.

Participants in the study are asked if they support stricter government control of microtargeted political advertisement. To incentivize truthful responses, their answers are aggregated and presented to members of the United States Senate and House of Representatives. The study reveals that a significant share of the public opposition to targeted political advertising is motivated by potentially biased beliefs about their impact on electoral outcomes. Partisans, who are informed that their preferred party has benefited from the use of targeted political advertising in the past become less supportive of stricter government regulation. This suggests that they are willing to accept violations of personal privacy as long as this this will help their party win. This outcome reveals the necessity of an informed and transparent public

debate about the impact of technological innovation on electoral processes. In this setting, information changed people's attitudes towards targeted political advertising, suggesting that current opinions on their stronger regulation might be biased due to the scarcity of insight about the effects of microtargeting and its intransparent nature.

Taken together, the results presented in the three chapters show how people act in the roles as recipients, providers, and regulators of information. The studies add to the growing behavioral and experimental literature that investigates how people's actions in these roles deviate from classical assumptions. By combining insight from the literature on social preferences and systematic biases, the findings show that people highly value information provision and transparency and that they strongly consider the outcomes of others that are the consequence of their actions. The three chapters emphasize the importance of information on economic and public life and open up new avenues for future research. The experimental paradigms that are presented in the three chapters enable future studies to further disentangle the motivations of people in their roles as recipients, providers, and regulators of information in an incentive-compatible, tightly-controlled manner. They make it possible to further build on the results that were presented here and to deepen our understanding of the role of information in people's lives.

Chapter 1

There is something you should know - the moral obligation to provide information

Stefan Meissner*

Abstract

Information provision can lead to gains for some but losses for other recipients. Because recipients might differ in their priors, or their abilities to update beliefs, information provision could negatively affect some recipients' decisions while helping others. This can create or contribute to inequality between recipients. This paper studies whether people are willing to provide information even when that leads to an unequal outcome, which they dislike. To do so, two large incentivized experiments are run in which Americans decide if they want to provide information. The findings of both experiments give robust causal evidence that people have the attitude that information provision is a

*Meissner: Department of Economics, NHH Norwegian School of Economics. I would like to thank Björn Bartling, Alexander W. Cappelen, Dirk Engelmann, Jana Friedrichsen, Alex Imas, Erik Sørensen, Bertil Tungodden, and seminar participants in Bergen and Berlin for helpful comments and discussions. Financial support from the Research Council of Norway through its Centre of Excellence Scheme (FAIR project No 262675) is gratefully acknowledged.

moral obligation and that they accept an outcome which they dislike when it is the consequence of information provision. The results further show that people who highly value others' autonomy are more likely to implement an unequal outcome they dislike if it is the consequence of information provision. This suggests that respect for autonomy is an important source of the moral obligation to provide information.

Keywords: Inequality acceptance, Information provision, Paternalism

JEL Classification: D63, D64, D82

1.1 Introduction

Companies, governments and people must frequently decide how much information they want to provide (Bar-Gill et al., 2019). Examples include sellers of financial services who decide how much information about their products they provide to potential buyers (Bolton et al., 2007), governments deciding how much information to give to citizens to achieve policy goals (Altmann et al., 2018), or doctors and teachers who have to choose how much information to reveal to patients and students (Caplin and Leahy, 2004). Information provision is often presented as an important obligation to respect people's "right to know" and their autonomy (Sunstein, 2020). However, information provision is not always helpful for recipients (Ben-Shahar and Schneider, 2011; Sunstein, 2019a). It is increasingly clear that decisions to provide information can sometimes negatively impact decision quality, thereby causing losses for some recipients (Lipnowski et al., 2020; Persson, 2018; Roberts, 2017).

Recent empirical work has documented negative consequences of information provision on decisions and welfare for some recipients in a wide range of settings, including financial information (Campbell, 2016; Frydman and Wang, 2020; Medina, 2018), or information about energy usage (Allcott and Kessler, 2019). Different explanations for some recipients' losses have been introduced in the literature. Most prominently, a large body of literature has shown that people who hold specific priors or have low abilities to correctly update beliefs based on new information can be hurt by information provision. For example, Andrade et al. (2019) present a case in which new information about monetary policy is beneficial to part of the recipients as it

helps them form more accurate beliefs about future economic circumstances. Yet, they show that the same information can cause others, who hold different prior beliefs, to become overly optimistic or pessimistic and adjust expectations about macroeconomic variables and investment behavior in a non-optimal way. Dohmen et al. (2009) further show that a large number of people are unable to draw correct inference from new information because they suffer from base-rate neglect, thereby not correctly accounting for the base-rate probability of an event. Base-rate neglect can shift some recipients' actions too strongly into the direction that the new information suggests, thereby contributing to worse decision-making (Benjamin, 2019; Hare et al., 2011; Trevino, 2020). For example, Hirshleifer (2001) presents evidence that some investors suffer from base-rate neglect and therefore overreact to new information, leading to worse trading-decisions and lower returns for them while other investors correctly adjust beliefs based on the new information.

Because information provision can negatively affect the outcome of some but create gains for others, it potentially creates or contributes to inequalities in wealth and income. Sunstein (2016) argues that it is important to consider the distributive consequences of information provision because they can influence decision-makers in their choice of how much information they want to provide. Previous research has shown that a majority of people are averse to inequality (e.g., Fehr and Schmidt (1999)) and that for a large share of the population, preferences for equal outcomes can outweigh efficiency considerations (e.g., Alesina and Giuliano (2011); Almås et al. (2020); Engelmann and Strobel (2004)). These preferences have an impact on individual choices, for example the decision how to distribute earnings between parties (e.g., Cappelen et al. (2007, 2013b); Charness and Rabin (2002); Konow (2003)). This suggests that people might be averse to the unequal outcome

that information provision produces. Consequently, they might prefer to withhold information even when provision would increase overall efficiency.

A broad literature in philosophy and ethics argues that people might want to provide information to others for non-consequentialist reasons. People could believe that withholding information from others infringes upon their autonomy to freely choose how to make use of it (Bok, 1989; Cox and Fritz, 2016; Epstein et al., 2010; Fallis, 2018; Harris and Keywood, 2001; Kymlicka, 2002; Schick, 1994). They could contend that people are freer with information (Sunstein, 2020). As previous research has shown, people highly value others' autonomy and prefer measures that are autonomy preserving or enhancing (Jung and Mellers, 2016; Reisch and Sunstein, 2016). Further, people might perceive withholding information as infringing upon others' right to know (Floridi, 2010; Holley, 1998; Zimmerman and Bradley, 2019). This attitude has already shaped important policy decisions, leading to the increasing prevalence of transparent information provision policies (Bennear and Olmstead, 2008; Berliner, 2014). For these reasons, people might see information provision as a moral obligation. As a result of this attitude, they could be willing to provide information and accept an unequal outcome which they dislike if it resulted from information provision.

This paper explores whether people are willing to provide information even when this produces an unequal outcome which they dislike. It reports from two large-scale incentivized studies which are conducted to investigate if people are willing to provide information even when that leads to an unequal outcome which they dislike. In total, more than 3000 participants from general population samples of the United States took part in the two studies. The participants act as impartial third-parties (spectators) and make a choice

that has real distributive consequences for two workers. In each study, there are two conditions in which spectators make a choice that results in either the equal but inefficient, or the unequal but efficient distribution of earnings for the two workers.[2,3] Spectators are randomly selected into one of the two experimental conditions. In the baseline condition, spectators can directly implement either the unequal or the equal outcome. Measuring which outcomes participants prefer to implement in this condition reveals the share of participants who prefer the equal outcome and the share who prefer the unequal one. In the treatment condition (information treatment), the participants can indirectly implement the exact same outcomes through their choice whether to provide or withhold information. Spectators are informed that information provision leads to the unequal outcome while withholding information results in the equal one. Spectators in the baseline condition and the information treatment receive the same information about the payout consequences of their actions. As participants are randomly assigned to baseline condition and information treatment, distributive preferences should on average be the same in the two treatments. Therefore, the comparison of the share of participants who decide to provide information with the share of participants who directly choose the same unequal outcome makes it possible to identify whether participants perceive information provision as a moral obligation and are willing to accept an unequal outcome which they dislike in order to respect that obligation. If a larger share of participants provide information than directly implement the unequal outcome, this would show that some participants who prefer the equal outcome are still willing to provide information that leads to the unequal outcome. This would suggest

[2]Throughout the paper, efficiency refers to the situation in which the sum of individual payouts is maximized(Güth et al., 2003).

[3]The equal but inefficient allocation is referred to as the equal outcome while the unequal but efficient allocation is called the unequal outcome.

that some participants accept the unequal outcome if it is the consequence of information provision, and thus gives evidence for the existence of the attitude that information provision is a moral obligation.

In both studies, spectators in the information treatment must decide how much information to provide. Each study investigates one crucial setting in which information can lead to losses for some recipients and gains for others. The first study focuses on *differences in priors*. Spectators are told that the decision to provide additional information will negatively impact one worker and increase the earnings of the other due to the differences in prior information that the workers had. It was random which prior information workers received; therefore the distribution of gains and losses depends on luck. The second study focuses on *differences in the ability to draw correct inference from new data*. Spectators are informed that the provision of additional information will harm the earnings of a worker who suffers from base-rate neglect and increase those of the other who does not. In the baseline conditions of both studies, spectators can directly implement the same payout consequences for the two workers who differ either in luck (study 1) or ability (study 2). Investigating whether people are willing to provide information in two important settings in which information provision leads to losses for some but gains for other recipients makes it possible to understand whether the attitude that information provision is a moral obligation robustly exists in different situations. The experimental designs and main hypotheses are pre-registered at the AEA RCT registry.[4]

The results of both studies show that the share of participants who implement the unequal outcome is significantly higher in the information treatment compared to the baseline condition. In the first study, 62 percent of partici-

[4]Study 1: AEARCTR-0003376 , Study 2: AEARCTR-0005018

pants choose to provide the workers with additional information, thereby implementing the unequal outcome. In the baseline condition, 40 percent of participants are willing to implement the unequal outcome for the two workers directly. In the second study, 40 percent of participants choose to provide additional information, producing the unequal outcome. 31 percent of participants choose the unequal distribution of earnings in the baseline condition. These findings suggest that people have the attitude that information provision is a moral obligation and that, to respect that obligation, they are willing to accept the unequal outcome which they dislike. However, in both studies, a substantial share of participants are willing to withhold information when that produces the equal outcome for the two workers.

The two studies provide evidence for demographic heterogeneity in the attitude that information provision is a moral obligation. Furthermore, the results suggest that there is a link between people's attitude towards autonomy and their attitude that information provision is a moral obligation. Participants, who state that they believe that the government should not restrict citizens' freedom even when that is in their best interest are more likely to provide information that leads to an outcome which they dislike than participants who believe the opposite. This result supports the interpretation that people's attitude that information provision is a moral obligation is linked to their intrinsic valuation of autonomy and suggests that this obligation might be a consequence of their respect for autonomy.

This paper contributes to the growing literature in economics and psychology (e.g., Benartzi et al. (2017); Sharot and Sunstein (2020); Sunstein (2016)) that investigates how distributive concerns influence information provision. This study is the first to show that people have the attitude that a moral obligation

for information provision exists and that, in order to respect that obligation, they accept an unequal outcome which they dislike if it is the consequence of information provision. By showing that a significant share of people are willing to accept an unequal outcome which they dislike when it is the consequence of information provision, this paper adds novel insight into this literature by offering a new foundation for the widespread use of information provision. This finding suggests that at least a fraction of people perceive it as a moral obligation to provide information. This might justify policy initiatives which focus on providing citizens with more information even when their distributive outcomes are undesirable (Bao and Ho, 2015; Loewenstein and Chater, 2017). However, a large share of participants are willing to withhold information when that leads to the equal outcome that they prefer for the potential recipients. A large literature has shown that distributive preferences are linked to policy preferences (Alesina et al., 2018; Almås et al., 2020; Fisman et al., 2017; Roth and Wohlfart, 2018). This finding should therefore caution policy makers as it could suggest that, even though a significant fraction perceives information provision as a moral obligation, a large share of people are willing to withhold information because they are averse to the outcome information provision would produce.

This paper further introduces a new approach to studying people's attitude towards information provision. While previous papers (e.g., Khalmetski et al. (2017); Serra-Garcia et al. (2011); Sheremeta and Shields (2017)) focused on situations in which people had a monetary incentive to provide or withhold information, this study presents a framework in which people's decision whether to provide information only influences other people's earnings (Rodriguez-Lara and Moreno-Garrido, 2012). This allows to cleanly identify whether information provision is at least partially driven by a moral

obligation instead of strategic or monetary considerations. Understanding this makes it possible to study situations in which outcomes of information providers are not directly influenced by their decision whether to provide information, for example, government agencies which provide information on energy usage.

The findings of this paper further contribute to the literature on social preferences. Overall, the results of this paper are in line with previous work that demonstrates that preferences for equal outcomes can outweigh efficiency considerations (e.g., Almås et al. (2020)). Furthermore, a large body of literature has shown that people's preferences over distributive outcomes for others influence their decisions (e.g., Almås et al. (2010); Andreoni and Miller (2002); Cappelen et al. (2013b); Durante et al. (2014); Fehr and Schmidt (1999)). This study provides evidence that for a substantial share of people, the moral obligation to provide information can influence their acceptance of distributive outcomes which they dislike. Comparing the share of participants who are willing to provide information with the share that are willing to directly implement the same outcome provides causal evidence that some people accept unequal outcomes which they dislike when they are the consequence of information provision. This finding reveals that people are willing to trade-off their aversion to inequality with the moral obligation to provide information. However, the studies also provide evidence that people's preference for an equal outcome can justify withholding information, emphasizing the importance of distributive preferences in the decision to provide information.

The findings further add to a small but growing literature in economics that shows how people's valuation of their own and other people's autonomy can

influence important decisions they make (Ambuehl et al., 2019; Bartling et al., 2014; Bobadilla-Suarez et al., 2017; Fehr et al., 2013). The result of the present study suggest that there is a link between the intrinsic valuation of autonomy and the moral obligation to provide others with information. People who value other's autonomy highly are more likely to accept outcomes which they dislike if they are the consequence of information provision. This suggests that respect for autonomy is an important source of the attitude that information provision is a moral obligation, thus suggesting that people use information provision as a way to give others autonomy. The findings of the present paper illustrate that the intrinsic valuation of other's autonomy can influence people's behavior in their role as information providers.

The paper proceeds as follows. The next section discusses the first study in more detail. In the third part of the paper, the second study is presented. Section four concludes the paper.

1.2 Study 1

1.2.1 Experimental design

The first study focused on a situation in which the information recipients hold different priors and, for this reason, information provision leads to monetary gains for one and losses for the other recipient. Two types of participants took part in the study, spectators and workers. Spectators were informed that in contrast to a hypothetical survey question, their decision might have real consequences for other people. Spectators were randomly assigned to one of two treatments and made an incentivized decision that,

with a certain probability, determined the earnings of two workers. In the baseline condition of the experiment, spectators had the choice to directly implement either the equal or the unequal earnings distribution for the two workers. In the information treatment, they had to decide whether they wanted to provide or withhold information from the two workers. Spectators were informed that their decision would, given that people follow the information they receive, have equivalent payout consequences as the decision in the baseline condition. Withholding information would result in the equal outcome for the two workers while providing information would lead to the unequal outcome. Experimental design and main hypotheses were pre-registered at the AEA RCT-Registry.

Spectators

Spectators were recruited in collaboration with the market research company Dynata. 2419 spectators from a general population sample of the United States took part in the study. Spectators were informed that their decision might be implemented.[5] Spectators were broadly representative of the adult population of the United States across age, income and location in census regions. The median age of participants was 45. About 52.5 percent of participants were women. Table 1.7 in the appendix provides an overview of the characteristics of the participants. The spectators were randomly assigned to one of two treatments. 1212 spectators were assigned to the baseline condition and 1207 participants were assigned to the information treatment. Table 1.8 in the appendix shows that treatment assignment was balanced on observable characteristics.

[5]The matching from spectators to pairs of workers was 1:20.

Workers

In total 242 workers were recruited on Amazon Mechanical Turk. 122 workers were assigned to the baseline condition and 120 workers were assigned to do the information task.[6] Two workers were paired with each other. In the baseline condition, spectators had to answer a few survey questions. Their earnings were determined by the decision that spectators made for the pair of workers. In the information task, workers had to complete a guessing task in which they had to estimate the share of red balls in an urn. Their earnings depended on the accuracy of their estimate.

Treatments

In the baseline condition, spectators could directly implement either an equal or an unequal outcome for the two workers. They were presented with one option that would lead to the pair of workers being paid $3.30 for their participation in an experiment. Alternatively, spectators could choose to change these payments to $5.00 for one and $2.50 for the other worker. Making that change resulted in an overall increase in payouts but introduced inequality between the two workers. Spectators in this treatment were explicitly informed that the allocation of gains and losses from the change was random and due to luck.

In the information treatment, spectators had to make the choice if they wanted to provide additional information. They were presented with a situation in which two workers had to estimate the share of red balls in an urn. They were informed that the earnings of the workers depended on the

[6]The median age of the workers was 35 years. 41.3 percent of workers identified as female.

accuracy of their estimate. If the worker correctly estimated the share of red balls in the urn, he or she would earn \$5.00. The information from the initial draws that the pair of workers received would put both workers equally far off the correct solution, by presenting either a share of red balls of 33 percent or 67 percent. Each percentage point off the correct 50 percent solution lowers the worker's earnings by \$0.10. Therefore, given that workers follow the information suggested by their prior information, both would earn \$3.30 because they are 17 percentage points away from the correct solution.

Spectators had the choice whether they wanted to present the initial draw or draw an additional ball from the urn and present it to the worker. The additional information from the draw would always result in one worker seeing a 50 percent share of red balls instead of the 33 percent or 67 percent they received in the initial draw while the other either sees a 25 percent or a 75 percent share.[7] Spectators were told that, given that they follow their information, one worker's earnings will increase from \$3.30 to \$5.00 (moving from 17 percentage points deviation to 0 percentage points) while the others' decrease from \$3.30 to \$2.50 (moving from 17 percentage points deviation to 25 percentage points).[8] Adding the additional ball would therefore lead to an increase in overall payouts but would also cause inequality between the two workers. The worker who gains from the additional information earns an additional \$1.70 if he or she follows the information while the other worker loses \$0.80. Spectators were further informed that, equivalent to the baseline condition, the color of the additional ball and therefore the distribution of gains and losses was random and not due to ability or effort.

[7] In the initial draw, one worker is presented with the draw that resulted in 2 blue and 1 red ball while the other receives information about the draw that resulted in 2 red and 1 blue ball. If the fourth ball is red, this would result in 3 blue and 1 red (25 percent) for the first while the other gets the result for the draw that resulted in 2 red and 2 blue ball (50 percent). The equivalent is true if the additional ball is blue.

[8] The modal responses of workers was equal to the information spectators received.

The appendix contains the full instructions to spectators. Further, information on age, gender, education, household income and attitudes towards paternalism of spectators were collected.

1.2.2 Empirical strategy

In both treatments, the share of spectators who implemented the unequal outcome was measured. Luck as a source of inequality was kept constant across treatments. Spectators in the information treatment were informed that, given that participants follow the information they receive, their decision would lead to the exact same payout consequences as the lottery in the baseline condition.[9] To identify whether people are willing to accept an unequal outcome which they dislike when it is the consequence of information provision, choices between the baseline condition and the information treatment are compared. Measuring the share of spectators who directly implement the unequal outcome in the baseline condition reveals the share of participants who prefer the efficient over the equal outcome. In the information treatment, the same outcomes are the result of the decision whether to provide information. Due to random assignment to treatment, distributive preferences of spectators should not vary between baseline condition and information treatment. This makes it possible to causally identify if a significant fraction of people are willing to accept an unequal outcome which they dislike when it is the consequence of information provision. If the share of spectators who provide information is higher than the share who directly implements the same outcome, that would show that some spectators who prefer the equal outcome are still willing to provide information that leads

[9] See instructions for exact description.

to the unequal outcome. This finding would suggest that these participants have the attitude that information provision is a moral obligation and that, to respect that obligation, they accept the unequal outcome which they dislike. This outcome provides a lower bound for the prevalence of people who have the attitude that information provision is a moral obligation. It might be that among participants who prefer the unequal outcome and participants who implement the equal outcome in the information treatment, some still have the attitude that information provision is a moral obligation.

To test whether the distributions of participants' choices between baseline condition and information treatment differed, a χ^2-test is performed. The results of the χ^2-test reveal whether the share of participants who chose to implement the equal and the share of participants who chose to implement the unequal outcome are different between baseline condition and information treatment. Further, an OLS-regression with the following specification is run:

$$Y_i = \beta_0 + \beta_1 \text{Information}_\text{i} + \beta X_i + \epsilon_i, \tag{1.1}$$

where Y_i is an indicator variable for the decision that leads to the unequal outcome and $Information_i$ is an indicator variable for assignment to the information treatment. All control variables, X_i are included as binary variables which are 1 if the participant is above the median in that demographic variable (or is female).

When excluding the control variables, β_0 measures the share of participants who implement the unequal distribution in the baseline condition. Therefore, β_0 identifies how large the share of spectators who prefer the unequal and efficient outcome over the equal one is. $1 - \beta_0$ measures the share of participants who prefer the equal outcome. The share of participants who are

willing to implement the unequal outcome in the information treatment is $\beta_0 + \beta_1$. $1 - \beta_0 - \beta_1$ measures the share of participants who withhold information to achieve an equal outcome. β_1 reveals the treatment effect. Given the assumptions that spectators believe that the payout consequences of information provision are as described to them and that preferences for outcomes on average do not differ between the two treatments due to randomization, β_1 allows to identify whether information provision causally impacts people's willingness to accept the unequal outcome which they dislike. β_1 is the difference between baseline condition and information treatment in the share of participants who implement the unequal outcome. If β_1 is positive and significantly different from 0, this provides causal evidence that a significant fraction of participants are willing to provide information even when they would prefer the equal outcome. This is interpreted as evidence for the existence of the attitude that information provision is a moral obligation as it would show that participants who prefer the equal outcome accept the unequal one when it is the consequence of information provision.

In addition, differences in treatment effects between demographic groups are analyzed. Analyzing whether the treatment effects vary across different demographic groups makes it possible to show if there are differences in people's willingness to provide information that leads to the unequal outcome which they dislike based on demographics. Understanding this could help to uncover whether certain demographic variables are linked to people's attitude that a moral obligation to provide information exists.

The following regression to analyze the heterogeneous treatment effects is run for each of the demographic background variables:

$$Y_i = \beta_0 + \beta_1 \text{Information}_\text{i} + \beta_2 x_i + \beta_3 x_i \text{Information}_\text{i} + \beta X_i + \epsilon_i, \qquad (1.2)$$

Y_i is an indicator variable for the decision to implement the unequal distribution of earnings. $\mathrm{Information_i}$ is an indicator variable for treatment assignment. x_i is the demographic variable which is tested and X_i is a vector of all other demographic variables. Variables x_i that are tested are (i) age, (ii) income, (iii) gender and (iv) education.

β_1 measures the treatment effect for participants for whom the tested variable is 0. The treatment effect on participants who are above median in the tested variable or female is measured as $\beta_1 + \beta_3$. β_3 gives insight into whether the size of the treatment effect differs between the low/high categories in that variable or between males and females respectively.

People's attitude that information provision is a moral obligation might also be driven by moral attitudes. One potential source of this might be respect for the autonomy of recipients. Restricting access to information could be perceived as a violation of this autonomy because withholding information takes other's freedom to decide how to respond to it (Kozlowski and Sweanor, 2016). Therefore, OLS regression 1.2 was run to test whether the size of the treatment effect is associated with participants' attitudes towards autonomy. This finding would suggest the attitude that information provision is a moral obligation is linked to people's attitude towards autonomy. Participants' attitudes towards autonomy were elicited with the question of whether they support government restrictions of freedom if that would be in the citizens' best interest. For participants that stated that they do not support government action in these cases, x_i is 1.

1.2.3 Results

Main analysis

Figure 1.2 provides an overview of spectator decisions. Overall, 51.1 percent

Figure 1.1: Study 1 - Share of spectators who choose the unequal distribution

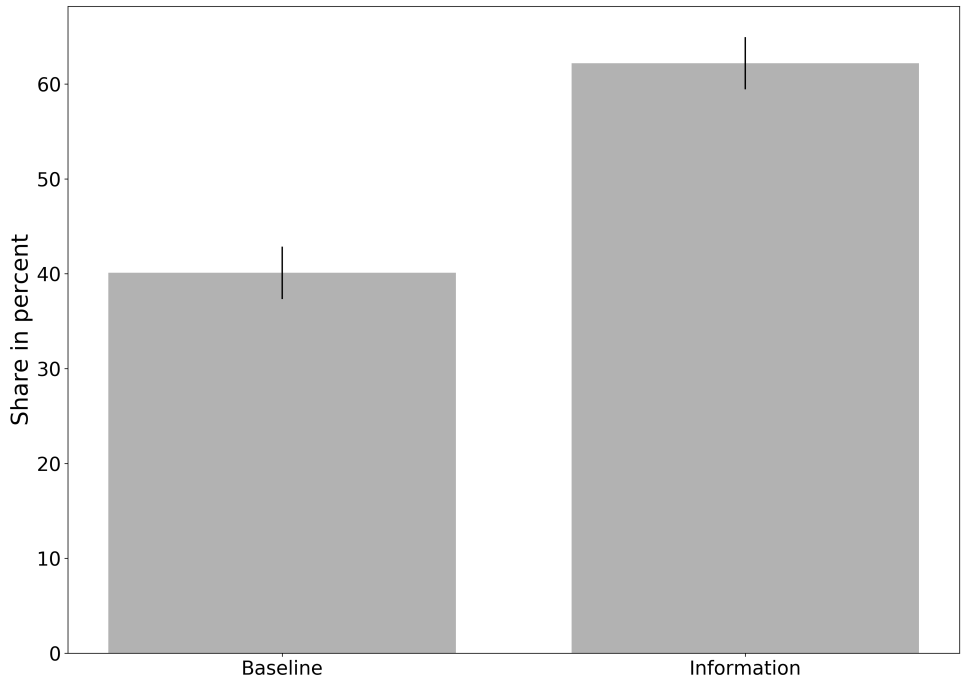

Note: The figure shows the share of spectators in the baseline condition and the information treatment who chose the alternative that leads to the unequal outcome for the two workers. The bars show 95% confidence intervals.

of participants chose to implement the unequal distribution of earnings for the workers. There are large differences in the share of spectators who chose the unequal outcome between treatments. 40.1 percent of participants chose to implement the unequal distribution in the baseline condition while 62.2 percent of participants did so in the information treatment.

The results of the main analysis reveal a significant difference in the participants' choices between the information treatment and the baseline condition.

The χ^2-test for differences in the distributions between the two conditions is highly significant (p < 0.001, $\chi^2 = 117.63$).

Table 1.1 shows the results of regression 1.1.

Table 1.1: Regression results for implemented inequality

	Inequality	Inequality	Inequality	Inequality	Inequality	Inequality
Information	0.221***	0.216***	0.220***	0.220***	0.220***	0.213***
	(0.020)	(0.020)	(0.020)	(0.020)	(0.020)	(0.020)
Age		-0.166***				-0.176***
		(0.020)				(0.020)
Income			0.026			0.015
			(0.020)			(0.022)
Female				-0.062***		-0.076***
				(0.020)		(0.020)
Education					0.019	0.012
					(0.020)	(0.022)
Constant	0.401***	0.485***	0.388***	0.434***	0.392***	0.517***
	(0.014)	(0.018)	(0.017)	(0.018)	(0.017)	(0.025)
Observations	2410	2410	2410	2410	2410	2410
R^2	0.049	0.070	0.050	0.053	0.049	0.076

Note: The table reports results from an OLS-regression of the share of spectators who implement the unequal outcome for the two workers. Information is an indicator variable which takes the value 1 if the participant was assigned to the information treatment. Age is an indicator variable with the value 1 if the participant was older than 45 years. Income is an indicator variable which is 1 if the participant was above the median household income ($60,000) in the sample. Female is an indicator variable for being female. Education is an indicator variable which is 1 when the participant has at least some college education. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Column 1 shows a large and significant (p < 0.001) difference in the willingness of spectators to implement the unequal distribution of earnings between baseline condition and information treatment. The result of the baseline condition reveals that approx. 60 percent of participants prefer the equal to the efficient outcome. The share that implements this outcome declines significantly when participants must withhold information to achieve equality. β_1

shows that the fraction of spectators who are willing to provide information that leads to inequality is 22.1 percentage points higher than the share of spectators who directly implement this outcome. This parameter is virtually unaffected by the inclusion of background variables as shown in columns 2-6. This leads to the following results:

Result 1: *More participants are willing to implement the unequal outcome in the information treatment than in the baseline condition.*

The increase in the share of participants who prefer the unequal outcome indicates that out of approx. 60 percent of participants who prefer an equal distribution of earnings, about one third is willing to accept the unequal and efficient outcome if it is the consequence of information provision. This result suggests that a substantial fraction of participants hold the attitude that information provision is a moral obligation.[10] However, 38.8 percent of participants in the information treatment withhold information to produce the equal distribution of earnings between the two workers. This demonstrates that a large share of participants are still willing to withhold information if that results in the equal distribution of income for the two workers that they prefer.

Result 2: *A large share of participants withholds information to produce the equal outcome for the two workers.*

Columns 2 – 6 further suggest that there are differences in the willingness to

[10]Participants had to complete two challenging attention checks in the experiment. All tests are also run with a sample which excludes participants who failed both attention checks in the survey. The instructions of the attention checks are available in the appendix part C. Appendix table 1.11 shows the regression results for this sample. The findings of this robustness check confirm the findings of this analysis.

implement the unequal outcome between groups. Columns 2 and 4 show that older participants and women are significantly less likely to implement the unequal distribution. The result for women is in line with previous research which shows that women are on average more inequality averse than men (Capraro, 2020; Croson and Gneezy, 2009).

Heterogeneity Analysis

Table 1.2 shows the results of the analysis for heterogeneity across demographic subgroups. Across all subgroups, the parameter for the assignment

Table 1.2: Regression results for heterogeneity across demographic groups

	Inequality			
	Age	Income	Gender	Education
	B = 1 if > 45	B = 1 if > $ 60,000	B = 1 if female	B = 1 if college
Information	0.139***	0.206***	0.196***	0.170***
	(0.028)	(0.028)	(0.028)	(0.028)
B*Information	0.151***	0.015	0.033	0.089**
	(0.039)	(0.039)	(0.039)	(0.039)
B	-0.251***	0.008	-0.092***	-0.032
	(0.027)	(0.029)	(0.028)	(0.029)
Constant	0.556***	0.521***	0.526***	0.537***
	(0.027)	(0.026)	(0.027)	(0.026)
$\beta_1 + \beta_2$	0.290***	0.221***	0.229***	0.259***
	(0.028)	(0.028)	(0.027)	(0.028)
R^2	0.083	0.076	0.076	0.078
Control variables	Yes	Yes	Yes	Yes
Observations	2410	2410	2410	2410

Note: The table reports results from an OLS-regression of the share of spectators who implement the unequal outcome for the two workers. Information is an indicator variable which takes the value 1 if the participant was assigned to the information treatment. B is an indicator variable which is 1 when the participant is older than 45 (Column 1), has an household income higher than $60,000 (Column 2), is female (Column 3) or has at least some college education (Column 4). B*Information are interactions between assignment to the information treatment and the respective demographic variable. $\beta_1 + \beta_2$ is the linear combination of the variables B*Information + Information. Control variables include all other demographic variables. Robust standard errors in parentheses. $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$

to the information treatment remains large and significant. Figure 1.3 in the appendix illustrates this finding. Across all groups, the share the treatment effect is significantly bigger than 0. That shows that in all groups, a substantial fraction are willing to accept the unequal outcome which they dislike when it is the consequence of information provision. The increase in shares of participants who implement the unequal outcome between baseline condition and information treatment ranges from 13.9 percentage points for participants who are younger than 45 to 29.0 percentage points for participants who are older than that age. This suggests that the finding that people have the attitude that information provision is a moral obligation is robust across demographic subgroups. Further, there is significant heterogeneity in treatment effects across some demographic subgroups. The increase in the share of spectators who implement the unequal distribution between baseline condition and the information treatment is significantly larger for older participants compared to younger individuals ($p < 0.001$). The same is true for better educated participants compared to less educated participants ($p = 0.022$). There is no evidence for differences in treatment effects between women and men and high- and low-earning participants.

Result 3: *The treatment effect is significantly stronger among older participants and college-educated participants compared to younger and non-college-educated participants.*

52.3 percent of the younger participants prefer to implement the efficient outcome in the baseline condition. This share increases to 66.6 percent in the information treatment. Only 28.3 percent of participants who are older than 45 implement the unequal outcome in the baseline condition.

This share more than doubles in the information treatment. 57.4 percent of older participants decide to provide information that leads to the unequal outcome. This finding suggests that there is a link between age and the attitude that information provision is a moral obligation. In this study, the increase in the share of participants who implement the unequal outcome increases significantly stronger for older participants, suggesting that they perceive a stronger obligation to provide information. However, the share of older participants who prefer the unequal outcome is still smaller than the equivalent share among younger participants.

Among college-educated participants, the share of participants who implement the unequal outcome increases from 38.6 percent to 65.6 percent between baseline condition and information treatment. 41.5 percent of non-college-educated participants implement the unequal outcome in the baseline condition. This share increases to 58.8 percent in the information treatment. This result suggests that higher education correlates with a stronger attitude that information provision is a moral obligation.

The same analysis is run with the sample that excludes participants who failed both attention checks in the study. Table 1.12 in the appendix shows the full results. The finding that people across all subgroups are willing to accept the unequal outcome which they dislike when it is the consequence of information provision is confirmed by these findings. However, the difference in the treatment effect between college-educated and non-college-educated participants becomes statistically insignificant.

The role of autonomy

Table 1.3 reports results for the analysis of heterogeneity in treatment effects depending on participants' attitudes towards autonomy. Column 1 shows

Table 1.3: Regression results for heterogeneity across different attitudes towards paternalism

	Inequality	Inequality
Information	0.210***	0.125***
	(0.019)	(0.031)
Non-Paternalist	-0.118***	-0.188***
	(0.020)	(0.028)
Non-Paternalist*Information		0.140***
		(0.040)
Constant	0.585***	0.631***
	(0.027)	(0.029)
β_2 + β_3		0.265***
		(0.025)
R^2	0.090	0.094
Control Variables	Yes	Yes
Observations	2410	2410

Note: The table reports results from an OLS-regression of the share of spectators who implement the unequal outcome for the two workers. Information is an indicator variable which takes the value 1 if the participant was assigned to the information treatment. Non-Paternalist is an indicator variable which is 1 if participants either strongly disagreed, disagreed or slightly disagreed with the statement "The federal government should restrict citizens' freedom if it is in the citizens' own best interests". In total 1470 participants disagreed with this statement. Non-Paternalist*Information is an interaction variable between assignment to the information treatment and participants' attitudes towards paternalism. $\beta_2 + \beta_3$ is the linear combination of the variables Non-Paternalist*Information + Information. Control variables include all demographic variables. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

that participants who state higher support for government paternalism are more likely to implement the unequal outcome. Column 2 shows the result of OLS-regression 1.2. It shows that both groups, people that support government intervention and people who oppose it, are willing to accept the

unequal outcome which they dislike when it is the consequence of information provision. The share of participants who are willing to implement an unequal outcome increases more strongly between baseline condition and information treatment for participants who highly value autonomy.

Result 4: *The treatment effect is significantly stronger among participants who highly value autonomy compared to participants who do not.*

The increase in the share who provide information compared to the share that directly implements the same payout consequences is approximately twice as large for them than for participants who indicate that they value outcomes over autonomy. Among people who support government restrictions of freedom if this is in the person's best interest, 52.8 percent implement the unequal outcome in the baseline condition. In the information treatment, 66.6 percent do so. Among participants who highly value autonomy, the share who implement the efficient outcome increases from 32.4 percent to 59.2 percent. Result 4 suggests that there is a connection between people's valuation of autonomy and their attitude towards information provision. The result is in line with the idea that providing others with information is perceived as a way to give them the opportunity to freely make decisions, indicating that respect for the autonomy of others could be an important source of the attitude that information provision is a moral obligation.

The next section presents the second study.

1.3 Study 2

1.3.1 Experimental design

While the first study investigated a situation in which information provision leads to inequality between recipients due to differences in their priors, this study focused on a situation in which information provision leads to inequality due to differences in the ability to correctly update beliefs between the two recipients. As in the first study, two types of participants were recruited to take part in the survey. The first type, spectators, made decisions that affected the monetary outcomes for the second type, workers. Spectators knew that they made choices that might have real consequences for two other people, a pair of workers. The study had two treatments and spectators were randomly assigned to one of them. As in the first study, spectators' choices in the baseline condition revealed their preferences for either the equal or the unequal outcome for the pair of workers. In the information treatment, spectators made the choice whether to provide or withhold information from the workers. In line with the first study, information provision would lead to the unequal outcome while withholding information resulted in the equal distribution of income. The experimental design and the main hypotheses were pre-registered at the AEA RCT-Registry.

Spectators

The study was conducted together with Norstat. 1004 participants from a general population sample of the United States were recruited for the study.

Spectators were informed that their decision might have real consequences.[11] Table 1.9 in the appendix presents demographics of the participants. Compared to the first study, participants in this study were slightly older. The median age of participants was 55. Further, the share of female participants was slightly higher in this study. 58.3 percent of participants were women. About 39.6 percent of participants identified as Republican. 501 participants were assigned to the baseline condition while 503 were assigned to the information treatment. Table 1.10 in the appendix shows that the treatment assignment was balanced on most observable characteristics.

Workers

Workers were recruited on Amazon Mechanical Turk. First, their ability to correctly estimate the probability of an event based on base-rate and additional information was pre-tested. 45 workers who suffered from base-rate neglect (incorrectly answer in the pre-test) and 45 Bayesian (correct answer in the pre-test) workers were identified and took part in the main task.[12] Low (base-rate neglecters) and high ability (Bayesians) workers were matched to pairs. 23 pairs were assigned to the information treatment and 22 pairs to the baseline condition. Workers' bonuses in both treatments depended on their pre-tested ability to correctly update their beliefs as they had to complete the same task. In the task, workers had to guess the probability of a lottery ticket paying out. The task that workers had to solve could be either simple or more complex. Workers were paid based on the accuracy of their estimate about the probability of the lottery ticket paying out.

[11] The matching from spectators to pairs of workers was 1:21.

[12] The median age of participants was 35. 46.7 percent of participants identified as female.

Treatments

In both treatments, the spectators were informed that they made a decision that affected a pair of workers in which one had a high ability to solve the complex task and one worker had a low ability to do so. Spectators in both treatments were told that abilities were pre-tested. In the two treatments, they decided which version of the task to assign to the pair of workers. In neither treatment did spectators know the exact task that the workers were asked to do.[13] Expected payout consequences of the choices of spectators were kept the same as in the first study.

In the baseline condition, spectators were presented with the choice which math task they want to assign to a pair of workers. Spectators could directly implement either the simple or the complex version of the math task.[14] Spectators were told that in the simple version of the math task, given workers pre-tested abilities, both would perform equally well while in the complex math task the high ability worker would be able to correctly solve the problem while the low ability worker would be unable to do so. Assigning the simple task would result in both workers earning $3.30 while the high ability worker would gain $1.70 to earn $5.00 in the complex task and the low ability worker would lose $0.80 to earn $2.50 based on their pre-tested abilities. Therefore, implementing the simple task leads to the equal outcome while implementing the complex task results in unequal earnings.

In the information treatment, spectators made the choice how much informa-

[13] See appendix for full instructions to the spectators.

[14] In the simple version of the task, workers only received information that lottery tickets like this have a base-rate probability to pay out that is 5 percent. In the more complex version, the workers additionally received information that a scanner with 90 percent sensitivity and specificity tested the ticket and predicted that it is a winning ticket.

tion they want to provide to the pair of workers. Spectators had to decide if they wanted to only provide the base-rate probability or provide base-rate and additional information to the two workers. Their decision whether to withhold or provide information lead to the simple or complex version of the math task. They were informed that providing additional information would, given the pre-tested abilities, enable high ability workers to correctly update their beliefs about the probability of the lottery ticket paying out while the low ability worker would drastically overestimate the probability of the ticket winning.[15] Given that both workers only receive information about the base-rate probability when the spectator decides to withhold information, this version of the task does not require updating. Spectators were therefore informed that the equal performance of the two workers would result in both of them earning $3.30. They were also told that their decision to provide information would, given the workers' abilities, increase the earnings of the high ability worker to $5.00, a gain of $1.70 while lowering them to $2.50, a loss of $0.80 for the worker with low abilities. Therefore, spectators were presented with two options in the information treatment that lead to the same payout consequences as the choice in the baseline condition.

The appendix contains full instructions to the spectators. Further, information on age, gender, household income and education were elicited. To understand whether political differences are linked to the attitude that information provision is a moral obligation, participants were asked to state their party preference.

[15] In the sample, only high ability workers were able to correctly solve the complex task.

1.3.2 Empirical strategy

The share of spectators who made a decision that lead to the unequal distribution was measured in both conditions. Ability as the source of inequality was kept constant across the two treatments. Measuring spectators' choices in the baseline condition made it possible to understand participants' preferences for outcomes in this scenario. Spectators' decision to assign the simple math task is interpreted as a preference for the equal over the efficient outcome while the opposite is true for spectators that assign the complex math task. Payout consequences of the decision to provide or withhold information and assign the simple or complex math task were described in the same way to the spectators in the information treatment and the baseline condition, respectively. Spectators were told that withholding information would lead to the equal outcome while information provision would result in the unequal outcome. Comparing whether more participants are willing to provide information than to directly implement the unequal outcome makes it possible to measure whether people are willing to accept the unequal outcome, which they dislike when it is the consequence of information provision. This finding would be interpreted as evidence that these people have the attitude that information provision is a moral obligation and that, to respect that obligation, they provide information even though they prefer the equal outcome. Therefore, this study makes it possible to identify the existence of the attitude that information provision is a moral obligation. However, the study only provides a lower bound for the prevalence of this attitude. Its prevalence among participants who already prefer the unequal outcome and participants whose aversion to inequality is sufficiently strong cannot be measured.

To study whether differences between choices in the two treatments exist, the distribution of choices in the baseline condition (simple or complex task) is compared to the distribution of choices in the information treatment (base-rate, or base-rate and additional information) with a χ^2-test. To further measure whether a difference between choices in the baseline condition and the information treatment exists, the following OLS-regression is run:

$$Y_i = \beta_0 + \beta_1 \text{Information}_\text{i} + \beta X_i + \epsilon_i, \tag{1.3}$$

where Y_i is 1 when the spectator makes a decision that leads to the unequal outcome. $\text{Information}_\text{i}$ is an indicator variable for assignment to the information treatment. All control variables, X_i, are included as binary variables which are 1 if the participant is (i) older than 45, (ii) has a yearly household income above \$60,000, (iii) is female and (iv) is college educated. That is equivalent to the groups in the first study.

In the regression without control variables, β_0 measures the share of spectators who implement the complex math task. This share is interpreted as participants who prefer the efficient over the equal outcome for the two workers. $1 - \beta_0$ is the share of participants who implement the simple math task which, based on pre-tested abilities, leads to the equal outcome for the two workers. The share of participants who decide to implement the unequal outcome in the information treatment is $\beta_0 + \beta_1$. The fraction of participants who are willing to withhold information, thereby implementing the equal outcome, is $1 - \beta_0 - \beta_1$. β_1 measures the treatment effect. If participants in the baseline condition and the information treatment believe that payout consequences of their choices are as described to them and that, due to random assignment, the share of participants who prefer the equal outcome

is the same in both treatments, β_1 makes it possible to measure if people's acceptance of an unequal outcome depends on whether it is the consequence of information provision. If β_1 is significantly larger than 0, this indicates that a significant fraction of participants who prefer the equal outcome are willing to provide information. This would suggest that these participants have the attitude that information provision is a moral obligation and that, to respect that obligation, they accept the unequal outcome which they dislike.

Furthermore, the robustness of the treatment effect across different subgroups is investigated. This analysis allows to better understand whether demographic variables are linked to participants' willingness to accept the unequal outcome which they dislike when it is the consequence of information provision. For this purpose, the following OLS-regression is run:

$$Y_i = \beta_0 + \beta_1 \text{Information}_\text{i} + \beta_2 x_i + \beta_3 x_i \text{Information}_\text{i} + \beta X_i + \epsilon_i, \qquad (1.4)$$

Y_i is an indicator variable for the implementation of the unequal distributions of earnings, $\text{Information}_\text{i}$ is an indicator variable for assignment to the information treatment and x_i is the demographic variable which is tested. X_i is a vector of all other demographic variables.

In the regression without control variables, β_1 offers a measure for the treatment effect among participants for whom x_i is 0. The linear combination of β_1 and β_3 measures the treatment effect for participants for whom x_i is 1. If β_3 is significantly different from 0, that provides evidence for differences in treatment effects between different demographic groups. This finding would suggest an association between this demographic variable and people's attitude that information provision is a moral obligation.

Study 1 provided evidence that people's intrinsic valuation of autonomy constitutes an important source of the moral obligation to provide information. This study analyzes whether this relationship is closely linked to party preferences. Previous research has shown that Republicans are more likely to accept outcomes of voluntary choices (Jost et al., 2009, 2003) and are opposed to interventions into these outcomes (Skitka, 1999; Skitka and Tetlock, 1993). Therefore, the OLS-regression 1.4 is run which tests whether Republicans are more willing to provide information that leads to the unequal outcome which they dislike than non-Republicans. In this regression, the variable x_i is an indicator variable for Republicans. The interaction effect between $\text{Information}_\text{i}$ and x_i allows to measure whether the willingness to provide information that leads to the unequal outcome which that they dislike varies between Republicans and non-Republicans.

1.3.3 Results

Main analysis

Figure 2 displays the shares of spectators who chose to implement the unequal outcome. Overall, 36.0 percent chose to implement the unequal distribution of earnings for the workers. The share of spectators who chose to provide additional information (40.8 percent) is higher than the share of spectators who assigned the complex math task (31.1 percent).

A χ^2-test provides strong evidence that the two distributions of choices are significantly different between information treatment and baseline condition (p = 0.001, $\chi^2 = 10.08$). Table 1.4 shows the results of OLS-regression 1.3.

Table 1.4: Regression results for implemented inequality

	Inequality	Inequality	Inequality	Inequality	Inequality	Inequality
Information	0.096***	0.103***	0.096***	0.097***	0.096***	0.104***
	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)	(0.030)
Age		0.089***				0.076**
		(0.031)				(0.031)
Income			0.081***			0.066**
			(0.030)			(0.033)
Female				-0.092***		-0.081**
				(0.031)		(0.031)
Education					0.010	-0.033
					(0.030)	(0.033)
Constant	0.311***	0.251***	0.272***	0.365***	0.307***	0.289***
	(0.021)	(0.029)	(0.024)	(0.028)	(0.024)	(0.039)
Observations	1004	1004	1004	1004	1004	1004
R^2	0.010	0.018	0.017	0.019	0.010	0.030

Note: The table reports results from an OLS-regression of the share of spectators who implement the unequal outcome for the two workers. Information is an indicator variable which takes the value 1 if the participant was assigned to the information treatment. Age is an indicator variable with the value 1 if the participant was older than 45 years. Income is an indicator variable which is 1 if the participant was above a household income of $60,000. Female is an indicator variable for being female. Education is an indicator variable which is 1 when the participant has at least some college education. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1.2: Study 2 - Share of spectators who choose the unequal distribution

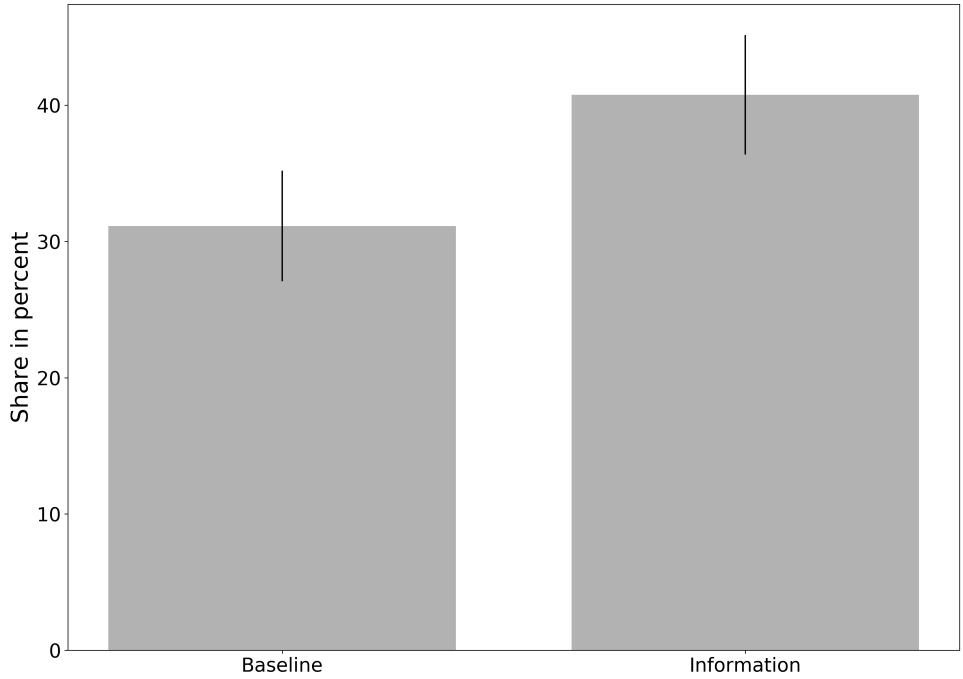

Note: The figure shows the share of spectators in the baseline condition and the information treatment who chose the alternative that leads to the unequal outcome for the two workers. The bars show 95% confidence intervals.

Overall the results show that a majority of participants in both treatments prefer the equal income distribution to the efficient one. However, there is a large and significant ($p < 0.001$) difference in the share of spectators who implement the unequal distribution between the baseline condition and the information treatment. The share of spectators who provide additional information is 9.6 percentage points higher than the share of participants who assign a complex math task. Columns 2 - 6 show that this parameter is almost unaffected by the inclusion of demographic control variables. Summarized, this leads to:

Result 5: *More participants are willing to implement the unequal outcome in the information treatment than in the baseline condition.*

In the baseline condition, about 70 percent of participants indicate that they prefer an equal to an efficient outcome. This share decreases to approx. 60 percent in the information treatment. This suggests that a significant fraction of people who are averse to inequality are willing to accept an unequal outcome if it is the consequence of information provision. This result provides robust evidence for the existence of the attitude that information provision is a moral obligation. Fewer people are willing to implement an equal outcome for two workers if that would require that they withhold information. However, in this study, a large majority is still willing to withhold information.

Result 6: *A large share of participants withholds information to produce the equal outcome for the two workers.*

This finding underlines the importance of distributional preferences in the decision to provide information. In the setting where inequality is due to differences in abilities, a majority of spectators prefer that less information is provided.

Together with the results of the first study, this shows that the moral obligation to provide information can influence people's decisions across different settings. Comparing the findings to the first study reveals a somewhat smaller treatment effect and overall lower levels of spectators implementing the unequal earnings for the two workers. This finding might be explained by differences in the source of inequality. Previous research (e.g., Almås et al. (2020)) has shown that the source of inequality matters for people's attitude towards it. In this setting, spectators might be more averse to inequality which is the result of innate ability compared to luck. This could be due to

beliefs which spectators hold about the affected groups. They might perceive people with low abilities as more vulnerable and are therefore averse to losses for this group. It is striking that a majority of spectators are willing to withhold information in the setting where information provision would negatively impact low ability workers. The overall lower levels might, however, also be driven by differences in samples between the first and the second study.

Column 2 further shows that older participants are significantly more likely (p = 0.004) to implement the unequal earnings distribution. Column 3 reveals that the share of high income participants who implement the unequal earnings is significantly (p = 0.007) higher than that of low income individuals. Column 4 shows that, in line with a larger literature, women are significantly (p = 0.003) less likely to implement the unequal distribution than men.

Heterogeneity Analysis

Table 1.5 shows the regression results for the analysis of heterogeneity in treatment effects across different demographic groups. Across all subgroups, the parameter for the assignment to the information treatment is positive. Yet, not all of the parameter are significant. Figure 1.4 in the appendix displays this finding. The smallest increase in the share of participants who implement the unequal outcome happens among college-educated participants. 35.7 percent prefer the unequal outcome in the baseline condition. This share increases to 37.6 percent in the information treatment. The highest increase is among non-college-educated participants for whom the share who prefer the unequal outcome increases from 27.4 percent to 43.8 percent. In all groups, a majority of participants prefer to withhold information however.

Table 1.5: Regression results for heterogeneity across demographic groups

	Inequality			
	Age	Income	Gender	Education
	B = 1 if > 45	B = 1 if income > $ 60,000	B = 1 if female	B = 1 if college
Information	0.072	0.162***	0.068	0.165***
	(0.049)	(0.041)	(0.048)	(0.041)
B*Information	0.049	-0.120**	0.062	-0.132**
	(0.062)	(0.060)	(0.062)	(0.061)
B	0.050	0.128***	-0.112***	0.035
	(0.043)	(0.044)	(0.043)	(0.045)
Constant	0.307***	0.257***	0.309***	0.257***
	(0.044)	(0.042)	(0.043)	(0.042)
$\beta_1 + \beta_2$	0.121***	0.043	0.130***	0.034
	(0.038)	(0.044)	(0.039)	(0.044)
R^2	0.031	0.034	0.031	0.035
Control Variables	Yes	Yes	Yes	Yes
Observations	1004	1004	1004	1004

Note: The table reports results from an OLS-regression of the share of spectators who implement the unequal outcome for the two workers. Information is an indicator variable which takes the value 1 if the participant was assigned to the information treatment. B is an indicator variable which is 1 when the participant is older than 45 (Column 1), has a household income higher than $60,000 (Column 2), is female (Column 3) or is college-educated (Column 4). B*Information are interactions between assignment to the information treatment and the respective demographic variable. $\beta_1 + \beta_2$ is the linear combination of the variables B*Information + Information. Control Variables include all other demographic variables. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

There is evidence for significant heterogeneity in treatment effects. Participants without college education and low income have a significantly higher treatment effect than participants with college education and with high income. There is no evidence for heterogeneity between younger and older and men and women.

Result 7: *The treatment effect is significantly stronger among non-college-educated and low-income participants compared to college-educated and high-income participants.*

This outcome suggests that among lower income and lower education participants, a higher share is willing to accept a distribution which they dislike when it is the consequence of information provision. This finding could suggest that these variables are associated with a stronger attitude that information provision is a moral obligation.

Taken together, the results of the two studies suggest that demographic variables might be linked to people's attitude that information provision is a moral obligation. This comparison also makes clear however that more research is needed to provide a conclusive picture which variables interact with this attitude.

Political preference

The first study revealed that participants' preferences for autonomy over outcomes can influence their willingness to provide information that leads to an unequal outcome which they dislike. Previous research has shown

that Republicans are more likely to emphasize the importance of personal autonomy and tend to reject restrictions that protect people from the consequences of their choices (Janoff-Bulman, 2009; Skitka et al., 2002). Therefore, party preferences might be associated with people's willingness to accept the unequal outcome which they dislike when they are the consequence of information provision. Column 2 of Table 1.6 shows if this association exists.

Table 1.6: Regression results for implemented inequality

	Inequality	Inequality
Information	0.106***	0.078**
	(0.030)	(0.038)
Republican	0.067**	0.032
	(0.031)	(0.042)
Republican*Information		0.071
		(0.062)
Constant	0.265***	0.280***
	(0.041)	(0.043)
$\beta_1 + \beta_3$		0.148***
		(0.048)
Observations	1004	1004
Control Variables	Yes	Yes
R^2	0.036	0.037

Note: The table reports results from an OLS-regression of the share of spectators who implement the unequal outcome for the two workers. Information is an indicator variable which takes the value 1 if the participant was assigned to the information treatment. Republican is an indicator variable which is 1 when the participant self-identifies as Republican. Republican*Information is an interaction variable which is 1 when the participant is Republican and assigned to the information treatment. $\beta_1 + \beta_3$ is the linear combination of Republican*Information and Information. Control Variables include all other demographic variables. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Column 1 provides evidence that Republicans are more likely to implement the unequal outcome for the two workers (p = 0.036). While across treatments 32.8 percent of non-Republicans implement the unequal distribution of income, this share is 40.7 percent among Republicans. This finding matches

the existing literature as it suggests that Republicans are more concerned about efficiency compared to equality than non-Republicans (Cappelen et al., 2018a,b; Kerschbamer and Müller, 2020). In the baseline condition, 29.3 percent of non-Republicans are willing to make a decision that leads to the unequal outcome. This share significantly increases to 36.2 percent in the information treatment. For Republicans, the increase between baseline condition and information treatment is almost twice as large. The share of participants who implement the unequal outcome increases from 33.8 percent to 48.2 percent. However, column 2 shows no significant difference in the treatment effect between Republicans and non-Republicans (p = 0.250).

Result 8: *There is no evidence for a difference in the treatment effect between Republicans and non-Republicans.*

Further, even among Republicans, a majority are still willing to withhold information when that results in the equal outcome for the two workers.

1.4 Concluding remarks

This paper reports results of two large-scale incentivized studies with two demographically diverse samples which measure whether people provide information even when that leads to an unequal outcome which they dislike. Comparing the shares of spectators who provide information that leads to the unequal outcome with the share of spectators who directly implement the same outcome makes it possible to causally identify if participants are willing to implement the unequal outcome which they dislike when it is the consequence of information provision. Across studies, spectators are

significantly more likely to implement the unequal outcome when it is the consequence of information provision rather than implementing it directly. Thus, the findings of the two studies support the hypothesis that at least some people have the attitude that information provision is a moral obligation and that, to respect that obligation, they provide information even when that leads to the unequal outcome which they dislike. This outcome robustly holds for two different mechanisms with which information provision can lead to gains for some and losses for other recipients. However, important differences between the results of the two studies exist. In particular, the relatively high willingness of participants to withhold information in the second study where that protects a low ability worker shows that, in certain situations, people are still opposed to the inequality that is the result of information provision. This result shows that aversion to inequality can shape people's decision whether to provide or withhold information. Future research is needed to better understand how situational factors determine the impact of people's aversion to inequality on their willingness to provide information.

The results of the two studies suggest that the attitude that information provision is a moral obligation is robust across demographic groups. There is some heterogeneity between groups in the willingness to provide information when that leads to the unequal outcome which they dislike. However, the comparison between studies does not make it possible to draw general conclusions about which groups are more willing to do so. Given that the two studies do not consistently allow to link demographic variables to the attitude that information provision is a moral obligation, it is unclear whether such links exists. This could be in line with the idea that this attitude is not only heterogeneous across but also within individuals and could depend

on the parameters of the situation (Gibson et al., 2013; Mazar et al., 2008). More research is needed to understand how demographic and situational factors interact in determining whether people perceive a moral obligation to provide information.

Exploratory results show that the intrinsic valuation of autonomy might be one source of the attitude that information provision is a moral obligation. The present results suggest that people perceive a moral obligation to provide information because they see it as a means to enable others to make their own decisions based on this information. This finding is in line with a large literature in philosophy and ethics that posits a close link between access to information and autonomy. This extends the result that people value policies which provide information in part because they are autonomy-preserving (e.g., Camerer et al. (2003)) as it suggests that they perceive these policies as autonomy-enhancing.

To study whether preferences for information provision in the study are associated with the participants' attitudes towards information provision more generally, spectators in the second study were presented with a hypothetical scenario in which a firm had to decide whether to provide simple or detailed information to consumers.[16] As shown in tables 1.13 and 1.14 in the appendix, a χ^2-test reveals that participants' choices in the information treatment are closely linked to their preferences in the hypothetical scenario while their choices in the baseline condition are not. Table 1.15 in the appendix presents results from an OLS-regression for participants in the information treatment. This analysis shows that participants who provided information in the study are more likely to support detailed information provision than participants

[16]The description of the hypothetical choice is available in the appendix and is similar to the scenario in Newell and Siikamäki (2014).

who withheld information in the study. These findings suggest that information provision in the study is tightly associated with more general attitudes of people towards information provision in different circumstances.

The findings of the present paper might inform the literature on lying costs. As suggested in Cappelen et al. (2013a) and reflected in a wider literature in philosophy, people might see withholding information as akin to lying and could therefore value the provision of information as an important moral obligation. The results of this paper support the idea that people have the attitude that information provision is a moral obligation. As a consequence of this moral obligation, people could be unwilling to falsify or withhold information they have regardless of the outcome which that behavior produces (Abeler et al., 2019; Cappelen et al., 2013a; Erat and Gneezy, 2012; Gneezy et al., 2018; Kerschbamer et al., 2019; López-Pérez and Spiegelman, 2013; Shalvi and Leiser, 2013). Yet, in both studies, a significant share of participants are willing to withhold information from others. This suggests that some people are willing to withhold information to produce the distributive outcome they prefer. This opens an interesting future avenue for research to understand how situational factors influence people's attitude that information provision is a moral obligation and potentially their willingness to lie. Future studies could investigate whether people are willing to lie if that results in a distributive outcome which they prefer. The present paper introduces a suitable experimental framework to approach this research question.

While this paper establishes the existence of the attitude that information provision is a moral obligation, several questions remain unanswered. Future research could try to link other moral attitudes with this moral obligation (Cappelen et al., 2013a). As the literature in philosophy posits, people might

perceive information provision as a contribution to a wider public good (Floridi, 2013). Understanding if this motive is a source of people's attitude that information provision is a moral obligation could provide an important insight into people's decision-making with regards to information provision. Further, while this study does not show a significant link between participants' attitude that information provision is a moral obligation and party affiliation, studying whether other political views influence this attitude could offer important insight. To gain a deeper understanding of the factors determining people's attitude towards information provision, it could further be necessary to understand if cultural factors are linked to the moral obligation to provide information.

1.A Overview

The following appendix presents demographic variables about the participants in the two studies (part A), additional results (part B), the instructions that were provided to the participants (part C) and shows and explains deviations from the pre-analysis plan (part D).

1.B Demographics

Table 1.7: Descriptive Statistics - Study 1

	Number of participants	Share
Household income less than or equal to $30,000	618	0.26
Household income between $30,001 and $60,000	608	0.25
Household income between $60,001 and $100,000	531	0.22
Household income between $100,001 and $150,000	337	0.14
Household income above $150,000	325	0.13
High School Education or below	727	0.30
Some College Education	515	0.21
Bachelor or equivalent	801	0.33
Master or equivalent	376	0.16
Between 18 and 34 years old	752	0.31
Between 35 and 44 years old	441	0.18
Between 45 and 54 years old	401	0.17
Between 55 and 64 years old	400	0.17
Older than 64 years	425	0.18
Female	1270	0.53
Male	1149	0.47
Observations	2419	

Note: The table shows demographic information about all participants in the first study.

Table 1.8: Descriptive Statistics - Treatments Study 1

	Baseline	Information	Difference
Household income less than or equal to $30,000	0.25	0.26	-0.01
Household income between $30,001 and $60,000	0.26	0.24	0.02
Household income between $60,001 and $100,000	0.21	0.23	-0.02
Household income between $100,001 and $150,000	0.14	0.14	0.00
Household income above $150,000	0.14	0.13	0.01
High School Education or below	0.32	0.29	0.03*
Some College Education	0.21	0.21	-0.00
Bachelor or equivalent	0.32	0.34	-0.02
Master or equivalent	0.15	0.16	-0.01
Age in years	46.46	45.87	0.59
Female	0.53	0.52	0.02
Non-Paternalist	0.62	0.59	0.03
Observations	1212	1207	

Note: The table shows the share of participants in the demographic categories and the mean for the category age in years. Differences were tested with t-tests. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.9: Descriptive Statistics - Study 2

	Number of participants	Share
Household income less than or equal to $30,000	264	0.26
Household income between $30,001 and $60,000	249	0.25
Household income between $60,001 and $100,000	226	0.23
Household income between $100,001 and $150,000	164	0.16
Household income above $150,000	101	0.10
High School Education or below	232	0.23
Some College Education	303	0.30
Bachelor or equivalent	261	0.26
Master or equivalent	208	0.21
Between 18 and 34 years old	199	0.20
Between 35 and 44 years old	143	0.14
Between 45 and 54 years old	146	0.15
Between 55 and 64 years old	238	0.24
Older than 64 years	278	0.28
Female	585	0.58
Observations	1004	

Note: The table shows demographic information about all participants in the second study.

Table 1.10: Descriptive Statistics - Treatment Study 2

	Baseline	Information	Difference
Household income less than or equal to $30,000	0.25	0.27	-0.02
Household income between $30,001 and $60,000	0.26	0.24	0.02
Household income between $60,001 and $100,000	0.21	0.24	-0.03
Household income between $100,001 and $150,000	0.18	0.15	0.03
Household income above $150,000	0.10	0.10	-0.00
High School Education or below	0.24	0.22	0.01
Some College Education	0.32	0.29	0.03
Bachelor or equivalent	0.25	0.27	-0.02
Master or equivalent	0.20	0.22	-0.02
Age in years	53.40	49.91	3.50***
Female	0.58	0.59	-0.01
Republican	0.41	0.38	0.03
Observations	501	503	

Note: The table shows the share of participants in the demographic categories and the mean for the category age in years. Differences were tested with t-tests. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

1.C Additional results

Table 1.11: Regression results for implemented inequality - Reduced sample Study 1

	Inequality	Inequality	Inequality	Inequality	Inequality	Inequality
Information	0.286***	0.285***	0.286***	0.287***	0.286***	0.285***
	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)	(0.026)
Age		-0.067**				-0.077***
		(0.027)				(0.027)
Income			0.009			-0.003
			(0.026)			(0.028)
Female				-0.045*		-0.055**
				(0.026)		(0.027)
Education					0.019	0.017
					(0.026)	(0.028)
Constant	0.244***	0.284***	0.239***	0.268***	0.234***	0.312***
	(0.017)	(0.025)	(0.022)	(0.023)	(0.022)	(0.036)
Observations	1299	1299	1299	1299	1299	1299
R^2	0.086	0.088	0.086	0.088	0.086	0.091

Note: The table reports results from an OLS-regression of the share of spectators who implement the unequal outcome for the two workers. Information is an indicator variable which takes the value 1 if the participant was assigned to the information treatment. Age is an indicator variable with the value 1 if the participant was older than 45 years. Income is an indicator variable which is 1 if the participant was above a household income of $60,000. Female is an indicator variable for being female. Education is an indicator variable which is 1 when the participant has at least some college education. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1.3: Treatment effects for different subgroups - Study 1

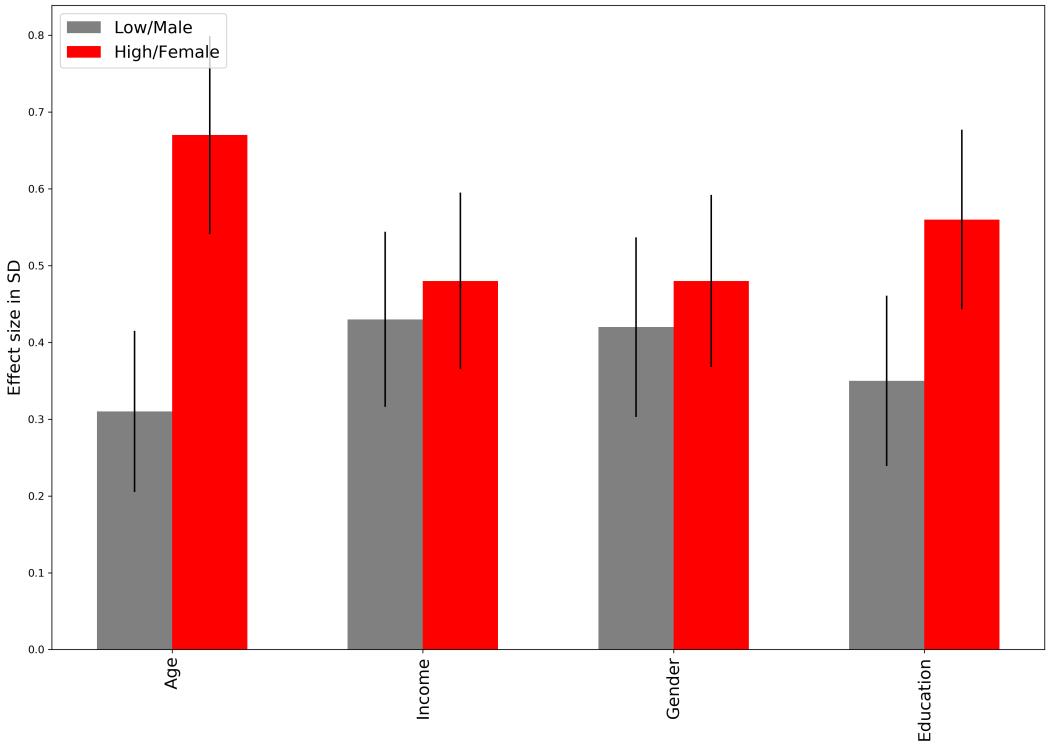

Note: The figure shows the effect size for the information treatment across different demographic subgroups. Effect size is measured in standard deviations. Grey bars indicate group younger than 45 (age), below $60,000 household income (income) and non-college educated (education) and male (gender). Red bars are older than 45 (age), above $60,000 household income (income) and college educated (education) and female (gender). The bars show the 95% confidence interval.

Table 1.12: Regression results for heterogeneity across demographic groups - reduced sample Study 1

	Inequality			
	Age	Income	Gender	Education
	B = 1 if > 45	B = 1 if income > $ 60,000	B = 1 if female	B = 1 if college
Information	0.220***	0.274***	0.244***	0.251***
	(0.042)	(0.037)	(0.039)	(0.037)
B*Information	0.109**	0.021	0.075	0.067
	(0.053)	(0.052)	(0.052)	(0.052)
B	-0.132***	-0.013	-0.092***	-0.016
	(0.036)	(0.035)	(0.034)	(0.035)
Constant	0.347***	0.318***	0.332***	0.327***
	(0.040)	(0.038)	(0.038)	(0.037)
$\beta_1 + \beta_2$	0.329***	0.295***	0.319***	0.318***
	(0.032)	(0.036)	(0.034)	(0.036)
R^2	0.095	0.091	0.092	0.092
Control Variables	Yes	Yes	Yes	Yes
Observations	1299	1299	1299	1299

Note: The table reports results from an OLS-regression of the share of spectators who implement the unequal outcome for the two workers. Information is an indicator variable which takes the value 1 if the participant was assigned to the information treatment. B is an indicator variable which is 1 when the participant is older than 45 (Column 1), has a household income higher than $60,000 (Column 2), is female (Column 3) or has at least some college education (Column 4). B*Information are interactions between assignment to the information treatment and the respective demographic variable. β_1 + β_2 is the linear combination of the variables B*Information + Information. Control variables include all other demographic variables. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 1.4: Treatment effects for different subgroups - Study 2

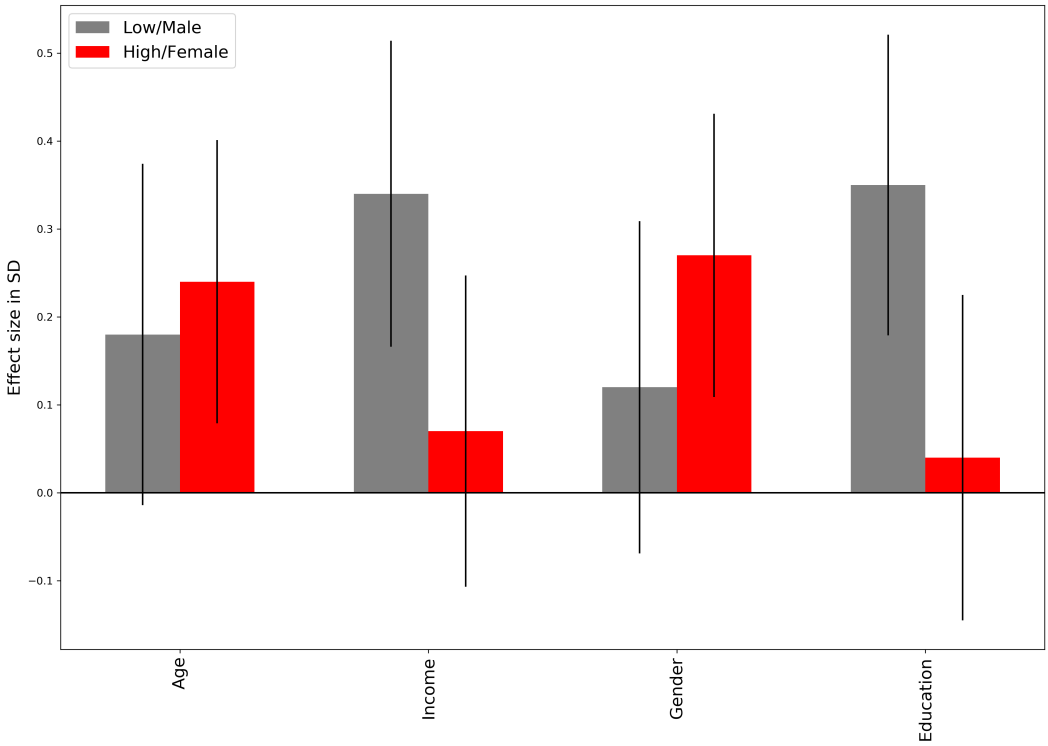

Note: This figure shows the effect size for the information treatment across different demographic subgroups. Effect size is measured in standard deviations. Grey bars indicate group younger than 45 (age), below $60,000 household income (income) and non-college educated (education) and male (gender). Red bars are older than 45 (age), above $60,000 household income (income) and college educated (education) and female (gender). The bars show the 95% confidence interval.

Table 1.13: Overview of choices in the information treatment and the hypothetical choice

		Hypothetical choice	
		Simple information	Detailed information
Information treatment	Withhold information	50 (16.8%)	248 (83.2%)
	Provide information	16 (7.8%)	189 (92.2%)

Note: The table shows an overview of people's choices in the information treatment and the choices they make in the hypothetical scenario. The share of participants who make a given decision in the information treatment in parentheses. χ^2-test for equality of distributions: p = 0.003.

Table 1.14: Overview of choices in the baseline condition and the hypothetical choice

		Hypothetical choice	
		Simple information	Detailed information
Baseline condition	Simple math task	38 (11.0%)	307 (89.0%)
	Complex math task	20 (12.8%)	136 (87.2%)

Note: The table shows an overview of people's choices in the baseline condition and the choices they make in the hypothetical scenario. The share of participants who make a given decision in the baseline condition in parentheses. χ^2-test for equality of distributions: p = 0.559.

Table 1.15: Regression results for choice in hypothetical scenario

	Detailed provision	Detailed provision
Information provider	0.090***	0.093***
	(0.029)	(0.030)
Constant	0.832***	0.787***
	(0.022)	(0.038)
R^2	0.017	0.027
Control variables	No	Yes
Observations	503	503

Note: The table provides results from an OLS-regression for people's support for detailed information provision of a firm. Only participants in the information treatment were included. Information provider is an indicator variable for the choice to provide information in the experiment. Control variables include all other demographic variables. Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

1.D Instructions

1.D.1 Study 1

Baseline condition

We now ask you to make a choice that might have consequences for a real life situation.

We are recruiting individuals via an international online market place to take part in a study for which they are paid. These individuals are randomly matched into pairs. You are now making a decision for one randomly selected pair of individuals. We plan to pay both 3.30 USD for their participation.

We ask you now to make the following decision: You can decide that we randomly select one of the two individuals to earn 5.00 USD, 1.70 USD more than before. The other individual then earns 2.50 USD, 0.80 USD less than before.

Your decision will be implemented with a certain probability and therefore might have real consequences.
I want to implement the alternative earnings
I do not want to implement the alternative earnings

Information treatment

We now ask you to make a choice that might have real consequences for a real life situation.

We are recruiting individuals via an international online market place to take part in a study in which they are asked to estimate the share of red balls in a container with a large number of red and blue balls. Their earnings in the study are determined by the precision of their estimate. The real share of red balls in the container is 50% and they earn 5.00 USD if their estimate is correct. Their earnings are reduced by 10 cents per percentage point their estimate differs from the correct number.

Two individuals, let us call them individual A and individual B, are matched in a pair. Each individual will see a result of three random draws from the container. One ball is drawn at a time, the color is reported to the individual and the ball is put back into the container.

The situation in which you are asked to make a decision is the following: Individual A has seen a draw of 2 red and 1 blue ball (a share of 67% red balls) while individual B has seen 2 blue and 1 red ball (a share of 33% red balls). Based on the three draws, Individual A estimates that the share of red balls is 67% (a 17 percentage point deviation from the correct number) while individual B estimates that the share of red balls is 33% (a 17 percentage point deviation from the correct number). If this is their final estimate, individual A earns 3.30 USD and individual B earns 3.30 USD.

We now ask you to decide whether or not you want to show the two individuals the color of one additional ball that is randomly drawn. If you do, one of the two individuals will have seen 2 red and 2 blue balls while the other individual will have seen 3 balls of one color and and 1 of the other color. Based on the four balls, one individual will estimate that the share of red ball is 50% (no deviation) and the other individual will estimate that the share of red balls is either 25% or 75% (a 25 percentage point deviation). One individual will therefore earn 5 USD, 1.70 USD more than before and the other individual will earn 2.50 USD, 0.80 USD less than before.

Your decision will be implemented with a certain probability and therefore might have real consequences.

I want to show them the additional ball

I do not want to show them the additional ball

Paternalism

To what extent do you agree to the following statement: The federal government should restrict citizens' freedom if it is in the citizens' own best interests?

Fully disagree
Disagree
Slightly disagree
Neither agree nor disagree
Slightly agree
Agree
Fully agree

Attention checks

Attention check 1:

Many people think it is a difficult and exhausting task to fill out online surveys with questions about a broad variety of situations. It is hard to stay focused and to keep your up your attention. To indicate that you still actively follow our survey, simply click on the arrow without answering the question to go to the next page.

Do you want to proceed to the next question?

Yes

No

Attention Check 2:

There always a big competition for our focus and our cognitive bandwidth in general is very limited. Studies report that the maximum time we can continuously pay attention is only about 7 minutes. To proceed to the next question and indicate your full level of attention please click on the only answer option constituting an even number, ignoring what the actual question is about.

According to research, what is the maximum time people can continuously pay attention?

5 minutes

6 minutes

7 minutes

9 minutes

1.D.2 Study 2

Baseline condition

We now ask you to make a choice that might have consequences for a real life situation. We have recruited two individuals, let us call them individual A and individual B, to take part in a study in which they have to solve math problems. They will both be paid a fixed amount for participating, but they can earn additional money depending on how well they did in the math task.

Previously, we have tested their math abilities and we want you to decide which math task they have to solve. They can either do a simple math task or a more complicated task. If you give them the first task then based on the result from the previous test, individual A and individual B will be able complete the task and both will earn $3.30. If you give them the second task then individual A, but not individual B, will be able to complete the task. Individual A will then earn $5.00, $1.50 more than in the first task and individual B will earn $2.50, $0.80 less than in the first task.

Please indicate your decision now:
I want to implement the **first task**
I want to implement the **second task**

Information treatment

We now ask you to make a choice that might have consequences for a real life situation. We have recruited two individuals, let us call them individual A and individual B, to take part in a study in which they have to assess the likelihood that a certain event happens based on the information given to them. They will both be paid a fixed amount for participating, but they can earn additional money depending on how well they did in the task.

Previously, we have tested their abilities and we want you to decide how much information to present to the participants. They can either be given basic information or basic and additional information. If you provide basic information on the likelihood of the event then based on the result from the previous test, individual A and individual B will make equally good predictions and both will earn $3.30. If you provide them with basic and additional information on the likelihood, individual A will be able to use this information to make a better prediction, but individual B will be unable to use the additional information and make a worse prediction. Individual A will then earn $5.00, $1.50 more than with basic information and individual B will earn $2.50, $0.80 less than with basic information.

Please indicate your decision now:
I want to provide **basic information**
I want to provide **basic and additional information**

Hypothetical Choice

Producers of electronic devices, such as refrigerators and air conditioners, have to present information about the energy efficiency of their products to consumers. Such information can help consumers select the product that is best for them. However, additional information might also result in some consumers getting confused.

Consider a situation in which a producer is deciding whether to provide more detailed information to its customers. This additional information will help the majority of consumers when they decide which product to purchase, but it will make this choice more difficult for a minority of consumers. In such a situation, do you think the producer should provide more detailed information?

Yes

No

1.E Deviations from the pre-analysis plan

The pre-analysis plan of the first study is available at AEARCTR-0003376. In the main analysis, an OLS-regression was run to better illustrate the main treatment effect. In the heterogeneity analysis, all control variables are included. This inclusion does not change the conclusions of this analysis.

The pre-analysis plan for the second study is available at AEARCTR-0005018. In this study, an OLS-regression was run for the main effect. Control variables were included for the heterogeneity analysis. In this case, the inclusion does not change the main conclusions either. For the external validity measure, another OLS-regression was run to present the main result.

All additional tests were run to test the pre-registered hypothesis. Hypotheses that were not pre-registered are made salient as "exploratory analysis" in the main body of the paper.

Chapter 2

Cancel the deal? An experimental study on the exploitation of irrational consumers

Alexander W. Cappelen Stefan Meissner

Bertil Tungodden*

Abstract

Consumers are sometimes exploited because they misunderstand the value of a product they buy. We present the results from a large-scale experimental study that examines whether third-party spectators from the general population in the US accept a deal where the irrational buyer pays more for a product than it is worth and the extent to which the willingness to cancel such deals depends on the seller's role in the process leading up to the deal. We find that the majority of the spectators cancel an unfair deal even when the seller's

*Affiliation of all authors: Department of Economics, NHH Norwegian School of Economics. We have received valuable comments and suggestions from Björn Bartling, Eleonora Freddi, Jana Friedrichsen, Alex Imas, George Loewenstein, and Erik Ø. Sørensen. We would also like to thank Sebastian Fest, for excellent research assistance. The project was financed by support from the Research Council of Norway through its Centres of Excellence Scheme, FAIR project no 262675, and the Research Council of Norway research grants no 236995, no 262636, and no 302145. The project was administered by FAIR - The Choice Lab, NHH Norwegian School of Economics.

role is limited to accepting a proposal made by the buyer. A large minority, however, are willing to uphold the deal even when the seller has obfuscated the value of the product and proposed the deal. The seller's involvement in obfuscating the buyer's information significantly increases the willingness to cancel the deal and to fine the seller. Our results shed new light on how people view fairness in market transactions and their attitudes to government regulation of businesses.

Keywords: Paternalism, Consumer naivety

JEL Classification: D63

2.1 Introduction

Consumers often make systematic mistakes in their valuation of products (Heidhues and Kőszegi, 2018) and such mistakes have been documented in numerous contexts, among them credit card contracts (Heidhues and Kőszegi, 2010), life- and health-insurance (Bhargava et al., 2017; Gottlieb and Smetters, 2016), cellular phone plans (Bar-Gill and Stone, 2012; Grubb, 2009), mortgages (Campbell et al., 2011), and financial investments (Choi et al., 2010). Companies can exploit such mistakes to take advantage of consumers (Agarwal et al., 2017; Akerlof and Shiller, 2015; Célérier and Vallée, 2017; Gabaix and Laibson, 2006; Grubb, 2015; Ru and Schoar, 2016). An important policy question is therefore whether the government should introduce policies aimed at limiting the opportunities businesses have to exploit ill-informed or irrational consumers (Agarwal et al., 2015; Bar-Gill and Warren, 2008; Campbell, 2016; Kőszegi, 2014; Nuñez, 2017; Pete, 2014; Warren and Wood, 2014). Such policies may be viewed as fair since they eliminate market transactions that would not happen if both parties were well informed and rational, but do at the same time reduce people's freedom to enter contracts (Epstein, 1975; Kronman, 1983). A large literature in economics, philosophy, and law, has argued that a person's autonomy is a valuable good in itself and needs to be respected even if the person's choices have adverse consequences for him- or herself (Bartling et al., 2014; Cohen, 1995; Mill, 1859; Nozick, 1974). Attitudes to such regulations thus depend on how people trade off different types of moral considerations.

In this paper, we present the results from a novel large-scale experiment designed to examine how the general population of the US views market

transactions involving irrational buyers. About 4000 participants are presented with real situations in which a buyer has misunderstood the true value of a product and agreed to pay more for the product than it is worth. The participants, acting as third-party spectators, are asked to decide whether the deal should be canceled. Spectators who decide to cancel the deal were also given the opportunity to penalize the seller financially. In a between-subject design, we vary the seller's role in the process leading up to the deal, while holding the consequences of canceling or not canceling the deal the same. In a base treatment with low involvement, the seller has only accepted a deal proposed by the buyer. In three additional treatments, the seller's involvement is increased either by the seller having proposed the deal, by the seller having obfuscated the information about the true value of the product, or both.

Our design allows us to cleanly identify people's moral views on contracts involving irrational buyers, and the extent to which such views depend on the involvement of the seller. We identify three types of spectators, who differ in the way they handle the trade-off between concern for the distributional consequences of market transactions and concern for people's autonomy: **Substantialists**, **Contractualists**, and **Proceduralists**. Substantialist are spectators who primarily care about the consequences of market transactions and therefore want to cancel deals that create an unfair distribution of gains and loses (Almås et al., 2020; Bolton and Ockenfels, 2000; Cappelen et al., 2007, 2013b; Fehr and Schmidt, 1999). Contractualists are spectators who primarily care about respecting individual autonomy and therefore do not want to cancel any voluntary deals (Arneson, 1980; Brownell et al., 2010; Epstein, 1975; Schwartz and Scott, 2003; Spector, 2006). Finally, Proceduralists are spectators who decide to cancel dependent on the seller's role in the process

leading up to the deal (Bolton et al., 2005; DiMatteo and Rich, 2005; Falk et al., 2008).

Our study provides three main findings. First, we show that the majority of spectators do not accept the exploitation of an irrational buyer, even when the contract reflects a voluntary agreement. Pooled across treatments, 60.7 percent of the spectators decide to cancel the contract. Second, we find that the involvement of the seller matters. In particular, the share of spectators canceling the deal increases significantly when the sellers have obfuscated the information given to the buyers, and it increases the share of spectators imposing a fine on the seller. Third, we estimate that the large majority of spectators are Substantialists, who focus on the unfair consequences of the market transaction and cancel the deal even when the involvement of the seller is minimal. We also find that a large minority are Contractualists, who focus on the contract reflecting a voluntary agreement and do not cancel the deal even when the seller has been actively involved in the exploitation of the buyer. We only find a small minority of spectators being Proceduralists, making the canceling of the deal conditional on the extent to which the seller has been actively involved. Taken together, our findings show that there is significant support for regulations of market transactions involving irrational actors in the US population, and that there might be significant risk for companies in pursuing strategies that exploit irrational consumers. However, our study also shows that a large minority accept voluntary contracts even when they have detrimental consequences for irrational buyers.

Our study relates to several literatures. It provides, to our knowledge, the first experimental study of people's attitudes to the exploitation of irrational consumers, which has recently been studied theoretically in several

important papers (DellaVigna and Malmendier, 2004; Eliaz and Spiegler, 2006; Gabaix and Laibson, 2006; Heidhues and Kőszegi, 2018; Kőszegi, 2014). These papers have demonstrated that the existence of irrational consumers in a marketplace creates opportunities for companies to exploit them. Our results suggest that the majority are in favor of initiatives that call for tighter restrictions on contract terms and more comprehensive demands for information disclosure (Pete, 2014; Seizov et al., 2019).

The results of this study also contribute to the literature on fairness in market transactions. Previous research has shown that perceptions of unfair pricing strategies can constrain a firm's opportunity to maximize profits (Bolton et al., 2003; Haitao Cui et al., 2007; Kahneman et al., 1986; Leibbrandt, 2020; Xia et al., 2004; Zbaracki et al., 2004). Our study demonstrates that outcome-based fairness considerations are of great importance for people's view on market transactions; a majority of the spectators decide to cancel a voluntary deal that benefits one side at the expense of the other. We also show that procedural-based fairness matters, by causally identifying that an increase in seller involvement increases the share canceling the deal and the share fining the seller. This is in line with previous research showing the importance of both outcomes and processes for people's assessment of fairness (Andreoni et al., 2020; Brock et al., 2013; Falk et al., 2008).

Finally, our findings further add to the growing literature that investigates the value of autonomy (Bartling et al., 2014; Fehr et al., 2013; Jacobsson et al., 2007). Empirical evidence has shown that people value their own and other people's autonomy and are reluctant to infringe upon it (Bartling et al., 2014; Fehr et al., 2013; Gagné and Deci, 2005; Iyer et al., 2012; Jung and Mellers, 2016; Pikulina and Tergiman, 2020; Reisch and Sunstein, 2016).

Consistent with these studies, our results show that a substantial share of the spectators are unwilling to intervene when doing so can be seen as infringing on people's autonomy.

The rest of the paper is organized as follows: Section 2.2 describes the experimental design and the sample, and Section 2.3 outlines the empirical strategy. Section 2.4 presents the results, while Section 2.5 concludes.

2.2 Experimental Design

The experiment involved 5587 participants, recruited using two different platforms. We recruited 3991 subjects from the general population in the United States through a leading international survey provider (Dynata) to act as third-party spectators, and we recruited 1596 subjects from the Amazon Mechanical Turk (AMT) online labor market platform to serve as stakeholders. The stakeholders were only recruited to create real economic situations, and in the analysis, we focus on the decisions made by the spectators.

The spectators were largely representative of the adult population (+ 18 years old) in the United States on observable characteristics (geographical location, age, and gender). The median age of spectators was 48 years, 52 percent of the sample were women, and the average level of education was somewhat higher than in the general population. The median household income in our sample was $60,000. Table 2.6 provides an overview of the demographic characteristics of the spectators. In line with the pre-analysis plan, the survey provider only included individuals in the sample if they passed an attention check.[2]

[2]The pre-analysis plan is available at 10.17605/OSF.IO/N5HVB.

Spectators were paid a fixed compensation for taking part in the study, independent of their spectator decision. The spectators answered a set of general questions about their background characteristics, including political affiliation. We also asked the spectators about their beliefs about the behavior of consumers and firms, and about their attitudes to government policies aimed at restricting businesses' opportunities to profit from customers who misunderstand the value of a product or a service.

2.2.1 Spectator decision

Each spectator was matched to a buyer and a seller who had agreed to a deal that would result in a loss for the buyer and a gain for the seller.[3] All spectators received the same basic description of the situation. They were informed that a seller had been endowed with a product that had no value for the seller but had a value of $2 for the buyer. Since there are no transaction costs, this means that there is a small efficiency gain if the product is sold to the buyer.

Both the seller and the buyer knew that product had no value for the seller. The spectators were, however, told that there was an information asymmetry between the seller and the buyer: only the seller knew that the value of the product for the buyer was $2. The spectators were furthermore told that the buyer had received complex information about the value of the product for him- or herself and that, if the buyer had used the information correctly, he or she could have calculated the correct value of the product. However,

[3]The spectators were told that there was a one-in-five chance that their choice would be implemented for the pair they were matched with. The spectators also knew that the pair they were matched with was aware that their earnings depended on the decision of another person.

using this information, the buyer had made a mistake when calculating the value of the product and therefore believed that the value was $20. Finally, the spectators were told that the seller knew that the buyer believed that the value of the product was $20 and that the buyer and the seller had agreed on a deal in which the buyer paid $10 for the product. As a result of the deal, the buyer would lose $8 and the seller would gain $10.

Spectators were asked whether they wanted to cancel the deal. If they decided to cancel the deal, the seller's gain and the buyer's loss would be eliminated. Those spectators who decided to cancel the deal were asked whether they wanted to fine the seller. If they decided to fine the seller, the payment to the seller would be reduced by $2.

The spectators were randomly assigned to one of four treatments that only differed in the how the deal between buyer and seller came about. Section 2.C in the appendix presents the full instructions for the spectators in each of the four treatments.

2.2.2 Treatment variations

The four treatments varied with respect to the level of involvement of the seller in proposing the deal and in obfuscating the information received by the buyer. In the Low treatment, the seller was not active and only accepted a proposal from the buyer. In this treatment, the spectators were told that the buyer had proposed to buy the product for $10 and that the seller had accepted this proposal. Furthermore, the spectators were informed that the seller had no role in obfuscating the information the buyer received.

In the other treatments, we increased the seller's involvement along two

Table 2.1: Overview of experimental treatments

		Seller obfuscated	
		No	Yes
Seller proposed	No	Low	Obfuscate
	Yes	Propose	High

Note: This table provides an overview of the experimental treatments. In the treatment *Low*, the seller's involvement is limited to having accepted a proposal made by the buyer. In the treatment *Proposed*, the seller has proposed the deal, but not obfuscated the information. In the treatment *Obfuscated*, the seller has obfuscated the information received by the buyer, but not proposed the deal. In the treatment *High*, the seller has obfuscated the information and proposed the deal.

dimensions. The Proposed treatment is identical to the Low treatment except that it was the seller who proposed the deal and the buyer who accepted it. The Obfuscated treatment is identical to the Low treatment except that the seller had decided to disclose information about the value of the product to the buyer in a complex manner rather than in an easily understandable manner. The seller knew that doing so would lead the buyer to overestimate the value of the product. Finally, in the High treatment, the seller both obfuscated the information about the value of the product and proposed the deal. Table 2.1 provides an overview of the four treatments.

2.2.3 The stakeholders

To create the situations for which the spectators make decisions, we recruited individuals on an online labor market platform to act as stakeholders. These individuals were randomly assigned to the role of either a *buyer* or a *seller*.

The sellers received a product, a lottery ticket, and were informed that the ticket had no value for themselves, but that it had a value of $2 for a buyer.

They were asked whether they wanted to obfuscate the information about the value of the ticket to the buyer. Sellers were informed that buyers who received the complex information might make a mistake when calculating the value and therefore overvalue the product.

The sellers who wanted to obfuscate the information given to the buyer were then asked whether they wanted to offer to sell the lottery ticket for $10 to a buyer who had overestimated the value of the product and whether they would accept an offer of $10 from such a buyer. The sellers who wanted to present in an easily understandable manner were asked whether they wanted to offer the ticket for $10 to a buyer who had received complex information from the experimenter and whether they would accept the same offer from such a buyer.

The buyers were informed that the product had no value for the seller but were given a complex description of the product that required them to make some calculations to correctly evaluate its value. They were told that the value of the product was either $20 or $2, and the description was given in such a way that buyers who suffered from base-rate neglect would evaluate the product at $20, while buyers who were Bayesian would correctly estimate the value to be $2 (Benjamin, 2019). Buyers who estimated the value to be $20 were asked whether they wanted to make an offer of $10 to a seller who owned the product and whether they would accept an offer to purchase the product for $10 from such a seller. Buyers who were unwilling to buy the product for $10 were not matched with any seller.

Buyers and sellers who were willing to trade when the price was $10, were matched in pairs to create the situations presented to the spectators. The stakeholders were told that they could receive their bonus payments within

a few days and that the bonus payment would be determined by their own decisions and the decisions of a third party. All participants got a bonus of $2, but the fine was subtracted from this amount for those sellers who were fined. The main stages of the experiment are summarized in Figure 2.1.

Figure 2.1: Flow of the experiment

- Sellers are endowed with a product and told the product has a value of $2 to a buyer, but no value to them. Asked whether they want to obfuscate the information and propose/accept to sell the product for $10.
- Buyers who mistakenly believe the product is worth $20 are identified and asked whether they want propose/accept to buy the product for $10
- Sellers and buyers are matched to create the four treatments

- Spectators are informed about a deal between a buyer and a seller. They are randomly assigned to treatments that differ with respect to whether the seller has obfuscated the information and proposed the deal.
- Spectators are given the opportunity to cancel the deal between buyer and seller

- Spectators who canceled the deal are given the opportunity to fine the seller $2.

- The spectator decisions are matched to pairs of buyers and sellers (5:1 matching rate)
- Payments are paid out to buyers and sellers

2.3 Empirical Strategy

The empirical strategy was pre-registered at the repository of the Open Science Foundation (OSF) before the data collection began.[4] Pre-registration included a power analysis, exclusion criteria for spectators, pre-specification of the main hypotheses, and an outline of the empirical strategy.

2.3.1 Main specifications

To examine how the involvement of the seller affects the spectators' willingness to cancel the deal, we use the following empirical specification

$$C_i = \beta_0 + \beta_1 P_i + \beta_2 O_i + \beta_3 P_i O_i + \gamma X_i + \epsilon_i, \tag{2.1}$$

where C_i is an indicator variable for whether the spectator canceled the deal, P_i is an indicator variable for whether the seller proposed the deal, O_i is an indicator variable for whether the spectator obfuscated the information, $P_i O_i$ is an interaction variable that captures a potential interaction effect between the two domains of the seller's involvement, and X_i is a vector for control variables, including age, education, income, gender, and political affiliation. As pre-registered, the main results are reported with and without the inclusion of control variables. We expected that the willingness to cancel would increase with the involvement of the seller, i.e. that $\beta_1 > 0$ and that $\beta_2 > 0$, but had no prior hypothesis about the interaction effect β_3.

We introduce the following classification of spectators, based on when they

[4]The pre-analysis plan is available at 10.17605/OSF.IO/N5HVB.

would want to cancel the deal:

- **Substantialists:** Always cancel the deal
- **Proceduralists:** Do not cancel the deal if the seller is not active, but cancel the deal if the seller has obfuscated information and proposed the deal
- **Contractualists:** Never cancel the deal

We assume that all spectators are of one of these three types and that the spectator type is independent of treatment. We can then estimate the share of each type using (2.1):

- **Substantialists:** $S = \beta_0$
- **Proceduralists:** $P = \beta_1 + \beta_2 + \beta_3$
- **Contractualists**: $C = 1 - \beta_0 - \beta_1 - \beta_2 - \beta_3$

We study heterogeneity in the willingness to cancel the deal using the background data collected in the survey. The heterogeneity analysis is conducted by estimating the following regression:

$$C_i = \beta_0 + \beta_1 B_i + \beta_2 P_i + \beta_3 B_i P_i + \beta_4 O_i + \beta_5 B_i O_i + \beta_6 P_i O_i + \beta_7 B_i P_i O_i + \gamma X_i + \epsilon_i, \tag{2.2}$$

where B_i is an indicator variable for spectator i either having high education, being female, having a household income that is higher than $60,000, being

at least 47 years old or being a Republican. In this regression, $\mathbf{X}_i$ includes all background variables except the variable captured by B_i. Using this regression, we can also estimate the share of Substantialists, Proceduralists, and Contractualists for the different subgroups.

To examine how the treatments affect the willingness to fine the seller, we run 2.1 where we replace the indicator variable for the decision to cancel with an indicator variable for whether the spectator fined the seller. The heterogeneity analysis is conducted in the same way as for the decision to cancel.

Treatment effects on the decision to fine can be driven by both the share of spectators who cancel the deal and the share who fine conditional on having canceled the deal. To shed light on the relative importance of these mechanisms, we will also examine how the treatments affect the decision to fine only using the respondents who decided to cancel.

2.4 Results

We first provide an analysis of the spectators' decision to cancel the deal between the buyer and the seller before we analyze the subsequent decision to fine. Finally, we examine the spectators' beliefs and policy attitudes.

2.4.1 The decision to cancel

Across treatments, we find that 60.7 percent of the spectators choose to cancel the voluntarily agreed-upon deal between the buyer and seller. Figure 2.2 reports the share of spectators who choose to cancel the deal in each of the

Figure 2.2: Share of spectators that cancel the deal

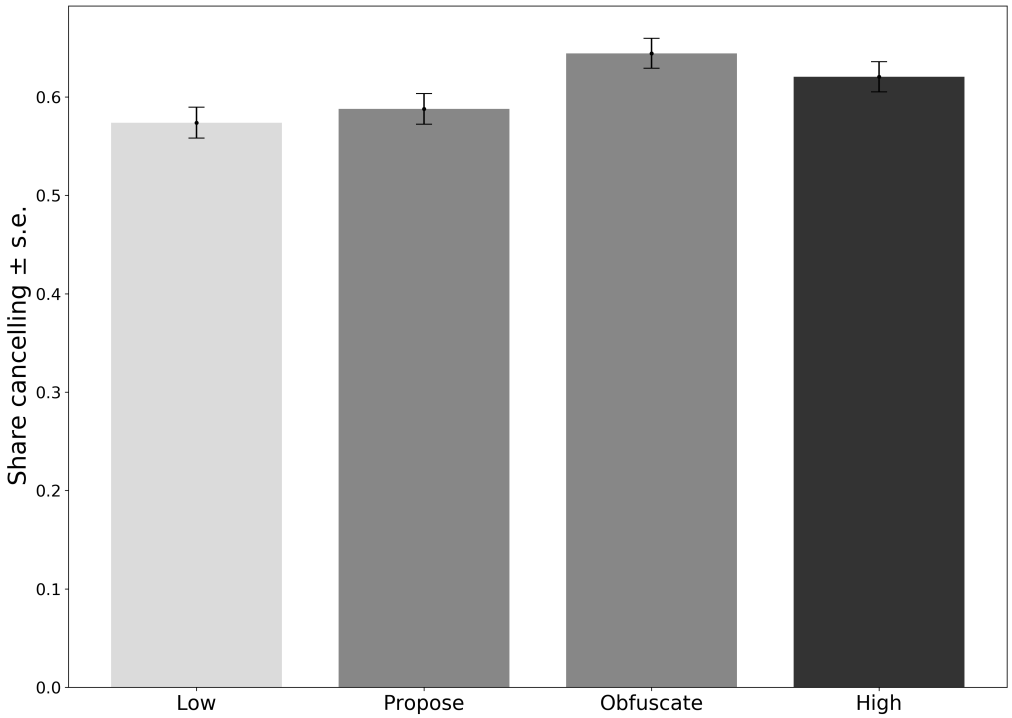

Note: The figure shows the share of participants who cancel the transaction between the buyer and the seller in each of the four treatments. The standard errors are indicated.

four treatments. We first observe that even in the Low treatment, where the seller's involvement is minimal, a majority of the spectators, 57.4 percent, cancel the deal. This share increases somewhat, to 59.0 percent, in the Propose treatment where the seller has proposed the deal. In the Obfuscate treatment where the seller has obfuscated the information given to the buyer, the share of spectators who cancel the deal further increases to 64.5 percent. Finally, we observe that 62.3 percent of the spectators cancel the deal in the High treatment, which implies that 37.7 percent do not cancel the deal even when the seller has both obfuscated the information and proposed the deal. We can summarize the first results as follows:

Result 1: *A majority of spectators cancel a voluntary deal between an irrational*

buyer and a minimally involved seller. A large minority do not cancel a voluntary deal between an irrational buyer and a seller even when the seller has obfuscated the information to the buyer and proposed the price that exploits the irrationality of the buyer.

Columns (1) and (2) in Table 2.2 report regressions on the spectators' decision on whether to cancel the deal between buyer and seller. We observe that the effect of having proposed the deal is not statistically significant ($p = 0.464$). However, the spectators' willingness to cancel the deal is significantly increased when the seller has obfuscated the information ($p = 0.001$), and this holds independently of whether the seller has proposed the deal or not. From column (2), we observe that these results hold when we control for background characteristics. Table 2.3 furthermore shows that these results hold for all subgroups. We can summarize these results as follows:

Result 2: *Obfuscation of information by the seller increases the share of spectators who cancel the deal. Spectator behavior is not affected by whether it is the seller or buyer who proposes the deal.*

From column (2) in Table 2.2, we also observe that female spectators and older spectators are significantly more likely to cancel the deal ($p < 0.001$ and $p < 0.001$), while spectators with high income $p < 0.001$ and Republican spectators ($p < 0.030$) are less likely to cancel the deal. The finding for Republicans is in line with previous research that suggests that conservatives are more likely to hold favorable views of the outcomes of free markets (Goren, 2005; Jost et al., 2009, 2003; Malka et al., 2014) and are opposed to interventions into these outcomes (Feldman and Johnston, 2014; Jost, 2017; Skitka, 1999; Skitka and Tetlock, 1993).

Table 2.2: Regression results for decision to cancel the deal and fine the seller

	Cancel	Cancel	Fine	Fine
Proposed	0.016	0.013	0.033**	0.033**
	(0.022)	(0.022)	(0.015)	(0.015)
Obfuscated	0.070***	0.072***	0.080***	0.080***
	(0.022)	(0.022)	(0.016)	(0.016)
Proposed*Obfuscated	-0.038	-0.037	-0.026	-0.024
	(0.031)	(0.031)	(0.023)	(0.023)
Female		0.076***		0.002
		(0.016)		(0.012)
Old		0.076***		-0.037***
		(0.015)		(0.012)
College		-0.019		-0.017
		(0.017)		(0.013)
Income		-0.037**		-0.005
		(0.017)		(0.013)
Republican		-0.061***		-0.030**
		(0.017)		(0.012)
Constant	0.574***	0.541***	0.112***	0.149***
	(0.016)	(0.022)	(0.010)	(0.016)
$\beta_1 + \beta_2 + \beta_3$	0.049**	0.048**	0.088***	0.089***
	(0.022)	(0.022)	(0.016)	(0.016)
$\beta_2 + \beta_3$	0.032	0.036*	0.055***	0.056***
	(0.022)	(0.022)	(0.017)	(0.017)
R^2	0.003	0.022	0.009	0.015
Observations	3991	3991	3991	3991

Note: The table reports results from an OLS-regression of the share of spectators who decides to cancel the deal (columns 1-2) and the share of spectators who decides to fine the seller conditional on having canceled the transaction (columns 3-4). "Proposed" is an indicator variable for the spectator being assigned to the treatment where the seller proposed the transaction. "Obfuscated" is an indicator variable for the spectator being assigned to the treatment in which the seller obfuscated the information. "Propose*Obfuscate" is the interaction between "Proposed" and "Obfuscated". "Female" is an indicator variable for being female. "College" is an indicator variable for the spectator having a bachelor degree or higher. "Old" is an indicator variable for the spectator being older than 47 years. "Income" is an indicator variable for the spectator having a household income of more than $60,000 per year. "Republican" is an indicator variable for the spectator voting for the Republican party. $\beta_1 + \beta_2 + \beta_3$ is the linear combination of the parameters for Proposed, Obfuscated and the interaction term between the two. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

Table 2.3: Heterogeneity regressions

	Cancel				
	Political	Gender	Education	Income	Age
	B=1 if Republican	B=1 if Female	B=1 if College	B=1 if higher than $60.000	B=1 if older than 47
Proposed	0.028	0.015	0.026	0.022	-0.004
	(0.026)	(0.032)	(0.031)	(0.029)	(0.032)
Obfuscated	0.074***	0.093***	0.061**	0.073**	0.084***
	(0.026)	(0.032)	(0.030)	(0.029)	(0.031)
Proposed*Obfuscated	-0.067*	-0.033	-0.034	-0.013	-0.014
	(0.037)	(0.045)	(0.043)	(0.041)	(0.044)
B	-0.057*	0.112***	-0.035	-0.037	0.079**
	(0.034)	(0.031)	(0.031)	(0.031)	(0.031)
B*Proposed	-0.046	0.005	-0.021	-0.014	0.036
	(0.049)	(0.044)	(0.044)	(0.045)	(0.044)
B*Obfuscated	-0.012	-0.040	0.019	-0.007	-0.027
	(0.047)	(0.044)	(0.044)	(0.044)	(0.044)
B*Proposed*Obfuscated	0.101	-0.014	-0.006	-0.055	-0.045
	(0.067)	(0.062)	(0.062)	(0.062)	(0.062)
Constant	0.591***	0.514***	0.591***	0.590***	0.535***
	(0.019)	(0.023)	(0.022)	(0.021)	(0.022)
R^2	0.007	0.012	0.005	0.008	0.009
Observations	3991	3991	3991	3991	3991

Note: The table reports results from an OLS-regression of the share of spectators who decides to cancel the transaction. "Proposed" is an indicator variable for the spectator being assigned to the treatment where the seller proposed the transaction. "Obfuscated" is an indicator variable for the spectator being assigned to the treatment in which the seller obfuscated the information. "Proposed*Obfuscated" is the interaction between "Proposed" and "Obfuscated". B is an indicator variable that takes the value 1 when the spectator is Republican (Column 1), female (Column 2), college-educated (Column 3), has a household income that is higher than $60.000 (Column 4) or is older than 47 years (Column 5). Robust standard errors in parentheses. $^{*}\ p < 0.10$, $^{**}\ p < 0.05$, $^{***}\ p < 0.01$. Standard errors in parentheses.

Turning to the analysis of the share of the different types of spectators, reported in Figure 2.3, we find that a majority of the spectators, 57.4 percent, are Substantialists who cancel the deal even when the seller only accepts an offer made the buyer. We interpret these spectators as primarily caring about the consequences of the deal and considering the outcome of the deal unfair. A large minority, 37.7 percent, are Contractualists who do not cancel the deal even when the seller has obfuscated the information and proposed the deal. We interpret these spectators as primarily wanting to respect individual autonomy. We cannot rule out the possibility that these spectators are also motivated by a concern for efficiency, but given the small efficiency gain associated with the deal, we find this less likely (Almås et al., 2020). Only 4.9 percent of the spectators are Proceduralists whose decision to cancel is dependent on the extent to which the seller has been involved in the process leading up to the deal. We interpret these spectators as mainly caring about whether the process leading up to the deal was fair.

Figure 2.3 also displays the distribution of Substantialists, Proceduralists, and Contractualists across the different subgroups. In all subgroups, the majority of spectators are Substantialists and Contractualists make up about a third of the spectators. There are some significant variations in the share of the types, with the share of Substantialists significantly higher among women and older spectators and spectators with low income ($p = 0.001$ and $p = 0.06$), but the distribution of spectators types is strikingly similar across subgroups.

Result 3: *We find that the majority of the spectators are Substantialists and that a large minority are Contractualists. Only a small minority are Proceduralists. The distribution of spectator types is largely robust across subgroups.*

Figure 2.3: Classification of spectators

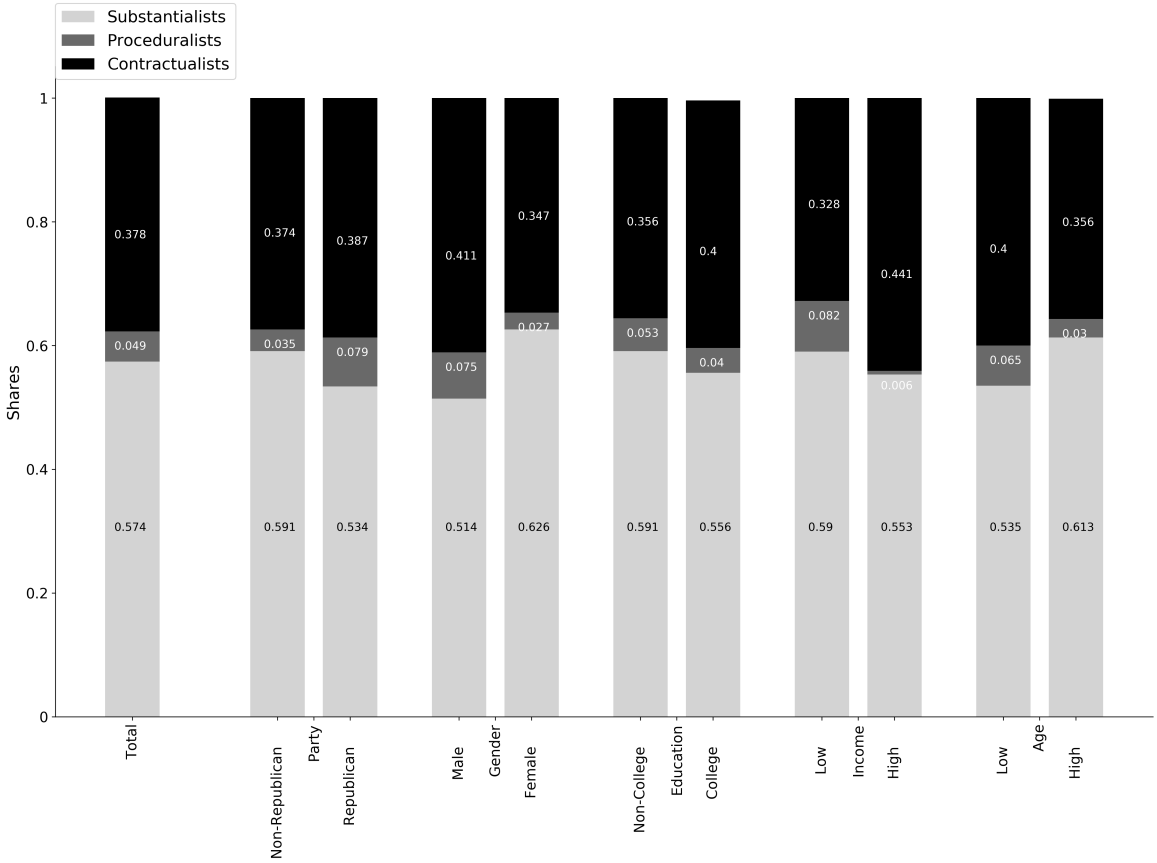

Note: The figure reports the estimated shares of spectators who are Substantialists, Proceduralists, and Contractualists, for the full sample and for different subgroups.

Figure 2.4: Share of spectators who fine the seller

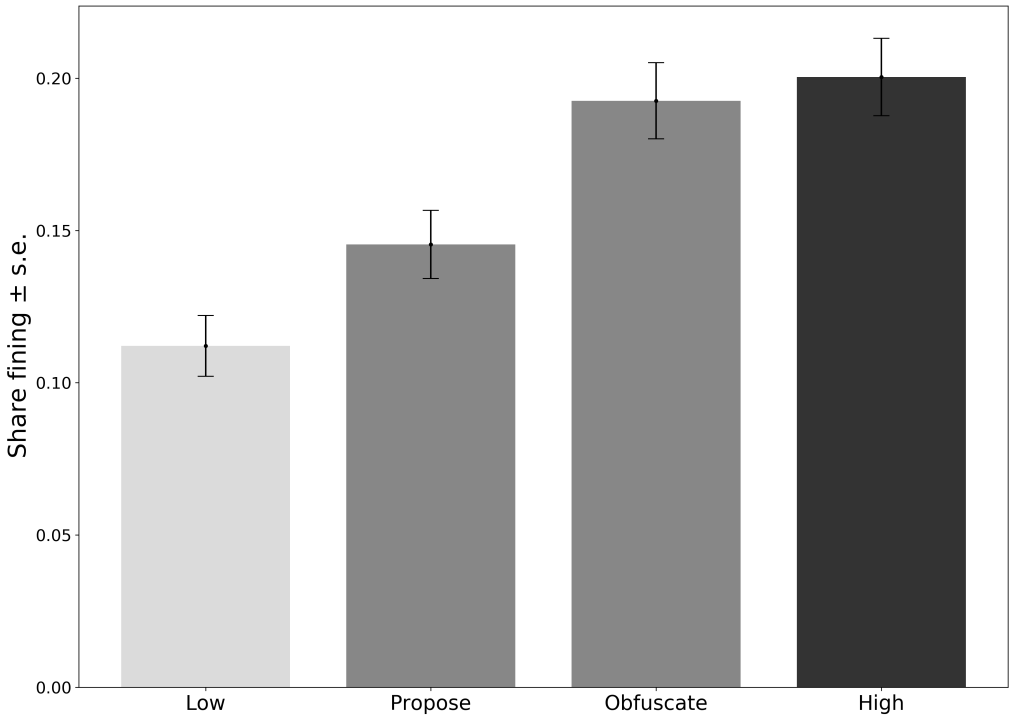

Note: The figure reports the share of spectators who decides to fine the seller for each of the four treatments. The standard errors are indicated.

2.4.2 The decision to fine

Spectators who canceled the deal may think that the seller deserves to be punished for the role the he or she played in the process leading up to the deal. The spectators who canceled the deal were therefore asked whether they also wanted to impose a fine on the seller. Figure 2.4 reports the share of spectators who, in addition to canceling the deal, decide to fine the seller. Across treatments, 16.3 percent of spectators decide to fine the seller, but there is considerable variation between treatments. In the Low treatment, only 11.2 percent of the spectators fine the seller, while 20.0 percent of the spectators decide to do so in the High treatment.

Columns (3) and (4) in Table 2.2 report regressions on the decision to fine the seller. From Column (3) we observe that the share of spectators increases by 3.3 percentage points when the seller proposed the deal ($p = 0.026$) and it increases by 8.0 percentage points when the seller has obfuscated the information ($p < 0.001$). Table 2.4 shows that these results hold for all subgroups. We furthermore observe that older spectators and Republican spectators are significantly less likely to fine the seller ($p = 0.002$ and $p = 0.015$ respectively).

Result 4: *Across treatments, only a minority decide to fine the seller. The share of spectators who fine the seller is higher when the seller has been active in the process leading up to the deal, by having proposed the transaction or having obfuscated the information.*

Since the share who canceled the deal is increasing with the involvement of the seller, the share of spectators who fine the seller will be increasing with the seller's involvement even when the share who fine conditional on canceling is the same across treatments. It could, however, also be the case that those who have canceled the deal are more likely to fine if the seller has been involved in obfuscating the information and proposing the deal. To shed light on this mechanism, we examine the share of spectators who fine the seller among those spectators who canceled the deal. Table 2.7 in the appendix shows the same overall pattern of results for the share of spectators who fine the seller conditional on having canceled the transaction. The share of spectators who fine the seller among the spectators who canceled the transaction increases from 19.5 percent in the Low treatment to 32.2 percent in the High treatment. This shows that the treatment effect on the share who

Table 2.4: Heterogeneity regressions

	Fine				
	Political	Gender	Education	Income	Age
	B=1 if Republican	B=1 if Female	B=1 if College	B=1 if higher than $60.000	B=1 if older than 47
Proposed	0.051***	0.016	0.059***	0.063***	0.023
	(0.018)	(0.022)	(0.022)	(0.021)	(0.022)
Obfuscated	0.088***	0.073***	0.085***	0.085***	0.109***
	(0.020)	(0.023)	(0.022)	(0.021)	(0.024)
Proposed*Obfuscated	-0.045	-0.007	-0.038	-0.042	-0.034
	(0.029)	(0.033)	(0.033)	(0.032)	(0.034)
B	-0.008	-0.011	0.003	0.013	-0.026
	(0.021)	(0.020)	(0.020)	(0.020)	(0.020)
B*Proposed	-0.063**	0.034	-0.052*	-0.068**	0.020
	(0.031)	(0.030)	(0.030)	(0.030)	(0.030)
B*Obfuscated	-0.023	0.015	-0.009	-0.010	-0.059*
	(0.034)	(0.032)	(0.032)	(0.032)	(0.032)
B*Proposed*Obfuscated	0.068	-0.036	0.027	0.039	0.019
	(0.048)	(0.047)	(0.047)	(0.047)	(0.046)
Constant	0.114***	0.118***	0.111***	0.106***	0.125***
	(0.012)	(0.015)	(0.014)	(0.013)	(0.015)
R^2	0.012	0.010	0.011	0.011	0.014
Observations	3991	3991	3991	3991	3991

Note: The table reports results from an OLS-regression of the share of spectators who decides fine the seller. "Proposed" is an indicator variable for the spectator being assigned to the treatment where the seller proposed the transaction. "Obfuscated" is an indicator variable for the spectator being assigned to the treatment in which the seller obfuscated the information. "Proposed*Obfuscated" is the interaction between "Proposed" and "Obfuscated". B is an indicator variable that takes the value 1 when the spectator is Republican (Column 1), female (Column 2), college-educated (Column 3), has a household income that is higher than $60.000 (Column 4) or is older than 47 years (Column 5). Robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

fine is not only driven by the higher number of spectators who cancel the transaction.

2.4.3 Beliefs and policy attitudes

Our study also provides evidence that people believe that the type of situation presented to the spectators, in which sellers exploit the irrationality of consumers, is common. At the end of the experiment, the spectators were asked about their beliefs about the behavior of consumers and firms, as well as their policy attitudes. The results are reported in Figure 2.5. We observe that a large majority of the participants agree with the statement that consumers often make mistakes when evaluating a product or service (78.0 percent strongly agree, or agree). A majority of spectators also agree with the statement that companies often profit from consumers' mistakes (78.5 percent agree, or strongly agree) and with the statement that companies actively target consumers who are likely to make mistakes (58.8 percent agree, or strongly agree).

We also measured attitudes to government regulations by asking the participants about the extent to which they agree with the statement that "The government should restrict businesses' opportunity to make profit from customers who misunderstand the value of a product or service". More people agree with the statement (42.6 percent strongly agree, or agree) than disagree (30.2 percent strongly disagree, or disagree).

We find a strong correlation between people's decision in the experiment and their support for government regulation. Table 2.5 reports regressions on whether the spectators either strongly agree or agree that the government

Figure 2.5: Beliefs and policy attitudes

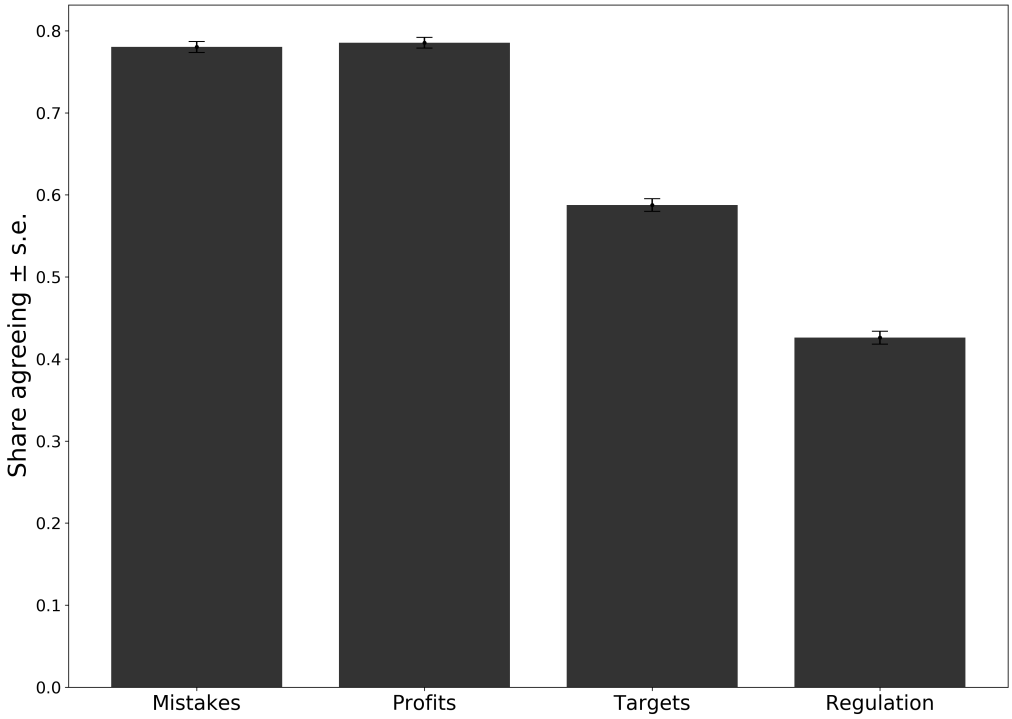

Note: The figure reports the share of spectators who strongly agree or agree to the following statements: "People often have the wrong beliefs about how valuable a product or service would be for them" (Mistakes), "Businesses often make profit from customers who misunderstand the value of a product or service" (Profits), "Businesses actively target customers who are likely to overestimate the value of their product or service" (Targets), and "The government should restrict businesses' opportunity to make profit from customers who misunderstand the value of a product or service" (Regulate).

Table 2.5: Government regulation

	Support for regulation	Support for regulation	Support for regulation	Support for regulation	Support for regulation	Support for regulation
Cancel	0.248***					0.205***
	(0.015)					(0.016)
Fine		0.243***				0.130***
		(0.021)				(0.022)
Mistakes			0.021			-0.016
			(0.019)			(0.019)
Profits				0.162***		0.106***
				(0.018)		(0.019)
Targets					0.152***	0.113***
					(0.016)	(0.016)
Female	-0.000	0.018	0.019	0.011	0.019	-0.003
	(0.015)	(0.016)	(0.016)	(0.016)	(0.016)	(0.015)
Age	-0.102***	-0.074***	-0.084***	-0.083***	-0.071***	-0.084***
	(0.015)	(0.015)	(0.016)	(0.015)	(0.016)	(0.015)
College	-0.024	-0.025	-0.030*	-0.034**	-0.033**	-0.028*
	(0.016)	(0.017)	(0.017)	(0.017)	(0.017)	(0.016)
Income	-0.023	-0.031*	-0.033*	-0.029*	-0.032*	-0.021
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.016)
Republican	-0.088***	-0.096***	-0.103***	-0.100***	-0.093***	-0.078***
	(0.016)	(0.016)	(0.017)	(0.017)	(0.016)	(0.016)
Constant	0.376***	0.470***	0.502***	0.396***	0.421***	0.234***
	(0.019)	(0.018)	(0.023)	(0.022)	(0.020)	(0.025)
Observations	3991	3991	3991	3991	3991	3991
R^2	0.081	0.055	0.022	0.040	0.045	0.117

Note: The table reports OLS regressions on an indicator for whether the spectator strongly agreed or agreed with the statement "The government should restrict businesses' opportunity to make profit from customers who misunderstand the value of a product or service". "Mistakes" is an indicator variable for the spectator strongly agree or agree with the statement "People often have the wrong beliefs about how valuable a product or service would be for them", "Profits" is an indicator variable for the spectator strongly agreeing or agreeing with the statement "Businesses often make profit from customers who misunderstand the value of a product or service", "Targets" is an indicator variable for the spectator strongly agreeing or agreeing with the statement "Businesses actively target customers who are likely to overestimate the value of their product or service". "Female" is an indicator variable for being female. "College" is an indicator variable for the spectator having a bachelor degree or higher. "Age" is an indicator variable for the spectator being . "Income" is an indicator variable for the spectator having a household income of more than $60.000 per year. "Republican" is an indicator variable for the spectator voting for the Republican party. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

should restrict businesses' opportunity to make profit from customers who misunderstand. From columns (1) and (2) we observe that those who cancel the deal or fine the seller are much more likely to agree with the need for government regulation ($p < 0.001$ and $p < 0.001$). From column (6) we observe that this holds even when we control for beliefs and background characteristics. We summarize these findings in the following result:

Result 5: *Behavior in the experiment is strongly predictive of attitudes to government regulation. Spectators who cancel the deal or fine the seller are more likely to agree that the government should restrict businesses' opportunity to profit from irrational consumers.*

2.5 Concluding remarks

We have presented the first set of evidence showing that a majority of Americans prefer to cancel voluntary deals in which a consumer has misunderstood the true value of a product and pays more for the product than it is worth. This is the case even when the seller's role is limited to accepting a proposal made by the buyer. This suggests that the main concern people have with deals in which a seller exploits the irrationality of a buyer are the consequences which these deals produce.

The willingness to cancel deals in which consumers are exploited sheds light on the widespread support for stricter government regulation of businesses, such as the US CARD act (Agarwal et al., 2015) or European regulations about "unfair" contract features (Heidhues et al., 2018). However, the fact that a large minority choose to uphold such transactions even in settings where the seller has actively contributed to the buyer's confusion and then

proposed the deal also sheds light on why such regulatory efforts often meet fierce resistance.

Only a small minority of the spectators consider the seller's involvement to be critical for the decision to cancel or not. In contrast, the seller's involvement, particularly whether they have obfuscated the information the buyer received, is important for the decision to fine. This suggests that procedural concerns play a more important role in people's willingness to fine the seller than in their willingness to cancel the deal. Further, the finding that procedural concerns are strongly driven by opposition to complex information disclosure could offer insight into widespread acceptance of policy initiatives that improve access to information from firms (Reisch and Sunstein, 2016; Sunstein, 2019b). This result underlines the importance of firms' decisions about how to disclose information about their products and it provides evidence that complex information can negatively influence consumer perception of a firm (Bao and Ho, 2015; Kozup et al., 2008; Nguyen and Mutum, 2012).

The present study also suggests that fairness considerations need to be taken into account by firms when considering how to handle their customers. In line with previous research (Bhattacharjee et al., 2017; Kahneman et al., 1986), we find that people are willing to punish firms that earn a profit from behavior that they perceive as unacceptable. We also show that simply accepting to be paid more for a product than it is worth can be seen as unacceptable.

In this study, we have examined people's attitudes to deals where consumers who make mistakes when calculating the value of a product are exploited. An interesting question for future research is to examine people's attitudes

to the exploitation of other types of irrationality among consumers. Finally, the present study is conducted with participants recruited from the general population in the US. An important question for future research is how attitudes in the US compare with attitudes in other countries.

2.A Demographics

Table 2.6: Descriptive Statistics

	Number of participants	Share
Income < $30,000	1130	0.28
Income 30,001 - $60,000	1136	0.28
Income 60,001 - $100,000	925	0.23
Income 100,001 - $150,000	523	0.13
Income > $150,000	277	0.07
High School Education or below	1158	0.29
Some College Education	1342	0.34
Bachelor or equivalent	956	0.24
Master or equivalent	535	0.13
18 - 34 years old	1035	0.26
35 - 44 years old	740	0.19
45- 54 years old	812	0.20
55 - 64 years old	682	0.17
65+ years old	722	0.18
Female	2084	0.52
Republican	1232	0.31
Observations	3991	

2.B Additional results

Table 2.7: Regression results for decision to fine the among the sellers that canceled the deal

	Fine	Fine
Proposed	0.051**	0.053**
	(0.024)	(0.024)
Obfuscated	0.104***	0.101***
	(0.025)	(0.024)
Proposed*Obfuscaed	-0.028	-0.027
	(0.036)	(0.035)
Female		-0.027
		(0.019)
Age		-0.003***
		(0.001)
College		-0.018
		(0.019)
Income		0.010
		(0.020)
Republican		-0.023
		(0.020)
Constant	0.195***	0.365***
	(0.017)	(0.035)
$\beta_1 + \beta_2 + \beta_3$	0.127***	0.128***
	(0.025)	(0.025)
$\beta_2 + \beta_3$	0.075***	0.073***
	(0.026)	(0.026)
R^2	0.012	0.026
Observations	2424	2424

Note: The table reports results from an OLS-regression of the share of spectators that fine the seller among spectators that canceled the deal. "Proposed" is an indicator variable for the spectator being assigned to the treatment where the seller proposed the transaction. "Obfuscated" is an indicator variable for the spectator being assigned to the treatment in which the seller obfuscated the information. "Proposed*Obfuscated" is the interaction between "Proposed" and "Obfuscated". "Female" is an indicator variable for being female. "College" is an indicator variable for the spectator having a bachelor degree or higher. "Age" is an indicator variable for the spectator being older than 47 years. "Income" is an indicator variable for the spectator having a household income of more than $60,000 per year. "Republican" is an indicator variable for the spectator voting for the Republican party. $\beta_1 + \beta_2 + \beta_3$ is the linear combination of the parameters for Proposed, Obfuscated and the interaction term between the two. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Standard errors in parentheses.

2.C Instructions

Low Treatment

On the next screen we ask you to make a decision that might have real consequences for two other individuals. We recruited these two individuals on an international online labor market to participate in a study. Both individuals received a participation fee.

We randomly assigned the two individuals to either the role of seller or the role of buyer. The seller was given a product that he or she could sell to the buyer.

We informed both the seller and the buyer that the product had no value for the seller. The seller, but not the buyer, was informed that the value of the product for the buyer was $2.
The seller could not disclose this information to the buyer. Instead, we gave the buyer information about the value of the product in a complex manner.

The buyer made a mistake when interpreting this information and wrongly believed that the value of the product for him or her was $20. The seller knew that the buyer made this mistake.

The buyer offered to buy the product for $10 from the seller. The seller accepted this offer.
- The seller gained $10 on the deal.
- The buyer lost $8 on the deal.

We now want you to decide whether this deal should be upheld or not.

o I want to uphold the deal. The seller gains $10 and the buyer loses $8.
o I want to cancel the deal. Neither the seller nor the buyer gain or lose.

There is a one-in-five chance that your decision will be implemented. If your decision is implemented, the seller and the buyer will receive payments according to your decision within a few days. The seller and the buyer are informed that a third-party will make a decision that determines their payments.

Propose treatment

On the next screen we ask you to make a decision that might have real consequences for two other individuals. We recruited these two individuals on an international online labor market to participate in a study. Both individuals received a participation fee.

We randomly assigned the two individuals to either the role of seller or the role of buyer. The seller was given a product that he or she could sell to the buyer.

We informed both the seller and the buyer that the product had no value for the seller. The seller, but not the buyer, was informed that the value of the product for the buyer was $2.
The seller could not disclose this information to the buyer. Instead, we gave

the buyer information about the value of the product in a complex manner.

The buyer made a mistake when interpreting this information and wrongly believed that the value of the product for him or her was $20. The seller knew that the buyer made this mistake.

The seller offered to sell the product to the buyer for $10. The buyer accepted this offer.
- The seller gained $10 on the deal.
- The buyer lost $8 on the deal.

We now want you to decide whether this deal should be upheld or not.

o I want to uphold the deal. The seller gains $10 and the buyer loses $8.
o I want to cancel the deal. Neither the seller nor the buyer gain or lose.

There is a one-in-five chance that your decision will be implemented. If your decision is implemented, the seller and the buyer will receive payments according to your decision within a few days. The seller and the buyer are informed that a third-party will make a decision that determines their payments.

Obfuscate Treatment

On the next screen we ask you to make a decision that might have real consequences for two other individuals. We recruited these two individuals on an international online labor market to participate in a study. Both indi-

viduals received a participation fee.

We randomly assigned the two individuals to either the role of seller or the role of buyer. The seller was given a product that he or she could sell to the buyer.

We informed both the seller and the buyer that the product had no value for the seller. The seller, but not the buyer, was informed that the value of the product for the buyer was $2.
The seller had the opportunity to disclose information about the value of the product to the buyer in an easy-to-understand manner, but decided not to do so. Instead, the seller decided to disclose the information to the buyer in a complex manner.

The buyer made a mistake when interpreting this information and wrongly believed that the value of the product for him or her was $20. The seller knew that the buyer made this mistake.

The buyer offered to buy the product for $10 from the seller. The seller accepted this offer.
- The seller gained $10 on the deal.
- The buyer lost $8 on the deal.

We now want you to decide whether this deal should be upheld or not.

o I want to uphold the deal. The seller gains $10 and the buyer loses $8.
o I want to cancel the deal. Neither the seller nor the buyer gain or lose.

There is a one-in-five chance that your decision will be implemented. If your decision is implemented, the seller and the buyer will receive payments according to your decision within a few days. The seller and the buyer are informed that a third-party will make a decision that determines their payments.

High Treatment

On the next screen we ask you to make a decision that might have real consequences for two other individuals. We recruited these two individuals on an international online labor market to participate in a study. Both individuals received a participation fee.

We randomly assigned the two individuals to either the role of seller or the role of buyer. The seller was given a product that he or she could sell to the buyer.

We informed both the seller and the buyer that the product had no value for the seller. The seller, but not the buyer, was informed that the value of the product for the buyer was $2.
The seller had the opportunity to disclose information about the value of the product to the buyer in an easy-to-understand manner, but decided not to do so. Instead, the seller decided to disclose the information to the buyer in a complex manner.

The buyer made a mistake when interpreting this information and wrongly

believed that the value of the product for him or her was $20. The seller knew that the buyer made this mistake.

The seller offered to sell the product to the buyer for $10. The buyer accepted this offer.
- The seller gained $10 on the deal.
- The buyer lost $8 on the deal.

We now want you to decide whether this deal should be upheld or not.

o I want to uphold the deal. The seller gains $10 and the buyer loses $8.
o I want to cancel the deal. Neither the seller nor the buyer gain or lose.

There is a one-in-five chance that your decision will be implemented. If your decision is implemented, the seller and the buyer will receive payments according to your decision within a few days. The seller and the buyer are informed that a third-party will make a decision that determines their payments.

Chapter 3

Partisan self-interest is an important driver for people's support for the regulation of targeted political advertising

Katharina Baum Stefan Meissner

Hanna Krasnova*

Abstract

The rapid emergence of targeted political advertising has sparked a heated public debate over what the government's response should be, and has led to public pressure advocating stricter regulation. To date, the regulatory debate has centered around public concerns about the collection and use of

*Meissner: Department of Economics, NHH Norwegian School of Economics; Baum and Krasnova: Weizenbaum Institute for the Networked Society and Faculty of Economic and Social Sciences, University of Potsdam. We thank Björn Bartling, Alexander W. Cappelen, Alex Imas, George Loewenstein, Erik Ø. Sørensen and Bertil Tungodden and seminar participants at the Weizenbaum Institute and the Norwegian School of Economics for helpful comments. This work was partially supported by the Research Council of Norway through its Centres of Excellence Scheme, FAIR project No 262675. This work has been partially funded by the Federal Ministry of Education and Research of Germany (BMBF) under grant no. 16DII116 ("Deutsches Internet-Institut"). This work has been partially funded by the Centre for Ethics and Economics at the Norwegian School of Economics.

citizens' private data. This paper tested and confirmed the hypothesis that public attitudes toward stricter regulation of targeted political advertising are also motivated by partisan self-interest. We conducted an experiment using an online survey of 1549 Americans who identify as either Democrats or Republicans. Our findings show that Democrats and Republicans believe that targeted political advertising benefits the opposing party. This belief is based on their conviction that their political opponents are more likely to mobilize by targeted political advertising than are supporters of their own party. We exogenously manipulated the beliefs of a random subset of participants by truthfully informing them that, in the past, targeted political advertising has benefited Republicans. This enabled us to establish a causal link between beliefs about partisan advantage and attitudes toward stricter regulation. Our findings show that Republicans informed about this benefit had less favorable attitudes toward regulation than did their uninformed co-partisans. This suggests that participants' attitudes regarding stricter regulation are based not solely on concerns about privacy violations, but also, in part, on beliefs about whether regulation would benefit their party. This result implies that people are willing to accept violations of their privacy if their preferred party benefits from the use of targeted political advertising.

Keywords: Electoral competition | Targeted political advertising | Government regulation | Third-person effect | Privacy concerns

3.1 Introduction

Recent advances in technology and the availability of vast amounts of personal data online have dramatically altered a key element of the electoral process: political campaigning (Fowler et al., 2020). Political parties and

campaigns can now microtarget specific messages to narrow groups of voters based on granular personal data (De Corniere and De Nijs, 2016) (see Appendix for a discussion). Targeted political advertising as a new method of political campaigning is quickly becoming a major tool for political actors (Hager, 2019) and has been publicly implicated as a factor causing unanticipated outcomes in a number of elections (Benkler et al., 2018). According to public opinion polling, the large majority of Americans considers the use of personal data for targeted online political ads unacceptable (Smith, 2018), and a heated public debate calling for stricter regulations has accompanied the emergence of such ads (Aral and Eckles, 2019; Dobber et al., 2019a; Kim et al., 2018). In response to public pressure (Isaac, 2019), Twitter and Google have already instituted self-imposed measures that either ban the use of targeted political advertising outright or limit the technical abilities of campaigns to use these platforms (Lerman, 2019; Wong, 2019). Facebook has responded by establishing an archive that stores all political ads that have been run on the platform (Leathern, 2020) and an option to opt-out from seeing political advertisement for people in the United States (Isaac, 2020). Despite the potentially far-reaching consequences posed by targeted online political ads and the mounting public pressure to regulate them, the political response has been slow. Consequently, targeted online political ads are still largely unregulated (Beyersdorf, 2019; Dommett and Power, 2019; Weintraub, 2019). Both public calls for regulation and private sector directives address primarily a lack of protection and transparency regarding the use of personal data for targeted political ads (Burkell and Regan, 2019; Dobber et al., 2019b; Dommett, 2019).

In fact, the recent debate over stricter regulation has focused on restrictions on the use of personal data (Sihvola, 2019) (see Appendix for a discussion).

Previous research has established that people value the privacy of their data and that privacy concerns are an important factor in determining people's attitudes toward the regulation of targeted advertising in general (Acquisti et al., 2016; Bellman et al., 2004; Milberg et al., 2000; Okazaki et al., 2009) (see Appendix for a discussion). Unease about the use of private data appears particularly pressing in the context of targeted political advertising, as such advertising requires the collection, storage and use of large amounts of sensitive data about people's political attitudes (Baum et al., 2019; Rubinstein, 2014). Furthermore, people seem especially worried about the use of their private data by political actors (Tan et al., 2018). Therefore, the public debate about stricter regulation of targeted political advertising has focused largely on the data security and privacy consequences of such data collection, as well as on the lack of transparency in its use (Boerman et al., 2017; Burkell and Regan, 2019; Dobber et al., 2019b; Dommett, 2019; Magalhães et al., 2018; Wood and Ravel, 2017).

In this paper, we argue that attitudes toward the regulation of targeted online political advertising are driven not only by concerns over the misuse of private data. While targeted commercial advertising influences only individuals' purchasing choices, targeted political advertising has the potential to influence voting decisions and, as a result, elections (Zuiderveen Borgesius et al., 2018; Magalhães et al., 2018). This has consequences for broader societal outcomes, affecting far more than individual data protection. We posit that people take these consequences into account when forming preferences regarding the regulation of targeted political ads. Research on public opinions about other aspects of the electoral process indicates that self-interest is an important factor in people's positions about the electoral effects of regulations (Alvarez et al., 2011; Biggers, 2019; Boix, 1999). Attitudes on gerrymander-

ing, voter ID laws, or same-day voter registration all seem to be driven by partisan self-interest, or by the concern for ensuring advantages for one's preferred party (Ansolabehere, 2009; Chen and Rodden, 2013; McCarthy, 2019). This study seeks to explore whether, in addition to privacy concerns, partisan self-interest is an important determinant of people's attitudes towards stricter regulation of targeted political advertising. For that purpose, we ran an experiment in the United States using an online survey. Working with a sample of Republican and Democratic participants, we investigated participants' beliefs about the consequences of using targeted political advertising aimed at voters of both parties in order to determine whether there is a link between partisan self-interest and attitudes towards targeted political advertising.

Understanding people's beliefs about the effects of political advertising on electoral outcomes is critical for ascertaining the underlying drivers of public attitudes towards stricter regulation. Political parties use targeted political ads mainly to mobilize their own voters. Therefore, people's perceptions as to whether targeted political advertising benefits or harms their party depend on whether they believe that voters of their own party are mobilized more strongly than are voters of the opposing party or vice versa. If people are motivated by partisan self-interest, they would oppose regulation in the first case, based on their perception that targeted political advertising would give their party an advantage in mobilization. In the latter case, people would demand regulation in order to mitigate the opposing party's mobilization advantage. We hypothesize that supporters of both parties believe targeted political ads yield an advantage for the opposing party.

Due to the potential difficulty people may have in correctly estimating the

actual effects of targeted political advertising on others, it is plausible that they could hold biased or unfounded beliefs about the issue. As a result, to assess the ads' effects, Democrats must guess how Republicans react to mobilizing messages and vice versa. However, campaign messages that are delivered to targeted recipients remain largely unavailable to others (Magalhães et al., 2018). Given the limited transparency of targeted political advertising (Zuiderveen Borgesius et al., 2018; Wood and Ravel, 2017), as well as the paucity of information about its effects on voters (Aral and Eckles, 2019), it seems likely that people do indeed have difficulties arriving at accurate estimates (FeldmanHall and Shenhav, 2019).

Academic work on the extent to which a person will be influenced by targeted online political ads, while still scarce, does suggest that demographics, place of residence, and political ideology all play a role in determining this phenomenon (Liberini et al., 2020). There is, however, a large body of research on people's beliefs about the effect of undesirable persuasive mass communication on others, documenting that people generally believe that others are influenced by it to a larger extent than they are themselves. This phenomenon is known as the third-person effect (Davison, 1983; Perloff, 1993) (See Appendix for a discussion). Past studies have shown that the strength of the third-person effect increases with social distance to the "other" (Jang and Kim, 2018; Perloff, 1999; White, 1997). Furthermore, the third-person effect predicts that people not only believe that others are more influenced by undesirable mass communication, but that these people also take action to rectify the consequences of such persuasive messages (Xu and Gonzenbach, 2008). High levels of polarization and mistrust between Democrats and Republicans in the United States suggest that the social distance between partisans is large (Ahler and Sood, 2018; Bordalo et al., 2016; Finkel et al., 2020;

Iyengar et al., 2019; Lees and Cikara, 2020; Martherus et al., 2019; Mason, 2018; Moore-Berg et al., 2020; Nyhan, 2020). Hence, the potential presence of the third-person effect, combined with a large social distance between the parties, suggests that both Republicans and Democrats may believe that opposing partisans are influenced by targeted political advertising to a larger extent than are supporters of their own party. Crucially, this means that the opposing party is perceived as gaining more from the use of mobilizing messages directed at their own electorate than is one's own party. As a result, according to the literature on the third-person effect, it can be inferred that people who believe voters of the other party are more influenced by targeted political ads than are voters of their own party will also support regulation of this advertising.

We therefore further hypothesize that supporters of each party believe that supporters of the opposing party are more strongly influenced by targeted political advertising than they themselves are. As a consequence, they believe that the other party experiences an advantage from targeted online political ads and, therefore, favor stricter regulation, perceiving it to be in their partisan self-interest. To test these hypotheses in our experiment, we measured participants' beliefs about the effect of targeted political advertising on both co-partisans and supporters of the opposing party, as well as their respective attitudes toward regulation. To establish the existence of a causal link between partisan self-interest and attitudes toward regulation, we exogenously manipulated participants' beliefs about the effect of targeted political advertising.

This study is composed of a correlational and an experimental part. The correlational part provides evidence that participants believe that supporters

of the opposing party are more influenced by targeted political advertising than are supporters of their own party. Importantly, we also show that beliefs about the effect of targeted political advertising on supporters of the other party relative to supporters of one's own party are positively correlated with a stronger demand for regulation. As a consequence, support for stricter regulation is linked not only to concerns about individuals' privacy, but also to participants' beliefs about partisan self-interest. In the experimental part of the study, we truthfully informed a randomly selected sample of participants that the Republican party benefited more than the Democratic party from the use of targeted political advertising in the 2016 presidential election. Thereby, we changed Republicans' perceptions of partisan self-interest without altering their concerns about privacy. Republican recipients of this information were less supportive of regulation than were their co-partisans who had not been given this information. This finding reveals a causal link between beliefs about partisan self-interest and people's attitudes toward stricter regulation. Our results reveal the challenges posed by new technological advances in the political domain and the ensuing need for new regulation. We show that some partisans are willing to oppose regulation if they believe that targeted political advertising benefits their preferred party, even at the expense of concerns about privacy violations and massive data collection. Our findings further reveal that attitudes toward regulation are partially driven by biased beliefs about the effect of targeted political advertising on others, since participants from both parties believe that regulation is in their own partisan self-interest.

3.2 Experimental methods

3.2.1 Experimental design

We conducted a pre-registered, incentivized online survey experiment with a sample of adult Americans identifying either as Democrats or Republicans. The pre-registration is available at the AEA RCT Registry AEARCTR-0005296. The study received an Institutional Review Board (IRB) approval from the Norwegian School of Economics, and all participants gave informed consent before taking part in the study. Part C of the appendix contains detailed information about instructions and measurements. Figure 3.3 in the appendix provides an overview of the structure of the experiment.

There were three phases to this study. In the first phase, we informed participants about targeted online political advertising and measured their beliefs about its effect on supporters of both the Republican and Democratic parties. In the second phase, the experimental manipulation was conducted by informing a random subset of participants about the beneficial effects of targeted political ads for Republicans. In the third phase, we measured all participants' attitudes toward the regulation of targeted political ads, performed a manipulation check, and measured respondents' demographics along with a number of other control variables. The following describes each phase in detail.

To ensure that all participants had the same knowledge on the subject, in the first phase of the study, participants were asked to read a text about targeted online political advertising that explained its technical aspects and its typical usage. We then asked participants to consider a hypothetical scenario in

which both Republicans and Democrats competed in a close electoral race in which they spent equivalent sums on targeted online political advertising. We elicited participants' beliefs about the extent to which they thought that they personally, Republicans and Democrats alike, would be influenced by targeted political advertising, using a five-point Likert scale ranging from "not at all" to "to a very great extent." This measurement corresponds to previous findings from the literature on the third-person effect (Jang and Kim, 2018; Perloff, 1999). The order of the questions about Republicans and Democrats was randomized. To address concerns that participants could potentially want to give negative answers about the opposing side while not necessarily believing that such answers had a basis in fact (Bullock et al., 2015; Gerber and Huber, 2010), we emulated the approach of previous studies (Prior et al., 2015), and asked participants to commit to answering the questions to the best of their knowledge.

In the second phase of the survey, participants were randomly placed in either the treatment or the control group. Participants in the treatment group were informed that controlling for the number of ads people saw, targeted political advertising on Facebook significantly increased voter turnout for the Republicans in the 2016 presidential election, while having no effect on Democrats. With this wording, we ensured that participants did not look to different levels of campaign spending as a possible cause of the ads' effects. The complete wording of the information used with the treatment group can be found in the appendix part C. The results are based on a study by Liberini et al. (Liberini et al., 2020).

In the final phase of the study, we measured all participants' attitudes towards regulation of targeted political online political advertising on a four-

item, seven-point Likert scale ranging from "strongly disagree" to "strongly agree"(adapted from (Krasnova et al., 2009)). The possible responses were: (i) Targeted political advertising should be banned; (ii) I support legislation that requires targeted online political advertising to be clearly marked as targeted; (iii) More regulation is needed when it comes to targeted online political advertising; and (iv) The government is already doing enough to regulate targeted online political advertising (reverse coded). The order of these responses was randomized. We incentivized honest answers by informing participants that their responses would be sent to the United States Congress in an aggregated and anonymous form (Elías et al., 2019), stressing that there was no deception in the study.

To determine whether the information treatment succeeded in manipulating beliefs about the effects of targeted political advertising of participants in the treatment group, all subjects were then asked to make an estimation of the number of interactions (likes, shares, comments) that social media campaigns on Facebook of both Republicans and Democrats received relative to each other prior to the midterm elections in 2018. This enabled us to ascertain whether participants generalized from information about the 2016 Presidential election and applied it to other elections. We offered a monetary incentive for participants to answer the question to the best of their knowledge (De Quidt et al., 2018). Participants giving the correct answer received a bonus of $1 (Bullock et al., 2015; Flynn et al., 2017). The exact wording of the question can be found in the appendix part C. In order to control for the possibility that the intervention influenced only beliefs about targeted political advertising's persuasiveness, but not about other problematic aspects of such advertising, we also measured whether participants thought the advertising was: (i) socially desirable; (ii) harmful to society

(reverse coded); (iii) beneficial to cultural values; and (iv) unfavorable to societal norms (reverse coded) on a ten-point scale.

To assess the level of privacy concerns, we presented participants with a four-item, seven-point Likert scale questionnaire (adapted from (Krasnova et al., 2009)) in which we asked participants whether they were concerned that their data was: (i) collected and stored by third-parties; (ii) shared with third-parties; (iii) used to display targeted advertising to them; and (iv) used for commercial purposes. The order of the items was randomized. We further included a fifth item as an attention check to ensure that participants carefully read the items. In accordance with our pre-analysis plan, participants who failed this attention check and another attention check were not included in the final sample.

We further collected data for political attitudes in terms of political engagement, subjective political knowledge, participants' level of social and economic conservatism (Everett, 2013), an ANES-based feelings thermometer towards both the Republican and the Democratic parties (Iyengar et al., 2019), and participants' perceived political efficacy (Bowler and Donovan, 2002). The demographic control variables included age, gender, ethnicity, education, income, household size, use time on the internet, use of an ad-blocker and social media usage.

3.2.2 Sample characteristics

We collected the data for this survey between the January 15, 2020 and January 24, 2020. We collaborated with the survey company Dynata to recruit our participants. For that purpose, we used Dynata's political panel to recruit

Republicans and Democrats, as Dynata collaborates with L2, the largest voting tracking service in the United States. Therefore, we were able to recruit Democrats and Republicans for whom party affiliation was partially verified by their actual voting behavior. That further enabled us to avoid recruiting Independents for our study. In total we recruited a demographically diverse sample of 1549 American participants who were either Democrats or Republicans. On average, participants were 47.49 years old. Of the sample, 50.55 percent were female and 25.05 percent were non-white. The participants were better educated than the overall population of the United States. Appendix Table 3.3 provides an overview of the characteristics of our sample. Among the participants, 777 identified as Republicans and 772 as Democrats. Given the nature of the experimental design, Independents were not included in the study. We randomly assigned the participants to either the treatment group (755 participants: 369 Democrats, 386 Republicans) or the control group (794 participants: 403 Democrats, 391 Republicans). Treatment assignment was balanced taking into consideration observable characteristics and pre-treatment beliefs (Appendix Table 3.4).

3.3 Results

This section presents the results of the study. First, we will present evidence supporting the hypothesis that supporters of both parties believe that supporters of the opposing party are influenced to a larger extent by targeted political advertising than are supporters of their own party. This implies that they believe that the use of targeted political advertising undermines their partisan self-interest. We will then present correlational results regarding the link between these beliefs, privacy concerns and support for stricter

regulation. Last, we will present our findings about the causal role of beliefs about the effects of targeted political advertising on attitudes towards regulation. The analysis was performed using Stata SE 16.0. The data, full instructions for participants, analysis code and variable coding are available at 10.17605/OSF.IO/QM7DZ.

3.3.1 Beliefs about the differential effect of targeted political advertising on opposing versus fellow partisans

Figure 3.1 shows the participants' beliefs about the extent to which targeted political advertising influences Republicans and Democrats. We found that Republicans believed that Democrats ($\mu = 3.20$, SD = 1.18) are more influenced than are Republicans ($\mu = 2.83$, SD = 1.10, Wilcoxon-signed-rank-test, z = -8.67, p < 0.001, r = 0.41). In contrast, Democrats stated that they believed that Republicans ($\mu = 3.41$, SD = 1.17) were more influenced than were Democrats ($\mu = 2.94$, SD = 1.02, Wilcoxon-signed-rank-test, z = -11.336, p < 0.001, r = 0.31). Consistent with the third-person effect, these results show that Republicans as well as Democrats expressed the belief that supporters of the opposing party are more influenced by targeted political advertisement than are supporters of their own party. Exploratory data analysis reveals that the difference in participants' beliefs about the effect of targeted political advertisement on opposing party supporters relative to supporters of their own party is not significantly different between Republicans and Democrats (two-sided Welch t-test, t(1540), d = 0.08, p = 0.11). The belief gap between own and other party indicates that supporters of both parties believe that the opposing party benefits more from the use of targeted political advertising than does their own, and therefore they perceive such advertising as harmful

to their partisan self-interest. In line with the overall low perceived desirability of these advertisements ($\mu = 4.66$, SD = 2.01, measured on a ten-point scale), we found that participants believed that targeted political advertising had a small influence on themselves ($\mu = 2.39$, SD = 1.21).

Figure 3.1: Beliefs about the effect of targeted political advertising

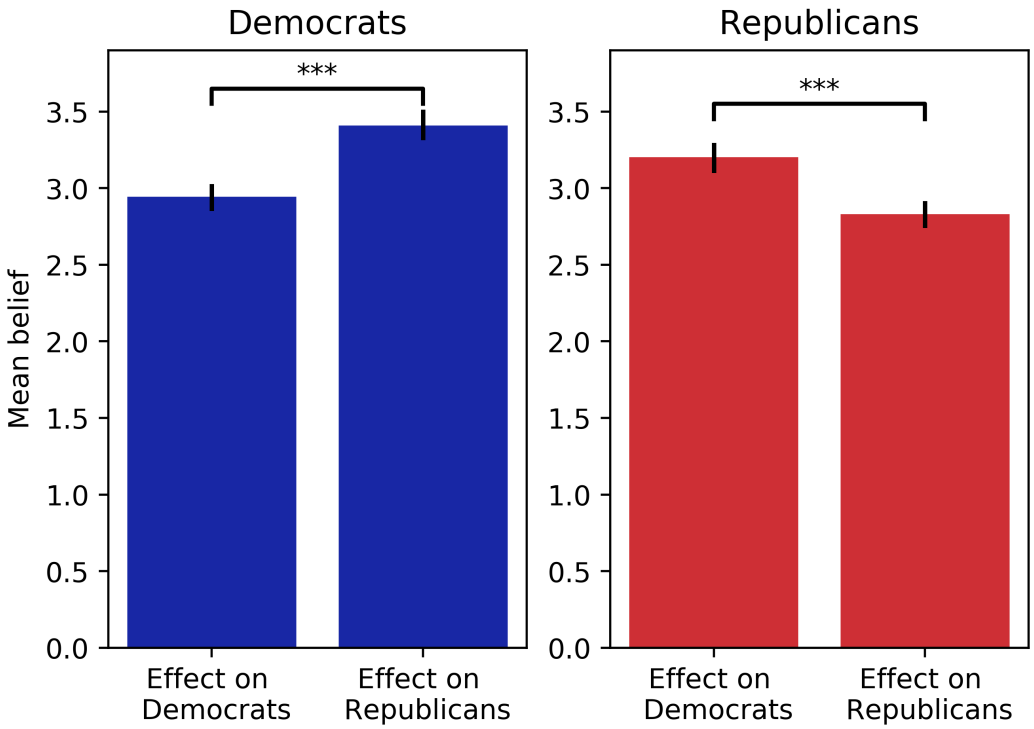

Note: The figure shows the beliefs of participants about the effect of targeted political advertising on Democrats and Republicans. Beliefs are measured on a five-point Likert scale (1 = "not at all", 5 = "to a very great extent"). The bars show 95% confidence intervals. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The size of the belief gap between one's own party and the other party is correlated to different attitudes that participants hold. We found a significant positive correlation between this gap and higher levels of affective and ideological polarization, perceived desirability of the advertising, and high subjective political knowledge. Participants holding a more negative view of the opposing party as measured on a feelings thermometer reported their belief in a larger difference in effects on supporters of the opposing parties

and supporters of their own party (Appendix Table 3.5, OLS-Regression, p < 0.001). We also found that the level of conservatism for Republicans and liberalism for Democrats as measured on a scale for social and economic conservatism (Everett, 2013) positively correlated with their beliefs about how strongly opposing party supporters are influenced by targeted political advertising (Appendix Table 3.5, OLS-Regression, p < 0.001). Participants who saw the advertising as more socially and culturally desirable reported a significantly smaller gap in beliefs between their own party and the other party (Appendix Table 3.5, OLS-Regression, p < 0.001). Taken together, these results suggest that people's belief that supporters of the opposing party are more influenced than supporters of their own party by political advertising is linked to a negative perception of the opposition and a more general dislike of targeted political advertising. This conclusion accords with previous literature on the third-person effect that suggests that people's belief about the influence of media messages on others relative to themselves correlates with the social distance to the other and a negative perception of the message. Moreover, participants who self-reported a high level of political knowledge reported a larger gap between their own party and the other party (Appendix Table 3.5, OLS-Regression, p = 0.04).

3.3.2 Support for government regulation

On average, we found that participants were in favor of regulation ($\mu = 4.82$, SD = 1.18 , Cronbach's-α $= 0.67$). Figure 3.4 in the Appendix shows the distribution of support for regulation. Overall, 70 percent of participants supported stricter regulation of targeted political advertisement. Support for stricter government regulation was higher in the baseline condition (two-

sided Welch t-test, t(782), Cohen's-d = 0.41, p < 0.001) among participants who identified as Democrats ($\mu = 5.06$, SD = 1.10) compared to Republicans ($\mu = 4.59$, SD = 1.21). We further found that, on average, participants were concerned about the use of their private data in targeted political advertising ($\mu = 5.63$, SD = 1.25, Cronbach's-α = 0.90). This concern was not significantly different (two sided Welch t-test, t(1529) = 0.10, Cohen's-d = 0.05, p = 0.31) between Democrats ($\mu = 5.67$, SD = 1.26) and Republicans ($\mu = 5.60$, SD = 1.25). Figure 3.5 shows the distribution of privacy concerns among participants. This concern was not significantly different (two sided Welch t-test, t(1529), Cohen's-d = 0.05, p = 0.31) between Democrats ($\mu = 5.67$, SD = 1.26) and Republicans ($\mu = 5.60$, SD = 1.25).

We ran an OLS-regression to test whether privacy concerns and beliefs about partisan self-interest were significantly correlated to participants' support for regulation. Partisan self-interest is measured as the difference between participants' beliefs about the effect targeted political advertising has on supporters of the other party and on supporters of their own party. Table 3.1 shows that support for stricter government rules is significantly linked to participants' belief about partisan self-interest (Belief other party - own party, p < 0.001). Column 1 shows that a 1 SD increase in the difference between the other party and one's own party is linked to a 0.12 SD increase in the support for government regulation. This parameter is virtually unaffected by the inclusion of control variables (Column 2). We further show that a 1 SD increase in privacy concerns of participants leads to a 0.28 SD increase in support for regulation (Column 1, p < 0.001). Column 2 shows that the inclusion of control variables does not significantly affect this parameter either. We find no significant link between participants' beliefs about the effect that targeted political advertising has on themselves and their support

for stricter regulation (Belief about effect on self, p = 0.19).

Table 3.1: Regression of determinants for the willingness to support stricter regulation of targeted advertising

	Support for regulation	Support for regulation
Belief other party-own party	0.119***	0.124***
	(0.034)	(0.035)
Belief about self	X	0.052
		(0.039)
Privacy concerns	0.283***	0.257***
	(0.043)	(0.045)
Observations	754	754
R^2	0.106	0.125
Demographics	No	Yes
Social Media use	No	Yes
Political Engagement	No	Yes

Note: Regressions only include participants that answered all questions of the survey. The table reports results from an OLS-regression in which people's support for regulation is the dependent variable. The value is standardized. Belief other party-own party is defined as the difference between people's belief about the effect on the other party and the effect on the own party. Belief about self is people's belief about the effect that targeted political advertising has on them. Privacy concerns are respondents' are measured on a seven-point 4 item-Likert scale. All three independent variables are standardized. Demographic information included age, education (dummy for above median in the sample), income (dummy for above median in the sample), household size (dummy for more than two members), gender (male dummy variable) and a dummy for being non-white. Social media use was a dummy variable for the use of social media, a continuous variable for the time people spent online in general (in hours), and the use of an ad-blocker (dummy for yes). Political engagement was a dummy variable for being politically active within the last year, external political efficacy, political knowledge (dummy for above median knowledge), and attitudes towards government regulation in general. Table 3.11 in the appendix provides an overview of all variables in the regression. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

To assess the robustness of our findings, we also ran an OLS-regression using participants' beliefs about the effect of targeted political advertising on the opposing party and their beliefs about the effect on their own party as independent variable. Table 3.6 in the Appendix shows that support for stricter government regulation is strongly positively correlated to participants' beliefs about the effect on the other party (p < 0.001) and negatively correlated to the effect on own party, but this effect is not significant (p = 0.26).

3.3.3 The causal effect of beliefs about voters' susceptibility to targeted political advertisement and support for its regulation

In the treatment condition, we informed a randomly selected subgroup of Republicans and Democrats that Republicans benefited more from the use of targeted political advertising in the 2016 presidential election than did Democrats. Figure 3.2 shows the effect that this information had on beliefs about social media interactions in the 2018 midterm election. We found that in this incentivized question, Republicans and Democrats who had not received that information reported beliefs that were qualitatively similar to the first measure of beliefs. Uninformed Republicans believed that Democrats received more interactions in the run-up to the 2018 midterm elections while uninformed Democrats believed that Republicans received more interactions. Responses to this question and to the more general question about the effects of targeted political advertising on Republicans and Democrats are well correlated (r = 0.24). Figure 3.6 in the appendix illustrates the relationship between the answers to the general belief question that was reported on in section 3.3.1 and this one asking about the 2018 midterm elections.

There is no difference in beliefs about the 2018 midterm elections between Democrats who received the information about the 2016 Presidential election and those who did not (χ^2-test, p = 0.65). This result is in accordance with the finding that Democrats already believed Republicans are more influenced by political online advertising than members of their own party and that the information they received was not new for them. Republicans who received that information reported that they believed that Republicans received more

interactions in 2018. This result represents a significant divergence in beliefs between informed and uninformed Republicans that corresponds to the information that they received (χ^2-test, p = 0.04). Figure 3.7 and Figure 3.8 in the appendix show the distribution of answers for this question.

Figure 3.2: Beliefs about social media engagement in the 2018 midterm elections

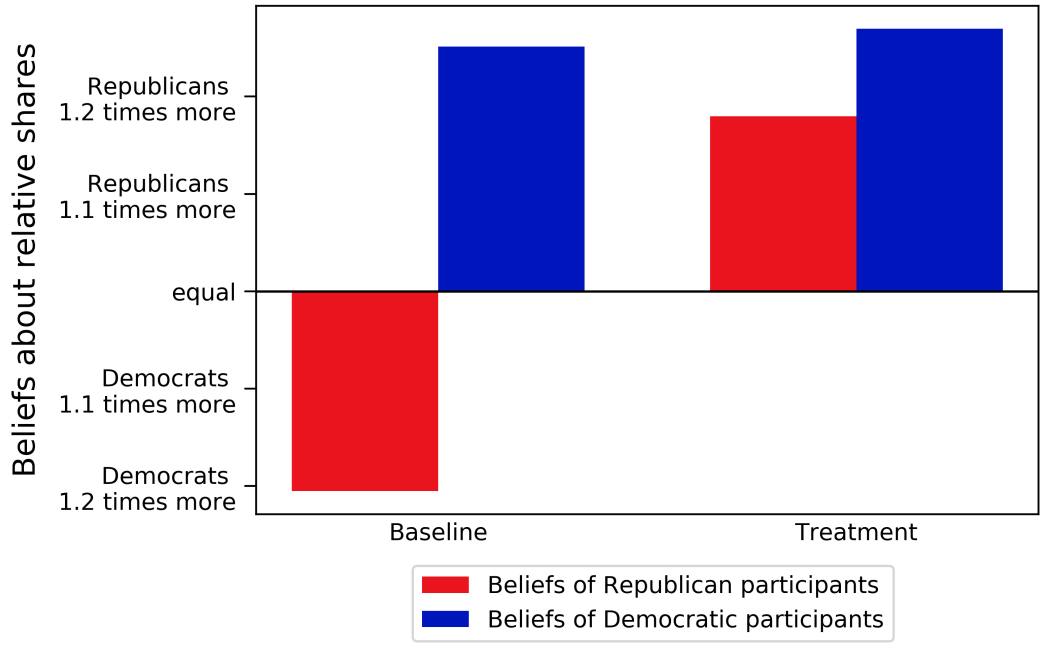

Note: The figure shows participants' beliefs about the ratio of interactions in the run-up to the 2018 Midterm election. This was measured on a scale that ran from "Democrats three times more than Republicans" to "Republicans three times more than Democrats" with "Equal" as the mid-point.

We next determined whether the information shared with participants shifted their support for stricter regulation of targeted political advertising. In accordance with the finding that beliefs of Democrats were not significantly influenced by the information, we found no effect on support for regulation between the treatment and the control group (two-sided Welch t-test, t(759), Cohen's d = 0.04, p = 0.58). Figure 3.9 in the appendix shows the distributions of answers for Democrats in the treatment and the control groups. With Republicans, we found significantly lower support for stricter regula-

tion of targeted political advertising between the treatment and the control groups (two-sided Welch t-test, t(776), Cohen's d = 0.15, p = 0.04). These effects remained qualitatively the same when examining only participants who wanted their opinions to be considered by Congress (98.7 percent of the sample) and participants who expressed trust in the information that they had received about the effect of targeted political advertisement (85.7 percent of the treatment group), although in the latter case, the effect became insignificant for Republicans (Appendix Table 3.7 and 3.8). Table 3.2 shows the magnitude of the shift for Republicans in a reduced form regression. Figure 3.10 shows the distribution of support for regulation for Republicans in the treatment and the control groups.

We found a downward shift in Republicans' support for regulation by 0.20 SD. That effect is approximately equivalent to a 1.68 SD increase in participants' belief regarding the extent to which Democrats are influenced by targeted political advertising relative to Republicans and a 0.71 SD downward shift in privacy concerns. This results in an approximately 50 percent increase in support for regulation between Republicans and Democrats in the treatment group compared to the control group ($\Delta_{control} = 0.47$, $\Delta_{treatment} = 0.70$).

To preclude the possibility that the information about the effect of targeted political advertising changed participants' perception of how desirable such advertising is or participants' privacy concerns, we tested for significant differences in these measures. We found that, in general, participants viewed the use of targeted political advertising as undesirable ($\mu = 4.66$, SD = 2.01). Comparing the ratings of the desirability of targeted political advertising for Republicans in the treatment ($\mu = 4.75$, SD = 2.00) and the control groups ($\mu = 4.85$, SD = 2.00), we found no statistically significant difference (two-

Table 3.2: Regression of Determinants for the willingness to regulate targeted ads

	Support for regulation	Support for regulation
Treatment*Republican	-0.203*	-0.192*
	(0.096)	(0.096)
Treatment	0.034	0.032
	(0.065)	(0.064)
Privacy concerns	0.297***	0.286***
	(0.030)	(0.030)
Republican	-0.362***	-0.411***
	(0.067)	(0.068)
Observations	1466	1466
R^2	0.150	0.166
Demographics	No	Yes
Social Media use	No	Yes
Political engagement	No	Yes

Note: Regressions only include participants that answered all questions of the survey. The table reports the results of an OLS-regression with the support for regulation as a dependent variable. The variable was standardized. Treatment is a dummy variable that is 1 when the participant was assigned to the treatment condition and 0 otherwise. Republican is a dummy variable that is 1 when the participant is a Republican and 0 if he or she is a Democrat. Treatment*Republican is the interaction of these two variables. Privacy concerns is a measure of people's privacy concerns. The value is standardized. Demographic information included age, education (dummy for above median in the sample), income (dummy for above median in the sample), household size (dummy for more than two members), gender (male dummy variable) and a dummy for being non-white. Social media use was a dummy variable for the use of social media, a continuous variable for the time people spent online in general (in hours), and the use of an ad-blocker (dummy for yes). Political engagement was a dummy variable for being politically active within the last year, external political efficacy, political knowledge (dummy for above median knowledge), and attitudes towards government regulation in general. Table 3.12 in the appendix provides an overview of all variables in the regression. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

sided Welch t-test, t(769), Cohen's d = 0.05, p = 0.49). The same result was found with Democrats in the treatment ($\mu = 4.42$, SD = 2.03) and the control groups ($\mu = 4.61$,SD = 2.00 two-sided Welch t-test, t(755), Cohen's d = 0.09, p = 0.20). We also found no significant differences in privacy concerns between the treatment and the control groups (two-sided Welch t-test, t(1526), d = 0.06, p = 0.32), for both Democrats (two-sided Welch t-test, t(759), Cohen's d = 0.05, p = 0.46) and Republicans (two-sided Welch t-test, t(768), Cohen's d = 0.05, p = 0.49).

Exploratory data analysis reveals that the effect of the information on Republicans was heterogeneous between different levels of conservatism. Figures 3.11 and 3.12 in the appendix illustrate the findings. We found that for those Republicans scoring below the median in social and economic conservatism, the information that their party benefited from the use of targeted political advertisement did not significantly change their support for regulation compared to the same group who did not receive this information (two-sided Welch t-test, t(403), Cohen's d = 0.02, p = 0.87). The support for stricter regulation of targeted political advertisement among Republicans scoring at or above the median in economic and social conservatism differed significantly between the treatment and the control groups (two-sided Welch t-test, t(373), Cohen's d = 0.27, p = 0.01). This effect can not be attributed to initial differences in the support for regulation in the baseline condition between above median and below median conservative Republicans (two-sided Welch t-test, t(347), p = 0.64).

3.4 Discussion

Our results provide evidence that the support for stricter regulation of targeted online political advertising is partially motivated by partisan self-interest. We show that both Republican and Democratic participants in our sample believed that supporters of the opposing party are influenced by targeted political advertising to a greater extent than are supporters of their own party. We found that both this belief and people's concern over privacy significantly drive people's support for policies limiting the use of such ads. Republicans who were informed about the beneficial effects of targeted online political ads for their party reported lower support for regulation than did Republicans in the control group. Therefore, we are able to show that the perception bias is causally linked to Republicans' support for stricter government regulation. This suggests that participants make a trade-off in favor of partisan self-interest and contrary to concerns about the violation of data privacy. We found that this effect is not present with all Republican participants, but is concentrated among those with the highest levels of conservatism. This finding concords with the idea that people trade-off personal costs, such as privacy concerns, against partisan self-interest. As more conservative Republicans gain more strongly from an electoral advantage for their party, they are more willing to accept violations of privacy if these violations provide their preferred party with a benefit in an election.

These results contribute to the findings of previous research examining motivations behind attitudes toward election laws. Previous work has shown that political party leaders are willing to use government regulation in ways that will increase the likelihood that they will get elected in the future (Alvarez

et al., 2011; Boix, 1999; Bol et al., 2019; Matakos and Xefteris, 2015). This behavior has been reported in the context of gerrymandering, voter ID laws or same-day registration laws for voting (Ansolabehere, 2009; Biggers and Hanmer, 2017; Chen and Rodden, 2013; McCarthy, 2019). To date, less is known about the way the public, as opposed to the political elite, forms their attitudes about electoral legislation (Biggers, 2019). While many scholars suspect that the public's strategic motivations resemble those of party elites, there is only scarce causal evidence to support this hypothesis (Alvarez et al., 2011; Stewart et al., 2016). For example, in the related domain of behavioral policy interventions, experimental evidence suggests that US adults hold more favorable opinions about policy interventions that are in line with their political beliefs (Tannenbaum et al., 2017). However, most studies cannot distinguish between when the public is pursuing strategic goals and when the public is simply following party leaders' cues (Biggers, 2019). Our findings support the idea that the broader public indeed pursues goals similar to those of party elites, favors regulation based on their partisan self-interest, and supports laws that contribute to the electoral success of their preferred party.

Our findings further add to an emerging body of literature that shows that some people are willing to make trade-offs between established democratic norms and partisan self-interest (Graham and Svolik, 2019; Kunst et al., 2019; Nyhan et al., 2020; Svolik, 2018, 2019). According to our results, participants holding the strongest policy views have the greatest reaction to the information that targeted political advertising benefits their party. This finding accords with previous findings that people are willing to accept the undermining of democratic principles if it benefits their policy goals. In our case, people's attitudes towards the regulation of targeted political advertising

are partially driven by the desire to set rules that benefit people's preferred party, even if they view targeted political advertising as harmful to societal norms. This behavior might be perceived as a threat to perceptions of the fairness of elections, which could then undermine peoples' support for a electoral system that relies on a shared understanding of democratic norms (Birch, 2010; Doherty and Wolak, 2012; Douglas, 2013; Levitsky and Ziblatt, 2018; Sunshine and Tyler, 2003; Welzel and Inglehart, 2007). We show that the rise of new technologies could potentially contribute to perceptions of "democratic backsliding" (Svolik, 2018), as people might be willing to use the newly-required rules for new technologies to pursue partisan self-interest.

We further show that beliefs about the impact that new technologies have on the electoral process are crucial to our understanding of public attitudes towards them. This finding contributes to a wider body of literature that investigates how potentially erroneous beliefs that people hold drive their behavior (Coibion et al., 2020; Esberg et al., 2020; Goldfayn-Frank and Wohlfart, 2019; Haaland et al., 2020; Lergetporer et al., 2018; Malmendier and Nagel, 2016; Roth and Wohlfart, 2019).

This study reveals that it is difficult to understand public preferences for certain policy measures without understanding the beliefs that people hold about key variables that are affected by these policies. Preferences for regulation of targeted political advertising are currently driven in part by third-person perceptions, leading to biased beliefs about their effect. This situation could lead to potentially sub-optimal policy decisions, as politicians might follow public preferences that are driven by biased beliefs. Our findings underscore the necessity of providing the public with truthful information about the effect of targeted political advertising. We show that support for

stricter regulation among Republicans would be significantly lower if they were correctly informed about the effect that it had on the 2016 Presidential election, because they underestimate the positive effect that targeted political advertising might have had or will have on their own party.

Previous research on the third-person effect found evidence for a gap between the perceived effect of persuasive mass communication on the self and on others (Davison, 1983; Perloff, 1993). Furthermore, correlational research supports the hypothesis that this gap motivates people in performing mitigating actions against the negative consequences of such persuasive communication (Xu and Gonzenbach, 2008). Our study adds to this literature in three ways. First, this study is the first to show that a perceptual gap exists in the context of targeted online political advertising. Second, this study is the first to establish a causal link between the perceptual gap described by the third-person effect and a behavioral measure for support for government regulation. By manipulating the perception gap of Republicans in our information treatment downward, and by showing that this decreases their support of the mitigating action, we were able to show causality between perception and behavior. Third, our results also add to previous studies reporting that the third-person perception increases with social distance, or between in-groups and out-groups (Jang and Kim, 2018; Perloff, 1999; White, 1997). To the best of our knowledge, this is the first study to show that the perceived gap between Democrats and Republicans in their perceptions of the influence of undesirable mass communication is strongly linked to affective as well as ideological polarization, and it is the first study to measure this outcome with an unincentivized and an incentivized measure.

Our results have some limitations. First, we were unable to show similar

causal results for Democratic supporters. We found a strong correlation between the beliefs that Democrats report about the effect that targeted political advertising has on Republicans and their support for stricter government regulation, but cannot claim causality for this group. Given that we needed to truthfully inform participants that we were not using deception in this study, we were unable to manipulate Democrats' beliefs in a way that was equivalent to that used with Republicans. Second, the main measure of interest, participants' support for stricter government regulation, indicates relatively low-scale reliability (Cronbach's $\alpha = 0.67$). In the Appendix Tables 3.9 and 3.10, we report exploratory results that show that a reduced scale (excl. the fourth item) has higher reliability (Cronbach's $\alpha = 0.75$) and that all of our main results are robust to the reduced scale.

This paper develops a new experimental paradigm to study people's attitudes towards technological change which has an influence on elections. We show that support for or opposition to the regulation of new technology that has implications for the political process is driven by potentially biased beliefs about how the use of this technology affects political outcomes for one's preferred party. Therefore, our findings add to a growing policy debate and underscore the necessity of making the effects of targeted political advertising transparent and of truthfully informing the public about the effects of the new technology so that the public can fully and knowledgeably realize their true attitudes. We believe that more research is necessary to fully understand the public's attitude towards these innovations, especially regarding beliefs about the spread and effect of false information and divisive messages. Further, our result indicating that people take into account the broader societal effects of targeted advertising might have implications for certain aspects of targeted commercial advertising. We would encourage future research to

investigate whether similar mechanisms would motivate people to oppose, for example, the use of targeted advertising to promote socially undesirable consumption, such as smoking, drinking or other unhealthy behavior.

3.A Overview

The appendix provides additional background information (part A), names and explains deviations from the pre-analysis plan (part B), provides an overview of the experimental design and instructions (part C) and presents additional results (part D).

3.B Additional text

3.B.1 Targeted political online advertising

Targeted online advertising refers to a kind of advertising where people's data is used to direct advertising content to them that maximizes the likelihood of them reacting to it (Chen and Stallaert, 2014). Thus, targeted advertisement is shown to people who most likely already have a preference for its content, which increases ad spending efficiency (Iyer et al., 2005). In case of political targeting, this means that parties show ads to voters who they are most likely to mobilize. Furthermore, political actors can customize their messages to the distinctive interests and concerns of these people, addressing issues that matter to them and employing language and form that appeals to them (Zuiderveen Borgesius et al., 2018). Data used for targeting can be demographic, or contain information about attitudes, interests, or personality traits, which individuals either revealed online or that can be deduced from their data (Boerman et al., 2017). This practice of tailoring advertisement based on personality variables is called "psychographic profiling" and often makes use of research on decision making to influence the

recipient's attitudes (Burkell and Regan, 2019).

While collecting voter data for advertising purposes is not new, the vast amount of personal data used for targeting and tailoring political ads is unprecedented and online behavioral data plays a major role in this (Rubinstein, 2014). With people generating more and more trackable information while surfing online, this data has gained significant importance for advertisers (Boerman et al., 2017). Online behavioral data can include search and purchase histories, websites visited, articles read or videos watched and what people communicate in e-mails and on social media (Boerman et al., 2017; Jansen et al., 2013). This data is gathered and then resold by so called "commercial data brokers", or companies whose business is "collecting personal information about consumers from a variety of sources and aggregating, analyzing, and sharing that information, or information derived from it, for purposes such as marketing" (FTC, 2014). These firms track online behavioral data by using tracking cookies, small text files that are stored on the user's computer after visiting a website. Tracking cookies that are placed on a website by another entity are also called third-party cookies, while cookies that originate from a website itself are called first-party cookies (Rubinstein, 2014). Third-party cookies can track users over several websites, with the data usually being used for advertising (Boerman et al., 2017). An individual website can have up to 350 built in cookies and the 100 most popular websites have more than 6000 cookies combined, more than 80 percent of them being third-party cookies (Altaweel et al., 2015). Tracking cookies can trace users' browsing history (including text entered or buttons clicked) and set-up long-term records of their online behavior (Altaweel et al., 2015). These personal user histories are then sold by the commercial data brokers to political strategists and are then combined with offline data from voter registration

databases, response data from door-to-door canvassing, telephone surveys or online behavioral data collected via candidates' websites. Subsequently, they are merged into detailed voter profiles (Rubinstein, 2014). Campaigners can then apply predictive modeling techniques to the data to make inferences on which users are most likely to vote for a party, and how to best address them (Burkell and Regan, 2019).

3.B.2 The role of privacy concerns in the regulation of targeted political advertising

A large body of literature has shown that individual privacy concerns are an important determinant for people's attitudes towards targeted advertising (Baek and Morimoto, 2012; Evans, 2009; Krasnova et al., 2009; Okazaki et al., 2009; Wang et al., 1998). The results of these studies also reveal that individual privacy attitudes are closely linked to individual preferences towards stricter privacy regulation (Milberg et al., 2000, 1995; Smith et al., 1996). This literature has identified several reasons for this connection. People's individual risk preferences can play a role in determining their stance on stricter regulation. If people perceive that the collection of their data can create large risks to them, for example due to data breaches, they are more likely to support stricter government regulation (Miltgen and Smith, 2015; Okazaki et al., 2009). Further, trust towards internet companies and the firms collecting data is an important factor in determining people's attitudes towards regulation (Edelman, 2011; Tang et al., 2008). If people have high levels of trust towards these companies they are less likely to support sweeping government regulation and believe that industry self-regulation is sufficient (Xu et al., 2011). Additionally, people might see their private data as valuable

commodity and demand stricter protection of their property right (Miltgen and Smith, 2015). They therefore demand stricter government laws to ensure that their property right is upheld (Acquisti et al., 2016; Jentzsch, 2003). Additionally, users might perceive targeted advertising as "creepy" when it relies on their private data too heavily (Moore et al., 2015; Tene and Polonetsky, 2013). Taken together, these different motives lead to a heightened sense of concern among people and can therefore motivate them to support stricter government regulation.

3.B.3 The third-person effect in communication

The third-person effect is broadly defined as people's perception that other people are influenced by undesirable mass communication to a larger extent than themselves (Davison, 1983). The emergence of the third-person effect is linked to a more general perception that people see biases and mistakes more strongly in others than in themselves (Pronin et al., 2002; Scopelliti et al., 2015). Previous research also suggests that people in general have too optimistic views of their own future outcomes and overestimate their abilities (Sharot, 2011).

Undesirableness of media content is defined as having either socially unacceptable content or intent (Gunther and Mundy, 1993) and the third-person effect is most likely to emerge when the topic of the media content is of personal importance and is perceived to be negatively biased against people's own interest (Perloff, 1993). Whether media content is perceived as undesirable can be dependent on societal factors and might be influenced by cultural traditions (Paul et al., 2000). Previous research has focused on media content that describes societal taboo topics like pornography (Lo and Wei,

2002), gambling (Youn et al., 2000), violence (McLeod et al., 1997), unhealthy behaviors (Henriksen and Flora, 1999) or, more recently, fake news (Jang and Kim, 2018) and has reliably found a third-person effect in that circumstances. The strength of the third-person effect also depends on the social distance to the "other" (Eveland et al., 1999). The more dissimilar, the more people think that that person will be influenced by undesirable media messages. Other work has focused on socially desirable topics (like ones promoting healthy behavior (Henriksen and Flora, 1999) or disaster preparedness (Atwood, 1994)) and have found the opposite effect, a first-person effect. A first-person effect is the perception that people themselves are more influenced by this sort of media content as compared to others. This is in line with people's belief that they themselves are more likely to make wise decisions and in general have higher abilities.

Previous research has further identified a second part of the third-person effect: The behavioral part. This is a direct consequence of people's perception of the media's effect on others. They adopt behaviors to rectify the consequences of undesirable media messages on others (Sun et al., 2008). The perception that others are strongly influenced by media messages can trigger different responses, for example adaptive behavior (Sun et al., 2008). This study focuses on behavior that is intended to rectify the consequences of undesirable media communication. This behavior involves the support for censorship of undesirable media content to shield more vulnerable groups from its influence (Rojas et al., 1996). Alternatively, it can also be driven by the intention to prevent negative influences for society, for example in the case of the promotion of unhealthy behavior (Shah et al., 1999). Rectifying behavior can also, in certain cases, lead to behavior that limits access of vulnerable groups to media (Hoffner and Buchanan, 2002) or increase support

for media literacy training (Jang and Kim, 2018).

3.C Deviations from the pre-analysis plan

The pre-analysis plan is available here: AEARCTR-0005296.

Data, full instructions, variable coding and analysis code are available at a public OSF-repository. We deviated from the pre-plan in minor ways. We excluded 41 responses on the survey because they were submitted from people that entered the survey multiple times. Because we did not expect the technical possibility that people can enter the survey several times, we did not consider that in our pre-analysis plan. None of our results is sensitive to the exclusion of these participants. We further had to exclude incomplete responses from some regressions. Our regression results show that the inclusion of control variables does not significantly affect our main parameters.

We further did not pre-register the regression we ran in table 3.1 of the main paper. This test does not deviate from the pre-registered analysis or hypothesis. Our main argument is built on the Welch t-test that is reported in the same section and that shows the same result as the regression. The regression we report in table 3.1 is pre-registered as secondary data analysis. The primary outcome is reported in the appendix in table 3.6. Both results are in line with the hypothesis that we pre-registered.

We further did not pre-register the comparison between privacy concerns in treatment and control condition explicitly. We pre-registered more generally that we will compare attitudes between treatment and control condition.

We further pre-registered some tests as one-sided tests. We decided to report two-sided results for all tests for ease of interpretation of the reader. Analyses that were not pre-registered are reported as exploratory data-analysis in the

main body of the paper.

3.D Instructions

Figure 3.3: Overview of the experimental design

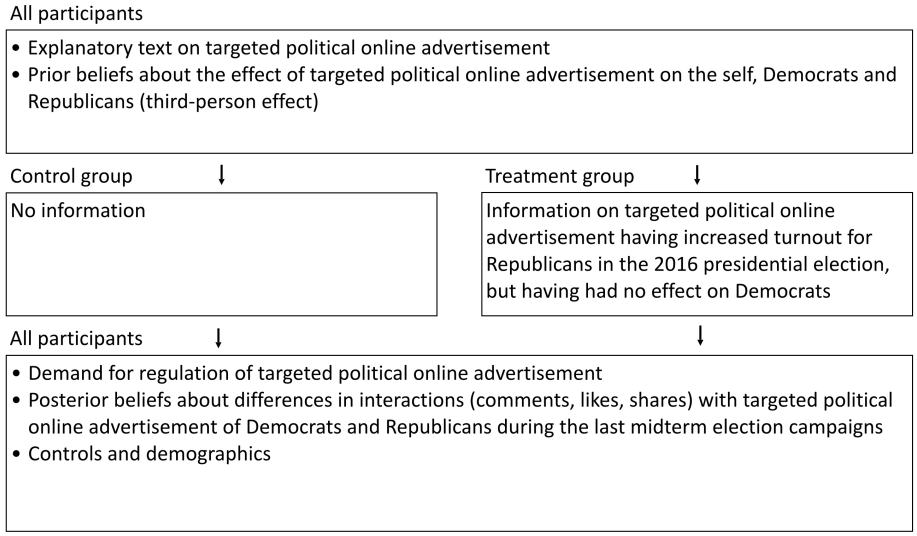

Information about targeted political advertising

Please read the following information carefully.

Targeted advertising is the practice of monitoring peoples' online behavior and using the collected information to show people individually targeted advertisements. Online behavioral data can include web browsing data, search histories, media consumption data (e.g., videos watched), app use data, purchases, click-through responses to ads, and communication content, such as what people post on social networking sites. This online data is often combined with demographic data like age, gender and location.

Political parties also use targeted advertising, for example before presidential elections. Targeted political advertisement involves creating messages targeted at narrow categories of voters based on data analysis gathered from individuals' demographic characteristics and their online behavior. This enables political campaigns to send very specific messages to certain groups of potential voters. These messages are selected to be the most appealing to this group. Political actors use targeted advertising, for example, to reach voters who are likely to vote for them with messages that will influence them.

I read and understood the information.

- ◯ Yes
- ◯ No

Question about the belief on the effect of targeted political advertising - Intro Screen

Both, Republicans and Democrats, use targeted political advertising on social media and other online platforms to show personalized messages to voters.

Imagine a political campaign in the run-up to a competitive election. Suppose that both parties, Republicans and Democrats, spend the same amount of money on social media and other online platforms to present targeted political advertisements to their likely voters. They use information on location, gender, education and political ideology to determine which messages they are displaying to these potential voters. Political ideology is inferred from peoples' online behavior like their "liked pages" on social media or their news consumption.

We now want to know what you believe about the effect that targeted political online advertising would have on potential voters.

Please answer the following questions to the best of your knowledge. This is important for the research project.

Question about the belief on the effect of targeted political advertising

To what extent do you think **targeted political online ads** have an influence on you?

- ◯ Not at all
- ◯ To a small extent
- ◯ To some extent
- ◯ To a great extent
- ◯ To a very great extent

To what extent do you think **targeted political online ads** have an influence on **Republicans?**

- ◯ Not at all
- ◯ To a small extent
- ◯ To some extent
- ◯ To a great extent
- ◯ To a very great extent

To what extent do you think **targeted political online ads** have an influence on **Democrats**?

- ◯ Not at all
- ◯ To a small extent
- ◯ To some extent
- ◯ To a great extent
- ◯ To a very great extent

Treatment-Information

A group of international researchers investigated the influence of targeted political ads on voters. For their analysis, they looked at how many ads people saw and related that to increases in votes for the Republican and the Democratic parties. The ads were targeted based on peoples' location, gender, education, age and political ideology.

The researchers found that the targeted political ads significantly increased the number of votes for the Republican party, but not for the Democratic party. Hence, targeted political ads influenced Republican voters, but did not influence Democratic voters.

Message to Congress

We now want to know your opinion on certain policy initiatives that can be launched by the US government.

We will inform the members of the US Congress with an anonymised summary of the attitudes towards these policy initiatives stated by all participants in the following questions. **There is no deception in this study. We will actually send a message with the results to US House Representatives and Senate members. However, your answers will remain anonymous. No one, not even the researchers involved in this study, will be able to match your answers to you.**

On the next page, you see a representation of what this message will look like.

Measure for support for regulation

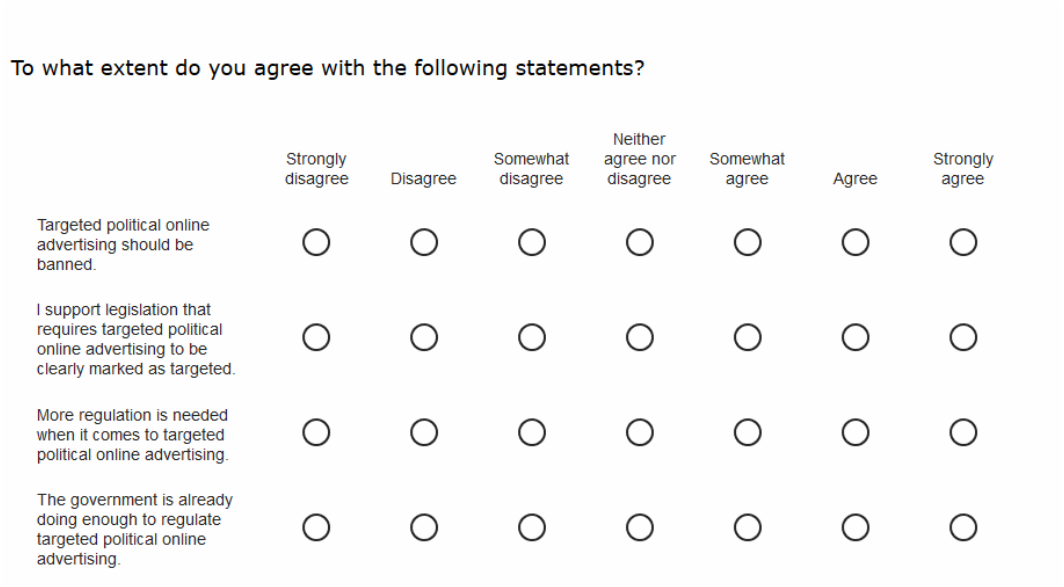

To what extent do you agree with the following statements?

	Strongly disagree	Disagree	Somewhat disagree	Neither agree nor disagree	Somewhat agree	Agree	Strongly agree
Targeted political online advertising should be banned.	○	○	○	○	○	○	○
I support legislation that requires targeted political online advertising to be clearly marked as targeted.	○	○	○	○	○	○	○
More regulation is needed when it comes to targeted political online advertising.	○	○	○	○	○	○	○
The government is already doing enough to regulate targeted political online advertising.	○	○	○	○	○	○	○

For participants in treatment: Confidence in Research findings

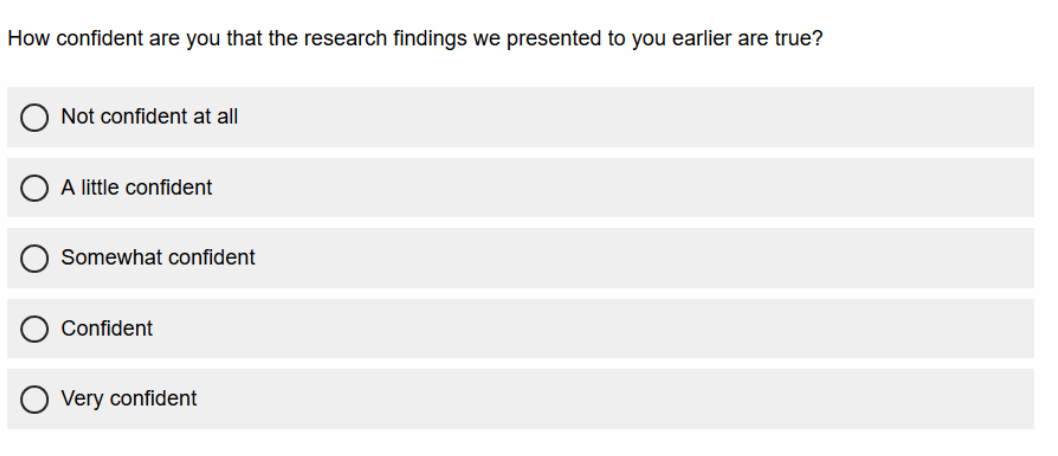

Question if Congress should take answers into account

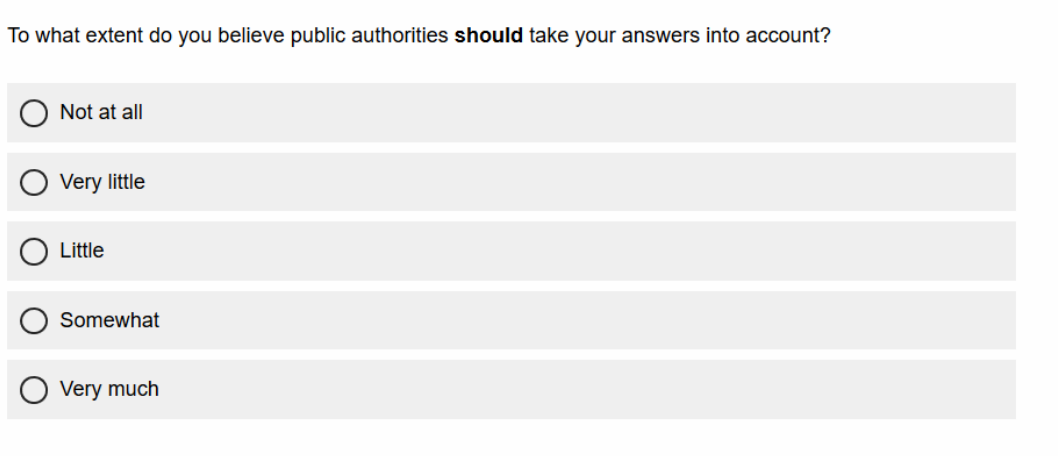

Manipulation check for change in beliefs

In this question, you can earn a small bonus based on the accuracy of your answer.

During the 2018 midterm elections, both major parties, Republicans and Democrats, used targeted advertising on social media platforms with the aim of influencing voters. One way to measure the success of this strategy is to see how much people interact (like, share, comment) with these advertising messages.

In the run-up to the 2018 elections, an international company measured interactions of people with Facebook posts of candidates for the House of Representatives. The measurement was taken between the 15. September and the 15. October 2018.

Please indicate now what you believe about the number of interactions that parties received relative to each other on these posts. If your answer is equal to the (rounded) correct answer, you will receive a bonus of $1.

- ◯ Democrats received four times the amount of Republicans
- ◯ Democrats received three times the amount of Republicans
- ◯ Democrats received twice the amount of Republicans
- ◯ Both received about the same amount
- ◯ Republicans received twice the amount of Democrats
- ◯ Republicans received three times the amount of Democrats
- ◯ Republicans received four times the amount of Democrats

Measurement of privacy concerns

To what extent do you agree with the following statements?
On the internet, I am concerned that my information:

	Strongly agree	Agree	Somewhat agree	Neither agree nor disagree	Somewhat disagree	Disagree	Strongly disagree
...can be collected and stored by third parties.	○	○	○	○	○	○	○
...can be shared with other third-parties (e.g. advertisers, employer, state).	○	○	○	○	○	○	○
Please choose "strongly agree" for this statement.	○	○	○	○	○	○	○
...can be used to display targeted advertising to me.	○	○	○	○	○	○	○
...can be used for commercial purposes (e.g. targeted advertising).	○	○	○	○	○	○	○

3.E Additional figures and results

Figure 3.4: Density plot of participants' support for stricter regulation of targeted political advertising, measured on four item, seven-point Likert Scale

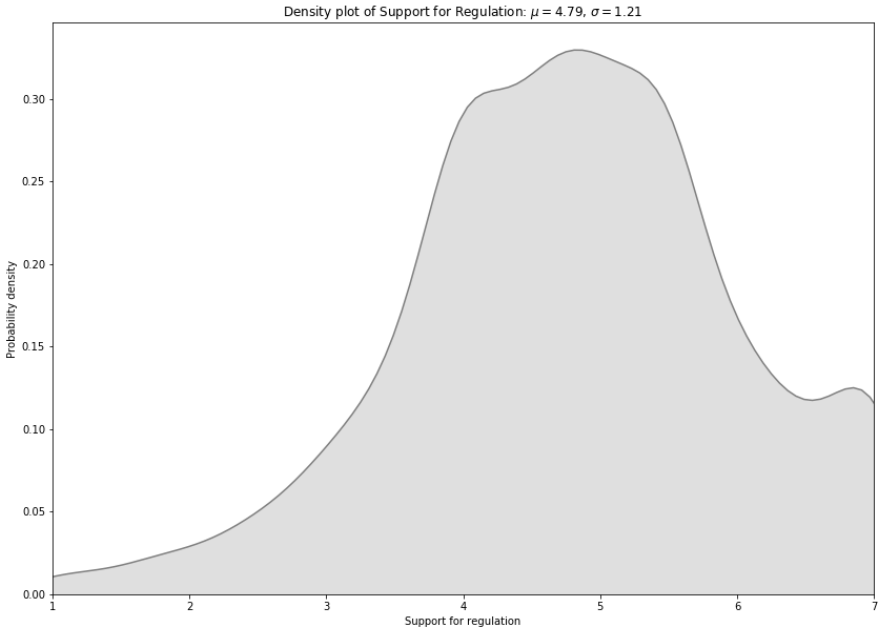

Note: Support for regulation was measured on a four item, seven-point Likert Scale. The plot shows the distribution of support for regulation for all participants

Figure 3.5: Density plot of participants' privacy concerns

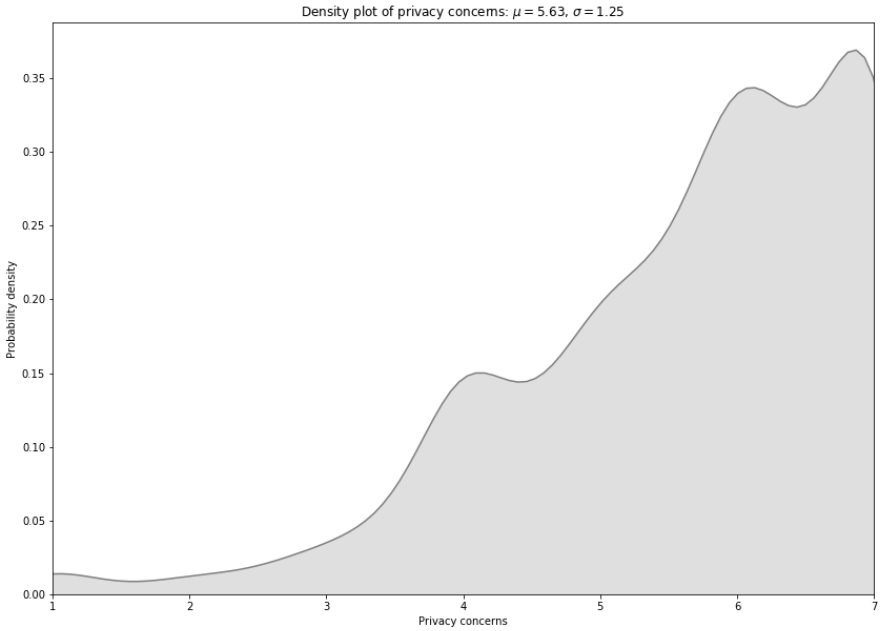

Note: Privacy concerns were measured on a four item, seven-point Likert Scale. The plot shows the distribution of privacy concerns for all participants

Figure 3.6: Relationship prior beliefs and manipulation check beliefs for participants in the control group

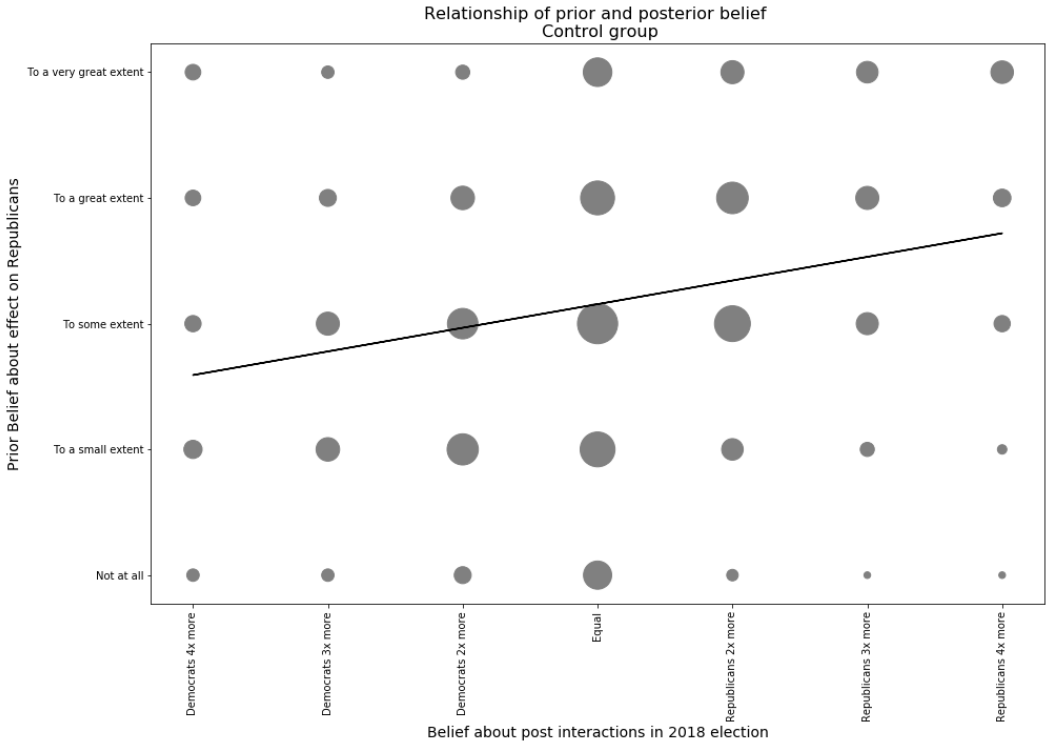

Note: The size of the circles indicate frequency of combinations. If the circle is bigger that indicates a higher frequency of both answers applying. The line shows a linear regression between the two measures.

Figure 3.7: Manipulation Check - Democrats

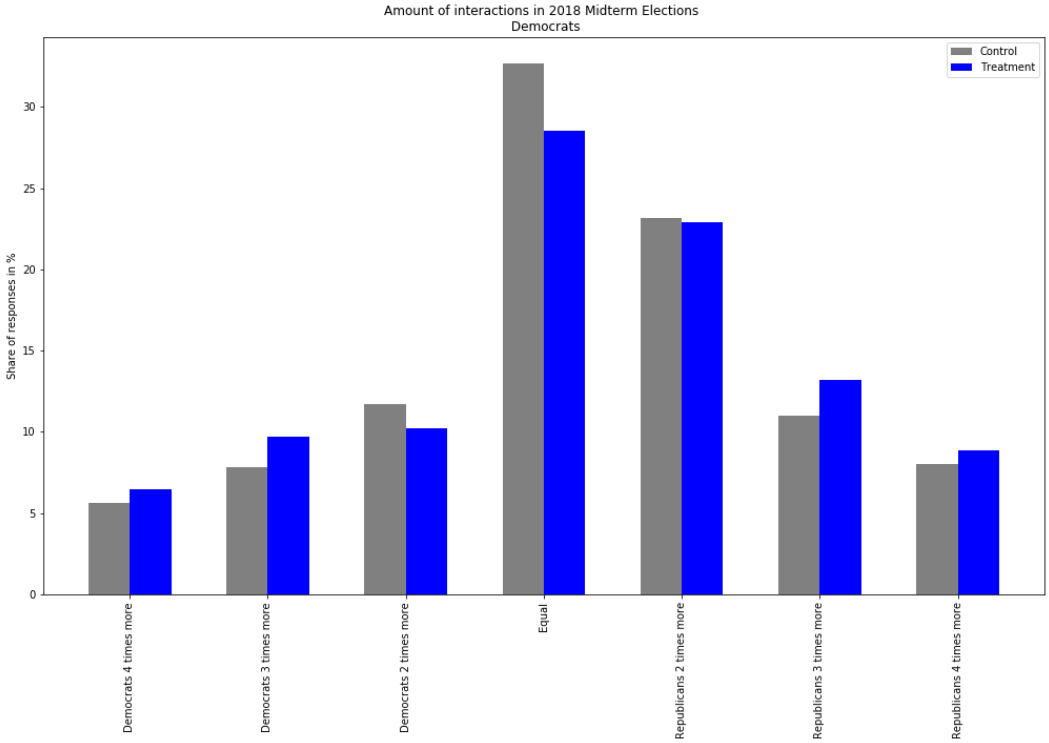

Note: The figure shows the responses of Democrats to the manipulation check for change in beliefs. The grey bars indicate responses for participants in the control condition. The blue bars indicate responses in the treatment condition.

Figure 3.8: Manipulation Check Republicans

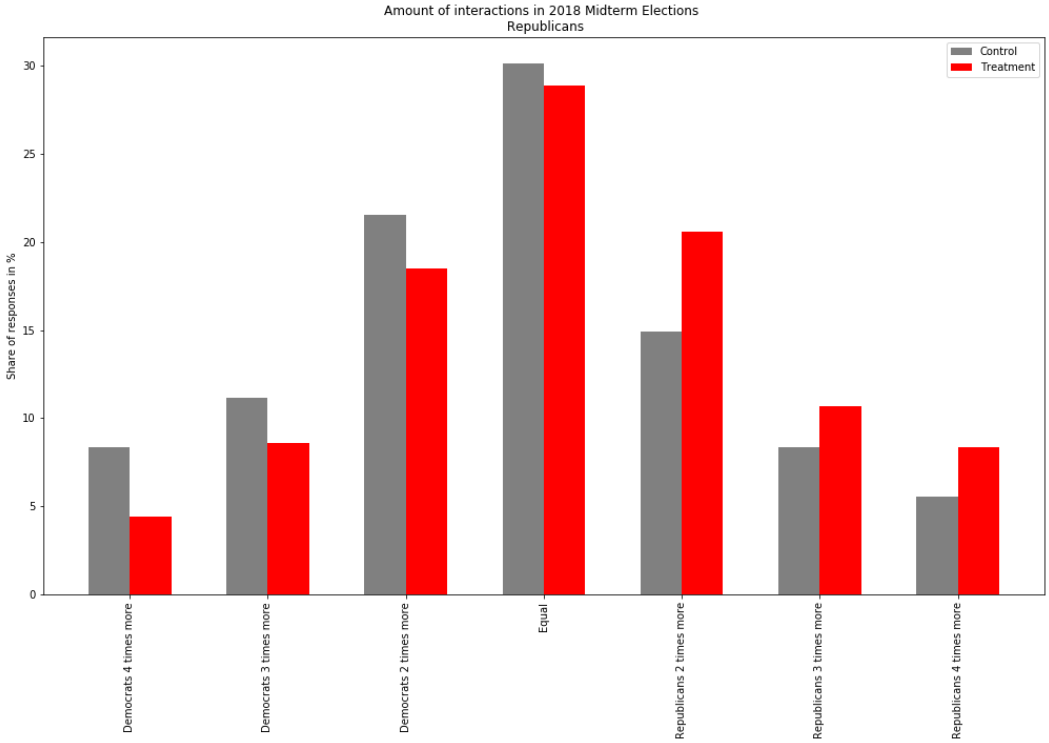

Note: The figure shows the responses of Republicans to the manipulation check for change in beliefs. The grey bars indicate responses for participants in the control condition. The red bars indicate responses in the treatment condition.

Figure 3.9: Treatment effect - Democrat

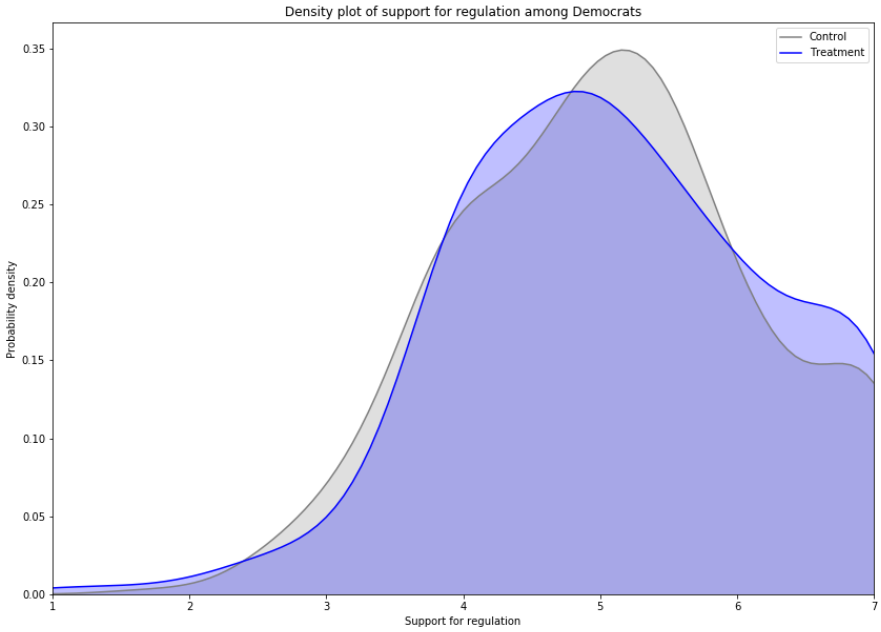

Note: The figure displays support for stricter government regulation among Democrats. Support for regulation was measured on a four item, seven-point Likert Scale. The blue area indicates the distribution for Democrats in the treatment condition. The grey area displays the distribution for Democrats in the control condition.

Figure 3.10: Treatment effect - Republicans

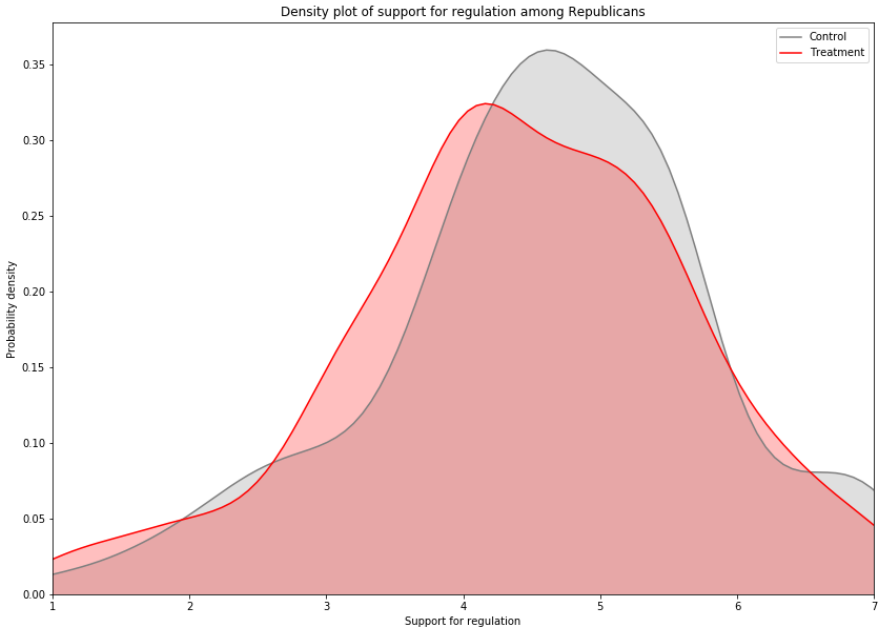

Note: The figure displays support for stricter government regulation among Republicans. Support for regulation was measured on a four item, seven-point Likert Scale. The red area indicates the distribution for Republicans in the treatment condition. The grey area displays the distribution for Republicans in the control condition.

Figure 3.11: Treatment effect - Republicans below median conservative

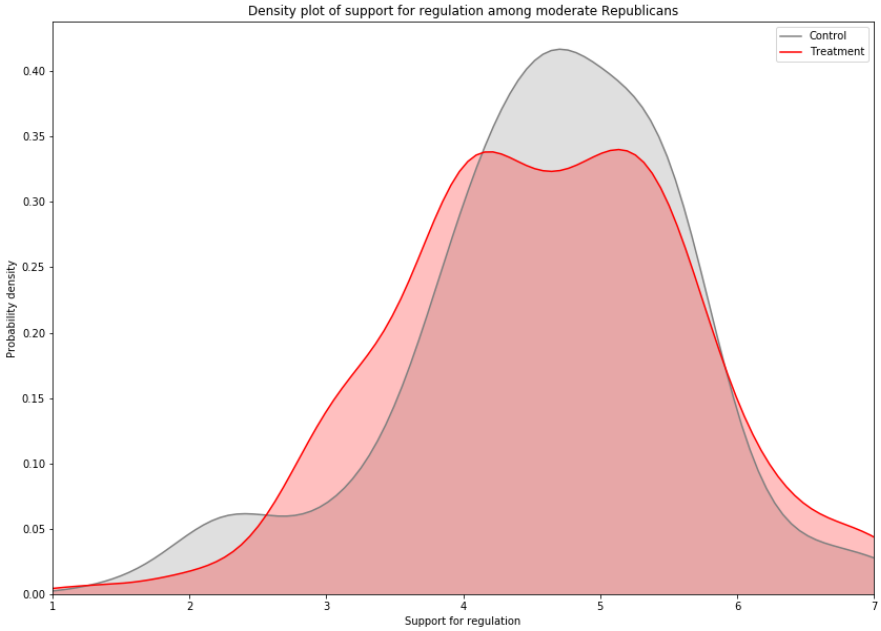

Note: The figure displays support for stricter government regulation among moderate Republicans. Support for regulation was measured on a four item, seven-point Likert Scale. Political ideology was measured on a scale for economic and social conservatism. Moderate Republicans were defined as being below median in this measure. The red area indicates the distribution for moderate Republicans in the treatment condition. The grey area displays the distribution for moderate Republicans in the control condition.

Figure 3.12: Treatment effect - Republicans at or above median conservative

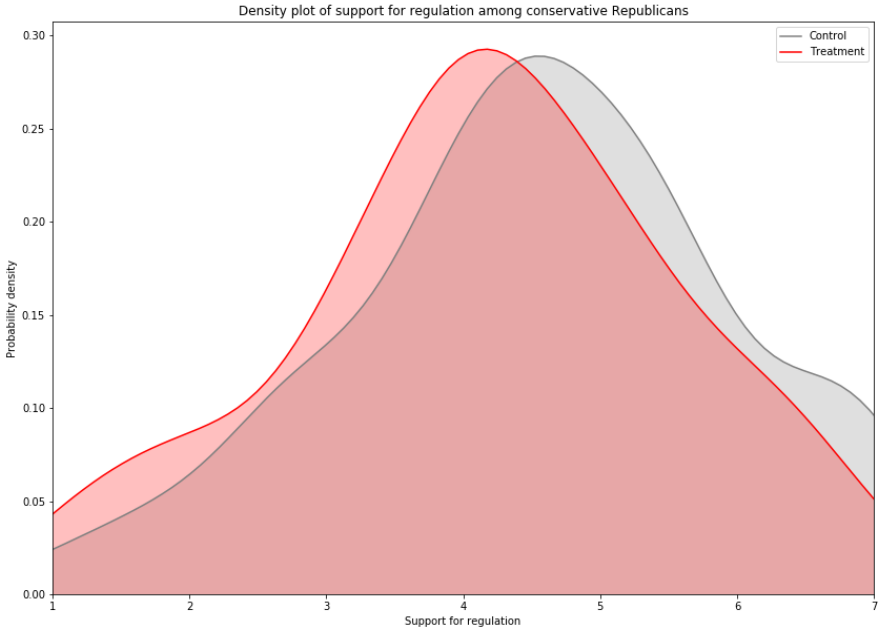

Note: The figure displays support for stricter government regulation among moderate Republicans. Support for regulation was measured on a four item, seven-point Likert Scale. Political ideology was measured on a scale for economic and social conservatism. Conservative Republicans were defined as being at or above median in this measure. The red area indicates the distribution for conservative Republicans in the treatment condition. The grey area displays the distribution for conservative Republicans in the control condition.

Table 3.3: Descriptive Statistics

	Amount of participants	Share
Income lower than $30,000	289	0.19
Income between $30,000 and $60,000	426	0.28
Income between $60,000 and $100,000	392	0.25
Income between $100,000 and $140,000	174	0.11
Income higher than $140,000	161	0.10
Less than high school	19	0.01
High school/GED	262	0.17
Some college	322	0.21
Associate degree	185	0.12
Bachelor degree	406	0.26
Postgraduate degree	287	0.19
Non-white participants	388	0.25
User of social media	1359	0.88
Observations	1549	

Note: The table provides an overview of demographic variables for all participants in the experiment.

Table 3.4: Descriptive Statistics

	Baseline	Treatment
Income lower than $30,000	0.18	0.19
Income between $30,000 and $60,000	0.28	0.27
Income between $60,000 and $100,000	0.26	0.24
Income between $100,000 and $140,000	0.11	0.11
Income higher than $140,000	0.10	0.11
Less than high school	0.01	0.01
High school/GED	0.16	0.18
Some college	0.23	0.19
Associate degree	0.11	0.13
Bachelor degree	0.27	0.25
Postgraduate degree	0.18	0.19
Non-white participants	0.25	0.25
Age	47.39	47.60
Belief about effect on self	2.38	2.41
Belief about the effect on Democrats	3.07	3.08
Belief about the effect on Republicans	3.17	3.07
Privacy Concerns	5.60	5.67
User of social media	0.89	0.87
Observations	794	755

Note: The table provides an overview of demographic variables for participants split into treatment and control condition in the experiment.

Table 3.5: Regression of correlates for the size of the difference between other party and own party

	Difference other party - own party
Affective polarization	0.170***
	(0.032)
Ideological polarization	0.112***
	(0.031)
Desirability of advertising	-0.149***
	(0.017)
High political knowledge	0.133*
	(0.065)
High income	-0.087
	(0.065)
Education	0.076
	(0.064)
Male	-0.064
	(0.062)
Non-white participants	-0.091
	(0.074)
Age	0.004
	(0.002)
Household size	-0.009
	(0.065)
Use of internet in hours	0.003
	(0.005)
Use of ad-block	0.018
	(0.029)
User of social media	0.000
	(0.085)
Attitude towards government regulation	-0.013
	(0.016)
External efficacy	-0.002
	(0.001)
Politically active	0.078
	(0.060)
Observations	1464
R^2	0.148

Note: The table reports the results for an OLS-regression with the difference between people's belief about the effect of targeted political advertising on the other - the own party as dependent variable. The dependent variable is standardized. Affective and Ideological polarization are standardized. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.6: Regression of determinants for the willingness to support stricter regulation of targeted advertising

	Support for regulation	Support for regulation
Belief about own party	-0.043	-0.010
	(0.038)	(0.048)
Belief about other party	0.169***	0.172***
	(0.039)	(0.039)
Belief about self	X	-0.019
		(0.046)
Privacy concerns	0.260***	0.239***
	(0.044)	(0.045)
Observations	754	754
R^2	0.115	0.135
Demographics	No	Yes
social media use	No	Yes
Political Engagement	No	Yes

Note: Regressions only include participants that answered all questions of the survey. The table reports results from an OLS-regression in which people's support for regulation is the dependent variable. The value is standardized. Belief other party is measured as participants' belief about the effect that targeted political advertisement has on other supporters of the other party. Belief own party is defined as the belief that participants have about the effect targeted political advertising has on supporters of their own party. Belief about self is people's belief about the effect that targeted political advertising has on them. Privacy concerns are respondents' are measured on a seven-point 4 item-Likert scale. All three independent variables are standardized. Demographic information included age, education (dummy for above median in the sample), income (dummy for above median in the sample), household size (dummy for more than two members), gender (male dummy variable) and a dummy for being non-white. Social media use was a dummy variable for the use of social media, a continuous variable for the time people spent online in general (in hours), and the use of an ad-blocker (dummy for yes). Political engagement was a dummy variable for being politically active within the last year, external political efficacy, political knowledge (dummy for above median knowledge), and attitudes towards government regulation in general. Table 3.11 in the appendix provides an overview of all variables in the regression. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.7: Effect of the information treatment on Republicans - Only Republicans that want that Congress takes their answers into account

	Baseline	Treatment	Difference
Observations	388	381	
Demand for regulation	4.58	4.41	-0.17*

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Demand for regulation measured on a four item 7 point Likert scale. p-value = 0.03 for one-sided t-test

Note: The table summarizes a Welsh-t-test that is run to compare Republicans' support for regulation. Support for regulation was measured on a four item 7 point Likert scale. Only Republicans who want their answers to be taken into account by Congress are included. The p-value of that test was 0.03 for a one-sided t-test. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.8: Effect of the information treatment on Republicans - Only Republicans that believe treatment information

	Baseline	Treatment	Difference
Observations	391	315	
Support for regulation	4.59	4.47	-0.12

Demand for regulation measured on a four item 7 point Likert scale. p-value = 0.10 for one-sided t-test

Note: The table summarizes a Welsh-t-test that is run to compare Republicans' support for regulation. Support for regulation was measured on a four item 7 point Likert scale. Only Republicans who trust the presented research finding were included. The p-value of that test was 0.10 for a one-sided t-test. $^{*}\ p < 0.05$, $^{**}\ p < 0.01$, $^{***}\ p < 0.001$

Table 3.9: Regression of determinants for the willingness to support stricter regulation of targeted advertising - Alternative regulation measure

	Support for Regulation	Support for Regulation	Support for Regulation	Support for Regulation
Belief other party - own party	0.066	0.112**	0.112**	0.107**
	(0.034)	(0.035)	(0.035)	(0.036)
Privacy concerns	0.305***	0.281***	0.274***	0.262***
	(0.042)	(0.042)	(0.042)	(0.043)
Belief about self	X	0.153***	0.145***	0.116**
		(0.038)	(0.038)	(0.039)
Observations	754	754	754	754
R^2	0.105	0.138	0.144	0.161
Demographics	No	Yes	Yes	Yes
social media use	No	No	Yes	Yes
Political Engagement	No	No	No	Yes

Note: Regressions only include participants that answered all questions of the survey. The table reports results from an OLS-regression in which people's support for regulation is the dependent variable. The value is standardized. Belief other party-own party is defined as the difference between people's belief about the effect on the other party and the effect on the own party. Belief about self is people's belief about the effect that targeted political advertising has on them. Privacy concerns are respondents' are measured on a seven-point 4 item-Likert scale. All three independent variables are standardized. Demographic information included age, education (dummy for above median in the sample), income (dummy for above median in the sample), household size (dummy for more than two members), gender (male dummy variable) and a dummy for being non-white. Social media use was a dummy variable for the use of social media, a continuous variable for the time people spent online in general (in hours), and the use of an ad-blocker (dummy for yes). Political engagement was a dummy variable for being politically active within the last year, external political efficacy, political knowledge (dummy for above median knowledge), and attitudes towards government regulation in general. Table 3.11 in the appendix provides an overview of all variables in the regression. Robust standard errors in parentheses. $^{*}\ p < 0.05$, $^{**}\ p < 0.01$, $^{***}\ p < 0.001$

Table 3.10: Regression of Determinants for the willingness to regulate targeted ads - Alternative regulation measure

	Support for Regulation	Support for Regulation
Treatment*Republican	-0.241*	-0.241*
	(0.098)	(0.097)
Treatment	0.081	0.081
	(0.063)	(0.063)
Privacy concerns	0.307***	0.293***
	(0.030)	(0.030)
Republican	-0.188**	-0.229***
	(0.068)	(0.069)
Observations	1466	1466
R^2	0.124	0.158
Demographics	No	Yes
social media use	No	Yes
Political engagement	No	Yes

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: Regressions only include participants that answered all questions of the survey. The table reports the results of an OLS-regression with the support for regulation as a dependent variable. The variable was standardized. Treatment is a dummy variable that is 1 when the participant was assigned to the treatment condition and 0 otherwise. Republican is a dummy variable that is 1 when the participant is a Republican and 0 if he or she is a Democrat. Treatment*Republican is the interaction of these two variables. Privacy concerns is a measure of people's privacy concerns. The value is standardized. Demographic information included age, education (dummy for above median in the sample), income (dummy for above median in the sample), household size (dummy for more than two members), gender (male dummy variable) and a dummy for being non-white. Social media use was a dummy variable for the use of social media, a continuous variable for the time people spent online in general (in hours), and the use of an ad-blocker (dummy for yes). Political engagement was a dummy variable for being politically active within the last year, external political efficacy, political knowledge (dummy for above median knowledge), and attitudes towards government regulation in general. Table 3.12 in the appendix provides an overview of all variables in the regression. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3.11: Regression of determinants for the willingness to support stricter regulation of targeted advertising - All controls

	Support for regulation
Belief other party - own party	0.124***
	(0.035)
Privacy concerns	0.257***
	(0.045)
Belief about self	0.052
	(0.039)
High income	0.049
	(0.073)
Education	-0.058
	(0.075)
Male	-0.112
	(0.070)
Non-white participants	-0.142
	(0.076)
Age	0.003
	(0.002)
Household size	0.063
	(0.073)
Use of internet in hours	-0.002
	(0.006)
Use of ad-block	-0.062
	(0.035)
User of social media	0.149
	(0.128)
Attitude towards government regulation	0.017
	(0.018)
High political knowledge	0.017
	(0.079)
External efficacy	-0.001
	(0.001)
Politically active	-0.054
	(0.085)
Constant	-0.132
	(0.235)
Observations	754
R^2	0.125

Note: The table reports results for the OLS-regression with the dependent variable Support for regulation. The dependent variable was standardized. Belief other party-own party is defined as the difference between people's belief about the effect on the other party and the effect on the own party. Belief about self is people's belief about the effect that targeted political advertising has on them. Privacy concerns are respondents' are measured on a seven-point 4 item-Likert scale. All three independent variables are standardized. Robust standard errors in parentheses. $^{*}\ p < 0.05$, $^{**}\ p < 0.01$, $^{***}\ p < 0.001$

Table 3.12: Regression of determinants for the willingness to regulate targeted ads - all controls

	Support for Regulation
Treatment*Republican	-0.192*
	(0.096)
Treatment	0.032
	(0.064)
Privacy concerns	0.286***
	(0.030)
Republican	-0.411***
	(0.068)
High income	0.072
	(0.053)
Education	-0.080
	(0.053)
Male	-0.084
	(0.051)
Age	0.001
	(0.002)
Household size	-0.027
	(0.053)
Non-white participants	-0.200***
	(0.058)
User of social media	0.047
	(0.085)
Use of internet in hours	0.001
	(0.004)
Use of ad-block	-0.037
	(0.024)
External efficacy	-0.001
	(0.001)
Politically active	-0.100
	(0.059)
Attitude towards government regulation	0.019
	(0.013)
High political knowledge	0.059
	(0.056)
Constant	0.239
	(0.174)
Observations	1466
R^2	0.166

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note: The table reports results for the OLS-regression with the dependent variable Support for regulation. The dependent variable was standardized. Treatment is a dummy variable that is 1 when the participant was assigned to the treatment condition and 0 otherwise. Republican is a dummy variable that is 1 when the participant is a Republican and 0 if he or she is a Democrat. Treatment*Republican is the interaction of these two variables. Privacy concerns is a measure of people's privacy concerns. The value is standardized. Robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

References

Abeler, Johannes, Daniele Nosenzo, and Collin Raymond, "Preferences for truth-telling," *Econometrica*, 2019, *87* (4), 1115–1153.

Acquisti, Alessandro and Christina Fong, "An experiment in hiring discrimination via online social networks," *Management Science*, 2020, *66* (3), 1005–1024.

_ **, Curtis Taylor, and Liad Wagman**, "The economics of privacy," *Journal of Economic Literature*, 2016, *54* (2), 442–92.

Agarwal, Sumit, Changcheng Song, and Vincent W Yao, "Banking competition and shrouded attributes: evidence from the US mortgage market," *Georgetown McDonough School of Business Research Paper*, 2017, (2900287).

_ **, Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel**, "Regulating consumer financial products: Evidence from credit cards," *The Quarterly Journal of Economics*, 2015, *130* (1), 111–164.

Ahler, Douglas J and Gaurav Sood, "The parties in our heads: Misperceptions about party composition and their consequences," *The Journal of Politics*, 2018, *80* (3), 964–981.

Akerlof, George A, "Sins of Omission and the Practice of Economics," *Journal of Economic Literature*, 2020, *58* (2), 405–18.

_ **and Robert J Shiller**, *Phishing for phools: The economics of manipulation and deception*, Princeton University Press, 2015.

Alesina, Alberto and Paola Giuliano, "Preferences for redistribution," in "Handbook of Social Economics," Vol. 1, Elsevier, 2011, pp. 93–131.

_ **, Stefanie Stantcheva, and Edoardo Teso**, "Intergenerational mobility and preferences for redistribution," *American Economic Review*, 2018, *108* (2), 521–54.

Allcott, Hunt and Judd B. Kessler, "The Welfare Effects of Nudges: A Case Study of Energy Use Social Comparisons," *American Economic Journal: Applied Economics*, jan 2019, *11* (1), 236–276.

_ **and Matthew Gentzkow**, "Social media and fake news in the 2016 election," *Journal of Economic Perspectives*, 2017, *31* (2), 211–36.

_ **, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow**, "The welfare effects of social media," *American Economic Review*, 2020, *110* (3), 629–76.

Almås, Ingvild, Alexander W Cappelen, and Bertil Tungodden, "Cutthroat capitalism versus cuddly socialism: Are Americans more meritocratic and efficiency-seeking than Scandinavians?," *Journal of Political Economy*, 2020, *128* (5), 1753–1788.

_ **,** _ **, Erik Ø Sørensen, and Bertil Tungodden**, "Fairness and the development of inequality acceptance," *Science*, 2010, *328* (5982), 1176–1178.

Altaweel, Ibrahim, Nathan Good, and Chris Jay Hoofnagle, "Web privacy census," *Technology Science*, 2015.

Altmann, Steffen, Armin Falk, Simon Jäger, and Florian Zimmermann, "Learning about job search: A field experiment with job seekers in Germany," *Journal of Public Economics*, aug 2018, *164*, 33–49.

Alvarez, R Michael, Thad E Hall, Ines Levin, and Charles Stewart III, "Voter opinions about election reform: Do they support making voting more convenient?," *Election Law Journal*, 2011, *10* (2), 73–87.

Ambuehl, Sandro, B Douglas Bernheim, and Axel Ockenfels, "Projective paternalism," Technical Report, National Bureau of Economic Research 2019.

Andrade, Philippe, Gaetano Gaballo, Eric Mengus, and Benoit Mojon, "Forward guidance and heterogeneous beliefs," *American Economic Journal: Macroeconomics*, 2019, *11* (3), 1–29.

Andreoni, James and John Miller, "Giving according to GARP: An experimental test of the consistency of preferences for altruism," *Econometrica*, 2002, *70* (2), 737–753.

_ , Deniz Aydin, Blake Barton, B Douglas Bernheim, and Jeffrey Naecker, "When Fair Isn't Fair: Understanding Choice Reversals Involving Social Preferences," *Journal of Political Economy*, 2020, *128* (5), 1673–1711.

Angrist, Joshua and Jörn-Steffen Pischke, "The credibility revolution in empirical economics: How better research design is taking the con out of econometrics," *Journal of Economic Perspectives*, 2010, *24* (2), 3–30.

_ , Pierre Azoulay, Glenn Ellison, Ryan Hill, and Susan Feng Lu, "Economic research evolves: Fields and styles," *American Economic Review*, 2017, *107* (5), 293–97.

_ , _ , _ , _ , and _ , "Inside job or deep impact? Extramural citations and the influence of economic scholarship," *Journal of Economic Literature*, 2020, *58* (1), 3–52.

Ansolabehere, Stephen, "Effects of identification requirements on voting: Evidence from the experiences of voters on Election Day," *PS: Political Science & Politics*, 2009, *42* (1), 127–130.

Aral, Sinan and Dean Eckles, "Protecting elections from social media manipulation," *Science*, 2019, *365* (6456), 858–861.

Arneson, Richard J, "Mill versus paternalism," *Ethics*, 1980, *90* (4), 470–489.

Atwood, L Erwin, "Illusions of media power: The third-person effect," *Journalism Quarterly*, 1994, *71* (2), 269–281.

Baek, Tae Hyun and Mariko Morimoto, "Stay away from me," *Journal of Advertising*, 2012, *41* (1), 59–76.

Bao, Jiayi and Benjamin Ho, "Heterogeneous Effects of Informational Nudges on Pro-social Behavior," *The B.E. Journal of Economic Analysis & Policy*, jan 2015, *0* (0).

Bar-Gill, Oren and Elizabeth Warren, "Making Credit Safer, 157 U," *Pa. L. Rev*, 2008, *1* (11).

_ **and Rebecca Stone**, "Pricing misperceptions: explaining pricing structure in the cell phone service market," *Journal of Empirical Legal Studies*, 2012, *9* (3), 430–456.

_ **, David Schkade, and Cass R Sunstein**, "Drawing false inferences from mandated disclosures," *Behavioural Public Policy*, 2019, *3* (2), 209–227.

Barberis, Nicholas, "Richard Thaler and the rise of behavioral economics," *The Scandinavian Journal of Economics*, 2018, *120* (3), 661–684.

Bartling, Björn, Ernst Fehr, and Holger Herz, "The intrinsic value of decision rights," *Econometrica*, 2014, *82* (6), 2005–2039.

_ , **Roberto A Weber, and Lan Yao**, "Do markets erode social responsibility?," *The Quarterly Journal of Economics*, 2015, *130* (1), 219–266.

Baum, Katharina, Stefan Meissner, Olga Abramova, and Hanna Krasnova, "DO THEY REALLY CARE ABOUT TARGETED POLITICAL ADS? INVESTIGATION OF USER PRIVACY CONCERNS AND PREFERENCES," in "Proceedings of the 27th European Conference on Information Systems (ECIS)" 2019.

Bellman, Steven, Eric J Johnson, Stephen J Kobrin, and Gerald L Lohse, "International differences in information privacy concerns: A global survey of consumers," *The Information Society*, 2004, *20* (5), 313–324.

Ben-Shahar, Omri and Carl E Schneider, "THE FAILURE OF MANDATED DISCLOSURE," *University of Pennsylvania Law Review*, 2011, pp. 647–749.

Bénabou, Roland and Jean Tirole, "Mindful economics: The production, consumption, and value of beliefs," *Journal of Economic Perspectives*, 2016, *30* (3), 141–64.

Benartzi, Shlomo, John Beshears, Katherine L Milkman, Cass R Sunstein, Richard Thaler, Maya Shankar, Will Tucker-Ray, William J Congdon, and Steven Galing, "Should governments invest more in nudging?," *Psychological Science*, 2017, *28* (8), 1041–1055.

Benjamin, Daniel J, "Errors in probabilistic reasoning and judgment biases," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 2, Elsevier, 2019, pp. 69–186.

Benkler, Yochai, Robert Faris, and Hal Roberts, *Network propaganda: Manipulation, disinformation, and radicalization in American politics*, Oxford University Press, 2018.

Bennear, Lori S and Sheila M Olmstead, "The impacts of the "right to know": Information disclosure and the violation of drinking water standards," *Journal of Environmental Economics and Management*, 2008, *56* (2), 117–130.

Berliner, Daniel, "The political origins of transparency," *The journal of Politics*, 2014, *76* (2), 479–491.

Beyersdorf, Brian, "Regulating the Most Accessible Marketplace of Ideas in History: Disclosure Requirements in Online Political Advertisements after the 2016 Election," *Calif. L. Rev.*, 2019, *107*, 1061–1100.

Bhargava, Saurabh, George Loewenstein, and Justin Sydnor, "Choose to lose: Health plan choices from a menu with dominated option," *The Quarterly Journal of Economics*, 2017, *132* (3), 1319–1372.

Bhattacharjee, Amit, Jason Dana, and Jonathan Baron, "Anti-profit beliefs: How people neglect the societal benefits of profit.," *Journal of Personality and Social Psychology*, 2017, *113* (5), 671.

Biggers, Daniel R, "Does Partisan Self-interest Dictate Support for Election Reform? Experimental Evidence on the Willingness of Citizens to Alter the Costs of Voting for Electoral Gain," *Political Behavior*, 2019, *41* (4), 1025–1046.

_ **and Michael J Hanmer**, "Understanding the adoption of voter identification laws in the American states," *American Politics Research*, 2017, *45* (4), 560–588.

Birch, Sarah, "Perceptions of electoral fairness and voter turnout," *Comparative political studies*, 2010, *43* (12), 1601–1622.

Bobadilla-Suarez, Sebastian, Cass R Sunstein, and Tali Sharot, "The intrinsic value of choice: The propensity to under-delegate in the face of potential gains and losses," *Journal of Risk and Uncertainty*, 2017, *54* (3), 187–202.

Boerman, Sophie C, Sanne Kruikemeier, and Frederik J Zuiderveen Borgesius, "Online behavioral advertising: A literature review and research agenda," *Journal of Advertising*, 2017, *46* (3), 363–376.

Boix, Carles, "Setting the Rules of the Game: The Choice of Electoral Systems in Advanced Democracies," *American Political Science Review*, 1999, *93* (3), 609–624.

Bok, Sissela, *Secrets: On the ethics of concealment and revelation*, Vintage, 1989.

Bol, Damien, Konstantinos Matakos, Orestis Troumpounis, and Dimitrios Xefteris, "Electoral rules, strategic entry and polarization," *Journal of Public Economics*, 2019, *178.*

Bolton, Gary E. and Axel Ockenfels, "A Theory of Equity, Reciprocity, and Competition," *American Economic Review*, 2000, *90* (1), 166–193.

Bolton, Gary E, Jordi Brandts, and Axel Ockenfels, "Fair procedures: Evidence from games involving lotteries," *The Economic Journal*, 2005, *115* (506), 1054–1076.

Bolton, Lisa E, Luk Warlop, and Joseph W Alba, "Consumer perceptions of price (un) fairness," *Journal of Consumer Research*, 2003, *29* (4), 474–491.

Bolton, Patrick, Xavier Freixas, and Joel Shapiro, "Conflicts of interest, information provision, and competition in the financial services industry," *Journal of Financial Economics*, 2007, *85* (2), 297–330.

Bond, Robert M, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler, "A 61-million-person experiment in social influence and political mobilization," *Nature*, 2012, *489* (7415), 295–298.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer, "Stereotypes," *The Quarterly Journal of Economics*, 2016, *131* (4), 1753–1794.

Borgesius, Frederik Zuiderveen, Judith Möller, Sanne Kruikemeier, Ronan Ó Fathaigh, Kristina Irion, Tom Dobber, Balazs Bodo, and Claes H de Vreese, "Online political microtargeting: promises and threats for democracy," *Utrecht Law Review*, 2018, *14* (1), 82–96.

Bowler, Shaun and Todd Donovan, "Democracy, institutions and attitudes about citizen influence on government," *British Journal of Political Science*, 2002, *32* (2), 371–390.

Brock, J Michelle, Andreas Lange, and Erkut Y Ozbay, "Dictating the risk: Experimental evidence on giving in risky environments," *American Economic Review*, 2013, *103* (1), 415–37.

Brooking, Emerson T and PW Singer, "War goes viral," *The Atlantic*, 2016, *318* (4), 70–83.

Brownell, Kelly D, Rogan Kersh, David S Ludwig, Robert C Post, Rebecca M Puhl, Marlene B Schwartz, and Walter C Willett, "Personal responsibility and obesity: a constructive approach to a controversial issue," *Health affairs*, 2010, *29* (3), 379–387.

Bullock, John G, Alan S Gerber, Seth J Hill, and Gregory A Huber, "Partisan Bias in Factual Beliefs about Politics," *Quarterly Journal of Political Science*, 2015, *10*, 519–578.

Burkell, Jacquelyn and Priscilla M Regan, "Voter preferences, voter manipulation, voter analytics: policy options for less surveillance and more autonomy," *Internet Policy Review*, 2019, *8* (4), 1–24.

Camerer, Colin F and Robin M Hogarth, "The effects of financial incentives in experiments: A review and capital-labor-production framework," *Journal of Risk and Uncertainty*, 1999, *19* (1-3), 7–42.

Camerer, Colin, Samuel Issacharoff, George Loewenstein, Ted O'donoghue, and Matthew Rabin, "Regulation for Conservatives: Behavioral Economics and the Case for" Asymmetric Paternalism"," *University of Pennsylvania Law Review*, 2003, *151* (3), 1211–1254.

Campbell, John Y, "Restoring rational choice: The challenge of consumer financial regulation," *American Economic Review*, 2016, *106* (5), 1–30.

_ **, Howell E Jackson, Brigitte C Madrian, and Peter Tufano**, "Consumer financial protection," *Journal of Economic Perspectives*, 2011, *25* (1), 91–114.

Caplin, Andrew and John Leahy, "The supply of information by a concerned expert," *The Economic Journal*, 2004, *114* (497), 487–505.

Cappelen, Alexander W., Astri Drange Hole, Erik Ø Sørensen, and Bertil Tungodden, "The pluralism of fairness ideals: An experimental approach," *American Economic Review*, 2007, *97* (3), 818–827.

_ **, Erik Ø Sørensen, and Bertil Tungodden**, "When do we lie?," *Journal of Economic Behavior & Organization*, 2013, *93*, 258–265.

_ , **Ingar K Haaland, and Bertil Tungodden**, "Beliefs about behavioral responses to taxation," Technical Report, Working Paper 2018.

_ , **James Konow, Erik Ø Sørensen, and Bertil Tungodden**, "Just luck: An experimental study of risk-taking and fairness," *American Economic Review*, 2013, *103* (4), 1398–1413.

_ , **Sebastian Fest, Erik Ø Sørensen, and Bertil Tungodden**, "Choice and personal responsibility: What is a morally relevant choice," *NHH Department of Economics Discussion Paper*, 2018, *27*.

Capraro, Valerio, "Gender differences in the trade-off between objective equality and efficiency," *Judgment and Decision Making*, 2020, *15* (4), 534–544.

Célérier, Claire and Boris Vallée, "Catering to investors through security design: Headline rate and complexity," *The Quarterly Journal of Economics*, 2017, *132* (3), 1469–1508.

Charness, Gary and Matthew Rabin, "Understanding social preferences with simple tests," *The Quarterly Journal of Economics*, 2002, *117* (3), 817–869.

Chen, Jianqing and Jan Stallaert, "An economic analysis of online advertising using behavioral targeting," *MIS Quarterly*, 2014, *38* (2), 429–450.

Chen, Jowei and Jonathan Rodden, "Unintentional Gerrymandering: Political Geography and Electoral Bias in Legislatures," *Quarterly Journal of Political Science*, 2013, *8* (3), 239–269.

Choi, James J, David Laibson, and Brigitte C Madrian, "Why does the law of one price fail? An experiment on index mutual funds," *The Review of Financial Studies*, 2010, *23* (4), 1405–1432.

Coffman, Lucas C, Clayton R Featherstone, and Judd B Kessler, "Can Social Information Affect What Job You Choose and Keep?," *American Economic Journal: Applied Economics*, 2017, *9* (1), 96–117.

Cohen, Gerald Allan, *Self-ownership, freedom, and equality*, Cambridge University Press, 1995.

Coibion, Olivier, Yuriy Gorodnichenko, and Tiziano Ropele, "Inflation expectations and firm decisions: New causal evidence," *The Quarterly Journal of Economics*, 2020, *135* (1), 165–219.

Corniere, Alexandre De and Romain De Nijs, "Online advertising and privacy," *The RAND Journal of Economics*, 2016, *47* (1), 48–72.

Cox, Caitriona L and Zoe Fritz, "Should non-disclosures be considered as morally equivalent to lies within the doctor–patient relationship?," *Journal of Medical Ethics*, 2016, *42* (10), 632–635.

Crawford, Vincent P and Joel Sobel, "Strategic information transmission," *Econometrica*, 1982, pp. 1431–1451.

Croson, Rachel and Uri Gneezy, "Gender differences in preferences," *Journal of Economic Literature*, 2009, *47* (2), 448–74.

Cui, Tony Haitao, Jagmohan S Raju, and Z John Zhang, "Fairness and channel coordination," *Management Science*, 2007, *53* (8), 1303–1314.

Davison, W Phillips, "The third-person effect in communication," *Public Opinion Quarterly*, 1983, *47* (1), 1–15.

DellaVigna, Stefano, "Psychology and economics: Evidence from the field," *Journal of Economic Literature*, 2009, *47* (2), 315–72.

_ **and Ulrike Malmendier**, "Contract design and self-control: Theory and evidence," *The Quarterly Journal of Economics*, 2004, *119* (2), 353–402.

_ **, John A List, Ulrike Malmendier, and Gautam Rao**, "Voting to tell others," *The Review of Economic Studies*, 2016, *84* (1), 143–181.

DiMatteo, Larry A and Bruce Louis Rich, "A consent theory of unconscionability: An empirical study of law in action," *Fla. St. UL Rev.*, 2005, *33*, 1067.

Dobber, Tom, Damian Trilling, Natali Helberger, and Claes de Vreese, "Spiraling downward: The reciprocal relation between attitude toward political behavioral targeting and privacy concerns," *New Media & Society*, 2019, *21* (6), 1212–1231.

_ **, Ronan Ó Fathaigh, and Frederik Zuiderveen Borgesius**, "The regulation of online political micro-targeting in Europe," *Internet Policy Review*, 2019, *8* (4), 1–20.

Doherty, David and Jennifer Wolak, "When Do the Ends Justify the Means? Evaluating Procedural Fairness," *Political Behavior*, 2012, *34* (2), 301–323.

Dohmen, Thomas J, Armin Falk, David Huffman, Felix Marklein, and Uwe Sunde, "The non-use of bayes rule: representative evidence on bounded rationality," 2009.

Dommett, Katharine, "Data-driven political campaigns in practice: understanding and regulating diverse data-driven campaigns," *Internet Policy Review*, 2019, *8* (4).

_ **and Sam Power**, "The political economy of Facebook advertising: Election spending, regulation and targeting online," *The Political Quarterly*, 2019, *90* (2), 257–265.

Douglas, Joshua A, "Procedural fairness in election contests," *Ind. L.J.*, 2013, *88*, 1–82.

Durante, Ruben, Louis Putterman, and Joël Van der Weele, "Preferences for redistribution and perception of fairness: An experimental study," *Journal of the European Economic Association*, 2014, *12* (4), 1059–1086.

Edelman, Benjamin, "Adverse selection in online "trust" certifications and search results," *Electronic Commerce Research and Applications*, 2011, *10* (1), 17–25.

Elbel, Brian, Rogan Kersh, Victoria L Brescoll, and L Beth Dixon, "Calorie Labeling And Food Choices: A First Look At The Effects On Low-Income People In New York City: Calorie information on menus appears to increase awareness of calorie content, but not necessarily the number of calories people purchase.," *Health Affairs*, 2009, *28* (Suppl1), w1110–w1121.

Elías, Julio J, Nicola Lacetera, and Mario Macis, "Paying for Kidneys? A Randomized Survey and Choice Experiment," *American Economic Review*, 2019, *109* (8), 2855–2588.

Eliaz, Kfir and Ran Spiegler, "Can anticipatory feelings explain anomalous choices of information sources?," *Games and Economic Behavior*, jul 2006, *56* (1), 87–104.

Engelmann, Dirk and Martin Strobel, "Inequality aversion, efficiency, and maximin preferences in simple distribution experiments," *American Economic Review*, 2004, *94* (4), 857–869.

Enikolopov, Ruben, Alexey Makarin, and Maria Petrova, "Social media and protest participation: Evidence from Russia," *Econometrica*, 2020, *88* (4), 1479–1514.

Enke, Benjamin, "What you see is all there is," *The Quarterly Journal of Economics*, 2020, *135* (3), 1363–1398.

Epstein, Richard A, "Unconscionability: A critical reappraisal," *The Journal of Law and Economics*, 1975, *18* (2), 293–315.

Epstein, Ronald M, David N Korones, and Timothy E Quill, "Withholding information from patients—when less is more," *New England Journal of Medicine*, 2010, *362* (5), 380.

Erat, Sanjiv and Uri Gneezy, "White lies," *Management Science*, 2012, *58* (4), 723–733.

Esberg, Jane, Jonathan Mummolo, and Sean J Westwood, "The Stability of Criminal Justice Policy Views: Evaluating the Effects of Factual Corrections and Appeals to Social Identity," 2020.

Evans, David S, "The online advertising industry: Economics, evolution, and privacy," *Journal of Economic Perspectives*, 2009, *23* (3), 37–60.

Eveland, William P, Amy I Nathanson, Benjamin H Detenber, and Douglas M McLeod, "Rethinking the social distance corollary: Perceived likelihood of expsoure and the third-person perception," *Communication Research*, 1999, *26* (3), 275–302.

Everett, Jim AC, "The 12 item social and economic conservatism scale (SECS)," *PloS One*, 2013, *8* (12), e82131.

Falk, Armin and Nora Szech, "Morals and markets," *Science*, 2013, *340* (6133), 707–711.

_ , **Ernst Fehr, and Urs Fischbacher**, "Testing theories of fairness—Intentions matter," *Games and Economic Behavior*, 2008, *62* (1), 287–303.

Fallis, Don, "Shedding light on keeping people in the dark," *Topics in Cognitive Science*, 2018.

Fehr, Ernst and Klaus M Schmidt, "A theory of fairness, competition, and cooperation," *The Quarterly Journal of Economics*, 1999, *114* (3), 817–868.

_ **and Urs Fischbacher**, "The nature of human altruism," *Nature*, 2003, *425* (6960), 785–791.

_ **, Holger Herz, and Tom Wilkening**, "The lure of authority: Motivation and incentive effects of power," *American Economic Review*, 2013, *103* (4), 1325–59.

Feldman, Stanley and Christopher Johnston, "Understanding the determinants of political ideology: Implications of structural complexity," *Political Psychology*, 2014, *35* (3), 337–358.

FeldmanHall, Oriel and Amitai Shenhav, "Resolving uncertainty in a social world," *Nature Human Behaviour*, 2019, *3* (5), 426–435.

Fernandes, Daniel, John G Lynch Jr, and Richard G Netemeyer, "Financial literacy, financial education, and downstream financial behaviors," *Management Science*, 2014, *60* (8), 1861–1883.

Finkel, Eli J., Christopher A. Bail, Mina Cikara, Peter H. Ditto, Shanto Iyengar, Samara Klar, Lilliana Mason, Mary C. McGrath, Brendan Nyhan, David G. Rand, Linda J. Skitka, Joshua A. Tucker, Jay J. Van Bavel, Cynthia S. Wang, and James N. Druckman, "Political sectarianism in America," *Science*, 2020, *370* (6516), 533–536.

Fisman, Raymond, Pamela Jakiela, and Shachar Kariv, "Distributional preferences and political behavior," *Journal of Public Economics*, 2017, *155*, 1–10.

Floridi, Luciano, *Information: A very short introduction*, OUP Oxford, 2010.

_ , *The ethics of information*, Oxford University Press, 2013.

_ , *The fourth revolution: How the infosphere is reshaping human reality*, OUP Oxford, 2014.

Flynn, DJ, Brendan Nyhan, and Jason Reifler, "The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics," *Political Psychology*, 2017, *38*, 127–150.

Fourcade, Marion, Etienne Ollion, and Yann Algan, "The superiority of economists," *Journal of Economic Perspectives*, 2015, *29* (1), 89–114.

Fowler, Erika Franklin, Michael M. Franz, Gregory J. Martin, Zachary Peskowitz, and Travis N. Ridout, "Political Advertising Online and Offline," *American Political Science Review*, 2020, p. 1–20.

Frydman, Cary and Baolian Wang, "The impact of salience on investor behavior: Evidence from a natural experiment," *The Journal of Finance*, 2020, *75* (1), 229–276.

FTC, Federal Trade Commission, "Data brokers: a call for transparency and accountability," *Report, FTC, Washington, DC.*, 2014.

Fuchs, Christian, *Internet and society: Social theory in the information age*, Routledge, 2007.

Gabaix, Xavier and David Laibson, "Shrouded attributes, consumer myopia, and information suppression in competitive markets," *The Quarterly Journal of Economics*, 2006, *121* (2), 505–540.

Gagné, Marylène and Edward L Deci, "Self-determination theory and work motivation," *Journal of Organizational Behavior*, 2005, *26* (4), 331–362.

Gerber, Alan S and Gregory A Huber, "Partisanship, political control, and economic assessments," *American Journal of Political Science*, 2010, *54* (1), 153–173.

Gibson, Rajna, Carmen Tanner, and Alexander F Wagner, "Preferences for truthfulness: Heterogeneity among and within individuals," *American Economic Review*, 2013, *103* (1), 532–48.

Gleick, James, *The information: A history, a theory, a flood*, Vintage, 2011.

Gneezy, Uri, "Deception: The role of consequences," *American Economic Review*, 2005, *95* (1), 384–394.

_ , Agne Kajackaite, and Joel Sobel, "Lying Aversion and the Size of the Lie," *American Economic Review*, 2018, *108* (2), 419–53.

Goldfayn-Frank, Olga and Johannes Wohlfart, "Expectation Formation in a New Environment: Evidence from the German Reunification," *Journal of Monetary Economics*, 2019.

Golman, Russell, David Hagmann, and George Loewenstein, "Information avoidance," *Journal of Economic Literature*, 2017, *55* (1), 96–135.

Goren, Paul, "Party identification and core political values," *American Journal of Political Science*, 2005, *49* (4), 881–896.

Gottlieb, Daniel and Kent Smetters, "Lapse-Based Insurance," *Working Paper*, 2016.

Graham, Matthew and Milan Svolik, "Democracy in America? Partisanship, Polarization, and the Robustness of Support for Democracy in the United States," *American Political Science Review*, 2019, *114* (2), 392–409.

Grubb, Michael D, "Selling to overconfident consumers," *American Economic Review*, 2009, *99* (5), 1770–1807.

_ , "Overconfident consumers in the marketplace," *Journal of Economic Perspectives*, 2015, *29* (4), 9–36.

Gunther, Albert C and Paul Mundy, "Biased optimism and the third-person effect," *Journalism Quarterly*, 1993, *70* (1), 58–67.

Güth, Werner, Hartmut Kliemt, and Axel Ockenfels, "Fairness versus efficiency: An experimental study of (mutual) gift giving," *Journal of Economic Behavior & Organization*, 2003, *50* (4), 465–475.

Haaland, Ingar, Christopher Roth, and Johannes Wohlfart, "Designing information provision experiments," *CEBI Working Paper Series, Working Paper 20/20*, 2020.

Hager, Anselm, "Do online ads influence vote choice?," *Political Communication*, 2019, *36* (3), 376–393.

Hare, T. A., J. Malmaud, and A. Rangel, "Focusing Attention on the Health Aspects of Foods Changes Value Signals in vmPFC and Improves Dietary Choice," *Journal of Neuroscience*, jul 2011, *31* (30), 11077–11087.

Harris, John and Kirsty Keywood, "Ignorance, information and autonomy," *Theoretical Medicine and Bioethics*, 2001, *22* (5), 415–436.

Heidhues, Paul and Botond Kőszegi, "Exploiting Naïvete about Self-Control in the Credit Market," *American Economic Review*, dec 2010, *100* (5), 2279–2303.

_ and _ , "Behavioral Industrial Organization," in "Handbook of Behavioral Economics: Applications and Foundations 1," Elsevier, 2018, pp. 517–612.

_ , Johannes Johnen, and Botond Kőszegi, "Browsing versus studying: A pro-market case for regulation," *unpublished MS*, 2018, *27*.

Henriksen, Lisa and June A Flora, "Third-person perception and children: Perceived impact of pro-and anti-smoking ads," *Communication Research*, 1999, *26* (6), 643–665.

Hirshleifer, David, "Investor psychology and asset pricing," *The Journal of Finance*, 2001, *56* (4), 1533–1597.

Hoffner, Cynthia and Martha Buchanan, "Parents' responses to television violence: The third-person perception, parental mediation, and support for censorship," *Media Psychology*, 2002, *4* (3), 231–252.

Holley, David M, "Information disclosure in sales," *Journal of Business Ethics*, 1998, *17* (6), 631–641.

Isaac, Mike, "Why Everyone Is Angry at Facebook Over Its Political Ads Policy," *NY Times*, 2019. `https://www.nytimes.com/2019/11/22/technology/campaigns-pressure-facebook-political-ads.html`, Last accessed on 2020-03-15.

_ , "Now You Can Opt Out of Seeing Political Ads on Facebook," *NY Times*, 2020. `https://www.nytimes.com/2020/06/16/technology/opt-out-political-ads-facebook.html`, Last accessed on 2020-07-18.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood, "The origins and consequences of affective polarization in the United States," *Annual Review of Political Science*, 2019, *22*, 129–146.

Iyer, Ganesh, David Soberman, and J Miguel Villas-Boas, "The targeting of advertising," *Marketing Science*, 2005, *24* (3), 461–476.

Iyer, Ravi, Spassena Koleva, Jesse Graham, Peter Ditto, and Jonathan Haidt, "Understanding libertarian morality: The psychological dispositions of self-identified libertarians," *PloS One*, 2012, *7* (8).

Jacobsson, Fredric, Magnus Johannesson, and Lars Borgquist, "Is altruism paternalistic?," *The Economic Journal*, 2007, *117* (520), 761–781.

Jang, S Mo and Joon K Kim, "Third person effects of fake news: Fake news regulation and media literacy interventions," *Computers in Human Behavior*, 2018, *80*, 295–302.

Janoff-Bulman, Ronnie, "To provide or protect: Motivational bases of political liberalism and conservatism," *Psychological Inquiry*, 2009, *20* (2-3), 120–128.

Jansen, Bernard J, Kathleen Moore, and Stephen Carman, "Evaluating the performance of demographic targeting using gender in sponsored search," *Information Processing & Management*, 2013, *49* (1), 286–302.

Jentzsch, Nicola, *The regulation of financial privacy: the United States Vs Europe*, ECRI RESEARCH REPORT NO. 5, 2003.

Jin, Ginger Zhe, Michael Luca, and Daniel Martin, "Is no news (perceived as) bad news? An experimental investigation of information disclosure," Technical Report, National Bureau of Economic Research 2015.

Jost, John T, "Ideological asymmetries and the essence of political psychology," *Political psychology*, 2017, *38* (2), 167–208.

_ , Christopher M Federico, and Jaime L Napier, "Political ideology: Its structure, functions, and elective affinities," *Annual Review of Psychology*, 2009, *60*, 307–337.

_ , Sally Blount, Jeffrey Pfeffer, and György Hunyady, "Fair market ideology: Its cognitive-motivational underpinnings," *Research in Organizational Behavior*, 2003, *25*, 53–91.

Jung, Janice Y and Barbara A Mellers, "American attitudes toward nudges.," *Judgment & Decision Making*, 2016, *11* (1).

Kahneman, Daniel, Jack L Knetsch, and Richard Thaler, "Fairness as a constraint on profit seeking: Entitlements in the market," *American Economic Review*, 1986, pp. 728–741.

Kerschbamer, Rudolf and Daniel Müller, "Social preferences and political attitudes: An online experiment on a large heterogeneous sample," *Journal of Public Economics*, 2020, *182*, 104076.

_ , Daniel Neururer, and Alexander Gruber, "Do altruists lie less?," *Journal of Economic Behavior & Organization*, 2019, *157*, 560–579.

Khalmetski, Kiryl, Bettina Rockenbach, and Peter Werner, "Evasive lying in strategic communication," *Journal of Public Economics*, 2017, *156*, 59–72.

Kim, Young Mie, Jordan Hsu, David Neiman, Colin Kou, Levi Bankston, Soo Yun Kim, Richard Heinrich, Robyn Baragwanath, and Garvesh Raskutti, "The stealth media? Groups and targets behind divisive issue campaigns on Facebook," *Political Communication*, 2018, *35* (4), 515–541.

Konow, James, "Which is the fairest one of all? A positive analysis of justice theories," *Journal of Economic Literature*, 2003, *41* (4), 1188–1239.

_ , "Is fairness in the eye of the beholder? An impartial spectator analysis of justice," *Social Choice and Welfare*, 2009, *33* (1), 101–127.

Korobkin, Russell, "Bounded rationality, standard form contracts, and unconscionability," *The University of Chicago Law Review*, 2003, pp. 1203–1295.

Kőszegi, Botond, "Behavioral contract theory," *Journal of Economic Literature*, 2014, *52* (4), 1075–1118.

Kozlowski, Lynn T and David Sweanor, "Withholding differential risk information on legal consumer nicotine/tobacco products: the public health ethics of health information quarantines," *International Journal of Drug Policy*, 2016, *32*, 17–23.

Kozup, John, Elizabeth Howlett, and Michael Pagano, "The effects of summary information on consumer perceptions of mutual fund characteristics," *Journal of Consumer Affairs*, 2008, *42* (1), 37–59.

Krasnova, Hanna, Oliver Günther, Sarah Spiekermann, and Ksenia Koroleva, "Privacy concerns and identity in online social networks," *Identity in the Information Society*, 2009, *2* (1), 39–63.

Kronman, Anthony T, "Paternalism and the Law of Contracts," *The Yale Law Journal*, 1983, *92* (5), 763–798.

Kruse, Lisa M, Dawn R Norris, and Jonathan R Flinchum, "Social media as a public sphere? Politics on social media," *The Sociological Quarterly*, 2018, *59* (1), 62–84.

Kunst, Jonas R, John F Dovidio, and Lotte Thomsen, "Fusion with political leaders predicts willingness to persecute immigrants and political opponents," *Nature Human Behaviour*, 2019, *3* (11), 1180–1189.

Kymlicka, Will, *Contemporary political philosophy: An introduction*, Oxford University Press, 2002.

Lazer, David MJ, Matthew A Baum, Yochai Benkler, Adam J Berinsky, Kelly M Greenhill, Filippo Menczer, Miriam J Metzger, Brendan Nyhan, Gordon Pennycook, David Rothschild et al., "The science of fake news," *Science*, 2018, *359* (6380), 1094–1096.

Leathern, Rob, "Expanded Transparency and More Controls for Political Ads," 2020. `https://about.fb.com/news/2020/01/political-ads/`, Last accessed on 2020-05-22.

Lees, Jeffrey and Mina Cikara, "Inaccurate group meta-perceptions drive negative out-group attributions in competitive contexts," *Nature Human Behaviour*, 2020, *4* (3), 279–286.

Leibbrandt, Andreas, "Behavioral constraints on price discrimination: Experimental evidence on pricing and customer antagonism," *European Economic Review*, 2020, *121*, 103303.

Lergetporer, Philipp, Guido Schwerdt, Katharina Werner, Martin R West, and Ludger Woessmann, "How information affects support for education spending: Evidence from survey experiments in Germany and the United States," *Journal of Public Economics*, 2018, *167*, 138–157.

Lerman, Rachel, "Google reins in political advertising," *Associated Press*, 2019. `https://apnews.com/6b57e86d91034b8c80a3c634a23bcfd8`, Last accessed on 2020-03-17.

Levitsky, Steven and Daniel Ziblatt, *How democracies die*, Broadway Books, 2018.

Liberini, Federica, Antonio Russo, Ángel Cuevas, Ruben Cuevas et al., "Politics in the Facebook Era-Evidence from the 2016 US Presidential Elections," 2020.

Lipnowski, Elliot, Laurent Mathevet, and Dong Wei, "Attention management," *American Economic Review: Insights*, 2020, *2* (1), 17–32.

List, John A, "Homo experimentalis evolves," *Science*, 2008, *321* (5886), 207–208.

_ , "Why economists should conduct field experiments and 14 tips for pulling one off," *Journal of Economic Perspectives*, 2011, *25* (3), 3–16.

Lo, Ven and Ran Wei, "Third-person effect, gender, and pornography on the lnternet," *Journal of Broadcasting & Electronic Media*, 2002, *46* (1), 13–33.

Loewenstein, George and Nick Chater, "Putting nudges in perspective," *Behavioural Public Policy*, 2017, *1* (1), 26–53.

_ **, Cass R Sunstein, Russell Golman et al.**, "Disclosure: Psychology Changes Everything," *Annual Review of Economics*, 2014, *6* (1), 391–419.

López-Pérez, Raúl and Eli Spiegelman, "Why do people tell the truth? Experimental evidence for pure lie aversion," *Experimental Economics*, 2013, *16* (3), 233–247.

Lyon, David, *The information society: Issues and illusions*, John Wiley & Sons, 2013.

Magalhães, Pedro C, John H Aldrich, and Rachel K Gibson, "New forms of mobilization, new people mobilized? Evidence from the Comparative Study of Electoral Systems," *Party Politics*, 2018, *26* (6), 605–618.

Malka, Ariel, Christopher J Soto, Michael Inzlicht, and Yphtach Lelkes, "Do needs for security and certainty predict cultural and economic conservatism? A cross-national analysis.," *Journal of Personality and Social Psychology*, 2014, *106* (6), 1031.

Malmendier, Ulrike and Stefan Nagel, "Learning from inflation experiences," *The Quarterly Journal of Economics*, 2016, *131* (1), 53–87.

Martherus, James L, Andres G Martinez, Paul K Piff, and Alexander G Theodoridis, "Party Animals? Extreme Partisan Polarization and Dehumanization," *Political Behavior*, 2019, pp. 1–24.

Mason, Lilliana, *Uncivil Agreement: How Politics Became Our Identity*, University of Chicago Press, 2018.

Matakos, Konstantinos and Dimitrios Xefteris, "Strategic electoral rule choice under uncertainty," *Public Choice*, 2015, *162* (3-4), 329–350.

Mazar, Nina, On Amir, and Dan Ariely, "The dishonesty of honest people: A theory of self-concept maintenance," *Journal of Marketing Research*, 2008, *45* (6), 633–644.

McCarthy, Devin, "Partisanship vs. Principle: Understanding Public Opinion on Same-Day Registration," *Public Opinion Quarterly*, 2019, *83* (3), 568–583.

McLeod, Douglas M, William P Eveland, and Amy I Nathanson, "Support for censorship of violent and misogynic rap lyrics: An analysis of the third-person effect," *Communication Research*, 1997, *24* (2), 153–174.

Medina, Paolina C, "Selective attention in consumer finance: Evidence from a randomized intervention in the credit card market," Technical Report, Working Paper 2018.

Milberg, Sandra J, H Jeff Smith, and Sandra J Burke, "Information privacy: Corporate management and national regulation," *Organization Science*, 2000, *11* (1), 35–57.

_ **, Sandra J Burke, Jeff H Smith, and Ernest A Kallman**, "Values, personal information privacy, and regulatory approaches," *Communications of the ACM*, 1995, *38* (12), 65–74.

Mill, John Stuart, "On Liberty," *Collected works of John Stuart Mill*, 1859, pp. 259–340.

Miltgen, Caroline Lancelot and H Jeff Smith, "Exploring information privacy regulation, risks, trust, and behavior," *Information & Management*, 2015, *52* (6), 741–759.

Moore-Berg, Samantha L, Lee-Or Ankori-Karlinsky, Boaz Hameiri, and Emile Bruneau, "Exaggerated meta-perceptions predict intergroup hostility between American political partisans," *Proceedings of the National Academy of Sciences*, 2020.

Moore, Robert S, Melissa L Moore, Kevin J Shanahan, and Britney Mack, "Creepy marketing: Three dimensions of perceived excessive online privacy violation," *Marketing Management*, 2015, *25* (1), 42–53.

Munafò, Marcus R, Brian A Nosek, Dorothy VM Bishop, Katherine S Button, Christopher D Chambers, Nathalie Percie Du Sert, Uri Simonsohn, Eric-Jan Wagenmakers, Jennifer J Ware, and John PA Ioannidis, "A manifesto for reproducible science," *Nature Human Behaviour*, 2017, *1* (1), 1–9.

Newell, Richard G and Juha Siikamäki, "Nudging energy efficiency behavior: The role of information labels," *Journal of the Association of Environmental and Resource Economists*, 2014, *1* (4), 555–598.

Nguyen, Bang and Dilip S Mutum, "A review of customer relationship management: successes, advances, pitfalls and futures," *Business Process Management Journal*, 2012.

Nosek, Brian A, George Alter, George C Banks, Denny Borsboom, Sara D Bowman, Steven J Breckler, Stuart Buck, Christopher D Chambers, Gilbert Chin, Garret Christensen et al., "Promoting an open research culture," *Science*, 2015, *348* (6242), 1422–1425.

Nozick, Robert, *Anarchy, state, and utopia*, Vol. 5038, New York: Basic Books, 1974.

Nuñez, Benjamin JR, "Developing behavioural economics as a practical tool for market authorities: Lessons learned from the first era of behavioural case work," *Journal of Antitrust Enforcement*, 2017, *5* (3), 375–406.

Nyhan, Brendan, "Facts and Myths about Misperceptions," *Journal of Economic Perspectives*, August 2020, *34* (3), 220–36.

_ , Gretchen Helmke, Mitchell Sanders, Katherine Clayton, John Carey, and Susan Stokes, "Who Will Defend Democracy? Evaluating Tradeoffs in Candidate Support Among Partisan Donors and Voters," *Journal of Elections, Public Opinion and Parties*, 2020, pp. 1–15.

Oehler, Andreas and Stefan Wendt, "Good consumer information: The Information Paradigm at its (dead) end?," *Journal of Consumer Policy*, 2017, *40* (2), 179–191.

Okazaki, Shintaro, Hairong Li, and Morikazu Hirose, "Consumer privacy concerns and preference for degree of regulatory control," *Journal of Advertising*, 2009, *38* (4), 63–77.

Paul, Bryant, Michael B Salwen, and Michel Dupagne, "The third-person effect: A meta-analysis of the perceptual hypothesis," *Mass Communication & Society*, 2000, *3* (1), 57–85.

Perloff, Richard M, "Third-person effect research 1983–1992: A review and synthesis," *International Journal of Public Opinion Research*, 1993, *5* (2), 167–184.

_ , "The third person effect," *Media Psychology*, 1999, *1*, 353–78.

Persson, Petra, "Attention manipulation and information overload," *Behavioural Public Policy*, 2018, *2* (1), 78–106.

Pete, Lunn, *Regulatory Policy and Behavioural Economics*, OeCD Publishing, 2014.

Pieters, Rik and Hans Baumgartner, "Who talks to whom? Intra-and interdisciplinary communication of economics journals," *Journal of Economic Literature*, 2002, *40* (2), 483–509.

Pikulina, Elena S and Chloe Tergiman, "Preferences for power," *Journal of Public Economics*, 2020, *185*, 104173.

Prior, Markus, Gaurav Sood, Kabir Khanna et al., "You cannot be serious: The impact of accuracy incentives on partisan bias in reports of economic perceptions," *Quarterly Journal of Political Science*, 2015, *10* (4), 489–518.

Pronin, Emily, Daniel Y Lin, and Lee Ross, "The bias blind spot: Perceptions of bias in self versus others," *Personality and Social Psychology Bulletin*, 2002, *28* (3), 369–381.

Quidt, Jonathan De, Johannes Haushofer, and Christopher Roth, "Measuring and bounding experimenter demand," *American Economic Review*, 2018, *108* (11), 3266–3302.

Reisch, Lucia A and Cass R Sunstein, "Do Europeans like nudges?," *Judgment and Decision Making*, 2016, *11* (4), 310–325.

Roberts, Jessica L, "Nudge-proof: distributive justice and the ethics of nudging," *Mich. L. Rev.*, 2017, *116*, 1045.

Rodriguez-Lara, Ismael and Luis Moreno-Garrido, "Self-interest and fairness: self-serving choices of justice principles," *Experimental Economics*, 2012, *15* (1), 158–175.

Rojas, Hernando, Dhavan V Shah, and Ronald J Faber, "For the good of others: Censorship and the third-person effect," *International Journal of Public Opinion Research*, 1996, *8* (2), 163–186.

Roth, Christopher and Johannes Wohlfart, "Experienced inequality and preferences for redistribution," *Journal of Public Economics*, 2018, *167*, 251–262.

_ **and** _ , "How do expectations about the macroeconomy affect personal expectations and behavior?," *Review of Economics and Statistics*, 2019.

Ru, Hong and Antoinette Schoar, "Do credit card companies screen for behavioral biases?," Technical Report, National Bureau of Economic Research 2016.

Rubinstein, Ira S, "Voter privacy in the age of big data," *Wis. L. Rev.*, 2014, *5*, 861–936.

Samuelson, Paul and William Nordhaus, "Principles of Economics," *McCraw-liill, New York*, 1985.

Schick, Thomas A, "Truth, accuracy (and withholding information)," *Public Relations Quarterly*, 1994, *39* (4), 7.

Schwartz, Alan and Robert E Scott, "Contract theory and the limits of contract law," *Yale LJ*, 2003, *113*, 541.

Scopelliti, Irene, Carey K Morewedge, Erin McCormick, H Lauren Min, Sophie Lebrecht, and Karim S Kassam, "Bias blind spot: Structure, measurement, and consequences," *Management Science*, 2015, *61* (10), 2468–2486.

Seizov, O., A. J. Wulf, and J. Luzak, "The Transparent Trap: A Multidisciplinary Perspective on the Design of Transparent Online Disclosures in the EU," *Journal of Consumer Policy*, 2019, *42* (1), 149–173.

Serra-Garcia, Marta, Eric Van Damme, and Jan Potters, "Hiding an inconvenient truth: Lies and vagueness," *Games and Economic Behavior*, 2011, *73* (1), 244–261.

Shah, Dhavan V, Ronald J Faber, and Seounmi Youn, "Susceptibility and severity: Perceptual dimensions underlying the third-person effect," *Communication Research*, 1999, *26* (2), 240–267.

Shalvi, Shaul and David Leiser, "Moral firmness," *Journal of Economic Behavior & Organization*, 2013, *93*, 400–407.

Sharot, Tali, "The optimism bias," *Current Biology*, 2011, *21* (23), R941–R945.

_ **and Cass R Sunstein**, "How people decide what they want to know," *Nature Human Behaviour*, 2020, pp. 1–6.

Sheremeta, Roman M and Timothy W Shields, "Deception and reception: The behavior of information providers and users," *Journal of Economic Behavior & Organization*, 2017, *137*, 445–456.

Sihvola, Elvira, "Privacy and Political Integrity: How European Data Protection Laws May Limit the Regulation of Foreign Political Interference in US Elections," *Colum. J. Eur. L.*, 2019, *25*, 135–166.

Skitka, Linda J, "Ideological and attributional boundaries on public compassion: Reactions to individuals and communities affected by a natural disaster," *Personality and Social Psychology Bulletin*, 1999, *25* (7), 793–808.

_ **and Philip E Tetlock**, "Providing public assistance: Cognitive and motivational processes underlying liberal and conservative policy preferences.," *Journal of Personality and Social Psychology*, 1993, *65* (6), 1205.

_ **, Elizabeth Mullen, Thomas Griffin, Susan Hutchinson, and Brian Chamberlin**, "Dispositions, scripts, or motivated correction? Understanding ideological differences in explanations for social problems.," *Journal of Personality and Social Psychology*, 2002, *83* (2), 470.

Smith, Aaron, *Public Attitudes Toward Computer Algorithms*, Pew Research Center, 2018.

Smith, H Jeff, Sandra J Milberg, and Sandra J Burke, "Information privacy: measuring individuals' concerns about organizational practices," *MIS Quarterly*, 1996, *20* (2), 167–196.

Spector, Horacio, "A Contractarian Approach to Unconscionability," *Chi.-Kent L. Rev.*, 2006, *81*, 95.

Stewart, Charles III, Stephen Ansolabehere, and Nathaniel Persily, "Revisiting Public Opinion on Voter Identification and Voter Fraud in an Era of Increasing Partisan Polarization," *Stan. L. Rev.*, 2016, *68*, 1455–1490.

Stigler, George J, "The economics of information," *Journal of Political Economy*, 1961, *69* (3), 213–225.

Sun, Ye, Lijiang Shen, and Zhongdang Pan, "On the behavioral component of the third-person effect," *Communication Research*, 2008, *35* (2), 257–278.

Sunshine, Jason and Tom R Tyler, "The role of procedural justice and legitimacy in shaping public support for policing," *Law & Society Review*, 2003, *37* (3), 513–548.

Sunstein, Cass R, *The Ethics of Influence*, Cambridge University Press, 2016.

_ , *# Republic: Divided democracy in the age of social media*, Princeton University Press, 2018.

_ , "Ruining popcorn? The welfare effects of information," *Journal of Risk and Uncertainty*, 2019, *58* (2-3), 121–142.

_ , "Which nudges do people like? A national survey," in "Handbook of Behavioural Change and Public Policy," Edward Elgar Publishing, 2019.

_ , *Too Much Information: Understanding What You Dont Want to Know*, MIT Press, 2020.

Svolik, Milan W, "When Polarization Trumps Civic Virtue: Partisan Conflict and the Subversion of Democracy by Incumbents," *Quarterly Journal of Political Science*, 2018, *15* (1), 3–31.

_ , "Polarization versus democracy," *Journal of Democracy*, 2019, *30* (3), 20–32.

Tan, Joshua, Mahmood Sharif, Sruti Bhagavatula, Matthias Beckerle, Michelle L Mazurek, and Lujo Bauer, “"Comparing hypothetical and realistic privacy valuations",” in “Proceedings of the 2018 Workshop on Privacy in the Electronic Society” 2018, pp. 168–182.

Tang, Zhulei, Yu Hu, and Michael D Smith, “Gaining trust through online privacy protection: Self-regulation, mandatory standards, or caveat emptor,” *Journal of Management Information Systems*, 2008, *24* (4), 153–173.

Tannenbaum, David, Craig R Fox, and Todd Rogers, “On the misplaced politics of behavioural policy interventions,” *Nature Human Behaviour*, 2017, *1* (7), 1–7.

Tene, Omer and Jules Polonetsky, “A theory of creepy: technology, privacy and shifting social norms,” *Yale JL & Tech.*, 2013, *16*, 59–102.

Thaler, Richard, “Toward a positive theory of consumer choice,” *Journal of Economic Behavior & Organization*, 1980, *1* (1), 39–60.

_ , *Misbehaving: The making of behavioral economics*, WW Norton New York, 2015.

_ **and Cass R Sunstein**, *Nudge: Improving decisions about health, wealth, and happiness*, Penguin, 2009.

Trevino, Isabel, “Informational channels of financial contagion,” *Econometrica*, 2020, *88* (1), 297–335.

Tversky, Amos and Daniel Kahneman, “Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment.,” *Psychological Review*, 1983, *90* (4), 293.

Wang, Huaiqing, Matthew KO Lee, and Chen Wang, "Consumer privacy concerns about Internet marketing," *Communications of the ACM*, 1998, *41* (3), 63–70.

Warren, Patrick L and Daniel H Wood, "The political economy of regulation in markets with naïve consumers," *Journal of the European Economic Association*, 2014, *12* (6), 1617–1642.

Webster, Frank, *Theories of the information society*, Routledge, 2014.

Weintraub, Ellen L., "Don't abolish political ads on social media. Stop microtargeting.," *Washington Post*, 2019. `https://www.washingtonpost.com/opinions/2019/11/01/dont-abolish-political-ads-social-media-stop-microtargeting/`, Last accessed on 2020-05-29.

Welzel, Christian and Ronald Inglehart, "Mass beliefs and democratic institutions," in Carles Boix and Susan C. Stokes, eds., *The Oxford handbook of comparative politics*, 2007.

White, H Allen, "Considering interacting factors in the third-person effect: Argument strength and social distance," *Journalism & Mass Communication Quarterly*, 1997, *74* (3), 557–564.

Wong, Julia Carrie, "Twitter's political ad ban to limit micro-targeting, putting pressure on Facebook," *The Guardian*, 2019. `https://www.theguardian.com/technology/2019/nov/15/twitter-facebook-ad-ban-micro-targeting`, Last accessed on 2020-05-22.

Wood, Abby K and Ann M Ravel, "Fool Me Once: Regulating Fake News and Other Online Advertising," *S. Cal. L. Rev.*, 2017, *91*, 1223.

Xia, Lan, Kent B Monroe, and Jennifer L Cox, "The price is unfair! A conceptual framework of price fairness perceptions," *Journal of Marketing*, 2004, *68* (4), 1–15.

Xu, Heng, Tamara Dinev, Jeff Smith, and Paul Hart, "Information privacy concerns: Linking individual perceptions with institutional privacy assurances," *Journal of the Association for Information Systems*, 2011, *12* (12), 798–824.

Xu, Jie and William J Gonzenbach, "Does a perceptual discrepancy lead to action? A meta-analysis of the behavioral component of the third-person effect," *International Journal of Public Opinion Research*, 2008, *20* (3), 375–385.

Youn, Seounmi, Ronald J Faber, and Dhavan V Shah, "Restricting gambling advertising and the third-person effect," *Psychology & Marketing*, 2000, *17* (7), 633–649.

Zbaracki, Mark J, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen, "Managerial and customer costs of price adjustment: direct evidence from industrial markets," *Review of Economics and Statistics*, 2004, *86* (2), 514–533.

Zimmerman, Michael J. and Ben Bradley, "Intrinsic vs. Extrinsic Value," in Edward N. Zalta, ed., *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, 2019.