NBER WORKING PAPER SERIES

SELFISH CORPORATIONS

Emanuele Colonnelli
Niels Joachim Gormsen
Timothy McQuade

Working Paper 30576
http://www.nber.org/papers/w30576

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2022

We thank Nicholas Barberis, Craig Brown, Leo Bursztyn, Nicola Gennaioli, Luigi Guiso, Sam Hartzmark, Mitch Hoffman, Ben Iverson, Steve Kaplan, Eben Lazarus, Yueran Ma, Elisa Maffioli, Alexey Makarin, Maria Marino, Filippo Mezzanotti, Claudio Michelacci, Lubos Pastor, Lasse Pedersen, Raghu Rajan, Chris Roth, Andrei Shleifer, Michelle Skinner, Stefanie Stantcheva, Amir Sufi, Edoardo Teso, Rob Vishny, David Yang, Florian Zimmermann, Lugi Zingales, and seminar participants at NBER Behavioral Finance, WFA, the Finance, Organization, and Markets conference, the Junior Entrepreneurial Finance and Innovation group, the IMF, BYU, EIEF, Bocconi, and the University of Siena for helpful comments and suggestions. Gabriel Lobato, Edoardo Marchesi, Sanhitha Jugulum, and Pranav Garg provided superb research assistance. We are grateful to The University of Chicago Booth School of Business, the Asness Junior Faculty Fellowship, the Fama Research Fund, and the Liew Family Junior Faculty Fellowship for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Selfish Corporations
Emanuele Colonnelli, Niels Joachim Gormsen, and Timothy McQuade
NBER Working Paper No. 30576
October 2022
JEL No. D63,D72,G34,G38,M14

ABSTRACT

Motivated by the public debate regarding corporate responsibility, we construct a memory-based model of decision-making to illustrate how corporate and political communication can impact policy preferences. We test the predictions of our model in a new large-scale survey of U.S. citizens on their support for economic policies such as corporate bailouts. We first establish that the public demands corporations to behave better within society, a sentiment we label "big business discontent." Then, using random variation in the order of survey sections and in the exposure to animated videos, we confirm the key predictions of our model. First, messages that prime respondents to think about policy through the lens of corporate responsibility make people more averse to bailouts, while reframing the issue in terms of economic trade-offs has opposite effects. Second, attempts to paint a positive public image of big business can actually backfire, as they focus attention on an aspect on which the public has well-established negative views.

Emanuele Colonnelli
Booth School of Business
University of Chicago
5807 S Woodlawn Ave
Chicago, IL 60637
and NBER
emanuele.colonnelli@chicagobooth.edu

Niels Joachim Gormsen
University of Chicago
5807 South Woodlawn Avenue
Office 401
Chicago, IL 60637
niels.gormsen@chicagobooth.edu

Timothy McQuade
Haas School of Business
University of California, Berkeley
Berkeley, CA 94720
tmcquade@haas.berkeley.edu

An online appendix is available at http://www.nber.org/data-appendix/w30576

1. Introduction

There is a fundamental debate in the United States about the role of large corporations in society, which has been accelerated by the social unrest and global pandemic of 2020. Now more than ever, regulators and the public argue about whether large corporations should foster diversity in the workplace, limit wage inequality, protect the environment, and care for local communities. This debate coincides with increasing calls by the public for less pro-business regulation, a change many argue is driven by a rising and widespread anti-corporate sentiment (Cowen, 2019). For instance, the recent regulatory scrutiny towards big tech (Apple, Amazon, Google, Facebook) might be seen as the result of a deterioration in their public image, which suddenly makes targeting big tech good politics.[1]

Corporate America appears to be reacting to this threat via extensive media communication that tries to paint an image of big business as friendly to society at large. A glaring example is the 2019 statement by the Business Roundtable—the association of chief executive officers of major US companies—which redefined the purpose of a corporation to promote "an economy that serves all Americans," marking a stark change from the famous statement by Friedman (1970) that "the social responsibility of business is to increase its profits."[2]

Motivated by this public debate, we ask how corporate and political communication strategies affect individual beliefs and policy preferences. We first present a simple model of memory-based decision-making in which preferences depend on which memories individuals recall when evaluating policies, and study how communication influences policy preferences through persuasion. We then test and validate the model's predictions in an experimental survey in the context of environmental, social, and governance (ESG) practices of large corporations.[3]

We develop a model of cognitive thinking inspired by the psychology model of associative memory recall by Kahana (2012). In contrast to standard economic models where rational agents have access to all relevant information, memory recall is limited and guided by contextual prompts. Within this framework, communications and messaging provide cues that prime the agent to recall experiences similar to the cue. However, the cue does not perfectly account for which experiences will be recalled. Similar, but non-cued experiences

[1]See Why does Washington suddenly seem ready to regulate Big Tech? (Vox, June 2019).

[2]See Green Gold: How Sustainability Became Big Business for Consumer Brands (Financial Times, November 2020) and Business Roundtable Redefines the Purpose of a Corporation to Promote 'An Economy That Serves All Americans' (Business Roundtable, August 2019).

[3]ESG is the leading model used in the investment world to measure the impact corporations have on society. Another typical term used in this literature is CSR, which stands for "corporate social responsibility." Both ESG and CSR are usually seen as "catchall" terms for several aspects of corporate responsibility policies, and they are extremely close to each other. Throughout the paper we use the terminology of ESG or "corporate responsibility."

will *interfere* with the recall of cued experiences. Policy preferences are thus dependent on the interplay between cue, similarity, and interference, since together they impact the set of experiences used to evaluate the policy.

To illustrate, consider the policy issue of corporate bailouts during a crisis. In our model, support for this policy is determined by which memories the agents recall when assessing the policy. This recall can then be influenced by cues that prime both a *policy domain*, which makes salient certain aspects of the policy decision, and a *valence framing* of the issues in that policy domain. For example, in an attempt to increase support for bailouts, big businesses might communicate positive examples of corporate social behavior. Such messaging would cue social responsibility as the policy domain, thereby leading agents to recall more memories about the social responsibility of large corporations when evaluating bailout policy. Moreover, the positive valence framing would lead agents to disproportionately draw positive memories about social responsibility. Crucially, though, there will be *interference* from memories that, while not directly cued, are similar to the cued memories. That is, there might be interference coming from (negative) memories of socially irresponsible behavior of large corporations. The larger the proportion of negative memories within a given policy context, the larger this interference will be. Concretely, in our example, positive messaging about social responsibility could backfire and lead to less support for bailouts if agents have largely negative views regarding the behavior of large corporations. In this context, corporate communication might be more effective if it were to draw attention away from social responsibility issues altogether, and instead focus the messaging on alternative policy domains, such as one where bailouts would save jobs.

We test the predictions of our model in a new broadly representative large-scale experimental survey of 6,727 U.S. citizens that we designed and conducted online. The survey is specifically set up to study the link between corporate responsibility and the public support for corporate bailouts and related policies during the 2020 coronavirus crisis. Focusing on bailouts at a time of crisis provides an apt setting for our analysis, because the stakes are high, the public is engaged in the policy debate, and media, politicians, and corporations all play an active role in shaping the debate via extensive communication efforts. In terms of measurement, we collect both perceptions about corporate responsibility and policy preferences. In terms of experimental variation, we are able to vary both the salience of corporate responsibility and its framing by embedding into the survey variation in the ordering of the survey's questions and by showing respondents different animated videos, respectively.

The survey begins by asking about the socioeconomic background of the respondents. We then show respondents professionally developed videos that are pitched as a way to explain the main topics we ask about in the survey. A key section of our survey then measures perceptions of corporate responsibility focusing on ESG policies of the 500 largest

U.S. corporations. We ask about some of the most important dimensions of ESG that respondents can easily relate to, such as executive pay, employee benefits, tax strategy, gender diversity, CO_2 emissions, and political donations. We measure perceptions by asking respondents both what they think specific corporate policies *are* as well as what they think the policies *should be*. For instance, we ask how much respondents think top executives are paid and how much they think they should be paid. By comparing what respondents think policies are and what they think they should be, we can measure whether corporate actions meet public expectations. Our section on policy preferences measures respondents' stated support for bailouts of large corporations and for policies aimed at helping small businesses.

Our first key contribution is to document a strong baseline "big business discontent" spanning the full socioeconomic range. That is, all respondents perceive corporations to not be doing enough for society, relative to what they think the benchmark should be.

Having established a pervasive big business discontent, we next test the first prediction of our model. According to our model, the support for policies such as corporate bailouts should depend on the cued policy domain, i.e., the aspect of the policy decision that is emphasized and made salient through media and communications. In particular, if individuals are primed to think about bailouts through the lens of corporate responsibility, and individuals are highly negative about corporate responsibility, then our model predicts that we should see support for bailouts decrease. To test this hypothesis, we experimentally vary the salience of corporate responsibility with a simple design choice, namely by varying the order of the perceptions and the policy preferences sections of our survey. Specifically, we treat half of the respondents to think about corporate responsibility—thus increasing its salience—before stating whether they support corporate bailouts. We do so by asking the perceptions questions *before* the question about support for bailouts. The other half of the respondents are asked about perceptions of large corporations only *after* they disclosed their support for bailouts and other policies.

This simple salience treatment strongly reduces the support for bailouts. This finding lends direct support to our key prediction, namely that increasing the salience of corporate responsibility decreases the support for bailouts and pro-corporations policies in a context in which the public has well-established negative views about big business. This result echoes a finding by Alesina et al. (2022) in the context of immigration, where they show that given the very negative baseline views that respondents have of immigrants, simply making them think about immigration makes them support less redistribution.

We next study the second prediction of the model, namely that providing a negative or positive framing around ESG issues can also impact support for economic policies which concern large corporations. This happens because the distribution of experiences which individuals draw from when evaluating the policy changes. We study the effects of framing through a second source of experimental variation introduced by means of animated videos. The videos aim to vary the policy domain and the valence framing described by our model. Each respondent is either presented with a treatment video or a control video. The control video consists of basic explanations of the concepts we ask about in the survey (e.g., bailouts), and it is shown to all respondents. Our two main treatment videos highlight large corporations' role in society in a negative or a positive light, respectively, while still providing accurate information. For example, in our negative treatment video, we emphasize that there are fewer women relative to men in executive and board positions, or that companies are reluctant to cut CO_2 emissions. Analogously, in our positive treatment video, we emphasize that in recent years there has been a rise in the number of women in executive and board positions, and that several companies are now voluntarily reducing and disclosing CO_2 emissions. As a result, these videos increase the salience of corporate responsibility while also priming respondents to think about negative and positive aspects of this topic.[4]

As expected, support for bailouts decreases when respondents are shown the video that frames corporate responsibility negatively. Consistent with framing having an impact on policy preferences, respondents shown the positive video are significantly more likely to support bailouts relative to those exposed to the negative video. More surprisingly, support for bailouts is lower upon seeing the video that frames corporate responsibility positively relative to respondents seeing the control video. This finding suggests that the salience effect of making respondents think about corporate responsibility outweighs any positive effect that might come from framing corporations in a good light. In the language of our model, this indicates that the good ESG cue in the positive video triggers *interference* by bad ESG memories. This happens if respondents have sufficiently many negative memories of corporate responsibility, which we established earlier. These findings of "backfiring" highlight a key point of our analysis: if respondents are not fully rational but have limited memory recall, providing positive information about corporate responsibility might actually reduce the support for corporation-friendly policies because it leads respondents to recall memories on an issue where most of their memories are negative.

While we largely apply our model by making salient issues related to corporate responsibility, the model can be applied more generally. Indeed, the model suggests that raising the salience of alternative policy contexts related to bailouts should also impact policy preferences, perhaps in a positive direction. To test this possibility, we include a final video treatment, which primes respondents to think about bailouts in the context of economic and financial stabilization. Consistent with predictions of the model, we find that individuals exposed to this alternative salient policy domain increase their support for bailouts relative

[4]This is similar to the experimental variation in Alesina et al. (2022) regarding either hard facts about immigrants or the stereotype of immigrants being hardworking.

to the control group. The evidence thus strongly indicates that shifting the cued policy domain, either to ESG or economic stabilization, can materially impact policy preferences.

We conduct several tests to pin down the memory channel highlighted by our model, whereby selective memory recall shapes individual beliefs and ultimately policy preferences. First, we find that respondents shown the negative video do in fact have worse perceptions of corporate behavior than respondents shown the positive or control video. Moreover, we also find that both positive and negative priming influence perceptions of how corporations actually behave, as opposed to how corporations should behave. As our theoretical model suggests, this provides strong support for the idea that the video treatments, i.e. the cues, work primarily by influencing selective memory recall. We also find that priming respondents to think positively about corporations also significantly increases the big business discontent relative to receiving no communication. This finding provides further evidence as to how providing a positive framing can actually backfire under certain circumstances due to interference. That is, the frequency of negative ESG experiences and their similarity to the cue *interferes* with the recall of positive ESG experiences. Second, we show that the backfiring effect of the good video treatment is significantly stronger when we focus on liberal and young respondents, whom we would expect to have even more negative views of corporate responsibility. Providing a positive framing for these individuals backfires quite strongly, leading to significantly less support for bailouts than when provided no communications at all. Third, we conduct an additional survey to more directly pin down memory effects and what comes to respondents' minds after watching our treatment videos. Consistent with a selective memory retrieval channel, we find that both the bad and the good treatment videos lead respondents to recall more memories related to bad corporate behavior (such as the Enron fraud scandal) opposed to other types of memories. Finally, we provide evidence that contextual prompts influence what comes to mind using a textual analysis of open-ended answers by survey participants.

We perform a battery of robustness checks that are standard in the literature on information experiments (Haaland et al., 2020). For example, while experimenter-demand effects are unlikely to be consistent with the effects uncovered by our salience treatment or by our positive video treatment, we further alleviate them by showing that our findings largely persist even after a week. We also note that our main survey focuses on self-reported individual preferences for government policies. One concern is that people's responses to our survey questions might not be fully reflective of their true policy preferences. We addressed this concern directly by conducting an additional large-scale experimental survey, where we find strong and consistent findings when using a number of different behavioral outcomes, namely signing a real petition urging a bailout of large corporations and addressed to the

U.S. Congress, emailing U.S. senators to express either support or opposition to bailouts, and donating money to the Business Roundtable.

1.1. **Literature Review.** Our paper highlights the presence of a dynamic relationship between corporate behavior and the public support for large corporations, thereby contributing to a large and growing body of work on corporate responsibility. Several recent studies in this literature argue that corporations might seek to maximize the welfare of all their stakeholders, not only shareholders (Bénabou and Tirole, 2010; Freeman et al., 2010; Edmans, 2011; Hart and Zingales, 2017; Broccardo et al., 2020). A related strand of papers provides evidence that social capital, trust, and culture matter for resource allocation and firm outcomes (Guiso et al., 2004, 2006, 2015a,b).[5] Yet, as outlined in the review by Kitzmueller and Shimshack (2012), the nexus between corporate behavior and policy outcomes has received little attention and remains generally "poorly understood." Thus, our paper makes a clear contribution to this literature by providing a direct link between the public perception of corporate behavior and the support for economic policies.

We further contribute to a burgeoning literature which uses online experimental survey designs to understand what shapes individual policy preferences, recently summarized by Haaland et al. (2020). Much of the work in this area has focused on identifying preferences for taxation and redistribution.[6] The two papers closest to ours are Alesina et al. (2022), who study how views over immigration impact support for redistribution, and Andre et al. (2022), who explore how views over the effects of macroeconomic shocks are impacted by prior experiences and contextual clues.[7] Both of these papers use experimental survey designs in conjunction with priming to conduct their analyses. Andre et al. (2022) use random variation in narrative vignettes on different economic shocks to prime individuals to think about alternative economic mechanisms, while Alesina et al. (2022) use random variation in the ordering of the survey's questions, similar to our salience treatment, to prime individuals to think about immigration.[8] Our contribution to this literature is threefold. First, we develop a memory-based framework to provide a formal basis for understanding how priming can be used to elucidate the relationship between domain-specific beliefs and policy

[5]For instance, Lins et al. (2017) find that high levels of public trust are particularly valuable for corporations during crises. See also Lins et al. (2020), Cororaton and Rosen (2020), and Albuquerque et al. (2020).

[6]Cruces et al. (2013), Kuziemko et al. (2015)), and Karadja et al. (2017) use randomized information treatments on the true income distribution to study how biased perceptions of income inequality impact demand for redistributive policies. Weinzierl (2017) argues that individuals resist the equalization of after-tax incomes that depend on luck. Alesina et al. (2018) argue that preferences over redistribution are impacted by views regarding intergenerational mobility, while Fisman et al. (2020) use a survey to study individual preferences over wealth taxation.

[7]Relatedly, Haaland and Roth (2017), Barrera et al. (2020), and Grigorieff et al. (2020) use randomized information treatments to study the determinants of support for immigration.

[8]For discussions of order effects and priming in surveys see, among others, Malhotra (2008), Lenz (2009), Johnston et al. (2017), and Krosnick (2018).

preferences. Second, we apply this conceptual framework towards a novel application area, showing that priming individuals through cues to think about corporate responsibility lowers support for corporate-friendly policies such as bailouts, due to robustly negative preexisting attitudes. Finally, we use our conceptual framework to show that when preexisting domain-specific attitudes are sufficiently negative, positive messaging and framing on those issues can actually backfire in terms of policy support, since it primes individuals to view the policy through that particular domain-specific lens and because negative memories interfere with the recall of positive memories.

In constructing our conceptual framework, we build our analysis on a growing literature regarding associative memory recall in economics. Kahana (2012) provides an overview of the theoretical frameworks for human memory in psychology. Our model follows closely the approach of Bordalo et al. (2020a), Bordalo et al. (2020b), and Bordalo et al. (2021), who develop models that highlight how selective memory recall and cues affect decision-making.[9] In particular, Bordalo et al. (2020a) show how cues and selective memory recall can impact probability judgments using the Tversky (1977) similarity metric, while Bordalo et al. (2020b) show how the selective recall of contextually similar past buying experiences impacts views on good quality and consumer choice. Enke et al. (2020) provide lab-based experimental support for these theories, showing that when identical signals about hypothetical companies are communicated with identical, but uninformative contexts, individuals systematically overreact to the signals. We apply a memory-based framework of selective recall to the setting of communications and show how the form and structure of a corporate or political message can impact policy preferences through the dual effects of priming a policy domain and priming a particular valence frame. We further contribute to this literature by using a randomized survey design to provide support for the effects of associative memory recall on individual decision-making and policy preferences.[10]

We finally contribute to a literature which explicitly looks at the economics of persuasion. The literature on persuasion has mostly focused on rational agents (see, e.g., Stigler, 1961; Crawford and Sobel, 1982, for early contributions) and uses "Bayesian persuasion" as the canonical model of communication (Kamenica and Gentzkow, 2011). However, a

[9]Other important recent contributions include da Silveira et al. (2020) and Dasgupta and Gershman (2021), who study how to optimally make decisions and process information in the presence of memory constraints. Wachter and Kahana (2019) show how memory-based recall can influence portfolio choice and asset pricing.

[10]Our work is also related to the case-based decision theory of Gilboa and Schmeidler (1995), who construct a decision theory in which, when facing complex decision problems in which all the outcomes and associated probabilities cannot be enumerated, individuals rely on past similar cases and the utility associated with the case's result. Analogously, agents in our model use past experiences in their memory database to evaluate whether to support a policy or not. Relative to this prior work, we emphasize how cues might influence selective memory recall, discuss how interference impacts the cue's effects, and apply our model in an experimental setting to study how perceptions of corporate social responsibility are shaped and influence policy preferences.

growing body of work, to which we contribute, focuses on the novel consequences of deviations from rationality within this context (Mullainathan, 2002; Mullainathan et al., 2008; DeMarzo et al., 2003). Mullainathan (2002) constructs a model featuring "associativeness", in which new events not only convey new information, but also impact what past information is recalled. Our work also features similarity-based recall, in that the details of the cue influence what experiences are recalled from the mental database. Building on these insights, and closest to our work is Mullainathan et al. (2008), who formalize notions of coarse thinking to highlight that communication can influence perceptions not only by providing information but also by changing the categorical lens through which the receiver of the information analyzes a given issue. Our behavioral model formalizes a similar theoretical result through memory-based microfoundations, and further shows how providing a positive framing or positive information can in fact backfire if it primes a policy domain in which agents have particularly negative views. We provide support for these predictions with novel experimental evidence.[11][12]

Our paper proceeds as follows. Section 2 introduces our model. Section 3 discusses the experimental survey. Section 4 provides a descriptive analysis of our data. Section 5 reports the main results from our experiments. Section 6 shows our analysis of behavioral outcomes and several robustness checks. Section 7 concludes.

2. Model

We study a setting where individuals form policy preferences on the basis of highly salient issues and where political and corporate communication strategies may shape such preferences through persuasion. We develop a simple memory-based model of cognitive thinking to conceptualize this setting, building on a widely adopted psychology model of associative memory recall by Kahana (2012). Our model recognizes that, in contrast to standard economic models in which a rational agent can retrieve all past relevant experiences, in practice memory recall is limited and guided by contextual prompts. In this framework, salient issues and messages are cues that retrieve recall of experiences based on similarity, subject to interference.

As a motivating example, suppose that an individual is asked to think about corporate bailouts, which will be the primary policy setting for our experimental survey. This is

[11]We also relate to a large literature on agenda setting in communications research, with the seminal study by McCombs and Shaw (1972) arguing that the media sets the agenda for what citizens focus on when thinking about policy. This work in the communications literature finds a counterpart in the work on media in economics. We refer to DellaVigna and Gentzkow (2010), Napoli (2014), DellaVigna and La Ferrara (2015), Strömberg (2015), and Enikolopov and Petrova (2015) for reviews of the literature.

[12]Our model is also related to a number of studies that focus on the specific role of attention (Enke and Zimmermann, 2019; Hartzmark et al., 2019; Enke, 2020) and the importance of prior experiences and emotions in financial decision-making (Kuhnen and Knutson, 2011; Rudorf et al., 2014; Kuhnen, 2015).

a complex policy and thus individuals' preferences over bailouts could be over multiple multifaceted aspects. For example, for reasons of fairness, individuals may not want to extend corporate bailouts if they believe large corporations act in a manner which hurts the interests of their various stakeholders, such as workers and society at large. In contrast, individuals may see the value in extending corporate bailouts to the extent they believe bailouts would stabilize the economy or financial system. Within this context, politicians and corporations, or media more broadly, can engage in persuasion by making certain aspects of the policy issue more salient or by framing certain aspects in a positive or negative light.

The way we formalize how individuals form policy preferences is through the agent recalling past experiences or news related to a given *policy domain*, and considering whether in that particular circumstance the policy would lead to a perceived positive outcome. For instance, if cued to think about bailouts through the lens of corporate responsibility, the agent may recall events such as the Enron accounting scandal, in which presumably most individuals would be less likely to support a bailout. In contrast, if cued to think about bailouts through the lens of economic stabilization, the agent may recall the Troubled Asset Relief Program (TARP) initiative during the financial crisis and the arguments surrounding it about restoring economic growth.

In addition to providing a policy domain, the cue can also offer what in the communications literature is known as framing, which we refer to as the *valence frame*. Other than setting the policy domain, the cue may include positive language or offer a narrative or specific examples for why a policy should be viewed favorably. This would constitute a positive valence frame. For example when cueing the agent to think about corporate bailouts through the lens of corporate responsibility and stakeholder capitalism, the prompt might also discuss these issues in a positive light by highlighting corporate social activism, environmentally aware corporate policies, or employee welfare programs. In contrast, the cue could instead offer a negative frame for the policy context. We parsimoniously model the effects of valence framing in a manner analogous to the policy domain. For a given policy domain, a positive valence framing will lead the agent to disproportionately recall those experiences from the mental database with a high valence. A negative valence framing will lead the agent to disproportionately recall those experiences with a low valence. For example, a positive frame will make it less likely that an agent recalls the Enron accounting scandal from the mental database.

2.1. **Mathematical Model.** We assume that individuals assess complex and multifaceted policy decisions by evaluating past memories. When confronted with a particular policy, agents recall past experiences and ask whether the given policy would have led to a high or low utility in that experience. If most of the recalled experiences are associated with a high utility, the agent will support the policy. The key for the following model is that agents are not fully rational and cannot recall all memories. Rather, the set of memories recalled can be manipulated through salience and messaging.

The memories of agent i are stored in a database of memories M_i. Different individuals can have different memory databases, but in what follows, since our focus will be on the memory-based recall of a single individual, we will suppress the dependence on i. Memory databases are comprised of a set of experiences $e_k \in \mathcal{E}$, where $1 \leq k \leq N$ indexes a particular experience and $\mathcal{E}$ denotes the universe of experiences. We take such experiences to be widely construed, reflecting either policy-relevant personal events or relevant pieces of information received through various forms of communication, e.g., through interacting with others or by engaging with news and media.

We will assume experiences have two relevant characteristics. The first is a policy domain $c_k \in \mathcal{C} \subset \mathbb{R}^m$, where $m > 0$ and $|\mathcal{C}| \in \mathbb{N}$ is finite. The second characteristic is a policy valence $u_{p,k} \in \{H, L\}$, which measures whether the policy p would have led to a high (H) or low (L) utility in the hypothetical case it was implemented in memory k. We measure the similarity between any two experiences in the mental database as:

$$S\left(e_k, e_{k'}; w_c, w_u\right) = \delta^{w_c|c_k - c_{k'}| + w_u 1\left[u_{p,k} \neq u_{p,k'}\right]},$$

where $0 < \delta < 1$ and $w_c > 0, w_u > 0$ determine the strength of associative memory recall.

Concretely, in our motivating example of corporate bailouts, relevant experiences could live either in a ESG policy domain or an "economic stabilization" policy domain. Experiences in the ESG policy domain could carry either a positive valence or negative valence. For example, experiences or memories of corporate charitable giving would likely have a positive valence and lead agents to view bailouts more favorably. In contrast, recall of harmful environmental practices or accounting scandals would carry a negative valence and lead agents to view bailouts less favorably.

When asked to think about a given policy, agents are also given a cue. This cue influences the probability that certain memories are recalled. In this framework, the cue influences the probability that a given memory is recalled through a similarity function, which is discussed further below. The more similar a cue is to a given memory, the more likely it is that this memory is recalled.

In particular, we think of a cue and communications as priming the recall of a set of experiences similar to the cue. Formally, we assume that a cue $\Gamma^* = (\Omega^*, \zeta^*)$ includes a set $\Omega^* \subset C \times \{\emptyset, L, H\}$, with each element of the set comprising a policy domain and (possibly) an associated valence framing. The empty set indicates that no valence framing was cued for a given constituent member. The cue also includes a parameter $\zeta^* \in \{0, 1\}$ which denotes the strength of the prime and impacts the degree of selective recall. This formulation of a cue is similar to Bordalo et al. (2021). Then, the similarity of an experience to the cue is defined as the average pairwise similarity between the experience and the cue's constituent members in Ω^*, with the degree of selective recall controlled by ζ^*. Specifically:

$$S\left(e_k, \Gamma^*\right) = \sum_{e_{k'} \in \Omega^*} S(e_k, e_{k'}; w_c^{\zeta^*}, w_u^{\zeta^*}) \pi(e'_k|\Omega^*), \tag{2.1}$$

with $w_j^1 > w_j^0$ for $j \in \{c, u\}$. That is, stronger cues lead to greater selective recall. Given a similarity between an experience e_k and the cue $\Gamma^* = (\Omega^*, \zeta^*)$, the probability $\Pi\left(e_k, \Gamma^*\right)$ that memory e_k is recalled is given by:

$$\Pi\left(e_k, \Gamma^*\right) = \frac{\pi\left(e_k\right) S\left(e_k, \Omega^*; w_c^{\zeta^*}, w_u^{\zeta^*}\right)}{\sum_{k'} \pi\left(e_{k'}\right) S\left(e_{k'}, \Omega^*; w_c^{\zeta^*}, w_u^{\zeta^*}\right)}, \tag{2.2}$$

where $\pi\left(e_k\right)$ is the true proportion of experience k in the mental database.

In what follows, when a cue primes a single policy domain c^* and no valence framing or a single valence framing through Ω^*, we denote the cue by $\Gamma^* = (c^*, u^*, \zeta^*)$ with $u^* \in \{\emptyset, H, L\}$. That the cue includes a strength ζ^* will be useful in our experimental work, in that some individuals are essentially primed twice to think about a given policy domain, and thus likely receive a stronger cue than those primed only once.

For instance, if media and communications prime agents to think about issues related to corporate responsibility, a specific policy domain, then agents will be more likely to recall experiences such as the Enron accounting scandal or corporate charitable initiatives, which are highly reflective of that domain. This framework captures the key psychological underpinnings of memory recall, as discussed by Kahana (2012). First, more frequent experiences are easier to recall than less frequent experiences. Second, it is easier to recall those experiences which are more similar to the cue than those experiences which are dissimilar to the cue. Finally, the denominator in equation (2.2) captures the idea of *interference*. Experiences $e_{k'}$ which are similar to the cue, but not in it, may intrude in memory recall and interfere with the ability to recall a cued experience e_k. In this way, agents cannot fully control what they recall. Experiences not in the cue may intrude in memory recall due to frequency and similarity. For instance, the cue may prime agents to recall positive experiences of ESG, such as charitable giving. However, examples of negative ESG, such as accounting scandals, may interfere with such recall since they live in the same policy domain and thus are similar, though not identical, to the cue. The greater the proportion such negative ESG memories constitute of the memory database, the larger such interference will be.

We assume that when accessing the mental database, the individual makes $T \geq 1$ draws, sampling with replacement.[13] Let $R_H\left(\Gamma^*\right)$ and $R_L\left(\Gamma^*\right)$ denote the number of draws of experiences with positive and negative policy valence, respectively. We assume the agent will

[13]If sampling occurs without replacement, then communications and priming can themselves impact the structure of the mental database. This is an interesting question with potential dynamic ramifications that we leave to future research.

support the policy if the number of positive valence draws exceeds the number of negative utility draws, that is if $R_H(\Gamma^*) > R_L(\Gamma^*)$.

In this way, by priming only a segment of the mental database, the cue can lead the agent to disproportionately recall experiences within a certain policy domain or of a certain valence. To the extent that the proportion of positive and negative valence experiences varies across policy domains, communications and priming can thus shift policy preferences. In the communications literature, this effect is known as second-order agenda setting, whereby agents are primed to think about complex and multifaceted policy issues through a specific lens, which in this case would be corporate responsibility. This is formalized in the following result:

THEOREM **2.1.** *Let $\pi(H|c^*)$ denote the fraction of positive valence experiences given domain c^*. Suppose that $\pi(H|c^*) < \pi(H|\tilde{c}^*)$ for $c^* \neq \tilde{c}^*$. Let $\Gamma^* = (c^*, \emptyset, \zeta^*)$ and $\tilde{\Gamma}^* = (\tilde{c}^*, \emptyset, \zeta^*)$. Then for $w_c \geq 0$ sufficiently large, $E[R_H(\Gamma^*)] < E\left[R_H\left(\tilde{\Gamma}^*\right)\right]$.*

Proof. See Appendix A.7. ∎

While mathematically simple, this result has significant economic content. Consistent with the intuitions provided above, it shows that by cueing the agent with a particular policy domain, the agent will draw a selected set of experiences from the mental database. If the fraction of positive valence experiences varies across different policy domains, then cueing the agent can impact policy preferences. This is in stark contrast to a model with rational agents, where priming an agent with a particular policy domain would not lead the agent to neglect or "forget" relevant information from another policy domain. For example, if agents have more positive views of bailouts in the economic stabilization policy domain than in the ESG policy domain, then a cue which primes agents to view the policy through the lens of economic stabilization would lead to greater support for bailouts than a cue which emphasizes issues related to corporate social responsibility.

The cue might also offer a positive or negative narrative or offer policy-relevant positive or negative examples, which we model as the cue being comprised of positive or negative valence experiences. This will prime the agent to disproportionately recall positive valence or negative valence experiences when they think about the policy. Thus, for example, if agents are primed to think about corporate responsibility, and moreover the cue provides a positive narrative surrounding these issues, then the agent will be more likely to recall instances of corporate charitable initiatives and less likely to recall the Enron accounting scandal more than if such a positive narrative had not been provided. This is because positive valence experiences will be more similar to the cue than negative valence experiences. In the communications literature, this is known as framing and is considered a distinct phenomenon from second-order agenda setting. Here, we show how both second-order agenda setting and framing can arise from the same cognitive basis.

Crucially, however, note that even with a positive framing, negative valence experiences within a policy domain can still interfere with the recall of positive valence experiences. The more frequent the negative valence experiences are in the mental database, the larger this interference will be. This psychological interference leads to a surprising result regarding communications. Providing a positive narrative in a specific policy domain can actually backfire, in the sense that agents would have been more supportive of the policy if no communications had been provided. We have the following formal result:

Theorem **2.2.** *Suppose that* $\pi(H|c^*) < \pi(H|\tilde{c}^*) < 1$ *for* $c^* \neq \tilde{c}^*$. *Let* $\Gamma^* = (c^*, H, \zeta^*)$ *and* $\tilde{\Gamma}^* = (\tilde{c}^*, \emptyset, \zeta^*)$. *Then for* $w_c \geq 0$ *sufficiently large, there exists* $\bar{w}_u > 0$ *such that* $E[R_H(\Gamma^*)] > E\left[R_H\left(\tilde{\Gamma}^*\right)\right]$ *for* $w_u > \bar{w}_u$ *and* $E[R_H(\Gamma^*)] < E\left[R_H\left(\tilde{\Gamma}^*\right)\right]$ *for* $w_u < \bar{w}_u$.

Proof. See Appendix A.7. ■

This result shows that—due to interference—when there are many negative valence experiences in the mental database within a policy domain, positive framing may be insufficient to drive individuals to view the policy more favorably than if the cue had not been provided. This result has significant implications for corporate and political communication strategies, especially if positive framing cannot be separated from setting a policy domain. Suppose, for example, that a corporation or political actor wanted to frame corporate behavior in a positive light as part of their pitch for bailouts. In a context where individuals consider corporations as uncaring of the needs of society at large, this could potentially backfire if the positive framing led agents to focus on fairness aspects, as opposed to economic aspects, of bailout policy.[14]

3. The Experimental Survey

In this section, we describe the empirical methodology we adopt. We focus our attention on the specific details of our main and largest survey in subsections 3.1, 3.2, and 3.3.

3.1. **Data Collection.** We launched our first experimental survey on May 5, 2020, in the midst of the policy discussion regarding how to implement corporate bailouts in response to the COVID-19 crisis. We received 91% of survey responses within one week, and the survey was closed after one month.

[14]In Appendix Section A.8, we provide an extension of the model in which we think of the cue as providing new *information* to individuals, relative to their existing memory database. In particular, we show that providing new positive information can still backfire in such a case. If positive information primes a policy domain in which individuals have a large stock of preexisting negative experiences, then the negative effects of priming the policy domain will dominate the effects of the new positive information, leading to lower policy support. On the other hand, if the memory database is sufficiently thin in that policy domain, new information can have a large effect and lead to increased policy support even if experiences are largely negative.

We designed the surveys using an online platform, and the survey links were then distributed by our data collection partner Dynata to a sample of U.S. citizens over 18 years old. Respondents are targeted to be representative of the U.S. population along the dimensions of gender, age, income, race and ethnicity, education, employment status, and political views. We collected a total of 6,727 survey responses. The median (average) time for completion of the survey was 11 (20) minutes. To test the persistence of the effects, we also conducted a follow-up survey—one week after the original survey—of approximately one-third of the sample for a total of 2,311 follow-up survey responses.

In Table 1, column 1, we report summary statistics on the socioeconomic background of our survey respondents. Going from top to bottom of the table, we can see that 51% of the sample are female, 30% are 35 years old or younger, 52% have a total household income of $70,000 or higher, 70% are white, 57% have completed a four-year college, or higher, degree, 61% are either business owners or employed full-time or part-time, and 31% see themselves as liberal or very liberal. In column 2 of Table 1, we report the same shares computed using the 2019 U.S. Current Population Survey (CPS). Our sample is largely representative of the U.S. population, with the exception of being more highly educated (57% vs 42%) and having a slightly lower percentage of individuals who are white (70% vs 78%).[15] We further report the geographical distribution of our respondents in Figure A1.

We discuss various techniques used to ensure we collect high-quality data in Section A.6.

3.2. **Survey Structure and Measurement.** We now provide a brief description of the survey, the structure of which is visually illustrated in Figure 1. We report the full text of the survey in Section A.2. Most questions in the survey are about large corporations and their shareholders and stakeholders, and the primary outcomes regard corporate bailouts. To make these concepts clear to all respondents, we define them in the survey. We ask the respondent to think of large corporations as the "*top 500 U.S. corporations*." We also provide a brief definition of the concepts of corporate bailouts, shareholders, and stakeholders.

After a brief introduction and consent form, the survey is organized into four main sections, covering the socioeconomic background of the respondent, the animated videos, the measurement of perceptions of large corporations, and the support for economic policies. We discuss each section of the survey in more details below.

3.2.1. *Socioeconomic Background.* The first section asks about the socioeconomic background of the respondent. We collect information on gender, age, total household income in

[15]Despite these being minor imbalances, we later show in the analysis that our results are essentially unchanged when we re-weight the sample so that it is representative along the education and race/ethnicity dimensions as well.

2019 (before taxes), race and ethnicity, the highest level of education, and current employment status. We additionally measure political orientation, by asking the following question: "*On economic policy matters, where do you see yourself on the liberal/conservative spectrum?*." The options given are: "Very Liberal, Liberal, Moderate, Conservative, Very Conservative."

3.2.2. *Animated Videos.* The second section of the survey consists of professionally animated videos we created to generate specific sources of experimental variation, as discussed in Section 3.3.

3.2.3. *Perceptions of Large Corporations.* A central part of our study consists of measuring individual perceptions of large corporations' impact on society at large. In particular, we measure perceptions of corporate policies related to environmental, social, and governance (ESG) issues. ESG covers a range of topics, from climate change, waste, pollution, and deforestation, to employee relations, working conditions, and engagement with local communities, as well as executive pay, tax strategy, political donations, corruption, and board diversity. Central to measurement is the conflicting tension between what is best for stakeholders as opposed to maximizing value for shareholders.

We measure perceptions by asking respondents both what they think specific corporate policies *are* as well as what they think the policies *should be*. The difference between such measures captures how "good" or "bad" large corporations are in the respondents' eyes from an environmental, social, and governance standpoint.

To keep the survey reasonably short so as to ensure high-quality data, we measure perceptions along some of the most important corporate responsibility dimensions that respondents can easily relate to and that we can reliably measure. We ask six topic-specific questions—executive pay, employee benefits, tax strategy, gender diversity, CO_2 emissions, and political donations—and one more abstract question—shareholders vs stakeholders. See Appendix A.2 for the specific questions and answer options.

The questions are phrased to be intuitive for the respondents. For this reason, we deliberately chose not to monetarily incentivize respondents's perceptions, as that would have required a considerably more complicated framing of the survey questions.[16] To further ease the readability of our analysis and results, we perform two basic transformations after we collect the data. First, we standardize all variables to be on a scale of 0-100.[17] Second, we

[16]As shown by Grewenig et al. (2020), monetary incentives may also trigger an individual to Google for a specific statistic, which might confound our measurement exercise, a point also raised by Roth et al. (2020).
[17]This transformation only affects the executive pay and the shareholders vs stakeholders questions. While the latter is simply multiplied by 10, the executive pay variable is standardized by assuming a linear increment with each higher value of the response. That is, the original variable takes value 1 if the response is "the same," it takes value 2 if "twice as high," and so on, up to taking value 6 if "500 times as high." We then standardize the variable by multiplying it by 100 and dividing by 6.

transform all responses so that a higher value can be interpreted as "worse" from an ESG perspective.[18]

Once all transformations are applied, we have separate measures of how bad (or good) individuals think large corporations *Are* and *Should Be* along several dimensions. Importantly, the difference between the *Are* and the *Should Be* responses (*Are* − *Should Be*) tells us how much large corporations fail to live up to the standards required by each individual respondent. This measure of difference in perceptions represents the main perception measure we use in our study. We label this measure the *big business discontent*.

3.2.4. *Support for Economic Policies.* We measure the support for economic policies with a focus on corporate bailouts and financial assistance to small businesses, which were both at the center of the policy debate at the time of our survey.

Our main question captures the support for bailouts of large corporations. Specifically, after defining once again the concept of corporate bailouts and re-emphasizing the focus on large corporations, we ask: "*On a scale from 0 to 10, where 0 means "do not support at all" and 10 means "strongly support," how would you rate your support for corporate bailouts?*"

Additionally, we ask a similar question to gauge respondents' support for similar policies aimed at helping small businesses, rather than large corporations. That is, before asking the questions, we state: "*The government also considers providing money directly to small businesses. By small businesses, we mean businesses with less than 100 employees, such as local retail stores, restaurants, and coffee shops.*"

3.3. **Experimental Variation.** Motivated by our conceptual framework developed in Section 2, we introduce into our survey two layers of randomization aimed at inducing experimental variation in the precise cue provided to survey participants. In this way, we seek to address how messaging and priming can influence policy preferences, as discussed in our Theorems 1 and 2. With the first layer, we prime a subset of the agents to think about bailout policy through the lens of a given policy domain, namely corporate responsibility or economic stabilization. With the second layer, in addition to setting a policy domain, participants are cued with either a positive or negative valence framing, primarily through a narrative which would prime agents to view corporate responsibility issues favorably or disfavorably.

We generate four main treatment groups and two main control groups. The same questions are asked to all respondents. We obtain variation by randomly varying the order of sections—layer 1—and by exposing respondents to different videos—layer 2. We illustrate

[18]As a result, the executive pay and political donations responses remain as they are, while all other responses are subject to the transformation 100-X. For example, a corporation that pays the CEO 500 times as much as the average worker is considered less-ESG friendly than a corporation with a lower CEO/worker pay ratio. However, a corporation that has more women in the top management is considered more ESG-friendly, and as a result the variable is transformed.

the experimental design, as well as the total number of observations in each treatment and control group, in Figure 1. Each of the eight possible groups is associated with a different cue $\Gamma^* = (c^*, u^*, \zeta^*)$ using the language of our theoretical framework in Section 2, in that each group is subject to a different combination of the policy domain c and valence framing u. This taxonomy showing the relationship between the treatments and the model specific cue is given in Table 2. In what follows, we detail each of these distinct treatments and the precise cue they coincide with.

3.3.1. *The Salience Treatment: Randomizing the Order of Questions.* The first layer of randomization consists of varying the order of Section 3 and Section 4 in the survey. Section 3—*Perceptions of Large Corporations*—questions respondents about corporate policies while Section 4—*Support for Economic Policies*—provides the policy preferences. Half of respondents see Section 3 *before* Section 4, while the other half see Section 3 *after* Section 4. The randomization is stratified by the second layer of the randomization, so that the half-half split holds within each of the video treatments, as shown in Figure 1.

The goal of this randomization is to randomly provide a subset of the survey participants with the cue $\Gamma^* = (\text{ESG}, \emptyset, 0)$. That is, this randomization allows us to prime a specific policy domain by focusing the attention of the respondents on corporate responsibility in a neutral way, i.e. without providing any valence framing. To the extent that individuals are indeed influenced by such communications, whereby the cue interferes with recall of experiences from alternative policy domains, providing such a cue may shift individuals' preferences regarding policies in which views about corporate responsibility play a role. Assuming individuals have deep-rooted beliefs about big business, this should be especially true in the context of large corporations' bailouts.

3.3.2. *The Animated Video Treatments.* The second layer of randomization consists of splitting the main sample into four groups, based on which video respondents are shown in Section 2 of the survey. All videos have been professionally scripted and developed, and they are similar to the animated videos seen in a variety of contexts, from marketing and advertisement campaigns to educational videos.[19] The full scripts of all videos are reported in Section A.4, and several screenshots are displayed in Figures A2, A3, and A4.

The baseline video is a *control video*, which consists of a brief introduction to the survey and to how to answer specific questions, such as those involving percentages and sliders. It also defines specific concepts that appear in the surveys, namely "large corporations," "corporate bailouts," and the difference between "shareholders" and "stakeholders" of a corporation. The control video is a subset of (and therefore slightly shorter than) all three treatment videos, which in turn start with the control video before adding the additional

[19]The full set of videos can be watched on the authors' websites (see http://emanuelecolonnelli.com).

content. We interpret this control video alone as providing no cue. Thus, survey participants who receive the salience treatment and the control video receive the cue $\Gamma^* = (\text{ESG}, \emptyset, 0)$. Survey participants who do not receive the salience treatment and receive the control video receive no cue, i.e., $\Gamma^* = \emptyset$.

The first treatment video—*T-Bad*—provides a (negative) valence framing of large corporations' behavior from an ESG standpoint. The video also naturally cues a policy domain as a byproduct by focusing the attention of survey participants on corporate responsibility issues. Specifically, we organize the animated video around the goals of corporations with a focus on the tension between maximizing value to shareholders or stakeholders. This is a standard way of thinking about corporations' impact on society. Corporations who only care about maximizing shareholders' profits are seen as the least friendly to society, while those who also care about their employees, society, the environment, and diversity and equality in the workplace, among other issues, are seen as having a more positive impact on society. For example, the video says: "Companies also have an obligation to promote a diverse and equal society. Yet they hire and promote very few women compared to men in executive and board positions. This will likely make it more difficult for other women to reach the top and reinforces the stereotype that men are better at doing business." Considering that executives are primarily men, the framing provided is accurate, yet the overall communication and language used place corporate behavior in a *negative* light and thereby prompts the recall of instances of bad corporate behavior. Thus, relative to the salience treatment, the T-Bad treatment primes the same policy domain but also crucially provides a negative valence framing.

In the context of the model, for survey participants who do not receive the salience treatment, but receive the T-Bad video, the cue is specifically $\Gamma^* = (\text{ESG}, L, 0)$. The cue primes both a policy domain and the valence of the experiences to draw. For survey participants who receive both the salience treatment and the T-Bad video, we denote the cue as $\Gamma^* = (\text{ESG}, L, 1)$. Here, the agent is primed twice to think about bailouts through the lens of corporate responsibility in addition to receiving the negative valence framing.[20]

Our second treatment video—*T-Good*—is similar to the T-Bad video in that it aims to provide a specific valence framing of corporate responsibility. However, while having the same structure and covering the same topics, T-Good places the approach of corporations in a *positive* light. For example, the video says: "Companies also have an obligation to promote diversity in the workplace. Over the last years, we have indeed seen a tremendous rise in the number of women in top management and in the boardroom." Again, the narrative provided in this video is accurate, given a recent trend towards increased diversity in executive

[20]Through the lens of the model, we interpret the effect of double priming the same cue as a potential increase in the parameters w_c and w_u. That is, it potentially increases the effects of associative memory recall.

positions, but the overall information delivery and language employed is designed to place corporate behavior in a positive light.

In the context of the model, for survey participants who do not receive the salience treatment, but receive the T-Good video, the cue is $\Gamma^* = (\text{ESG}, H, 0)$. For survey participants who receive both the salience treatment and the T-Good video, the cue is $\Gamma^* = (\text{ESG}, H, 1)$, since the ESG policy domain is primed twice.

Finally, we design a third and final treatment video—*T-Economy*—aiming to provide an altogether alternative policy domain for the policy decision. The main difference of this video relative to the control video is the addition of a scene conveying that corporate bailouts are likely needed for the economy to recover, a view many experts shared regarding the COVID-19 crisis at the time of the survey. The scene reads as follows: "Leading economists of all political views, from liberal to conservative, mostly agree that corporate bailouts will likely help the economy." In this way, the video primes survey participants to view bailout policy through the lens of economic stabilization, as opposed to corporate responsibility. With no salience treatment, the cue is $\Omega^* = (\text{Econ}, \emptyset, 0)$. With the salience treatment, the full cue is $\Omega^* = (\text{ESG+Econ}, \emptyset, 0)$, since two distinct policy domains are primed.

3.3.3. *Balance Checks.* A key assumption for our experimental design to be valid is that there is no statistical difference between treatment and control groups. We report the balance tests in Table 1, which shows that the characteristics of respondents in any of the treatment groups are essentially indistinguishable from those of respondents in their respective control groups. Columns 3-6 of Table 1 report the results from univariate regressions of an indicator variable for each treatment group on the main demographics we collect, namely gender, age, income, race and ethnicity, education, employment status, and political views. Columns 7-10 of Table 1 report a similar analysis where the demographic characteristics are included together in the same regression. The results in the table display the randomization was effective, as there are extremely few coefficients that are statistically significant—all of which are small and marginally significant—across the several specifications.

4. Descriptive Analysis

In this section, we provide a brief descriptive analysis of the data we collect on perceptions of large corporations and support for economic policies.[21] While the descriptive facts are interesting per se, the goal of this section is to establish the presence of a strong and widespread big business discontent, and to show there exists an association between what individuals think about large corporations and their policy preferences.

[21] All the tables and figures discussed in this section are constructed from the sample of respondents included in the control video group of the main study (May 2020) to ensure that our descriptive analysis in unaffected by the treatment. The only exception is the correlation table between perceptions and outcomes, which relies on the full sample of the main study (May 2020).

4.1. **A First Look at Individual Perceptions of Large Corporations.** Table 3 shows what respondents think environmental, social, and governance policies of large corporations currently are and what, in their mind, these policies should be. All numbers are reported after we apply the transformations discussed in Section 3.2, such that a higher value corresponds to less ESG-friendly policies.[22]

A clear pattern emerges, highlighting a key motivating finding of our paper: respondents think corporate policies are less friendly to society than they should be. This can be seen in column (7), which reports the difference between what respondents think corporate policies *are* and what they think corporate policies *should be*, which is our measure of big business discontent. The big business discontent is positive and highly significant for all measures, indicating that respondents think large corporations are not doing enough along a multitude of attributes.

We find the largest big business discontent in the questions about political donations and the environment. For example, respondents think 69.79% of large corporations donate money to politicians, but they think fewer than 30% of corporations should make political contributions. Similarly, respondents think that 40% of corporations disclose CO_2 gas emissions, but they believe 70% of companies ought to. All other ESG attributes also generate a discontent in the respondents, as they believe top executives and managers should be paid less (first row of Table 3), corporations should pay a larger fraction of employees' health care costs (second row) and more in federal income taxes (third row), and that there should be an equal gender distribution among top managers and executives (fourth row). The answer to the broader question of shareholder vs stakeholder maximization displays a similar pattern. We further report the full distributions of responses to our perceptions questions using histograms, as shown in Figures A5 (big business discontent), A6 (are), and A7 (should be) in the Appendix.

Figure 2 reports the big business discontent for different subgroups of survey respondents. We also separately report the two components of big business discontent—what respondents think corporate policies are and what they think they should be—in Appendix Figures A8 and A9 respectively. A key observation from Figure 2 is that all policies are perceived as being insufficiently friendly to society in each subgroup we consider.

4.2. **Individual Perceptions and Policy Preferences.** Our main hypothesis is that perceptions of corporations' role in society influence public support for government policies. In our survey, we place a special emphasis on policies related to corporate bailouts, which were at the center of public debate during the time of our surveys.

[22]At times, we implicitly mention the non-transformed variable, when it is descriptively more meaningful.

Figure 3 shows the differences in support for bailouts across subgroups of the population. The bottom panel of Figure 3 illustrates respondents' support for government initiatives providing money directly to small businesses. In general, the support for these "small business bailouts" is significantly higher than the support for bailouts of large corporations across all groups of respondents.[23]

We report in Table 4 the results of a regression of support for corporate bailouts and small businesses onto big business discontent along the different corporate policies we measure. In the first column, the table shows that respondents displaying a higher big business discontent also disapprove of corporate bailouts. Interestingly, in column (2), we find that the support for small business bailouts is positively correlated with the big business discontent. This is consistent with the idea that negative views regarding corporate responsibility are primarily about large corporations rather than all U.S. businesses.[24]

5. Big Business Discontent and Economic Policies: Experimental Evidence

In the descriptive analysis of the previous section, we showed that those with more negative views regarding corporate behavior are less likely to support corporate bailouts. In this section, motivated by our model of cognitive decision-making, we use random variation in the survey design, as previously described in Section 3.3, to evaluate whether priming and communications can shift policy preferences. Here, we recall the model's key predictions. First, having established that individuals do indeed have widespread negative views regarding corporate responsibility, cues which prime ESG as the policy domain should decrease the support for corporate bailouts. The model further predicts that policy preferences can be influenced by valence framing. In particular, a cue which primes ESG as the policy domain and offers a positive valence framing of corporate responsibility issues should lead to greater support for bailouts than a cue which primes ESG as the policy domain but offers a negative framing of corporate responsibility issues. Furthermore, a cue which primes ESG and provides a positive framing could both lower or increase the support for bailouts if the negative effect of the policy domain dominates the positive effect of the framing, which it will do if the big business discontent is sufficiently large. Finally, priming an alternative policy domain, namely the economic stabilization effects of bailouts should increase the support for bailouts.

We first present the results from our experimental design focused on the main effects of priming, utilizing random variation in the salience and animated video treatments (section

[23]In Appendix Figure A10, we also report the full distributions of responses.

[24]Indeed, recent Gallup polling shows that 33% of Americans have very little confidence in big business as an institution, whereas only 7% have a great deal of confidence. In contrast, 38% of Americans have a great deal of confidence in small business, while only 6% have very little confidence. See https://news.gallup.com/poll/5248/big-business.aspx.

5.1). We then discuss additional results in line with a mechanism of selective memory retrieval (section 5.2).

5.1. **Effects of Priming.**

5.1.1. *The Salience Treatment.* Our experimental design includes a randomization layer where we experimentally vary the question order so that half of the respondents (treatment group) are exposed to the corporate perceptions questions *before* the questions on policy preferences, while the other half (control group) first state their support for bailouts and only *after* answer perceptions questions, as shown in Figure 1. In short, varying the question order is a treatment which primes the ESG policy domain by focusing the attention of survey participants on corporate responsibility issues. We then study whether providing this policy context for the bailout decision is sufficient to change policy preferences. To the extent that individuals are indeed influenced by the cues they receive and, moreover, individuals have deep-rooted beliefs about big business, providing such a policy domain may have substantial effects.

Our results are reported in Table 5. In columns (1) and (2), we first examine the effects of the salience treatment exclusively within those individuals who received the control video. This corresponds to those survey participants who were randomly assigned the cues $\Gamma^* = \emptyset$ or $\Gamma^* = (\text{ESG}, \emptyset, 0)$, as described in Section 3.3. As column (1) shows, we find that the salience treatment has a negative and statistically significant effect on the support for bailouts of large corporations. Through the lens of our cognitive framework, this is because the salience treatment primes agents to recall policy-relevant experiences in the ESG policy domain and *interferes* with the recall of experiences in alternative policy domains. Interestingly, as illustrated in column (2), we find that providing the salience treatment has no impact on individuals' support for small business, consistent with our evidence in Section 4.

In columns (3)-(4) of Table 5, we examine the impact of the salience treatment in the full sample, controlling for the video the survey participants were assigned to. In particular, we estimate the following specification:

$$Y_i = \alpha + \beta T_i^{Salience} + \sum_{j=1}^{j=3} \beta^j T_i^j + \nu_i, \tag{5.1}$$

where Y_i are the outcome variables we observe for each respondent i, namely the support for bailouts of large corporations or small businesses. $T_i^{Salience}$ is an indicator variable equal to 1 if respondent i was subject to the salience treatment (and 0 otherwise). We control for the second randomization layer by adding an analogous indicator variable T_i^j for each of the video treatments.

As seen from column (3), using the full sample and controlling for the video treatments does not attenuate the effect. This suggests that even though the T-Good and T-Bad videos also prime the ESG policy domain, priming the ESG policy domain through the salience treatment continues to incrementally shift policy preferences. That is, double priming has an effect. Within the context of our model this is because double priming increases the strength of selective recall, leading the respondent to recall even more experiences within the primed policy domain. Indeed, if anything, the effect appears to be supermodular. We return to these considerations in Section 5.2. In terms of economic significance, the treatment effect represents an almost 10% increase relative to the mean support for bailouts in the control group. Using the full sample, we once again find that the salience treatment has little impact on individuals' support for small business, as shown in column (4).

5.1.2. *The Video Treatments: Framing of Corporate Responsibility and Support for Economic Policies.* In addition to showing how priming the policy domain can impact policy preferences, the model further shows how cueing a valence frame can impact policy preferences by leading agents to disproportionately recall experiences with positive or negative policy utility (see Theorem 2). As discussed in Section 3.3.2, we introduce positive and negative valence framing of corporate responsibility issues by randomizing the animated videos respondents are exposed to right before the main sections of our survey. We thus study how the different video treatments influence the set of policy preferences we observe by estimating a specification analogous to equation 5.1. We start with the comparison of the T-Bad and T-Good treatments to the control group, while controlling for the salience treatment, as reported in Table 6.

We find a negative and highly significant coefficient of T-Bad on the support for bailouts of large corporations in column (1) of Table 6. The effect is sizable, as the -0.720 coefficient corresponds to a decrease of more than 13% relative to the average support for bailouts in the control group. We once again find little impact of the T-Bad video on the support of policies that aid small businesses.

Consistent with the prediction that individuals do respond to valence framing, we find that survey respondents who watch the T-Good video are significantly more likely to support corporate bailouts than survey participants who receive negative framing through the T-Bad video, as shown in the test for the difference of the coefficients near the bottom of the table.[25] In column (2) of Table 6, we also find that watching the video presenting ESG activities of large corporations in a positive light increases the support for government initiatives aimed at helping small businesses by 0.289 relative to a mean of 7.641.

Table 6 thus confirms one of the key predictions of our model, namely that valence framing in communications can materially impact the support for economic policies. Interestingly,

[25]Later in the paper, we provide evidence suggesting this is because agents are recalling different experiences or memories of corporate responsibility from their mental database, as our cognitive model suggests.

however, Table 6, column (1) further shows that the T-Good treatment actually has a negative impact on the support for bailouts, relative to the no communication baseline, albeit the coefficient is statistically marginally insignificant (p-value of 0.118). This is consistent with the interference channel of the model discussed in Theorem 2. In particular, if the positive messaging surrounding corporate responsibility leads agents to evaluate policy through an ESG lens more than they otherwise would, and agents have many negative experiences regarding responsible corporate behavior in their mental databases on which to draw, then such positive framing can lead to less support for the policy than if agents had received no communication at all. The key point is that due to their frequency and similarity, the negative valence experiences interfere with the recall of the cued positive valence experiences.

5.1.3. *The Video Treatments: Framing Bailouts as Helping the Economy.* Our empirical results thus far strongly suggests that individuals respond to cues which prime the ESG policy domain. That is, when primed to think about bailouts through the lens of corporate responsibility, support for the policy decreases. However, as we have noted, bailouts are a complex and multifaceted policy. This raises the natural question that if agents were primed to think about bailouts in a different policy context, one in which the balance of their experiences might lead them to view bailouts in a more favorable light, support for the policy could be increased. We test this by showing a subset of survey participants a treatment video, i.e., T-Economy, that contains a scene where economists state that corporate bailouts are important for the economy to recover, but that does not mention anything regarding corporate responsibility.

Consistent with our predictions, column (1) of the third row of Table 6 shows that this video treatment leads to increased support for corporate bailouts. Through the lens of the model, priming agents to think about bailouts through the lens of economic stabilization interferes with the recall of relevant ESG information, leading to greater support. Column (2) shows that this video treatment also leads to increased support for small businesses.

5.1.4. *Unpacking the Treatment Effects.* Our results show that the policy preferences can be shifted by communication cues, both by priming the policy domain and by providing a valence frame. To get a better overview of the effect of the individuals cues, Table 7 reports the treatment effects on policy preferences of survey participants for all the byproducts of cues of our randomly assigned treatment layers, which are visually illustrated in Figure 1. The effects are all relative to the control group of respondents who receive the control video and who are not subject to the salience treatment. We make two comments on these results.

First, Table 7 shows that providing the salience treatment in addition to the videos always reduces support for bailouts, indicating that priming the policy domain multiple times leads to stronger results than priming the policy domain only once. In particular, we find that providing participants with the salience treatment and the T-Good treatment strongly and significantly reduces support for bailouts. These disaggregated findings further reinforce how positive messaging can backfire when it also sets a policy domain in which individuals have particularly negative views.

Second, and relatedly, we find that being exposed to the T-Economy video without the salience treatment leads to significantly higher support for bailouts. In contrast, when the salience treatment is provided in conjunction with the T-Economy video treatment, we find no statistically significant impact on policy preferences. This result is consistent with our theoretical framework, since the cue primes two different policy domains, one of which (Econ) leads to greater support for bailouts, while the other (ESG) leads to lower support.

5.2. **Discussion of Mechanisms and Additional Predictions.** Our memory-based model recognizes that there is a distinction between cues and databases and that memory is subject to interference, thus allowing us to rationalize puzzling results such as our finding that positive messaging can "backfire." If memory was only a cue or if interference was not at play, reminding that companies do good would simply increase the support for corporation-friendly policies. If memory was only databases, priming would have no effect. The interaction of these various forces explains why people primed with good ESG information are more supportive of bailouts than people primed with bad ESG information (*cue*) and why priming good ESG information can backfire (*interference*). In this subsection, we provide a set of additional results consistent with individual cues affecting beliefs and policy preferences via a selective memory retrieval channel, such as the fact that the backfiring effects are stronger for the liberals and young individuals than for conservative and older individuals (*database*).

5.2.1. *Treatment Effects on Perceptions of Large Corporations.* A central tenet of our model is that policy preferences are shaped by the individual recall of large corporations' behavior within society. In particular, the cue should influence which memories or experiences are recalled from the mental database. Our experimental survey design allows us to directly test for this, using the questions of section 3 of the survey as dependent variables in our baseline specification.

We find a strong effect of our animated video experiment on perceptions of corporate responsibility. The top row of Table 8, Panel A, shows that the T-Bad treatment, which provides a negative valence framing of corporate behavior, significantly increases the big business discontent along all dimensions. The magnitude of the T-Bad treatment is substantial, as it changes perceptions by around one-third of the mean and one-fourth of the standard deviation. For example, the smallest magnitude pertains to the question about corporations donating to politicians, and yet we find an increase of 21.71% relative to the mean. The largest relative effect is the one regarding health care benefits to employees, where we find an increase of 43.58% relative to the mean, which is similar to the increase of

38.93% we observe for the question on gender diversity. All coefficients are not only large in magnitudes, but also highly significant.

Panel A of Table 8 further shows that the big business discontent is not as large when the survey participants receive the T-Good treatment, which provides a positive valence framing of corporate behavior. This is consistent with the positive framing leading agents to recall more positive examples of corporate responsibility. To provide further evidence that framing is in fact impacting memory recall, in Panels B and C of Table 8, we report the results separately for the *Are* and *Should Be* components of the big business discontent, respectively. Consistent with our model's predictions of selective memory recall, we find that the videos primarily alter respondents' perceptions of how unsatisfactory corporate policies within society *are*. In contrast, the videos appear to have little impact on survey participants' perceptions of what they think large corporations *should* do.

While the T-Good treatment leads to less big business discontent than the T-Bad treatment, consistent with the effects of framing, we do find that the T-Good treatment does increase the big business discontent relative to survey participants who did not receive the video, as shown in the second row of Table 8, Panel A. In particular, while we find positive but statistically insignificant effects of the video on individual perceptions related to health care benefits, gender diversity, and shareholder vs stakeholder maximization, we see a positive and statistically significant impact on policies such as executive compensation, tax strategy, disclosure of CO_2 emissions, and political donations. These surprising results once again show how, due to priming the policy domain and interference, positive communications on corporate responsibility issues can actually lead to more negative views than would otherwise be with no messaging.

Finally, when looking at the T-Economy treatment, with the exception of our broad question on shareholders vs stakeholders in Panel A, all coefficients are close to zero, statistically insignificant, and precisely estimated, displaying a mix of negative and positive signs. This latter result provides a further important validation test for our experimental design.

5.2.2. *Heterogeneous Effects: Liberals versus Conservatives, Young versus Old.* Our model, with Theorem 2, predicts that the negative effects of positive framing should be larger for respondents with a higher proportion of pre-existing negative memories about corporations. We provide suggestive evidence in support of this prediction in Table 9, where we study how the effects of the treatment videos vary depending on the political orientation (Panel A) and age (Panel B) of the respondent.

In Panel A, the underlying assumption is that liberal respondents are more likely to have memories of large corporations behaving poorly within society in their database. Specifically, the table reports results from a regression where we augment our baseline regression with a set of interaction terms using indicator variables for both conservative and liberal respondents, respectively.[26] The excluded category consists of individuals who identify themselves as moderate, who make up 40% of the sample. We focus on the support for bailout of large corporations as dependent variable.

Across the full political spectrum, negative valence framing leads to less support for bailouts than positive framing. Thus, our basic conclusion that policy preferences are impacted by framing remains robust. Furthermore, we now find a significantly stronger (i.e., more negative) effect of our T-Good treatment among liberals, as shown in the top row of column (2). That is, consistent with our prediction above, liberals become significantly less likely to support bailouts when shown the T-Good video compared to receiving no communication at all. For this particular subset of respondents who have highly established negative views regarding corporations' role in society, attempting to provide a positive framing on corporate responsibility backfires. Through the lens of our model, this is because the positive framing cannot be separated from priming the policy domain. That is, by providing them with a positive narrative on ESG issues, liberals are primed to consider bailout policy through an ESG lens, which ultimately lowers the support for bailouts more than the positive framing increases it.

In Panel B, we focus on how the effects vary by age. We hypothesize that younger demographics—for whom the global financial crisis and subsequent economic hardship, accompanied by a rising anti-corporate sentiment, make up most of their recent experiences—have a database with a larger share of bad memories regarding large corporations relative to older respondents. As a result, positive ESG messaging should suffer from less interference for older individuals. Table 9, Panel B, shows results consistent with this hypothesis, as we see a strong backfiring of the T-Good treatment fully concentrated among the young.

We provide additional supportive evidence for memory effects in Appendix Table A1, where we report an analogous analysis of heterogeneous treatment effects using perceptions of large corporations as dependent variables, as in section 5.2.1. Once again, we find evidence of significant backfiring of the T-Good video among the liberal and young respondents.

5.2.3. *A Survey Measuring What Comes to Mind.* Our approach to identify a selective memory channel is to rely on a number of priming experiments, by varying the order of questions, the content of videos, and the framing of the policy. Priming experiments are a common technique to assess memory effects. An alternative approach to measure memory effects is to obtain direct evidence of selective memory retrieval. Following an approach similar to Andre et al. (2022), we therefore conducted an additional round of surveys in June 2022 to measure what comes to respondents' minds (Gennaioli and Shleifer, 2010).

[26]"Liberal" includes both Liberal and Very Liberal, while "Conservative" includes both Conservative and Very Conservative.

We collected a total of 608 responses from individuals who did not participate in any of our previous surveys with our online data collection partner. The survey begins in a way that is identical to our main survey, and with identical language and instructions. The respondents answer a set of sociodemographic questions and are then shown a video. We only focus on the T-Bad and T-Good treatment arms. Of the 608 total respondents, 212 see the T-Bad video, 210 see the T-Good video, and 186 see the Control video. The main difference with our main survey is that we do not ask questions about perceptions of ESG or support for economic policies. Instead, we are interested in measuring what comes to respondents' minds right after watching the experimental videos.

Specifically, right after the video, we ask respondents: "Could you tell us a bit more about what thoughts you had after watching the video? The following statements describe different thoughts you might have had on your mind after watching the video. Did you have any of these thoughts on your mind? Please tick all that you had on your mind." The respondents are then presented with a list of thoughts, which are illustrated in Appendix Table A2. We randomize the order in which the thoughts are shown to address potential order effects.[27] At the bottom of the list we always show the option "None of the above." The "thoughts" are short sentences aimed at capturing specific memories or experiences the respondents might recall when primed with a specific video. Therefore, we created three thoughts related to bad corporate behavior (e.g., "Large corporations are often involved in corruption scandals. For example, think of Enron!"), three thoughts related to good corporate behavior (e.g., "Facebook and other Big Tech companies are leading the way in diversity and inclusion."), and three other thoughts that are unrelated to corporate responsibility and that instead concern economic aspects of the policy discussion (e.g., "The economic consequences of mass layoffs during a crisis can be catastrophic."). We consider these thoughts as close conceptually to the memories or experiences present in people's mental databases.

We report our results in Table 10. We are interested in whether our videos prime the recall of specific memories in the database, and specifically whether individuals exposed to the T-Bad or T-Good videos are more likely to recall memories or experiences related to corporate responsibility relative to respondents in the control group. We create two types of dependent variables aimed at capturing the *Share* (relative to the number of thoughts selected) and the total *Count*, respectively, of thoughts related to bad corporate behavior, good corporate behavior, and to other economic aspects. The econometric specification is analogous to our baseline equation 5.1. We observe findings that largely echo our earlier results. First, individuals exposed to the T-Bad video are significantly more likely to recall bad memories of corporate responsibility, both with respect to the control group and to those exposed to the T-Good video. Second, and consistent with the presence of interference in

[27]We also randomize the number of thoughts respondents are shown. Approximately half the respondents are shown all nine thoughts, while the other half are shown six randomly drawn thoughts.

the memory recall process, the T-Good video backfires, as it leads individuals to recall more memories of bad corporate behavior relative to the control group. We do not find significant differences with respect to the recall of good memories of corporate behavior or of other memories.

In sum, while definitively measuring memory effects remains challenging, we believe these additional results on what comes to mind, combined with our previous empirical findings aligned with our model, provide important support for a selective memory retrieval channel explaining our findings.

5.2.4. *Text Analysis of Open-Ended Questions.* To provide additional support for a selective memory retrieval channel, we perform a textual analysis of the answers to open-ended questions. Specifically, in a follow-up survey we conducted in October 2020, we asked respondents: "Could you tell us a bit more about why you have these views on policies regarding large corporations? What makes you being friendly or unfriendly with respect to helping large corporations?"[28] Our hypothesis is that, relative to the control group, respondents who watched the T-Bad video would be more likely to use ESG-related reasoning, relative to reasoning based on economic aspects of the issue, when answering the question.

To test this hypothesis, following an approach similar to Andre et al. (2022), we ask two independent external reviewers, namely research assistants, to classify the responses based on the type of reasoning used by respondents. The reviewers only observe the answers to the open-ended questions and are unaware of the treatment status of the observations. First, the reviewers filter out answers that are not of a sufficiently high quality to assess. Those answers are categorized as "Invalid" and taken out of the sample. Second, reviewers group responses into an ESG category if the answer shows any reasoning related to ESG issues. Examples include responses addressing executive compensation, employees' health care, corporate tax evasion, and gender diversity in executive positions, among others. On the other hand, if the answer only focuses on economic aspects of the issue, reviewers are asked to classify that as economic reasoning. Examples include responses focused on the importance of bailouts for economic recovery and the need to save jobs. The reviewers can also assign responses to a "Other" category in those cases where the answer does not fit in neither the ESG nor in the economic category. The reviewers almost uniformly agree on whether the quality of the response is high enough and, importantly, they agree on around 80% of the classifications into ESG and Economic reasoning. This inter-rater reliability, shown in Table A3, is similar to the one in Andre et al. (2022).

Finally, reviewers also separately grouped answers into those expressive a negative or a positive sentiment. If responses presented a negative and unfriendly tone, say towards large corporations or toward policy making, they were placed in the *negative* category. Conversely,

[28] This survey, discussed in detail in Section 6.3, only had one treatment arm, namely the T-Bad video.

reviewers grouped answers in a *positive* category if responses had a positive and friendly tone. If responses used a neutral tone they were placed into a *neutral* category.

Figure 4 shows the results. The figure shows the percentage of respondents who rely on ESG and economic reasoning among both treated and control respondents. In Panel A, the figure shows that the prevalence of ESG-based reasoning is significantly higher for the treated than the control group. Among the treated respondents, around 57% uses ESG-based reasoning whereas only 47% of control respondents use ESG-based reasoning. We furthermore observe a significantly lower use of economic reasoning in the treatment group, consistent with our memory framework. Approximately 58% of participants in the control group use economic reasoning, while only 51% of participants in the treatment group do.[29] Overall, these results provide supporting evidence that the T-Bad video treatment influences policy preferences by priming individuals to disproportionately think about bailouts of large corporations through the lens of ESG-related issues. In further support of a memory channel, in Panel B of Figure 4, we also show that participants who received the T-Bad treatment used reasoning based on *negative* ESG views when accounting for their policy preferences, relative to participants in the control group. Specifically, 44% of participants in the control group use negative ESG views to rationalize their policy preferences, compared to 52% of participants in the treatment group.

6. Robustness Tests and Behavioral Outcomes

In this final section, we report several additional results aimed at testing the robustness of our findings. In particular, we first discuss alternative explanations and the persistence of the effects (Section 6.1). In Section 6.2 we show the robustness of our main findings to alternative specifications. Finally, in Section 6.3 we describe the findings from an additional survey designed to measure behavioral outcomes.

6.1. Alternative Explanations and Persistence of Effects.

A typical, mechanical alternative explanation in information experiments is that outcomes are driven by experimenter-demand effects.[30] However, our finding that the T-Good video has a negative effect on respondents' perceptions of large corporations is strongly inconsistent with such a story. In fact, if experimenter-demand effects were at play, we would have expected the opposite, as the treatment likely displays an intention of the researcher to shed a positive light on corporations' behavior towards society.

[29] For these analyses, in case of disagreement between the reviewers, we classify the answer within a given category if at least one reviewer does so. Also, in unreported results, we find that the fraction of respondents using "Other" reasoning is largely the same across the treatment and control group.

[30] Yet, recent evidence by De Quidt et al. (2018) indicates that such concerns are of rather limited quantitative importance in online surveys like ours.

If anything, a related but opposite concern arises when thinking about the finding that the T-Good video backfires. Specifically, it is plausible that respondents in the T-Good treatment group believe that our study is linked to a pro-corporations think-tank or policy institute. If that were the case, the T-Good findings would not be driven by the mechanisms of our model, but possibly by, for example, respondents' thinking large corporations are even worse than they thought. However, such a "reverse" experimenter-demand effect seems inconsistent with our other empirical findings, and in particular with the result that respondents exposed to the T-Economy treatment do not react negatively to the pro-bailouts pitch of the video.

Moreover, notice that the above concerns regarding individuals' beliefs of the researchers' intention are likely inconsistent with our salience treatment results, which largely do not rely on informational treatments and yet generate strong effects.

To further alleviate these concerns, we also conducted a follow-up survey one week after the original survey, which further allows us to test the persistence of our results. We chose to have approximately one week between the two surveys to test for persistence while also minimizing attrition.[31] We surveyed a total of 2,311 respondents in the follow-up, which consisted of asking respondents only Sections 3 and 4 of our survey, namely the questions on perceptions and on support for government policies.[32] Crucially, we do not show any video to anyone, so that the follow-up survey does not provide differential information to respondents, and answers are detached from our immediate treatments. This is a common test against experimenter-demand concerns in this type of study (Alesina et al., 2018; Fehr et al., 2019; Haaland et al., 2020).

We replicate our analysis of treatment effects on economic policies using the follow-up survey responses in Table A4. We find that respondents who were exposed to the T-Bad video still hold significantly different (negative) views on support for bailouts. The effect is smaller in magnitude, but it is still strongly statistically significant. The coefficient on T-Good shows no effect on support for bailouts.

We conduct a similar analysis in Table A5, where we replicate the analysis of Table 8, but using the measures of perceptions collected in the follow-up study. Focusing on the big business discontent, we find that respondents who were exposed to the T-Bad video still hold significantly different views on the policies adopted by large corporations, continuing to display a higher dissatisfaction with their behavior. The magnitudes of the effects are smaller, but only around one-third so, depending on the specific corporate policy we measure.

[31]The precise lag between the original survey and the follow-up one ranges between 3 and 13 days for all respondents. The average difference was 6.12 days.

[32]In Table A6, we show that respondents to the follow-up surveys are somewhat selected since, for example, young and employed people are less likely to respond to the follow-up while white respondents are more likely to do so. Therefore, in Section 6.2, we show that our results are largely unchanged by the inclusion of all possible sets of individual socioeconomic controls.

We find some persistence of the T-Good treatment as well, but to a lesser extent, consistent with the original effects being milder.

In sum, these follow-up surveys without information treatments have reassuring implications for the robustness and interpretation of our findings.[33]

6.2. **Robustness Checks.** We conduct several robustness checks. First, in Appendix Tables A7 and A8, we show the robustness of our findings to a re-weighting procedure to make our sample representative of the U.S. population along all socio-demographic dimensions. Relatedly, in Appendix Tables A9 and A10, we illustrate the robustness of our results to the inclusion of individual socioeconomic controls in our estimation. We further show that our results are unaffected by the inclusion of every possible combination of socioeconomic controls, by reporting coefficient stability plots in Appendix Figures A11 and A12.[34]

In Table A11 we report additional robustness results. First, in column (1), we show that our results hold when we drop respondents who say they have put little to no effort into the survey. Then, in column (2), we show our results remain largely unchanged when we control for the time spent on the survey. In column (3), we further show that the results are unaffected when we drop respondents who state they felt the survey was politically biased. Columns (4)-(6) report the analogous results using the support for small businesses as the dependent variable.

6.3. **Measuring Policy Preferences through Costly Behavioral Actions.** We conducted a new survey in October 2020, where we collected data from a sample of 1,683 new respondents who were never exposed to our original survey. The main objective of this additional survey is to collect behavioral outcome measures to complement the analysis of our initial survey based on self-reported policy preferences. By collecting behavioral measures, we can alleviate concerns that self-reported survey responses might not be fully reflective of true individual policy preferences since they do not require costly actions on behalf of the respondents. Moreover, by conducting the survey five months after the original survey and after the initial shock induced by the coronavirus crisis, we can maximize external validity and test for the robustness of our results over time.

The survey is identical in structure to our main survey, but focuses on our strongest treatment only, namely the negative treatment video. The sample is split into 855 respondents who are exposed to the control video and 828 respondents who are exposed to the negative

[33] A further interpretation of the persistence of the effects, within the context of our model, could be that of an instance in which our treatments implant new memories in the mental database as a byproduct of our "treatment as a cue." That is, when respondents are asked about their preferences on a given issue, they retrieve memories of the previous survey experience, thereby generating the persistence we observe.

[34] We follow the procedure by Bursztyn et al. (2020). Coefficient stability plots for the one-week follow-up are also reported, in Appendix Figures A13 and A14.

treatment video. The balance statistics for this study are reported in Table A12. The full questionnaire is reported in Appendix Section A.3.

We measure the public support for corporate bailouts, and large corporations in general, in three ways. First, we create a petition on the website Change.org to support a bailout of large corporations. The full page of the petition is shown in Appendix Section A.5 and is designed to be consistent with similar types of petition asking for various forms of economic support during the coronavirus crisis. The petition is addressed to the U.S. Congress and contains a concrete policy proposal arguing in favor of a bailout of large corporations at a time when a new economic stimulus plan was being discussed. Given the potential real policy consequences of signing the petition, external validity concerns are attenuated. We make this issue more salient to the respondents by stating: "Few citizens sign petitions, making policy makers take them all the more seriously." Since we are unable to track whether our respondents actually sign the petition, our analysis focuses on the responses to our survey question, and specifically whether the respondent indicates either *I will sign the petition* or *I will not sign the petition*.

Our second behavioral measure consists of asking respondents' permission to contact U.S. senators on their behalf. In practice, we create ready-to-send emails, and we give the option to send them to any senators of their choice. One version of the email is in clear support of bailouts of large corporations, while another version is in clear opposition to such bailouts. To make this action costly, we tell respondents that by giving the OK they agree to have their name included in the email to the U.S. Senators, together with the names of other survey respondents who also agreed. The full text of the question is shown in Q27 in Appendix Section A.3.

The third behavioral measure aims at capturing an individual's broader support for large corporations, rather than just corporate bailouts. To do so, we enroll respondents into a lottery for multiple $25 gift cards. We then ask them whether they would like to donate part of their winnings to the Business Roundtable, which we describe next as a "non-profit organization that represents chief executive officers of America's largest corporations and that advocates policies to strengthen the economy while protecting the business interests of corporations."[35] As a result, this question elicits another costly action, as respondents are asked to forego part of their compensation.[36]

[35]We minimize experimenter-demand concerns by truthfully telling respondents: "We will now randomly select one of two nonpartisan and nonprofit organizations: one advocates supporting workers and communities; the other advocates more support for large corporations and their executives." In practice, we randomized almost all of our respondents to the Business Roundtable.

[36]Donations to liberal and conservative non-profit organizations and initiatives are widely accepted in the literature as a way to measure policy preferences (Perez-Truglia and Cruces, 2017; Haaland and Roth, 2019; Grigorieff et al., 2020; Bursztyn et al., 2020; Haaland et al., 2020).

We report the results of our analysis in Table 11. We start in column 1 by estimating the impact of the treatment on the same self-reported measure we used for our main analysis, which asks about the support for bailouts on a scale from 0 to 10. It is reassuring to see that the point estimate for treatment effects in our October 2020 survey (-0.719) is both strongly statistically significant and nearly identical to the point estimate in Table 6.[37]

In columns 2-5, we move to the analysis of the behavioral outcome measures. Specifically, column 2 reports the impact of our treatment on an indicator variable taking value 1 if the respondent indicated she would sign the petition in support of bailouts for large corporations. Columns 3 and 4 use as dependent variables indicators for whether the respondent gave permission for her name to be included in an email to U.S. Senators in support of or opposition to bailouts, respectively. Finally, in column 5, we report the total amount of money (in U.S. dollars) individuals would agree to donate to the Business Roundtable in case they won one of the several $25 lotteries we enrolled them in.

The results in Table 11 provide strong support to our main findings, as we find that the self-reported public support for corporate bailouts is largely reflected in costly actions by the respondents. First, we report significant effects in terms of individual willingness to take real policy action, as we find that 42% of respondents indicated they would sign the petition, and 22.3% (28.2%) decided to communicate to U.S. Senators support for (opposition to) a bailout of large corporations. Similarly, respondents are willing to donate approximately one-third ($7.43) of potential winnings to the non-profit organization representing the interests of large U.S. corporations and their executives, i.e., the Business Roundtable.

The magnitudes of the treatment effects are significant. We find that treated respondents are 10.8 percentage points less likely to sign the petition, which is a 25.71% decline relative to the average in the control group. Treated respondents are also 8.9 percentage points less likely to email U.S. Senators to support bailouts (relative to a mean of 22.3). We find marginally statistically significant effect on the willingness to email U.S. senators to oppose bailouts, with the treatment video leading to an increase of 13.12%. Finally, we find the treatment induces a 27.11% ($2.015) decrease in the amount of money respondents are willing to donate in support of large corporations.

While external validity concerns are always present, we believe the findings in this section help alleviate them considerably, in line with a large experimental literature using largely similar measures for similar purposes (see Haaland et al. (2020) for a review).

[37]In Table A13 we report the results on perceptions of large corporations for this additional survey. We observe nearly identical measures of big business discontent in the control group (as shown in the mean of the dependent variable row of the table), and largely similar magnitudes of the treatment effects, which are all large and statistically significant at the 0.01 level.

7. Conclusion

Corporate America appears to be under more public scrutiny than ever before. From boycotts to protests and a number of social actions, large corporations' role in society is taking center stage in the public debate. Plausibly, societal perceptions of corporate responsibility could have an impact on governmental policies impacting the corporate sector. Corporations are reacting to this development with messaging and communication campaigns attempting to paint a positive image of their impact on the environment, their employees, and society at large.

In this paper, we provide some of the first evidence linking public perceptions of corporate behavior, as well as the messaging surrounding corporate responsibility, to the support for economic policies. By conducting large-scale surveys on public opinions of the policies of corporate America, we show that the public demands corporations to behave better within society, a sentiment we label "big business discontent." We also find a strong baseline link between big business discontent and the support for economic policies, with people dissatisfied with large corporations' behavior within society also opposing corporate bailouts.

To formalize a link between communications and policy preferences, we build a theoretical, cognitive model of decision-making featuring similarity-based, limited memory recall. In our framework, communications take the form of a cue, which can provide both a policy domain and a valence framing. The policy domain primes the agent to consider the policy decision through a particular lens. The valence framing primes the agent to disproportionately recall experiences which paint the policy in either a positive or negative light. Our behavioral model predicts the form of the cue can shift policy preferences. In particular, if the cue primes a policy domain where the views of agents are quite negative, then support for the policy decreases. Moreover, holding the policy domain fixed, a positive valence framing should lead to greater support for the policy than a negative valence framing.

We test these predictions using experimental variation in our survey design and find support for these effects. Our empirical findings also confirm another subtle prediction of the model: positive communications surrounding corporate behavior can actually lead to less support for corporation friendly policies than providing no communication if agents have sufficiently negative established beliefs regarding corporate responsibility. This result has significant implications for corporate and political communication strategies, especially if positive framing of an issue cannot be separated from priming the policy domain.

Our paper leaves many open questions for future research. To start with, more evidence, both experimental and non-experimental, is needed to better understand how views about corporations affect a range of economic policies in both good and bad economic times. Moreover, our study illustrates how actual policy preferences are impacted by communications. Much more work can be done on the long-term determinants of individual perceptions of large corporations. Digging deeper into the structure of how societal beliefs about corporations are formed, perhaps through direct experiences or social interactions, seems like a first-order issue. Finally, societies around the world are obviously different, and therefore we see studying the relationship between people and big business outside the U.S. as an immediate next step.

References

Albuquerque, R. A., Y. Koskinen, S. Yang, and C. Zhang (2020): "Love in the time of covid-19: The resiliency of environmental and social stocks," .

Alesina, A., A. Miano, and S. Stantcheva (2022): "Immigration and Redistribution," *The Review of Economic Studies*.

Alesina, A., S. Stantcheva, and E. Teso (2018): "Intergenerational mobility and preferences for redistribution," *American Economic Review*, 108, 521–54.

Andre, P., C. Pizzinelli, C. Roth, and J. Wohlfart (2022): "Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples," *The Review of Economic Studies*.

Barrera, O., S. Guriev, E. Henry, and E. Zhuravskaya (2020): "Facts, alternative facts, and fact checking in times of post-truth politics," *Journal of public economics*, 182, 104123.

Bénabou, R. and J. Tirole (2010): "Individual and corporate social responsibility," *Economica*, 77, 1–19.

Bordalo, P., K. Coffman, N. Gennaioli, F. Schwerter, and A. Shleifer (2020a): "Memory and representativeness." *Psychological Review*.

Bordalo, P., J. Conlon, N. Gennaioli, S. Kwon, and A. Shleifer (2021): "Memory and Probability," *Working paper*.

Bordalo, P., N. Gennaioli, and A. Shleifer (2020b): "Memory, attention, and choice," *The Quarterly journal of economics*.

Broccardo, E., O. D. Hart, and L. Zingales (2020): "Exit vs. Voice," Tech. rep., National Bureau of Economic Research.

Bursztyn, L., I. K. Haaland, A. Rao, and C. P. Roth (2020): "DISGUISING PREJUDICE: POPULAR RATIONALES AS EXCUSES FOR INTOLERANT EXPRESSION," .

Cororaton, A. and S. Rosen (2020): "Public firm borrowers of the US Paycheck Protection Program," *Available at SSRN 3590913*.

Cowen, T. (2019): *Big Business: A Love Letter to an American Anti-Hero*, St. Martin's Press.

Crawford, V. P. and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.

Cruces, G., R. Perez-Truglia, and M. Tetaz (2013): "Biased perceptions of income distribution and preferences for redistribution: Evidence from a survey experiment," *Journal of Public Economics*, 98, 100–112.

da Silveira, R. A., Y. Sung, and M. Woodford (2020): "Optimally imprecise memory and biased forecasts," Tech. rep., National Bureau of Economic Research.

Dasgupta, I. and S. J. Gershman (2021): "Memory as a computational resource," *Trends in Cognitive Sciences*.

De Quidt, J., J. Haushofer, and C. Roth (2018): "Measuring and bounding experimenter demand," *American Economic Review*, 108, 3266–3302.

DellaVigna, S. and M. Gentzkow (2010): "Persuasion: Empirical Evidence," *Annual Review of Economics*, 2, 643–669.

DellaVigna, S. and E. La Ferrara (2015): "Economic and social impacts of the media," in *Handbook of media economics*, Elsevier, vol. 1, 723–768.

DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003): "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly journal of economics*, 118, 909–968.

Edmans, A. (2011): "Does the stock market fully value intangibles? Employee satisfaction and equity prices," *Journal of Financial economics*, 101, 621–640.

Enikolopov, R. and M. Petrova (2015): "Media capture: empirical evidence," in *Handbook of media Economics*, Elsevier, vol. 1, 687–700.

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

Enke, B., F. Schwerter, and F. Zimmermann (2020): "Associative memory and belief formation," Tech. rep., National Bureau of Economic Research.

Enke, B. and F. Zimmermann (2019): "Correlation neglect in belief formation," *The Review of Economic Studies*, 86, 313–332.

Fehr, D., J. Mollerstrom, and R. Perez-Truglia (2019): "Your Place in the World: The Demand for National and Global Redistribution," Tech. rep., National Bureau of Economic Research.

Fisman, R., K. Gladstone, I. Kuziemko, and S. Naidu (2020): "Do Americans want to tax wealth? Evidence from online surveys," *Journal of Public Economics*, 188, 104207.

Freeman, R. E., J. S. Harrison, A. C. Wicks, B. L. Parmar, and S. De Colle (2010): *Stakeholder theory: The state of the art*, Cambridge University Press.

Friedman, M. (1970): ""The Social Responsibility of Business is to Increase Its Profits"," *New York Times Magazine*.

Gennaioli, N. and A. Shleifer (2010): "What comes to mind," *The Quarterly journal of economics*, 125, 1399–1433.

Gilboa, I. and D. Schmeidler (1995): "Case-based decision theory," *The quarterly Journal of economics*, 110, 605–639.

Grewenig, E., P. Lergetporer, K. Werner, and L. Woessmann (2020): "Incentives, search engines, and the elicitation of subjective beliefs: Evidence from representative online survey experiments," *Journal of Econometrics*.

Grigorieff, A., C. Roth, and D. Ubfal (2020): "Does Information Change Attitudes Toward Immigrants?" *Demography*, 57, 1117–1143.

Guiso, L., P. Sapienza, and L. Zingales (2004): "The role of social capital in financial development," *American economic review*, 94, 526–556.

——— (2006): "Does culture affect economic outcomes?" *Journal of Economic perspectives*, 20, 23–48.

——— (2015a): "Corporate culture, societal culture, and institutions," *American Economic Review*, 105, 336–39.

——— (2015b): "The value of corporate culture," *Journal of Financial Economics*, 117, 60–76.

Haaland, I. and C. Roth (2017): "Labor market concerns and support for immigration," *Available at SSRN 3017420*.

——— (2019): "Beliefs about racial discrimination and support for pro-black policies," .

Haaland, I., C. Roth, and J. Wohlfart (2020): "Designing Information Provision Experiments," .

Hart, O. and L. Zingales (2017): "Companies should maximize shareholder welfare not market value," *ECGI-Finance Working Paper*.

Hartzmark, S. M., S. Hirshman, and A. Imas (2019): "Ownership, Learning, and Beliefs," *Available at SSRN 3465246*.

Johnston, R. J., K. J. Boyle, W. Adamowicz, J. Bennett, R. Brouwer, T. A. Cameron, W. M. Hanemann, N. Hanley, M. Ryan, R. Scarpa, et al. (2017): "Contemporary guidance for stated preference studies," *Journal of the Association of Environmental and Resource Economists*, 4, 319–405.

Kahana, M. J. (2012): *Foundations of human memory*, OUP USA.

Kamenica, E. and M. Gentzkow (2011): "Bayesian persuasion," *American Economic Review*, 101, 2590–2615.

Karadja, M., J. Mollerstrom, and D. Seim (2017): "Richer (and holier) than thou? The effect of relative income improvements on demand for redistribution," *Review of Economics and Statistics*, 99, 201–212.

Kitzmueller, M. and J. Shimshack (2012): "Economic perspectives on corporate social responsibility," *Journal of Economic Literature*, 50, 51–84.

Krosnick, J. A. (2018): "Questionnaire design," in *The Palgrave handbook of survey research*, Springer, 439–455.

Kuhnen, C. M. (2015): "Asymmetric learning from financial information," *The Journal of Finance*, 70, 2029–2062.

Kuhnen, C. M. and B. Knutson (2011): "The influence of affect on beliefs, preferences, and financial decisions," *Journal of Financial and Quantitative Analysis*, 46, 605–626.

Kuziemko, I., M. I. Norton, E. Saez, and S. Stantcheva (2015): "How elastic are preferences for redistribution? Evidence from randomized survey experiments," *American Economic Review*, 105, 1478–1508.

Lenz, G. S. (2009): "Learning and opinion change, not priming: Reconsidering the priming hypothesis," *American Journal of Political Science*, 53, 821–837.

Lins, K. V., L. Roth, H. Servaes, and A. Tamayo (2020): "Sexism, Culture, and Firm Value: Evidence from the Harvey Weinstein Scandal and the# MeToo Movement." .

Lins, K. V., H. Servaes, and A. Tamayo (2017): "Social Capital, Trust, and Firm Performance: The Value of Corporate Social Responsibility during the Financial Crisis," *The Journal of Finance*, 72, 1785–1824.

Malhotra, N. (2008): "Completion time and response order effects in web surveys," *Public opinion quarterly*, 72, 914–934.

McCombs, M. E. and D. L. Shaw (1972): "The agenda-setting function of mass media," *Public opinion quarterly*, 36, 176–187.

Meade, A. W. and S. B. Craig (2012): "Identifying careless responses in survey data." *Psychological methods*, 17, 437.

Mullainathan, S. (2002): "A memory-based model of bounded rationality," *The Quarterly Journal of Economics*, 117, 735–774.

Mullainathan, S., J. Schwartzstein, and A. Shleifer (2008): "Coarse thinking and persuasion," *The Quarterly journal of economics*, 123, 577–619.

Napoli, P. M. (2014): "Automated media: An institutional theory perspective on algorithmic media production and consumption," *Communication Theory*, 24, 340–360.

Perez-Truglia, R. and G. Cruces (2017): "Partisan interactions: Evidence from a field experiment in the united states," *Journal of Political Economy*, 125, 1208–1243.

Roth, C., S. Settele, and J. Wohlfart (2020): "Beliefs about public debt and the demand for government spending," .

Rudorf, S., B. Weber, and C. M. Kuhnen (2014): "Stock Ownership and Learning from Financial Information," *reunião de.*

Stigler, G. J. (1961): "The Economics of Information," *Journal of Political Economy*, 69, 213–225.

Strömberg, D. (2015): "Media and politics," *economics*, 7, 173–205.

Tversky, A. (1977): "Features of similarity." *Psychological review*, 84, 327.

Wachter, J. A. and M. J. Kahana (2019): "A retrieved-context theory of financial decisions," Tech. rep., National Bureau of Economic Research.

Weinzierl, M. (2017): "Popular acceptance of inequality due to innate brute luck and support for classical benefit-based taxation," *Journal of Public Economics*, 155, 54–63.

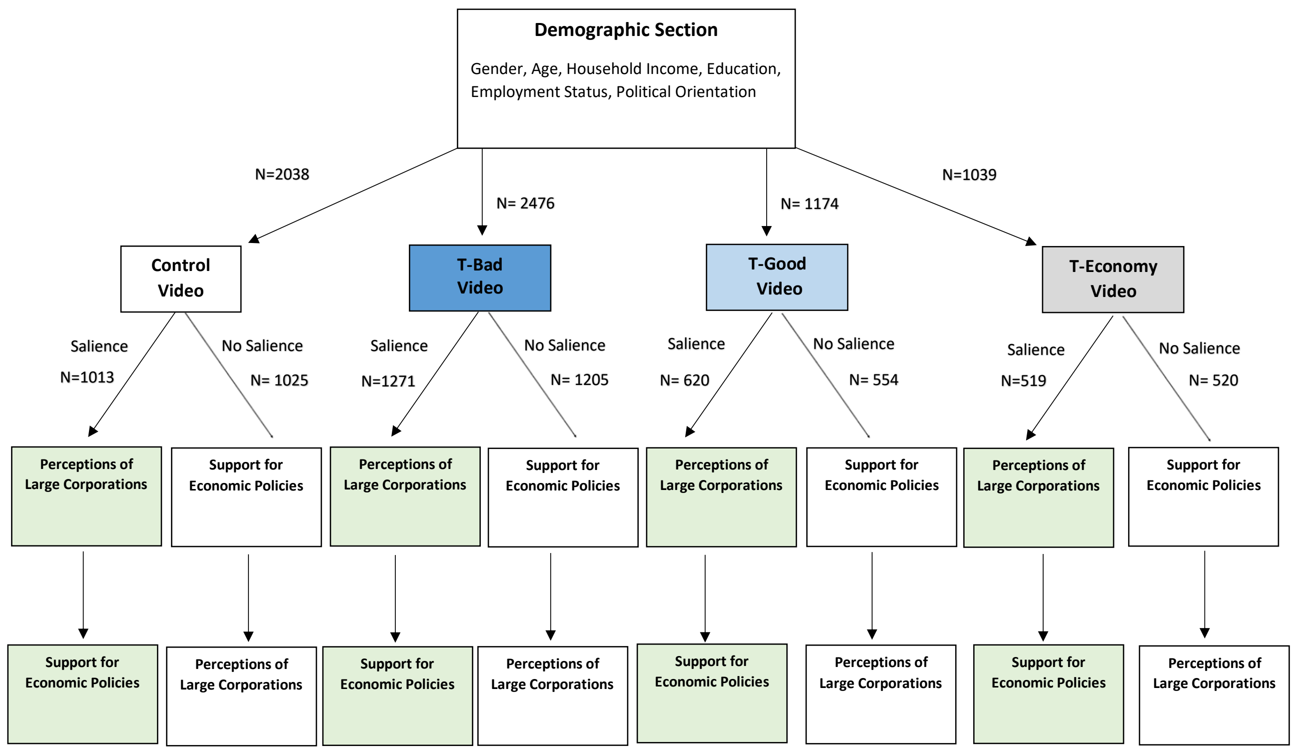

FIGURE 1. Experimental Design

Notes: This figure illustrates our experimental design, including the randomization layers and the sample sizes associated with each treatment and control group. The details of the design are discussed in Section 3.3.

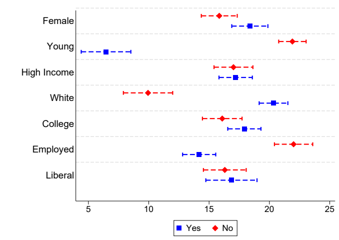

A. Executive Compensation

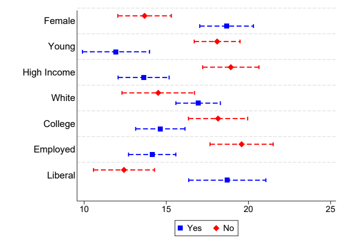

B. Health Care

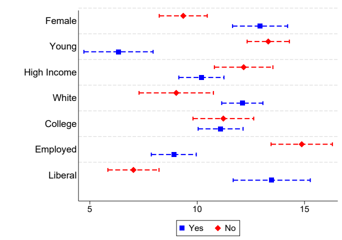

C. Taxes

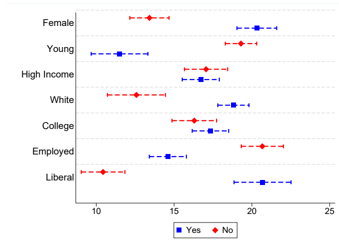

D. Women Executives

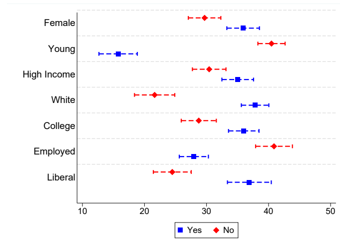

E. CO_2 Disclosure

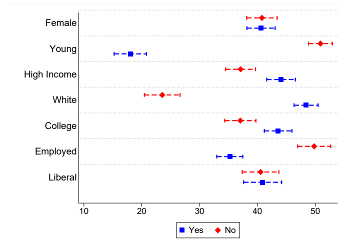

F. Political Donations

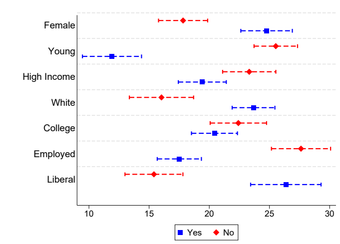

G. Shareholders

FIGURE 2. Heterogeneity in Big Business Discontent

Notes: This figure shows how our measure of big business discontent varies across socio-demographic characteristics of the respondents. The sample consists of respondents in the Control video group. *Yes* indicates the respondent belongs to the given group in the y-axis, while *No* indicates otherwise. See Section 3.2 for a definition of big business discontent and each specific measure, and see Table 1 for a definition of each specific socio-demographic indicator variable. The sub-figures display the average and the 95% confidence interval.

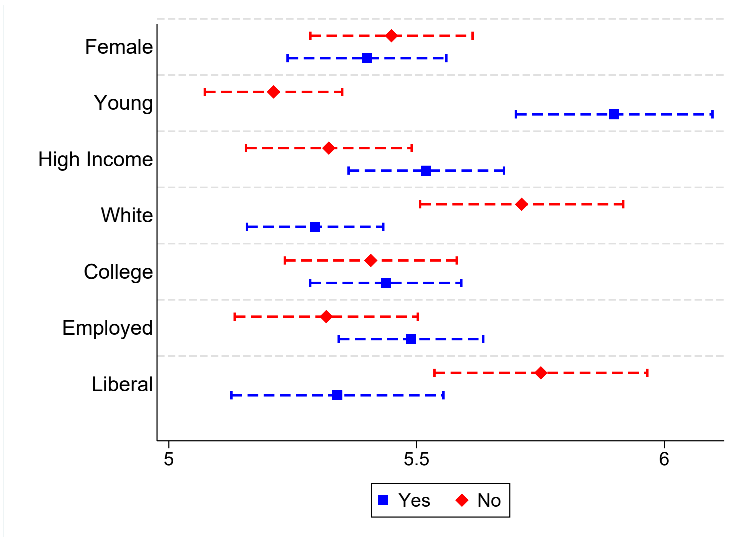

A. Support for Bailouts

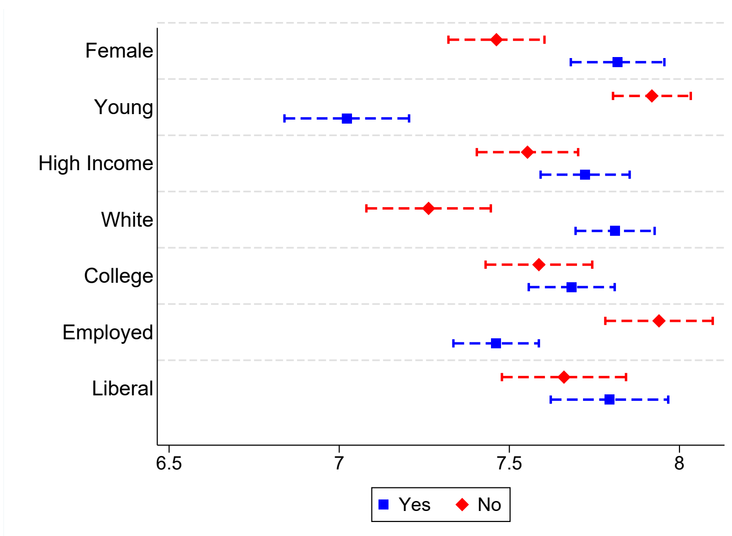

B. Support for Small Businesses

FIGURE 3. Heterogeneity in Support for Economic Policies

Notes: This figure shows how our support for economic policies measures vary across socio-demographic characteristics of the respondents. The sample consists of respondents in the Control video group. *Yes* indicates the respondent belongs to the given group in the y-axis, while *No* indicates otherwise. See Table 1 for a definition of each specific socio-demographic indicator variable. The sub-figures display the average and the 95% confidence interval. All outcomes are measured on a scale of 0 to 10, and they are defined in Section 3.2.

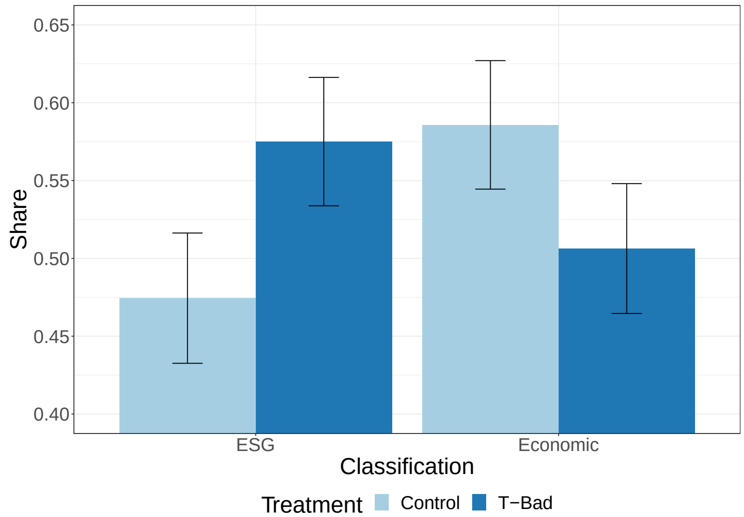

A. ESG vs. Economic Reasoning

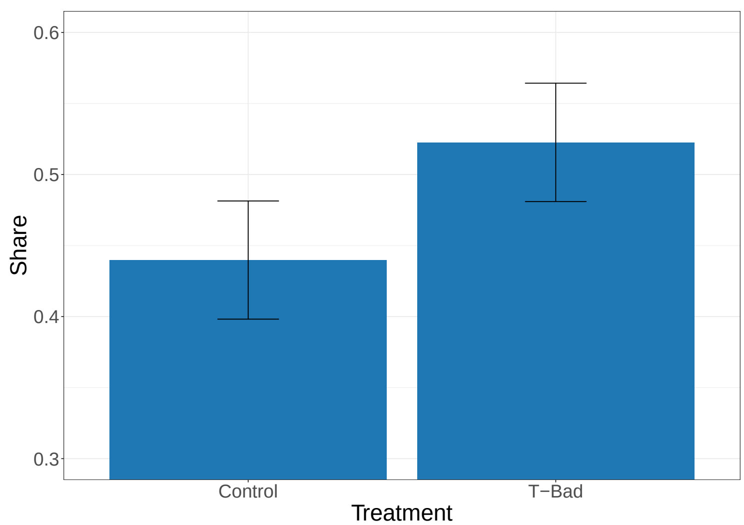

B. Negative Tone and ESG Reasoning

Figure 4. Type of Reasoning

Notes: In this figure we analyze the answers to the open-ended question that was included as part of our October 2020 survey, which reads: *"Could you tell us a bit more about why you have these views on policies regarding large corporations? What makes you being friendly or unfriendly with respect to helping large corporations?"* See Section 5.2.4 for details. The figure excludes responses that were classified as invalid or those that were placed in the "Other" category. The sample therefore consists of a total of 1,179 observations. Figure A shows the share of respondents in the T-Bad versus Control group who use ESG-based or Economic reasoning in their response. Figure B shows the share of responses in the T-Bad versus Control group that were classified as using ESG reasoning and negative sentiment.

Table 1. Sample and Balance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Share		***Univariate Balance***				***Joint Balance***			
Variables	***Data***	***CPS***	T-Salience	T-Bad	T-Good	T-Economy	T-Salience	T-Bad	T-Good	T-Economy
Female	0.51	0.52	-0.002	0.003	0.021	0.006	-0.002	0.005	0.022	0.005
			(0.894)	(0.847)	(0.215)	(0.744)	(0.891)	(0.738)	(0.205)	(0.753)
Young	0.30	0.32	-0.004	-0.026	-0.029	0.010	-0.004	-0.018	-0.023	0.024
			(0.738)	(0.112)	(0.115)	(0.594)	(0.771)	(0.309)	(0.277)	(0.232)
High income	0.52	0.54	0.002	0.003	-0.005	0.003	0.005	-0.005	-0.007	0.010
			(0.856)	(0.857)	(0.759)	(0.873)	(0.680)	(0.747)	(0.727)	(0.607)
White	0.70	0.78	0.001	0.012	0.028	0.021	-0.002	0.000	0.021	0.026
			(0.967)	(0.449)	(0.128)	(0.267)	(0.900)	(0.999)	(0.303)	(0.190)
College	0.57	0.42	-0.010	0.023	0.004	-0.007	-0.012	0.029*	0.009	-0.005
			(0.424)	(0.120)	(0.804)	(0.677)	(0.385)	(0.076)	(0.635)	(0.797)
Employed	0.61	0.61	-0.004	-0.027*	-0.027	-0.016	-0.002	-0.028*	-0.021	-0.017
			(0.777)	(0.078)	(0.123)	(0.377)	(0.883)	(0.082)	(0.245)	(0.371)
Liberal	0.31	-	-0.008	-0.009	0.011	-0.008	-0.006	-0.008	0.014	-0.008
			(0.550)	(0.561)	(0.550)	(0.654)	(0.650)	(0.644)	(0.446)	(0.671)
Observations	6,727	258,821,976	6,727	4,514	3,212	3,077	6,727	4,514	3,212	3,077
Joint significance: p-value	-	-	-	-	-	-	0.989	0.310	0.348	0.823

Notes: This table reports summary statistics on socio-demographic characteristics as well as the balance between treatment and control groups in our experiment. Column 1 reports the shares for our sample of survey respondents, while column 2 shows the same shares from the 2019 U.S. Current Population Survey (CPS). We check for balance in two ways: (i) through univariate regressions of an indicator variable equal to 1 if the individual is subject to a given treatment on each demographic characteristic separately (columns 3-6), and (ii) through multivariate regressions of an indicator variable equal to 1 if the individual is subject to a given treatment on all demographic characteristics jointly (columns 7-10). The sample for each column consists of all individuals in the specific treatment group and all individuals in the control group. *Female* is an indicator variable equal to 1 for females. *Young* is an indicator variable equal to 1 for individuals who are 35 years old or younger. *High income* is an indicator variable equal to 1 for individuals with a total household income of $70,000 or higher. *White* is an indicator variable equal to 1 for white or European American. *College* is an indicator variable equal to 1 for individuals who have completed a 4-year college or higher degree (Master's Degree, PhD, or Professional Degrees such as JD, MD and MBA). *Employed* is an indicator variable equal to 1 for individuals who are either business owners or are employed full-time or part-time. *Liberal* is an indicator variable equal to 1 for individuals identifying themselves as Very liberal or Liberal. *T-Bad* is an indicator variable equal to 1 for the sample of individuals subject to the Bad video treatment. *T-Good* is an indicator variable equal to 1 for the sample of individuals subject to the Good video treatment. *T-Economy* is an indicator variable equal to 1 for the sample of individuals subject to the Economy video treatment. *T-Salience* is an indicator variable equal to 1 for the sample of individuals subject to the Salience Treatment. p-values in parentheses.

Table 2. Treatments and Associated Cues

Treatment	Cue
No Salience + Control Video	$\emptyset$
No Salience + T-Bad Video	$(ESG, L, 0)$
No Salience + T-Good Video	$(ESG, H, 0)$
No Salience + T-Economy Video	$(ECON, \emptyset, 0)$
Salience + Control Video	$(ESG, \emptyset, 1)$
Salience + T-Bad Video	$(ESG, L, 1)$
Salience + T-Good Video	$(ESG, H, 1)$
Salience + T-Economy Video	$(ESG + ECON, \emptyset, 0)$

Notes: This table provides a taxonomy of the various treatments used in our experimental survey design. Each treatment is associated with a precise cue which primes a particular policy domain and valence framing, as described in our theoretical framework in Section 2.

TABLE 3. Perceptions of Large Corporations

Variables/Statistics	(1) mean	(2) *Are* median	(3) sd	(4) mean	(5) *Should Be* median	(6) sd	(7) *Big Business* Diff. in Means	(8) *Discontent* p-value
Executive Compensation (How Many Times Higher)	65.04	66.67	24.86	47.92	50.00	19.17	17.12	0.00
100-% Health Care Paid By Corporations	43.22	45.00	23.30	27.04	20.00	23.55	16.18	0.00
100-% Federal Income Tax Corporations Pay	68.02	80.00	26.97	56.87	65.00	24.35	11.15	0.00
100-% Women Executives	65.87	70.00	23.87	48.99	50.00	19.11	16.88	0.00
100-% Corporations Disclose CO_2	60.10	66.00	27.97	27.29	16.00	31.23	32.81	0.00
Political Donations (% of Corporations)	69.79	75.00	26.35	29.10	16.50	32.18	40.69	0.00
Care Only About Shareholders (% of Corporations)	51.32	50.00	27.74	30.01	30.00	24.16	21.31	0.00

Notes: This table provides summary statistics on perceptions of large corporations. The sample consists of respondents in the Control video group. We report perceptions of what individuals think large corporations' policies "Are" (columns 1-3) and "Should Be" (columns 4-6). Column 7 reports the difference between these two measures, i.e. the big business discontent. Column 8 tests for whether such difference is statistically significant. For each measure, a higher number indicates a less ESG-friendly corporation. All variables are measured on a scale of 0 to 100 and they are defined in details in Section 3.2.

Table 4. Correlation Between Perceptions and Support for Economic Policies

Variables	(1) Support for Bailouts	(2) Support for Small Businesses
Executive Compensation (How Many Times Higher)	-0.012***	-0.001
	(0.002)	(0.001)
100-% Health Care Paid By Corporations	-0.007***	0.000
	(0.001)	(0.001)
100-% Federal Income Tax Corporations Pay	-0.014***	-0.001
	(0.002)	(0.002)
100-% Women Executives	-0.004**	0.008***
	(0.002)	(0.002)
100-% Corporations Disclose CO_2	-0.004***	0.003***
	(0.001)	(0.001)
Political Donations (% of Corporations)	-0.009***	0.008***
	(0.001)	(0.001)
Care Only About Shareholders (% of Corporations)	-0.006***	0.006***
	(0.001)	(0.001)
Observations	6,727	6,727
Mean D.V. Control	5.424	7.641
SD D.V. Control	2.634	2.272

Notes: This table reports the correlation between the big business discontent measures (the regressors) and support for economic policies (the dependent variables). The specification is $Y_i = \delta + \sum_{j=1}^{j=7} \theta^j X_i^j + \epsilon_i$. *Support for Bailouts* represents how strongly individuals support corporate bailouts. *Support for Small Businesses* represents how strongly individuals support for small-business bailouts. All dependent variables (regressors) are defined in Section 3.2. At the bottom of the table we report mean and standard deviations of dependent variables measured using only information from the control group. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 5. Salience of Corporate Responsibility and Support for Economic Policies

VARIABLES	(1) Support for Bailouts	(2) Support for Small Businesses	(3) Support for Bailouts	(4) Support for Small Businesses
Treatment: T-Salience	-0.259**	0.021	-0.502***	-0.069
	(0.117)	(0.101)	(0.065)	(0.054)
Observations	2,038	2,038	6,727	6,727
Sample	Control	Control	Full	Full
Mean D.V Control	5.424	7.641	5.424	7.641
SD D.V Control	2.634	2.272	2.634	2.272

Notes: This table shows the treatment effects of our Salience experiment on support for economic policies. Columns 1-2 show the treatment effects of our Salience experiment on support for economic policies in the sample of respondents exposed to the Control video. The specification is $Y_i = \alpha + \beta T_i^{Salience} + \nu_i$. $T_i^{Salience}$ is an indicator variable equal to 1 for the sample of individuals subject to the Salience Treatment (and 0 otherwise). Columns 3-4 of this table show the treatment effects of our Salience experiment on support for economic policies in the full sample, while controlling for which video the respondent was exposed to. The specification is $Y_i = \alpha + \beta T_i^{Salience} + \sum_{j=1}^{j=3} \beta^j T_i^j + \nu_i$. $T_i^{Salience}$ is an indicator variable equal to 1 for the sample of individuals subject to the Salience Treatment (and 0 otherwise). *T-Bad* is an indicator variable equal to 1 for the sample of individuals subject to the Bad video treatment. *T-Good* is an indicator variable equal to 1 for the sample of individuals subject to the Good video treatment. *T-Economy* is an indicator variable equal to 1 for the sample of individuals subject to the Economy video treatment. *Support for Bailouts* represents how strongly individuals support corporate bailouts. *Support for Small Businesses* represents how strongly individuals support small-business bailouts. All dependent variables are measured on a scale in the range of 0 to 10 and are defined in Section 3.2. At the bottom of the table we report mean and standard deviations of dependent variables measured using only information from the control group. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 6. The Video Experiment: Framing of Corporate Responsibility and Support for Economic Policies

	(1)	(2)
Variables	Support for Bailouts	Support for Small Businesses
Treatment: T-Bad	-0.720***	0.084
	(0.079)	(0.067)
Treatment: T-Good	-0.152	0.289***
	(0.097)	(0.082)
Treatment: T-Economy	0.317***	0.268***
	(0.101)	(0.085)
Observations	6,727	6,727
Control for Salience	Yes	Yes
T-Bad vs T-Good	0.000	0.009
T-Bad vs T-Economy	0.000	0.026
T-Good vs T-Economy	0.000	0.825
Mean D.V Control	5.424	7.641
SD D.V Control	2.634	2.272

Notes: This table shows the treatment effects of our experiments on support for economic policies. The specification is $X_i = \lambda + \sum_{j=1}^{j=3} \phi^j T_i^j + S_i + \eta_i$. *T-Bad* is an indicator variable equal to 1 for the sample of individuals subject to the Bad video treatment. *T-Good* is an indicator variable equal to 1 for the sample of individuals subject to the Good video treatment. *T-Economy* is an indicator variable equal to 1 for the sample of individuals subject to the Economy video treatment. S_i is equal to 1 if the respondent was subject to the salience treatment (and 0 otherwise). *Support for Bailouts* represents how strongly individuals support corporate bailouts. *Support for Small Businesses* represents how strongly individuals support for small-business bailouts. All dependent variables are measured on a scale in the range of 0 to 10 and are defined in Section 3.2. The table also reports the p-value for the test of difference in the treatment effects across treatments. At the bottom of the table we report mean and standard deviations of dependent variables measured using only information from the control group. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

TABLE 7. Unpacking the Treatment Effects

		(1)
Treatment	Cue (i.e., Ω^*)	Support for Bailouts
Control + Salience	(ESG)	-0.259**
		(0.117)
T-Bad Only	(ESG, Bad)	-0.600***
		(0.112)
T-Bad + Salience	(ESG+, Bad)	-1.100***
		(0.111)
T-Good Only	(ESG, Good)	0.223
		(0.140)
T-Good + Salience	(ESG+, Good)	-0.760***
		(0.135)
T-Economy Only	(Economy)	0.406***
		(0.143)
T-Economy + Salience	(Economy+ESG)	-0.033
		(0.143)
Observations		6,727
Mean D.V. Control		5.553
SD D.V. Control		2.542

Notes: This table shows the treatment effects of all sub-layers of our experiment on support for economic policies. The specification is $Y_i = \alpha + \beta + \sum_{j=1}^{j=7} \beta^j T_i^j + \nu_i$. *Control + Salience* is an indicator variable equal to 1 for the sample of individuals in the control group subject to the Salience Treatment (and 0 otherwise). *T-Bad Only* is an indicator variable equal to 1 for the sample of individuals subject to the Bad video treatment only (and 0 otherwise). *T-Bad+Salience* is an indicator variable equal to 1 for the sample of individuals subject to the Bad video treatment and to the Salience treatment (and 0 otherwise). *T-Good Only* is an indicator variable equal to 1 for the sample of individuals subject to the Good video treatment only (and 0 otherwise). *T-Good+Salience* is an indicator variable equal to 1 for the sample of individuals subject to the Good video treatment and to the Salience treatment (and 0 otherwise). *T-Economy Only* is an indicator variable equal to 1 for the sample of individuals subject to the Economy video treatment only (and 0 otherwise). *T-Economy+Salience* is an indicator variable equal to 1 for the sample of individuals subject to the Economy video treatment and to the Salience treatment (and 0 otherwise). *Support for Bailouts* represents how strongly individuals support corporate bailouts and is measured on a scale in the range of 0 to 10 and is defined in Section 3.2. At the bottom of the table we report mean and standard deviations of the dependent variable measured using only information from the control group. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 8. Changing Perceptions with Animated Videos

Variables	(1) Executive Compensation	(2) Health Care	(3) Taxes	(4) Women Executives	(5) CO_2 Disclosure	(6) Political Donations	(7) Shareholders
Panel A: Big Business Discontent							
Treatment: T-Bad	5.390***	7.056***	3.499***	6.571***	12.427***	8.833***	7.389***
	(0.696)	(0.805)	(0.600)	(0.609)	(1.287)	(1.234)	(1.049)
Treatment: T-Good	2.415***	0.983	2.658***	1.103	4.989***	8.056***	0.663
	(0.852)	(0.987)	(0.735)	(0.746)	(1.577)	(1.512)	(1.285)
Treatment: T-Economy	1.360	0.840	-0.392	0.871	-0.267	0.109	2.758**
	(0.887)	(1.026)	(0.765)	(0.776)	(1.640)	(1.573)	(1.337)
Observations	6,727	6,727	6,727	6,727	6,727	6,727	6,727
Control for Salience	Yes	Yes	Yes	Yes	Yes	Yes	Yes
T-Bad vs T-Good	0.000	0.000	0.237	0.000	0.000	0.595	0.000
T-Bad vs T-Economy	0.000	0.000	0.000	0.000	0.000	0.000	0.000
T-Good vs T-Economy	0.287	0.901	0.000	0.789	0.004	0.000	0.161
Mean D.V Control	17.120	16.190	11.150	16.880	32.810	40.690	21.310
SD D.V Control	24.320	26.640	19.660	20.930	42.790	41.270	34.190
Panel B: Are							
Treatment: T-Bad	4.885***	6.036***	6.845***	8.896***	7.683***	2.494***	7.521***
	(0.715)	(0.694)	(0.747)	(0.656)	(0.818)	(0.791)	(0.816)
Treatment: T-Good	2.426***	0.309	5.986***	3.462***	-0.931	2.599***	0.811
	(0.875)	(0.850)	(0.915)	(0.803)	(1.002)	(0.969)	(1.000)
Treatment: T-Economy	1.447	0.049	0.222	0.636	0.619	0.355	1.509
	(0.911)	(0.885)	(0.951)	(0.836)	(1.042)	(1.008)	(1.040)
Observations	6,727	6,727	6,727	6,727	6,727	6,727	6,727
Control for Salience	Yes	Yes	Yes	Yes	Yes	Yes	Yes
T-Bad vs T-Good	0.004	0.000	0.332	0.000	0.000	0.911	0.000
T-Bad vs T-Econ	0.000	0.000	0.000	0.000	0.000	0.029	0.000
T-Good vs T-Econ	0.336	0.793	0.000	0.000	0.183	0.046	0.548
Mean D.V Control	65.040	43.220	68.020	65.870	60.100	69.790	51.310
SD D.V Control	24.860	23.300	26.970	23.870	27.970	26.350	27.740
Panel C: Should Be							
Treatment: T-Bad	-0.506	-1.019	3.346***	2.325***	-4.744***	-6.340***	0.131
	(0.547)	(0.689)	(0.703)	(0.528)	(0.928)	(0.925)	(0.717)
Treatment: T-Good	0.011	-0.674	3.329***	2.359***	-5.921***	-5.457***	0.149
	(0.670)	(0.844)	(0.861)	(0.647)	(1.137)	(1.133)	(0.878)
Treatment: T-Economy	0.086	-0.791	0.615	-0.235	0.886	0.246	-1.249
	(0.697)	(0.877)	(0.896)	(0.673)	(1.183)	(1.178)	(0.913)
Observations	6,727	6,727	6,727	6,727	6,727	6,727	6,727
Control for Salience	Yes	Yes	Yes	Yes	Yes	Yes	Yes
T-Bad vs T-Good	0.425	0.672	0.984	0.956	0.284	0.421	0.984
T-Bad vs T-Econ	0.381	0.788	0.002	0.000	0.000	0.000	0.119
T-Good vs T-Econ	0.923	0.905	0.007	0.001	0.000	0.000	0.171
Mean D.V Control	47.920	27.040	56.870	48.990	27.290	29.100	30.010
SD D.V Control	19.170	23.550	24.350	19.110	31.230	32.180	24.160

Notes: This table reports the estimates for the first stage, namely the impact of our treatments on some of our measures of perceptions, namely our primary measure of perception *"Big Business Discontent"* (Panel A), what individuals think large corporations policies *are* (Panel B) and *should be* (Panel C). The specification is $X_i = \lambda + \sum_{j=1}^{j=3} \phi^j T_i^j + S_i + \eta_i$. *T-Bad* is an indicator variable equal to 1 for the sample of individuals subject to the Bad video treatment. *T-Good* is an indicator variable equal to 1 for the sample of individuals subject to the Good video treatment. *T-Economy* is an indicator variable equal to 1 for the sample of individuals subject to the Economy video treatment. S_i is equal to 1 if the respondent was subject to the salience treatment (and 0 otherwise). All dependent variables are defined in details in Section 3.2. For each dependent variable measure, a higher number indicates a higher big business discontent, that is the respondent thinks large corporations are less ESG-friendly than they should be. The table also reports the p-value for the test of difference in the first stage coefficients across treatments. At the bottom of the table we report mean and standard deviations of dependent variables measured using only information from the control group. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 9. Heterogeneity Across Political Views and Age

	(1)	(2)
	Support for Bailouts	
Variables	T-Bad	T-Good
Panel A: Political Views		
Treatment x Liberal	-0.483**	-0.438*
	(0.190)	(0.231)
Treatment x Conservative	0.002	-0.012
	(0.192)	(0.239)
Treatment	-0.579***	-0.016
	(0.125)	(0.154)
Liberal	0.091	0.091
	(0.140)	(0.141)
Conservative	0.502***	0.502***
	(0.143)	(0.143)
Observations	4,514	3,212
Mean D.V. Control	5.424	5.424
SD D.V. Control	2.634	2.634
Panel B: Age		
Treatment x Young	-0.247	-0.472**
	(0.173)	(0.213)
Treatment	-0.641***	-0.015
	(0.095)	(0.116)
Young	0.687***	0.687***
	(0.127)	(0.127)
Observations	4,514	3,212
Mean D.V Control	5.424	5.424
SD D.V Control	2.634	2.634

Notes: This table shows heterogeneous effect of the treatments on support for bailouts, using as heterogeneity of interest the political orientation (Panel A) and age (Panel B) of the respondents. The specification in Panel A is: $Y_i = \alpha + \beta_L L_i \times T_i + \beta_C C_i \times T_i + \beta T_i + \alpha_L L_i + \alpha_C C_i + \nu_i$. The specification in Panel B is $Y_i = \alpha + \beta_Y Y_i \times T_i + \beta T_i + \alpha_Y Y_i + \nu_i$. *T-Bad* is an indicator variable equal to 1 for the sample of individuals subject to the Bad video treatment. *T-Good* is an indicator variable equal to 1 for the sample of individuals subject to the Good video treatment. We group respondents into three groups based on political orientation: Liberal (comprising Very liberal or Liberal), Moderate, and Conservative (comprising Very conservative or Conservative). *Young* is an indicator variable equal to 1 for individuals who are 35 years old or younger. *Support for Bailouts* represents how strongly individuals support corporate bailouts and it is measured on a scale in the range of 0 to 10. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 10. What Comes to Mind

Variables	(1) Share of Bad Thoughts	(2) Count of Bad Thoughts	(3) Share Good Thoughts	(4) Count of Good Thoughts	(5) Share Econ Thoughts	(6) Count of Econ Thoughts
Treatment: T-Good	0.063**	0.155**	−0.004	0.032	0.003	0.088
	(0.032)	(0.067)	(0.027)	(0.059)	(0.034)	(0.085)
Treatment: T-Bad	0.166***	0.352***	−0.039	−0.047	−0.042	0.004
	(0.032)	(0.067)	(0.027)	(0.059)	(0.034)	(0.085)
Observations	608	608	608	608	608	608
T-Bad vs T-Good	0.001	0.003	0.190	0.163	0.170	0.303
Mean D.V. Control	0.19	0.47	0.17	0.39	0.36	0.85
SD D.V. Control	0.30	0.66	0.29	0.60	0.37	0.87

Notes: This table measures what comes to respondents' minds after being exposed to our video treatments in our June 2022 survey. The specification is $Y_i = \alpha + \beta^{Good} T_i^{Good} + \beta^{Bad} T_i^{Bad} + \nu_i$. *T-Bad* is an indicator variable equal to 1 for the sample of individuals subject to the Bad video treatment. *T-Good* is an indicator variable equal to 1 for the sample of individuals subject to the Good video treatment. *Share of Bad/Good/Econ Thoughts* represents the share of thoughts from the Bad/Good/Econ group over the total number of thoughts the respondent selected. *Count of Bad/Good/Econ Thoughts* represents the count of thoughts from the Bad/Good/Econ group the respondent selected. The table also reports the p-value for the test of difference in the coefficients across treatments. At the bottom of the table we report mean and standard deviations of dependent variables measured using only information from the control group. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Table 11. Treatment Effects: Behavioral Outcome Measures

Variables	(1) Support for Bailouts	(2) Sign Petition	(3) Email Senators to Support Bailouts	(4) Email Senators to Oppose Bailouts	(5) Donation to Business Roundtable
Treatment: T-Bad	-0.719***	-0.108***	-0.089***	0.037*	-2.015***
	(0.140)	(0.023)	(0.019)	(0.022)	(0.408)
Observations	1,683	1,683	1,683	1,683	1,683
Mean D.V. Control	5.386	0.420	0.223	0.282	7.433
SD D.V. Control	2.830	0.494	0.417	0.450	8.864

Notes: This table shows the treatment effects of our experiments on the behavioral outcome measures we collect in our October 2020 survey. The specification is $Y_i = \alpha + \beta_i^{T-Bad} + \nu_i$. *T-Bad* is an indicator variable equal to 1 for the sample of individuals subject to the T-Bad treatment. *Sign Petition* is an indicator variable for whether the respondent indicated she would sign the petition to support corporate bailouts. *Email Senators to Support Bailouts* is an indicator variable for whether the respondent agreed to include her name in the message to the U.S. Senators to support corporate bailouts. *Email Senators to Oppose Bailouts* is an indicator variable for whether the respondent agreed to include her name in the message to the U.S. Senators to oppose corporate bailouts. *Donation to Business Roundtable* is the total amount of U.S. dollars the respondent indicated she would like to donate to the Business Roundtable. All dependent variables are explained in more details in Section 6.3. At the bottom of the table we report mean and standard deviations of dependent variables measured using only information from the control group. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.