Consuming Values

Jacob Conway, *Stanford University**

Levi Boxell

October 24, 2023

(Click Here for Latest Version)

Abstract

We study the extent to which individuals' consumption decisions are influenced by firms' stances on controversial social issues and the implied incentives for firms to take such stances. We use transactions from a major payment card company to predict cardholders' likely social alignment with firm stances and to quantify effects on consumption. The social stances taken by firms increase revenue on average, with significant heterogeneity across consumers and firm stances. Consumers most aligned with a firm's social stance increase their consumption at the firm by 19 percent in the month following widely-known social stance events, and consumers most opposed to the firm's stance decrease their consumption by 11 percent. These diverging consumption responses decrease in magnitude over time but persist even a year later. Firms tend to take stances that align with the social preferences of their consumer base. The direction of a firm's social stance is also correlated with the firm's ownership structure and its employees' social preferences.

*E-mail: jcconway@stanford.edu, leviboxell@gmail.com. We thank Matthew Gentzkow, Liran Einav, and Amit Seru for their invaluable mentorship and advice. We thank Marianne Bertrand, Pradeep Chintagunta, Emanuele Colonnelli, Jose Ignacio Cuesta, Michael Dinerstein, Kwabena Donkor, James Druckman, Jean-Pierre Dubé, Dana Foarta, Wesley Hartmann, Caroline Hoxby, Guido Imbens, Emir Kamenica, Yewon Kim, James Lattin, John List, Neale Mahoney, Sanjog Misra, Sridhar Narayanan, Muriel Niederle, Lubos Pastor, Claudia Robles-Garcia, Navdeep Sahni, Brad Shapiro, Jann Spiess, Avner Strulov-Shlain, Amir Sufi, Stefan Wager, Jason Weitze, David Yang, Constantine Yannelis, Ali Yurukoglu, Anthony Lee Zhang, Luigi Zingales, and seminar participants at Stanford University for valuable comments and suggestions. We gratefully acknowledge financial support from the National Science Foundation (grant number: DGE-1656518), the Institute for Research in the Social Sciences (IRiSS) at Stanford University, the Institute for Humane Studies, the Stanford Center on Philanthropy and Civil Society, the B.F. Haley and E.S. Shaw Fellowship for Economics (through a grant at the Stanford Institute for Economic Policy Research), and the George P. Shultz Dissertation Support Fund at SIEPR. We thank the data science team at a payment card company for generously providing data used in this paper. Our research was approved by the Stanford IRB (eProtocol 53506).

1 Introduction

Consumer surveys and firm behavior suggest that individuals care about firms' values when making consumption decisions. Most consumers state that they have made or avoided a purchase due to the social values of the company, that they are attracted by firm values that align with their own, and that they are willing to pay more for sustainable goods (Nielsen, 2015; ANES Data Center, 2016; Barton et al., 2018). Firms also increasingly espouse social goals beyond narrow profit-maximization and devote significant resources to improving their social impacts (Meier and Cassar, 2018; Business Roundtable, 2019; Rajan et al., 2022). Consumer-facing firms in particular spend more on their environmental, social, and governance (ESG) efforts in some contexts than firms that sell to businesses (Kitzmueller and Shimshack, 2012).

Despite this suggestive evidence, it is unclear to what extent individuals' consumption decisions respond to firms' values and how this response impacts firms' profits and decision-making. Some researchers hypothesize that consumers' self-reported demand may be "cheap talk" in that stated preferences may not reflect their actual purchase decisions (Auger and Devinney, 2007). The magnitude of this response matters for society as it determines the extent to which consumers can incentivize or discipline firm behavior. Some have argued that socially conscious consumerism can effectively cause firms to internalize their social externalities (Barboza et al., 2021; The Economist, 2021a), while others have argued that it is too weak to do so and may crowd out other forms of regulation (Csutora, 2012; Wicker, 2017; Sheffi, 2021).

One increasingly important setting for potential engagement occurs when firms take controversial social stances, such as advertising campaigns featuring divisive racial justice protesters, corporate policy on contraceptives and reproductive rights, or comments regarding sexual orientation and gender legislation (Lin, 2022). These controversial stances have become more frequent in recent years (Klostermann et al., 2021; The Economist, 2021b), while consumers have also become more socially divided (Iyengar and Westwood, 2015) and increasingly report caring about and seeking out information on firm stances (Global Strategy Group, 2018). One important hypothesis is that firms take social stances only to the extent that doing so increases their profits, and testing this hypothesis is important for our understanding of corporate governance and how consumers and other stakeholders influence firm behavior.

In this paper, we study the extent to which individuals' consumption decisions respond to their social alignment with firms around events in which firms took salient and controversial social stances, and we analyze the implied incentives for firms to take such stances. We use consumer transactions from a major payment card company to predict cardholders' likely social alignment with firm stances and to quantify effects on consumption. We estimate that observed firm stances

increase revenue on average, with considerable heterogeneity across consumers and events. We find that consumers whose social views are likely to be aligned with a firm's stance increase their consumption at the firm in the months following the social stance event and that consumers likely to be socially-opposed to the firm's stance decrease their consumption. Social stances thus typically have more positive revenue impacts when the stance better aligns with the views of the firm's customers, and revenue-maximizing social stance decisions therefore vary across firms. We also show that in practice, firms tend to take stances that align with their consumers' and employees' social preferences and that correlate with the firm's ownership structure.

We start by building a dataset of 117 events in which controversial social stances were taken by firms within our transaction data. We identify these events in part by searching systematically for unusual spikes combining the firm name with keywords indicative of social stances in either Google Trends searches or in news articles. We extend this list of events based on contemporaneous brand perception surveys and queries to a large language model (OpenAI's GPT-4). We measure consumer awareness of each event in contemporaneous surveys that asked respondents whether they had recently heard any good or bad news about the firm.

We then use the credit and debit card transaction data to measure the effect of each event on consumption at the firm, quantifying overall impacts as well as heterogeneity across consumer groups that are likely more or less aligned with the firm's stance. This disaggregation is valuable because demand for social alignment with firms would predict heterogeneous responses between these different consumer groups, and this would matter for the incentives of firms with different consumer bases to take stances.

To enable this disaggregation, we infer cardholders' likely alignment with firm stances from their transactions and demographics. We first identify more than 30 million consumers who have clearly expressed their likely alignment on social issues through their donations to PACs, charitable organizations, and other non-profits. We use these donors to train a model that predicts an individual's likely alignment based on a wider set of their other transactions and demographics, which we then use to predict likely social alignment among all non-donors. Using this predicted alignment, we quantify the distribution of revenues across more aligned/opposed consumers for each firm in the year preceding its social stance event.

We then account for changes in consumption at the firm unrelated to its stance by predicting the counterfactual consumption that would have occurred absent the firm's stance based on consumption at related firms. Following the synthetic difference-in-differences approach of Arkhangelsky et al. (2021), this prediction draws from contemporaneous consumption at each of the thousands of other firms in the economy and from past consumption at the firm taking a social stance. This synthetic series closely tracks consumption at the firm prior to its social stance.

Comparing actual consumption to this counterfactual, we estimate that observed firm social stances have positive revenue impacts on average, with considerable heterogeneity across consumers and events. To illustrate how the magnitude of these responses varies with consumer awareness, consider a stance that 25 percent of consumers report hearing about in contemporaneous surveys, which would be the fourth largest event and at the 97th percentile across events in our data. On average, we estimate that such a stance increases overall revenues by 3 percent in the month following the firm's stance. Our estimated impacts in subsequent months are weakly positive on average but not statistically distinguishable from zero (at a 95% significance level).

Disaggregating this overall response among consumers by alignment, we estimate starkly diverging responses that provide clear evidence of consumer demand for social alignment with firms. We find that consumers most aligned with highly salient firm stances increase their consumption at the firm in the following month by 19 percent, and that consumers most opposed to the firm's stance decrease their consumption at the firm by 11 percent following the firm's stance. These consumption responses attenuate over time but persist even a year later. Among consumers predicted to be more weakly aligned or opposed, we similarly see diverging consumption responses depending on alignment with the firm's stance, but with smaller response magnitudes than among consumers at the extremes of predicted alignment.

Turning to the supply-side implications of our consumer response estimates, we analyze when firm stances increase revenue for hypothetical firms facing different baseline consumption shares across consumer social alignment groups. For example, we show how revenue-maximizing social stance decisions vary for typical firms depending on the state or industry in which they operate.

Decision-makers taking firm social stances may care not only about impacts on revenues, but may also seek to align with their own preferences or the preferences of other stakeholders. For example, companies often face internal pressure from employees (Maks-Solomon and Drewry, 2021) as well as external pressure from their shareholders and owners (Baron, 2009) on social issues.[1] Combining our stances with measures of the preferences of a firm's different stakeholders, we analyze which stakeholders' preferences best predict the direction of a firm's stance and how this interacts with the firm's corporate governance structure. The direction of a firm's stance is best predicted by the preferences of its consumers and employees, as well as by its public vs. private ownership status. In contrast, the personal preferences of a firm's CEO and corporate board are

[1]Relatedly, there is an active normative debate regarding what the purpose and goals of a corporation *should* be. Normative theory in business ethics has long been dominated by the Friedman doctrine (Friedman, 1970), which argues that firms are beholden to their shareholders (i.e., "shareholder primacy"). More recently, however, the main business association for CEOs has argued that companies should also commit to benefiting their customers, employees, suppliers, and communities (Business Roundtable 2019, i.e., "stakeholder theory").

not strong predictors.

Our paper contributes to several existing literatures. The first analyzes socially-conscious consumerism, quantifying the extent to which consumers' preferences on social or environmental issues impact their purchase decisions. Closest to our own paper are studies which examine consumer responses to controversial firm social stances and the net impacts of these stances. For example, Liaukonyte et al. (2022) analyze a controversial social stance by Goya, finding evidence of increased consumption at store locations in counties home to many consumers likely aligned with the firm's stance ("aligned buycotts"). They do not find similar evidence of boycotts in more opposed areas, and thus estimate that Goya's stance had a positive net impact on its sales in the following weeks. Similarly, Painter (2021) uses smartphone-location data to quantify foot traffic responses to a Walmart statement on gun control. Painter again finds increases among locations in socially-aligned counties, consistent with aligned buycotts. In contrast to Liaukonyte et al. (2022), Painter finds decreases in foot traffic at stores in generally socially-opposed counties ("opposed boycott"), resulting in a negative overall impact on foot traffic to Walmart relative to local competitors. Klostermann et al. (2021) analyze the impact of controversial firm social stances on self-reported favorability towards the firm in YouGov BrandIndex data, finding negative overall impacts on favorability on average. Hydock et al. (2020) similarly find evidence of aligned buycotts, opposed boycotts, and negative overall impacts on average when providing information about firm stances in unincentivized survey experiments. Schoenmueller et al. (2023) find evidence that after some types of firm stances, the firm's Twitter following shifts towards users who are likely aligned with the firm's stance.[2] These papers provide mixed evidence on the existence and relative magnitudes of aligned buycott and opposed boycott responses to controversial firm social stances, and on the net revenue impacts caused by firm stances.[3]

Our paper makes several contributions to this existing literature on socially-conscious con-

[2]Schoenmueller et al. (2023) focus primarily on documenting increased polarization in consumer behavior following the 2016 U.S. presidential election.

[3]Related research on socially-conscious consumerism analyzes consumption responses to other forms of firm social engagement or social impact. For example, Panagopoulos et al. (2020) experimentally manipulate consumer beliefs about socially impactful behaviors and then observe subsequent choices between firm gift cards. They find evidence that consumers are more likely to choose gift cards from firms that align with personal social preferences. There is also a related literature examining consumer responses to foreign-policy, with some finding evidence of a significant consumer response (e.g., Chavis and Leslie 2009; Fuchs and Klann 2013; Heilmann 2016; Pandya and Venkatesan 2016; Chen and Zhong 2019; Fouka and Voth 2021) and others finding no such evidence (e.g., Ashenfelter et al. 2007; Davis and Meunier 2011) in different contexts. Another group of papers, such as Elfenbein et al. (2012), Bartling et al. (2014), Barrage et al. (2020), and Hart et al. (2022), investigates consumer responses to firm activities on less controversial social issues on which most consumers hold similar views.

sumerism and on consumer responses to firm social stances. First, we focus on actual consumption choices made by consumers representing a large and representative portion of firm revenues, rather than relying on self-reported survey responses or other proxies that might not reflect true consumer behavior. Second, we measure individuals' social alignment and heterogeneous consumption responses at high time frequencies, thereby strengthening our identification relative to papers analyzing data at lower frequencies or using less precise proxies of alignment based on the geography of store locations. Third, we systematically identify and analyze a larger number of social stance events. This allows us to provide robust evidence of aligned buycotts and opposed boycotts, to quantify the ensuing revenue tradeoff between these two countervailing effects, to explain heterogeneous impacts across different events, and to better reconcile the mixed evidence in the existing literature. Lastly, our analysis delves into the supply-side implications of consumer demand, examining the incentives of profit-maximizing firms to engage with controversial social issues.

We also build on a literature analyzing the impacts and drivers of firms' ESG (Environmental, Social, and Governance) or CSR (Corporate Social Responsibility) behavior. This includes work in a variety of research fields including economics (reviewed in Kitzmueller and Shimshack, 2012), marketing (e.g., Hydock et al., 2019), and management science (e.g., McWilliams and Siegel, 2001). This literature has analyzed such firm behavior in relation to other (non-consumer) stakeholders, including work on employees (e.g., Hedblom et al., 2019; Burbano, 2021), on financial performance (e.g., Dimson et al., 2015), on investors (e.g., Larcker and Watts, 2019; Bonnefon et al., 2022; Broccardo et al., 2022), and on local governments (e.g., Bertrand et al., 2020). This literature provides the strongest support for impacts on employees, with weaker and mixed evidence of impacts on financial performance and investors. Contemporaneous work in Barari (2023) looks at the preferences of different stakeholders' as potential predictors of firms' controversial speech online,[4] finding moderate correlations between the firms' choice of language and proxies for the preferences of potential consumers, employees, and elected officials, without quantifying connections to firm profits. We contribute to this literature by providing strong evidence on the impacts of consumers' social preferences on firm behavior, and by considering firms' social stances on particularly controversial and salient issues.

[4] Firms have increasingly used controversial language in their corporate communication online (Cassidy and Kempf, 2022) and have experienced increased homophily in their executive teams (Fos et al., 2021).

By considering the relative importance of different stakeholders' preferences as drivers of firm behavior, our paper also contributes to a literature on corporate governance and agency problems within the firm. For summaries of this literature, see Stein (2003), Becht et al. (2003), and Milosevic et al. (2015). Our analysis also relates to debates around stakeholder theory vs. shareholder primacy as firm objective functions (e.g., Friedman, 1970; Hart and Zingales, 2017). We contribute to this literature by analyzing realized firm behavior in a social stance context in which we can precisely quantify revenue impacts of firm actions and in which we can measure the (potentially competing) personal preferences of different stakeholders.

The remainder of the paper is structured as follows. Section 2 provides our conceptual framework modeling a revenue-maximizing firm's decision to take a social stance and highlights key empirical targets. Section 3 describes the data sources we use to estimate these empirical targets. Section 4 describes our event selection procedure and quantification of event size. In Section 5, we turn to analysis of our transaction dataset and describe our measurement of individual consumers' social preferences. Section 6 presents our synthetic difference-in-differences procedure for imputing no-event counterfactual consumption and our resulting event study estimates of overall and disaggregated consumption responses. With these empirical targets estimated, in Section 7 we return to our conceptual framework and discuss the supply-side implications of these consumer response estimates along with other potential drivers of firm behavior. Section 8 concludes.

2 Conceptual Framework

In this section, we provide a conceptual framework to illustrate the consumer demand elasticity we wish to estimate, the tradeoffs firms face when deciding whether and how to take social stances, and the key parameters that determine optimal firm behavior.

2.1 Firm Problem

We consider a single firm choosing to either take a stance on a binary social issue or choosing not to take a stance, denoting its stance decision as $s \in \{\mathbf{F}or, \mathbf{A}gainst, \mathbf{N}one\}$. There exists a continuum of consumers that can be partitioned into different groups (g). The net present value of revenues from each group may depend on the firm's stance and adds up to total revenue at the firm: $\sum_{g \in G} y_g(s) = y(s)$. Consumers are independently aware of the firm's stance with probability τ, and otherwise

believe that the firm has not taken a stance.[5] The firm seeks to maximize its revenue, $y(s)$. This is equivalent to maximizing the overall revenue growth induced by its social stance decision, which for estimation purposes can be split into the product of three terms that summarize our empirical targets:

$$\max_{s\in\{\mathbf{F},\mathbf{A},\mathbf{N}\}} \underbrace{\frac{y(s)-y(\mathbf{N})}{y(\mathbf{N})}}_{\substack{\text{Overall}\\ \text{Revenue Growth}}} = \sum_g \underbrace{\frac{y_g(\mathbf{N})}{y(\mathbf{N})}}_{\text{Baseline Share}} \times \underbrace{\tau}_{\text{Awareness}} \times \underbrace{\frac{[y_g(s)-y_g(\mathbf{N})]\tau^{-1}}{y_g(\mathbf{N})}}_{\substack{\text{Consumption Responsiveness}\\ \text{(Conditional on Awareness)}}} \tag{1}$$

The overall revenue growth induced by its stance is a weighted average of group-specific revenue growth responses, with weights given by the share of consumption dollars a firm would receive from a given group if it did not take a stance (which we refer to as baseline shares). The induced consumption growth of a given group can be split into the product of two terms: the share of consumers aware of a firm's stance, and the group's consumption response conditional on awareness. This split is useful when comparing responses to social stance events with varying levels of consumer awareness, as the induced consumption growth scales linearly with the share of consumers aware of the firm's stance.[6] We can therefore think of τ as a measure of treatment intensity or event size that varies across potential events.

In this stylized model, firms face a potential tradeoff in catering to the preferences of different groups when taking controversial social stances. For example, suppose that there are two consumer groups denoted by their social views on this issue (i.e., $G = \{\mathbf{F}or, \mathbf{A}gainst\}$). The firm will prefer taking an **F** stance to no stance if and only if the consumption increase among the aligned (**F**) group is at least as large as the decrease among the opposed group. The net revenue impact of an **F** stance by the firm is more positive if aligned consumers account for a larger baseline share and/or if aligned consumers have greater consumption responsiveness (conditional on awareness) than opposed consumers. Consumer awareness (τ) affects the magnitude of revenue impacts, but does not affect the firm's optimal social stance decision given the assumption that awareness is constant across consumer groups.

[5] Although consumer awareness could in principle vary across groups, we show in Section 4 that consumer awareness is empirically similar across groups. Our conceptual framework and subsequent empirical analysis therefore assumes that consumer awareness of a given event does not vary across groups.

[6] To see this linearity in τ, define $\tilde{y}_g(s)$ as the consumption by group g at the firm that would occur if all group members were aware of the firm's stance, thus $y_g(s) = \tau \times \tilde{y}_g(s) + (1-\tau) \times y_g(\mathbf{N})$. Then $[y_g(s) - y_g(\mathbf{N})]/y_g(\mathbf{N}) = \tau \times [\tilde{y}_g(s) - y_g(\mathbf{N})]/y_g(\mathbf{N})$.

2.2 Key Empirical Targets

Estimating the terms in Equation 1 requires: identifying salient social stance events (s); measuring consumer awareness of each firm stance (τ); separating consumers into different groups (g) with likely similar social alignment; and assembling data that allows us to measure consumption at firms by each group and at different times ($y_g(s)$). We can then reasonably proxy for baseline shares ($y_g(\mathbf{N})/y(\mathbf{N})$) using consumption shares during the year preceding the firm's social stance event. The final term left to then be estimated in the equation above is $y_g(\mathbf{N})$, the counterfactual consumption that would have occurred if the firm had not taken its stance. We can predict this value based on contemporaneous consumption at other firms as well as the firm's historical seasonal patterns. We estimate each of these targets in the subsequent sections. These estimated parameters allow us to quantify the typical strength of consumer demand responses, to test the optimality of existing firm social stances, and to analyze the optimal behavior of a firm facing consumers with arbitrary baseline shares.

3 Data

In this section, we summarize the data sources we use to estimate the key empirical targets highlighted by our conceptual framework.

3.1 Transaction Data

We primarily use credit- and debit-card data from a large payment card company, which allows us to measure individuals' actual consumption at particular firms over time. This dataset contains transactions in the U.S. from 2008 through March 2023, and covers approximately 20% of all U.S. consumption. The dataset is longitudinal and transactions can be linked at the card-level. For each transaction, we observe the date, dollar amount, and merchant (along with other information). The transaction data is depersonalized, so name, address, and other personal information about the cardholder is not observable, other than what can be inferred given the card's transaction history. The data also do not specify which goods or services were purchased from a particular merchant, nor the prices of those items. Transaction data, in the aggregate, may be combined with depersonalized demographic data from consumer credit reports.[7] This demographic information includes

[7] Aggregate demographic data from consumer credit reports is not available for cards that were only active during earlier years or for debit cards. This data instead covers only credit cards that were active in recent years, representing

the cardholder's home census block, gender, age, and estimated household income.[8] We use this transaction dataset to impute cardholders' likely alignment on social issues (forming groups g from our conceptual framework), to measure a firm's baseline shares across these groups ($y_g(\mathbf{N})/y(\mathbf{N})$), to predict the counterfactual consumption that would have occurred had a firm not taken a social stance ($y_g(\mathbf{N})$), and to measure actual consumption by these groups at the firm over time ($y_g(s)$).

3.2 Other Complementary Data Sources

Our analysis also relies on several other complementary data sources, which we use to identify social stance events and measure consumer awareness of each event (s and τ), and to analyze related outcomes associated with these events. One key such data source is YouGov's BrandIndex dataset of contemporaneous brand perception surveys of consumers, in which YouGov surveys a nationally representative sample of at least 5,000 people each day (from their panel of 4+ million U.S. respondents) about their perceptions of more than 2,000 brands operating in the U.S. Importantly for our analysis, YouGov has been collecting this data continuously since 2007, allowing us to analyze changes in respondents' contemporaneous perceptions of firms during the period surrounding their social stance event.[9] YouGov also collects a large number of profile variables for each respondent (including information about their demographics, party affiliation, location, attitudes, and behaviors), allowing us to separately analyze responses among consumers with likely different social alignment starting in November 2012. We use BrandIndex data primarily to measure consumer awareness of firms' social stances, as well as to analyze respondents' interpretation of social stance news and their self-reported consumption responses.[10]

We use data from Google Trends and from ProQuest's U.S. Newsstream primarily to identify salient firm social stances, and to construct alternative proxies of events' salience to consumers. Google Trends data consists of daily relative search frequencies for given keywords on Google, which can be compared over time, across search terms, and/or across geographies. The coverage

15% of cards in our sample.

[8] At no point do we analyze consumption responses at the level of an individual or card, instead aggregating cards into large groups prior to our analysis of consumption responses.

[9] Outcome metrics available in BrandIndex include the share of respondents who: have heard something positive and/or negative about the brand in the past two weeks; have a positive vs. negative impression of the brand; would consider purchasing from the brand; are most likely to choose that brand relative to its industry competitors; are current and/or former customers of that brand; would feel proud vs. embarrassed to work for the brand; and would recommend others to purchase from vs. avoid the brand.

[10] For additional information regarding BrandIndex data, including sampling methodology and complete text for all survey questions used, see Appendix Section A.3.

of this data begins in 2006. ProQuest's U.S. Newsstream dataset contains full-text news articles published by more than 350 U.S. print and online newspapers, and is intended as a comprehensive collection of U.S. news that is available throughout our analysis period. For each news article, we observe the full text, the publication date and outlet, and additional metadata including the names of firms mentioned as subjects in the article. For additional detail on these different data sources and the precise construction of variables used in our analyses, see Appendix Section A.

4 Event Selection and Consumer Awareness

In this section, we describe how we systematically identify events in which firms took controversial and salient social stances (s) and measure consumer awareness (τ) of each event based on Google Trends searches, news reports, contemporaneous consumer surveys, and queries to a large language model.

4.1 Identifying Candidate Social Stance Events

We construct a dataset of 117 salient social stance events that were associated with particular firms, had a clear event date, and which were likely to affect consumer perceptions of a firm's social values. We restrict our analysis to events that occurred between 2011 and 2022Q1, inclusive, to align with the coverage of our transaction data and empirical methods.[11] Examples of the social stance events we identify include a controversial advertising campaign related to racial justice, stances on widely debated LGBTQ rights and legislation, corporate policy regarding the provision of contraceptives or abortion access, and salient stances on gun control issues or voting legislation.

We combine several different methods to identify these candidate corporate social stance events, which we overview in this section and describe in more detail in Appendix Section A.1. We first implement a procedure to identify candidate events by searching systematically for spikes in daily Google Trends searches for a given firm name and for the firm name and keywords indicative of social stances, using keywords like "*transgender*" or "*gun control*" and repeating this search for each of the 10,000 largest U.S. firms by revenue.[12]

[11] The transaction dataset we use covers 2008–2023Q1, and our empirical method requires data three years prior to and one year after the event date.

[12] Additional detail on this procedure can be found in Appendix Section A.1, including a complete list of searched keywords and a description of how keywords were chosen. We use data from D&B Hoovers to identify the largest U.S. firms by revenue.

We also implement a similar approach to identify candidate events based on news coverage in ProQuest's U.S. Newsstream, looking for unusual spikes in the number of news articles that mention firm names together with keywords indicative of social stances. We complement our news-based approach using an existing list of firm stances from Klostermann et al. (2021), which identifies events by searching for any individual news articles that contain their own set of keywords indicative of corporate stances.

While the vast majority of events we analyze are selected by these keyword-based Google Trends and news methods, we complement these methods with two additional approaches based on brand perception surveys and queries to a large language model to ensure that we have not omitted salient events due to our choice of keywords.[13] In the first such complementary method, we identify candidate social stance events based on different favorability responses towards a firm among two groups of respondents who hold likely opposite views (based on their party affiliations) in contemporaneous brand perception surveys from YouGov's BrandIndex dataset. We further extend our list of candidate social stance events by querying OpenAI's GPT-4 large language model for a list of the most salient events in which firms took stances on controversial social issues in the U.S., considering the top 50 most salient events returned by GPT-4 as candidate social stance events.[14]

Taking the union of candidate firm-dates generated by the four methods above, we then manually filter this list by using news coverage to confirm the existence and exact timing of a social stance event while removing candidate firm-dates not associated with a social stance event. Using the consumer awareness event-size measure defined below, we also exclude rare candidate firm-dates that occur within two years of a larger event at the same firm, as well as three candidate events that were estimated to have a weakly negative event size. We typically choose the final event date based on the earliest news coverage of a given event. We note that events are often each selected by multiple methods, and that our main results are robust to dropping any one method from our event selection procedure.

We provide generic descriptions for each of the 117 selected firm social stance events (we denote this set J) in Appendix Table B1, also providing for each event the year, direction of alignment with our consumer clusters (as defined in Section 5.1), and the share of consumers we estimate

[13] Appendix Section A.1 lists the share of events identified by each of our event selection methods. In practice, 9.4% of events in our final sample were not identified by our keyword-based methods and were added by our brand perception and/or large language model methods.

[14] The full text of the prompt provided to GPT-4 via ChatGPT can be found in Appendix Section A.1.

were aware of the firm's stance.[15]

4.2 Quantifying Event Size Based on Awareness in Contemporaneous Consumer Surveys

Having selected a set of social stance events, we use contemporaneous surveys from YouGov's BrandIndex dataset to measure consumer awareness of firms' social stances. To do so, we first define the intermediate series a_{jt} as the share of BrandIndex respondents in event-time month t who report having heard something positive and/or negative about firm j in the past two weeks.[16] We then define our BrandIndex-based estimate of consumer awareness as $\hat{\tau}_j := \frac{a_{j0}-a_{j,-1}}{1-a_{j,-1}}$, i.e., the pre- vs. post-month change in the share of respondents who have heard good or bad news about the firm, scaled by the share of respondents who were not already reporting having heard news about the firm.[17] For the 15 percent of events that are not covered by the BrandIndex dataset, we impute consumer awareness of these events based on changes in news coverage and Google Trends searches for the firm.[18]

In Figure 1 Panel A, we show variation in average consumer awareness over time for our social stance firms. BrandIndex respondents report hearing good and/or bad news about firms at fairly constant rates in the months before their social stance. In the month of the firm's social stance, consumer awareness increases by 5 percent on average. This consumer awareness measure varies significantly across events as shown in Figure 1 Panel B, with consumer awareness around 40 percent for the most salient social stances, whereas the 75th percentile and median values are

[15] These events represent social stances taken by 96 unique firms. Results are similar if we restrict to the first or most salient event for each firm.

[16] When analyzing variation over time around firm events, we define event-time "months" as four-week periods relative to the firm's event. Month $t = 0$ is defined as the four-week period starting with the day of the firm's event, with month $t = -1$ then denoting the four weeks directly preceding the firm's event.

[17] This scaling accounts for the fact that responses will not change among respondents who would have already reported hearing news about the firm in the absence of its stance. It is theoretically justified if news about the firm other than its social stance is constant over time and independent of news about the firm's social stance (see Appendix Section A.2 for detail). This scaling is also closely related to the literature on persuasion rates (see DellaVigna and Gentzkow, 2010), which similarly scales changes by the share of respondents who could possibly be converted by a treatment.

[18] See Appendix Section A.2 for detail on this imputation procedure. We choose our imputation method via cross-validation. We prefer BrandIndex-based measures of consumer awareness (when available) to measures based on news coverage and Google Trends searches, as this preferred measure is most closely related to the empirical target τ highlighted by our conceptual framework. It also appears more stable over time in the absence of an event (e.g., exhibits less seasonality and noise) than alternative measures, and avoids potential issues when comparing the news mentions or searches of firms that vary in their name's commonality or potential for variants. Estimates based on Google Trends also face a concern that searches for a given firm could in part reflect purchase intent (i.e., searching for their website in order to buy a product) and might therefore directly reflect changes in consumer demand. Our results are robust to excluding events for which consumer awareness was imputed by news coverage and Google Trends.

much smaller at 6 percent and 3 percent respectively. This variation highlights the need to scale observed consumption responses relative to consumer awareness, as discussed in Section 2.

In Appendix Figure B1, we report consumer awareness separately among respondents depending on their likely alignment with the firm's stance (based on their self-reported party affiliations). In Panel A, we see similar magnitude spikes in consumer awareness among consumers who are likely aligned with the firm's stance, opposed to the firm's stance, or who are less strongly socially aligned/opposed. These spikes in the month of the firm's social stance are slightly larger in magnitude on average among opposed vs. aligned consumers, but this difference is not statistically significant at standard significance levels (5.4% vs. 4.2% awareness respectively, p-value=0.17). Similarly plotting the mean of a_{jt} itself over time by respondent alignment in Panel B, we note that aligned and opposed respondents report hearing good and/or bad news about the firm at nearly identical rates on average in the month of the firm's stance (20.7% vs. 20.7% among aligned and opposed respondents, respectively).[19] In the empirical analysis below, we assume that awareness does not vary with alignment.

5 Measurement of Consumer Social Alignment and Baseline Shares

In this section, we describe how we use transaction data to impute cardholders' likely social alignment with firm stances, to aggregate consumers into groups with similar imputed social alignment (groups g), and to measure the baseline consumption share each of our social stance firms receives from these different groups ($y_g(\mathbf{N})/y(\mathbf{N})$). We will use these imputed social alignment groups and baseline shares in our analysis of consumption responses in subsequent sections.

5.1 Imputing Social Alignment and Consumer Groups

The longitudinal nature of the transaction dataset allows us to impute a cardholder's likelihood of alignment with the firm's social stance based on the other transactions they make throughout the card's history, as well as their demographics (when available). Prior work has demonstrated how consumption histories can predict a myriad of demographic characteristics including income, education, gender, ideology, and race (Bertrand and Kamenica, 2018). To impute social alignment in our context, we start with a subset of consumers with donations to PACs, charitable organizations,

[19] The insignificant difference in estimated consumer awareness is driven by the fact that aligned consumer are slightly more likely to report having heard good or bad news about the firm throughout most of the ten months *prior* to the firm's event. See Appendix A.3 for detail on the construction of these series.

and other non-profits that clearly indicate that these donors are likely socially aligned with or opposed to one or more of the 117 social stances in our analysis. For computational purposes, we then partition these donations into two clusters based on correlated donation patterns, which we arbitrarily label "For" and "Against." We define a donor as aligned with a stance if they donate to a cause that is associated with a similar position to the firm's social stance (or to other causes in the same cluster as these aligned causes), and as opposed if they donate to a cause that is associated with or shares a cluster with an opposing position on this issue.[20]

We use these donors as a labeled dataset (including more than 30 million cards) on which to train a machine learning model to predict social alignment with firm stances, defined as the probability of likely sharing the same For/Against position. We include the following as predictors: indicators for ever purchasing at each of the 1,000 merchants in the data most predictive of donor alignment on their own by χ^2 (excluding our set of firms with social stance events and the donations directly used for tagging consumer social preferences);[21] the demographics of inferred home counties;[22] and other general demographics (when available from credit reports).

In our prediction exercise, we first randomly split the dataset into a training sample (70 percent of cards) and a holdout sample (30 percent). We use XGBoost (Chen and Guestrin, 2016), a tree-based ensemble method, to classify donors in our training sample, fitting this model via weighted maximum likelihood estimation.[23] We empirically tune the parameters of the XGBoost algorithm

[20] We would ideally like to estimate the likelihood of a cardholder's alignment or opposition to each firm social stance separately, but doing so would increase the computation burden of these predictions by a factor of 117 (the number of social stance events). We implement our clustering approach to minimize this computational burden. All causes that are associated with a position direction on a given social issue are included in the same cluster, and we group together position directions across distinct issues when there is a higher relative co-occurrence of donations to those associated causes than to causes associated with the opposite position on this social issue (e.g., clustering "pro-LGBTQ+" donors with "support for stricter gun control" donors in the "For" cluster). While this clustering is motivated primarily by computational constraints, this can be justified by the fact that individuals' views are strongly and increasingly correlated across distinct social issues (Fiorina, 2016). We exclude from our definition of donors a small share of cardholders who donate to causes in both clusters.

[21] We construct this value for a given firm j as $\chi^2_j = \frac{(O_{jF}-E_{jF})^2}{E_{jF}}$, where O_{jF} is defined as the observed consumption by donors to the "For" cluster at firm f at any point in time, and E_{jF} as the expected consumption that would have occurred at the firm by For-cluster donors under the assumption that the relative consumption share of each donor cluster is constant across firms (i.e., $E_{jF} = (O_{jF} + O_{jA}) \times \frac{\Sigma_{j'} O_{j'F}}{\Sigma_{f'}(O_{j'F}+O_{j'A})}$). This χ^2 statistic is highest for firms that have particularly skewed consumption shares from donors by cluster (relative to their consumption shares in the entire economy) and high overall dollar spending by donors, which together make them useful for differentiating between a large number of donors by cluster.

[22] We infer an individual's home county as the modal county of their in-person transactions throughout time. County characteristics include population distributions across age groups, race, voting, urbanicity, and other similar characteristics.

[23] We observe a smaller number of donors in the Against cluster, and so we uniformly upweight donors from this cluster so that both clusters are given the same total weight when evaluating the likelihood of our prediction.

using five-fold cross-validation on the 70 percent training sample. We then fit XGBoost to the full 70 percent training sample of donors using the parameters selected by this cross-validation, and we make predictions for our holdout sample to evaluate the model's out-of-sample performance.

We evaluate the performance of our predictive model among our donor sample in Figure 2 Panel A, which shows the density of predicted alignment probabilities by true alignment status among donors. We see that our predicted class probabilities effectively separate donors socially aligned with For vs. Against positions, in that our predicted probabilities of being aligned with For positions are high among consumers truly aligned with For stances (based on their observed true donations, not used as predictors) and are low for consumers truly aligned with Against positions on those stance issues. When making out-of-sample predictions on our holdout sample, we achieve 81 percent balanced accuracy, which can be interpreted as the likelihood with which a randomly drawn donor would be assigned to the correct class (for which a random coin flip would be expected to produce 50 percent). We similarly achieve an area under the ROC-curve (i.e., ROC-AUC) of 0.89, which can be interpreted as the probability that a randomly drawn consumer truly aligned with a given class will receive a higher predicted probability of membership for that class than would a randomly drawn donor truly aligned with the other class (for which a random coin flip would be expected to produce 0.5). We show a generic list of the most influential predictors selected by this algorithm (as defined by the gain in predictive accuracy from including this predictor) in Appendix Table B2, noting that these predictors are highly interpretable and intuitive. While we are unable to identify individual merchants under our data agreement, these include media subscriptions, donations to other non-profit organizations associated with clear social alignments, purchases at merchants with industries and geographic distributions that are particularly correlated with likely social alignment, and other similar transactions and demographics that are plausibly predictive of an individual's alignment on social issues.

Having estimated this predictive model on our training dataset of donors, we now have a function mapping an individual's transactions and demographics to a measure of their social alignment (likelihood of alignment with a given For vs. Against stance among donors) which we apply to the transactions and demographics of non-donors to impute their individual social alignment. In doing so, we are making the assumption that the relationships between an individual's social alignment and their transactions and demographics among observed donors are similar to the relationships among non-donors. We are more confident of this external validity when analyzing the relative ordering of consumers' social alignment as opposed to the levels of predicted class probabilities

due to the fact that our training data contain more donors from one of the two clusters and due to our computationally-motivated clustering approach (which ignores that true alignment levels likely vary across stances and issues). As a result, we focus primarily on the relative ordering of consumers in terms of social alignment (rather than the predicted likelihood of alignment as levels) in our subsequent analysis.

We partition consumers into 12 groups that contain individuals with relatively similar social alignments and intensities. These ordered groups consist of: donors aligned with For causes (and likely opposed to the opposite Against positions); card-weighted deciles among non-donors decreasingly ordered by their likelihood of alignment with For causes; and donors aligned with stances associated with Against causes (and likely opposed to the opposite For positions). To give an intuition for how these predicted social aligned vary, in Figure 2 Panel B we map each state's median imputed social preference decile across cards with imputed home counties in that state.[24] Aggregating individual cards into fixed groups is computationally efficient in our subsequent analysis of consumption, as the number of individual cards is very large.

5.2 Measuring Baseline Consumption Shares by Group

To quantify the relative importance of each group to a given firm, we would ideally like to know the share of consumption (in $) that the firm would receive from that group (among all groups) if it did not take a stance (i.e., $y_g(\mathbf{N})/y(\mathbf{N})$). While we cannot directly observe this counterfactual quantity given that the firms in our analysis did in fact take social stances, we can reasonably proxy for this value as the share of consumption that the firm received from a given group in the year *preceding* its stance.

We show our estimated baseline shares in Figure 3, i.e., the share of consumption at a firm coming from each consumer alignment group in the year preceding its stance. We show the τ_j-weighted averages of these baseline shares across events, respectively, as well as each group's share of consumption (in $) aggregating across all U.S. firms in the transaction data throughout the period studied (2008-2023Q1).[25] We see that on average firms take stances that are aligned with the social preferences of their existing customer base. Firms that take stances received more pre-existing consumption from groups that are likely socially-aligned with the direction of their stance than from groups that are likely opposed to the firm's stance.

[24] See Appendix Figure B2 for an analogous map by county.
[25] We show these shares by position direction in Appendix Figure B3.

6 Consumption Responses

We now use our consumer awareness measures, consumer alignment groups, and transaction data to estimate potentially heterogenous consumption responses to the set of firm social stance events we have selected. Our empirical target throughout this section is consumption responsiveness conditional on awareness, i.e., $\frac{[y_g(s)-y_g(\mathbf{N})]\tau^{-1}}{y_g(\mathbf{N})} \approx \frac{log(y_g(s))-log(y_g(\mathbf{N}))}{\tau}$, and we will use this log approximation in our empirical estimation.[26] We note that, given the pieces we have already estimated, we can directly observe log consumption at the social stance firm over time given its observed social stance ($log(y_g(s))$) and measured consumer awareness of this stance (τ), leaving only no-event counterfactual log consumption ($log(y_g(\mathbf{N}))$) to be estimated. We first show consumption changes by group normalized by a simple counterfactual: each group's log consumption changes at all other firms in the economy. We then provide our preferred causal estimates of consumption responsiveness by imputing our no-event consumption counterfactuals via a synthetic difference-in-differences design that uses contemporaneous consumption at related firms and past historical patterns at the social stance firm.

6.1 Consumption Changes by Group

Before diving into our more sophisticated counterfactual, we first show changes in log consumption by each group at the event-study firm in the months surrounding the firm's social stance event. We normalize this log consumption relative to month $t = -1$, relative to changes in a given group's log consumption at all other firms in the data, and scale by consumer awareness. We therefore show $[\Delta log(y_{gjt}(s)) - \Delta log(y_{g,-j,t})] \times \tau^{-1}$, where $y_{gjt}(s)$ denotes observed consumption by group g at firm j in event-time month t, $y_{g,-j,t}(s)$ similarly denotes consumption for this group and month at all *other* firms in the economy, where Δ denotes changes from month -1 to month t, and using our measure of consumer awareness (τ) estimated using contemporaneous brand perception surveys. Taking changes over time controls for pre-existing differences across firms and groups (i.e., removing any group$\times$firm effects) and removing group-specific trends in consumption at all other firms controls for group-specific trends in consumption that affect all firms (i.e., removing any group$\times$time effects). This simple counterfactual does not control for group time-trends that are specific to the social stance firm or shocks to consumption at the social stance firm itself that

[26] We make this log approximation to consumption growth (coming from a first-order Taylor expansion) in our empirical estimation due to log consumption's increased robustness to outliers and for increased tractability, as this avoids the appearance of the unknown $log(y_g(\mathbf{N}))$ in the denominator.

vary over time.

We plot these values in Appendix Figure B4, taking a weighted-average across events.[27] We note that in the ten months preceding the firm's event, the normalized log consumption of relatively aligned and opposed consumer groups varies over time but generally moves similarly. This suggests that any time-specific shocks (such as seasonal changes in demand) generally have similar effects across these groups.

In sharp contrast, in the month of a firm's social stance event we see large and sharply diverging changes in normalized consumption across groups consistent with a demand for social alignment with firms. Donors aligned with the firm's stance see a sharp increase in consumption of about 18 percent (per 25 percent consumer awareness) in the month of a firm's stance, while donors opposed to the firm's stance see a sharp decline in consumption of about 13 percent in this same month. We see similar divergent sorting among non-donors by imputed social alignment decile; non-donors predicted to be most aligned with the firm's stance increase their consumption by about 8 percent while the most opposed decile of non-donors decreases their consumption by about 5 percent. All 12 groups' consumption changes are ordered exactly as predicted by a preference for social alignment, with smaller change magnitudes among non-donors less clearly aligned with or opposed to the firm's social stance. We see that these different consumption differences decrease slightly in magnitude (especially for the large responses among donors) but largely persist even a year after the firm's stance.

The sharp timing and unusual divergence in these consumption changes provide clear evidence of different relative consumption responses consistent with consumer demand for alignment with the perceived social values of firms. Because the magnitude of this change among donors is large relative to typical variation in pre-event months, we can also reasonably conclude that these events had an overall positive revenue impact among aligned donors and an overall negative revenue impact among opposed donors. However, these consumption responses also reflect other shocks to the firm (e.g., seasonality or discounts) that affect consumption unrelated to the firm's social stance. Indeed, we see some month-to-month changes in consumption that similarly affect all consumers during both pre-event and post-event months and are therefore most consistent with these confounds. While these results provide some of the best evidence to date regarding consumer demand for alignment with firms' social values, we need to account for these potential confounds

[27] For comparability with Figure 4, we use the same precision weights used in that analysis. These weights are proportional to τ_j^2, but also scale inversely with the estimated precision of our estimates of no-event counterfactual consumption. See Section 6.2 for detail on the construction of and motivation for these weights.

in order to quantify the causal effects of the firm's stance on consumption (for each group and overall) and to analyze optimal behavior for revenue-maximizing firms.

6.2 Imputing No-Event Counterfactual Consumption

To identify causal effects on consumption, we impute the counterfactual consumption that would have occurred at our event-study firms if they had not taken a social stance using a synthetic difference-in-differences (henceforth synthetic DiD) design (Arkhangelsky et al., 2021). We can think of estimating consumption responses to the firm's stance as an imputation problem, in that we observe the actual consumption that occurred after the firm's stance but do not directly observe the consumption that would have occurred if the firm had not taken this stance.

We train this counterfactual by predicting, for each event, the log weekly consumption series for the event-study firm ($log(y_{j\tilde{t}})$) in the two-years before the firm's event ($-104 \leq \tilde{t} < 0$, where $\tilde{t}$ denotes weeks relative to the date of the firm's social stance).[28] We then use this model to forecast consumption at the firm in the absence of an event.

Our synthetic DiD estimator can be expressed as follows:

$$\min_{\omega_0,\omega} \sum_{\tilde{t}=-104}^{-1} \Bigg[\underbrace{\left(\omega_0 + \sum_k \omega_k log(y_{k\tilde{t}}) \right)}_{\widehat{log(y_{j\tilde{t}})}} - log(y_{j\tilde{t}}) \Bigg]^2 + \zeta \left\| \omega \right\|_2^2$$

$$\text{s.t. } \sum_k \omega_k = 1;\ \omega_k \geq 0$$

In words, for a given event-firm's log weekly consumption series ($log(y_{j\tilde{t}})$) we seek to form a synthetic series ($\widehat{log(y_{j\tilde{t}})}$) as an ω_k-weighted average of control units (k) such that this synthetic series moves in parallel with the target social stance firm, while allowing for a fixed difference (ω_0) between the two.[29] We use the following as a superset of possible control units: log consumption at the firm in the same week of the previous year ($log(y_{j,\tilde{t}-52})$); log consumption at each of the thousands of other U.S.-based firms in the economy for the same week ($log(y_{k\tilde{t}})\ \forall k \neq j$), also excluding as predictors any other firms with a social stance event; and contemporaneous log total con-

[28] The higher frequency of week-level data (relatively to monthly data shown elsewhere) helps to inform the weights placed on different control units in our synthetic DiD counterfactual by ensuring that this counterfactual moves similarly to the firm week-by-week during the pre-event training period, while still aggregating across day-of-week patterns that are less important for our synthetic DiD to match.

[29] Our intercept ω_0 is chosen to normalize average log consumption in the pre-event month to zero.

sumption across all other firms in the same n-digit NAICS industry as j ($log(\sum_{j' \in F_n(j) \setminus j} y_{j'\tilde{t}})$ $\forall n \in \{0, 1, \cdots, 6\}$, where $F_n(j)$ denotes the set of firms with the same n-digit industry as our event-study firm). We note that when making predictions after the firm's event, this choice of possible control units ensures that we only use pre-event data at the social stance firm and/or contemporaneous data at other firms that did not take a social stance. As in synthetic controls, the weights placed on these control units are constrained to be non-negative and sum to 1, which helps to avoid regularization bias. The synthetic DiD estimator also penalizes the squared sum of weights ($\|\omega\|_2^2$) in order to spread out weights across units, with the amount of this regularization determined by the hyper-parameter ζ.

Two decisions need to be made when setting up this estimation. Most importantly, we need to decide how to restrict the set of possible control units that can be used to form this counterfactual. We consider restricting to the set of units chosen as controls by a first-stage Lasso regression (regressing the social stance firm's log weekly consumption series on all potential controls), restricting to the largest firms by revenue, and/or restricting to firms in the same industry as the event-study firm. We also need to choose the regularization hyper-parameter ζ. We tune hyper-parameters governing these control selection and regularization decisions in a data-driven way for each firm-event by maximizing our out-of-sample forecast accuracy in pre-event data.[30] More specifically, we look at a series of three-year periods which occur entirely before the firm's social-stance event. For each such three-year period and possible combination of hyper-parameters, we use the first two years as training data on which we fit a synthetic DiD estimator, and then use this synthetic control to forecast (out-of-sample) log weekly consumption at the event-study firm. We choose hyper-parameters that minimize the root mean squared error of these pre-event out-of-sample forecasts, as in these pre-event windows we do actually directly observe "no stance" consumption as the firm has not yet taken a stance and therefore we want to forecast this series as accurately as possible. We also use this root mean squared error as an estimate of the average variance of our estimator for a given firm, which we denote $\hat{\sigma}_j^2 := \widehat{Var}[log(y_{j\tilde{t}}) - \widehat{log(y_{j\tilde{t}})}]$.

Having thus selected our model's hyper-parameters in a data-driven way, we fit our synthetic DiD estimator to the two years of data preceding a firm's social stance event. This fit determines our choice of weights and intercept (ω_k and ω_0), which determines our synthetic control

[30]The hyper-parameters we tune are: the L1 penalty λ in a first stage Lasso regression used to select predictors by restricting to firms given non-zero Lasso coefficients; the share of firms ν to keep when selecting the largest firms by revenue as predictors; the number of digits m to use when restricting to firms in the same m-digit industry as predictors; the synthetic DiD weight regularization hyper-parameter ζ. We include values of ν and m that allow for no filtering on these size and industry characteristics.

series ($\widehat{log(y_{j\tilde{t}})}$) and estimated consumption responsiveness estimate ($[log(y_{j\tilde{t}}) - \widehat{log(y_{j\tilde{t}})}]/\tau_j$) in our training period ($\tilde{t} \in [-104, -1]$) and in our out-of-sample forecasts following the firm's event ($\tilde{t} \in [0, 51]$). We aggregate these treatment effect estimates of consumption responsiveness by taking a precision-weighted average of event-specific estimates, with precision weights given by $w_j := \widehat{Var}([log(y_{j\tilde{t}}) - \widehat{log(y_{j\tilde{t}})}]/\tau_j)^{-1} = \frac{\tau_j^2}{\hat{\sigma}_j^2}$. We plot our actual and predicted consumption series by event-week in Appendix Figure B5.

6.3 Estimates of Consumption Responsiveness

Comparing this actual scaled log consumption at the firm ($log(y_{j\tilde{t}})/\tau_j$) relative to our synthetic DiD control ($\widehat{log(y_{j\tilde{t}})}/\tau_j$) allows us to estimate causal effects of the firm's stance on consumption at the firm, which we plot in Figure 4. Panel A shows estimated overall consumption responsiveness along with a 95% confidence interval.[31] We estimate a statistically significant increase in overall consumption of about 3 percent (per 25 percent consumer awareness) on average in the month of the firm's event. This decreases on average in the following months to values which are generally weakly positive but not statistically distinguishable from zero (at a 95% significance level).

To provide group-specific estimates of consumption responsiveness, we assume that the group-specific consumption series in Appendix Figure 4 are on parallel trends except for the firm's event (after having already differenced out groups' consumption trends at all other firms and initial differences in consumption levels at $t = -1$). We therefore shift each group's series for a given firm by the same amount in a given month such that their average (weighted by baseline shares) aggregates up to our estimated treatment effect on overall consumption for that social stance firm (based on the synthetic DiD estimator above).[32] We then take a (w_j) precision-weighted average

[31] When performing statistical inference on our consumption response estimates, it is important to account for uncertainty in our synthetic DiD control ($\widehat{log(y_{j\tilde{t}})}/\tau_j$). We do so using a wild cluster bootstrap approach that incorporates residuals from our past forecasts on pre-event data. In each bootstrap iteration, we randomly sample firm-events with sampling probabilities equal to their precision weights. Once we have drawn the time-series for a given firm-event within a bootstrap iteration, we then uniformly sample one past forecast series from the set of all possible past forecasts for that firm-event, and add the residuals from the drawn past forecast series multiplied by a Rademacher weight (± 1 each with 50% probability) to the estimated overall consumption for that firm-event. We then average across the sampled firm-event+residual series within a bootstrap iteration to produce an estimated overall consumption time series for that bootstrap iteration. We conduct ten thousand such independent bootstrap iterations. We then construct a 95% confidence interval for our overall consumption response estimates in Figure 4 Panel A by using as our bounds the 2.5% and 97.5% quantiles across bootstrap iterations for each month.

[32] More formally, we shift each group-level response for a given firm-event-month by the constant $c_{jt} = \left[log(y_{jt}) - \widehat{log(y_{jt})}\right] - \sum_g \frac{\widehat{y_{gj}(\mathbf{N})}}{y_j(\mathbf{N})} [log(y_{gjt}) - log(\tilde{y}_{gjt})]$, where $\widehat{log(y_{jt})}$ is our synthetic DiD counterfactual for overall log consumption, $\frac{\widehat{y_{gj}(\mathbf{N})}}{y_j(\mathbf{N})}$ are our estimated baseline shares, and $[log(y_{gjt}) - log(\tilde{y}_{gjt})]$ is the group-specific consump-

across events. This shift incorporates our synthetic DiD counterfactual to control for seasonality and other confounding shocks unrelated to the firm's social stance, allowing us to estimate consumption responses under our parallel-trends assumption.

We show these shifted, group-specific estimates of consumption responsiveness in Figure 4 Panel B. We again observe sharply diverging consumption responses, with $+19$ percent and -11 percent effects among aligned and opposed donors, respectively, during the month of the firm's social stance event. As expected given this constant shift, we again see ordering of consumption responses according to social value alignment. Responses are generally positive among aligned donors and most non-donor deciles, and are mostly negative among opposed donors and the decile of non-donors most opposed to the firm's stance. Consumption responses are remarkably persistent, with some gradual decreases in magnitudes.

In Figure 5, we compare each group's estimated consumption response to the mean probability of alignment with the firm's stance among all cards in that group, using the prediction probabilities estimated in Section 5.1. We refer to this mean probability as the estimated share of that group aligned with the firm's stance ($ShareAligned_{gj}$). In Panel A of this figure, we plot each group's estimated consumption response in the event-month against its share aligned. In addition to the ordering of consumption responses by alignment and the weakly positive response point estimates among most donors noted previously, we can also see from this figure how the consumption response gradient with respect to the share aligned changes at different points in the $ShareAligned_{gj}$ distribution. Aligned and opposed donor responses (and to a lesser extent the responses of the most aligned and opposed non-donor deciles) are more extreme than we might expect were we to linearly extrapolate how consumption changes with $ShareAligned_{gj}$ based on the non-donor deciles. The second most extreme deciles (i.e., the 10-20th and 80-90th percentiles) differ in their share aligned by about 0.51 and in their consumption response by about 6 percent. The aligned and opposed donor groups themselves presumably differ in their aligned share by 1 (less than doubling the difference vs. these deciles) but differ in their consumption responses by 30 percent (roughly quintupling the difference vs. these deciles). This suggests that the consumption of aligned and opposed donors is likely more responsive to social alignment than that of non-donors even after accounting for the fact that they hold a greater share of aligned/opposed consumers, which makes intuitive sense as these donors have already shown a willingness to pay financially for their social

tion responses shown in Appendix Figure 4 (prior to this synthetic DiD adjustment).

views through their donations.[33]

In Figure 5 Panel B, we estimate the gradient of consumption responses with respect to a group's share aligned, comparing across groups within a month. We measure these gradients as the coefficients β_t in the following regression specification:

$$\frac{log(y_{gjt}) - \widehat{log(y_{gjt})}}{\tau_j} = \gamma_t + \beta_t ShareAligned_{gj} \times \psi_t + \varepsilon_{gjt}$$

We regress our estimated consumption response for a given group g, event j, and event-month t on the alignment share of that group (allowed to vary flexibly by month), controlling for month fixed effects and using the same precision weights described in Section 6.2. Our estimated consumption response differences out our synthetic difference-in-differences no-stance counterfactual ($log(y_{gjt})$), and is scaled by consumer awareness (τ_j), as described in Sections 6.2 and 4 respectively. We initially see a gradient of about 21 percent in the event-month, which gradually attenuates to about 9 percent after ten months. Clustering our standard errors by event, this gradient estimate is statistically significant at the 95% level in each post-event month, and is not statistically significant in any pre-event month.

In Appendix Figure B6, we also visualize and provide 95% confidence intervals for the difference in consumption responses between the aligned and opposed donors (Panel A) and between the most aligned and most opposed non-donor deciles (Panel B). These differences are again small and not statistically significant in any month prior to the firm's social stance. We then see a sharp and statistically significant jump in this consumption response in the month of the firm's social stance, which gradually attenuates by about half after ten months. This gap remains statistically significant among donors at this endpoint of our analysis and through eight months for the most aligned and opposed non-donor deciles.

Appendix Figures B7-B9 show analogous results by group from a placebo exercise in which we shift actual social stance event dates forward in time by one year and rerun all analysis (including synthetic DiD training and estimation) using these one-year-earlier placebo dates. We do not see the same pattern of sharply diverging consumption responses by alignment with the firm's stance in this placebo exercise.

We note that our estimates rely on our parallel-trends assumption, which we evaluate during

[33]The weaker relationship between consumption responses and the share aligned among non-donors vs. donors could also be driven in part by attenuation given that the share aligned is measured with noise among non-donor groups. Given that the share aligned is calculated fairly precisely when aggregating across the many millions of cards within each group as we do here, attenuation is unlikely to be the primary driver of this result.

the pre-period by reproducing our group-level consumption responses with three years of pre-event data in Appendix Figure B10. Having already differenced out group-specific trends at all other firms in the economy, we see generally similar trends across groups in log consumption at the event study firms throughout the three years preceding an event, and in particular we never observe sharply diverging responses among consumers with different social preferences. This lends support to our parallel trends assumption.

We similarly show consumption responses over a longer 2-year post-event horizon in Appendix Figure B11.[34] We find that among non-donors, consumption by alignment deciles reconverges after roughly one year, with somewhat greater persistence among the most opposed decile until two years after the firm's stance. In contrast, aligned−opposed donor differences persist strongly even two years later and show little sign of abating.

Appendix Figure B12 shows our overall and group-specific estimates of consumption responsiveness when separately aggregating events in which "For" cluster donors are likely aligned with vs. opposed to the firm's stance. We observe similar divergent responses for both cluster For-aligned and -opposed stances, with aligned consumption increases and opposed consumption decreases in both cases on average. On average, stances aligned with Against donors seem to induce somewhat larger consumption responses (both positive and negative) for a given level of consumer awareness, and these events drive the increase in overall consumption during the event-month.

6.4 Stance Impacts on Related Outcomes

We also use BrandIndex data to analyze respondents' interpretation of social stance news, as well as their self-reported consumption responses.[35] Appendix Figure B13 shows that respondents who are socially aligned with the firm's stance (based on the respondents' self-reported demographics) interpret this news positively and increase their favorability towards the firm, while socially opposed respondents feel more negatively about this news and about the firm following its social stance event.[36] This change in favorability translates into self-reported purchase behavior in Appendix Figure B14, as socially-aligned consumers more frequently say that they would consider

[34] We show these responses without our synthetic DiD adjustment, as our use of lagged data at the event-study firm as a predictor prevents us from making valid synthetic DiD predictions over horizons longer than one year.

[35] See Appendix A.3 for detail regarding the construction and analysis of these related outcomes.

[36] We note that we used divergence in favorability towards the firm or in interpretation of news about the firm as one of several criteria for selecting possible events. We observe similarly diverging BrandIndex favorability and news interpretation responses when restricting to events chosen by our other event selection methods, alleviating the potential concern that this result could mechanically reflect our BrandIndex-based event selection procedure, and suggesting that this is instead a real effect typical of controversial social stances taken by firms.

purchasing from or that they intend to purchase from the firm, while socially opposed consumers say that they are likely to do so less following the firm's stance. These favorability and self-reported consumption responses persist even a year after the firm's stance. This finding corroborates and complements the more incentivized responses we estimate in our analysis of transaction data.[37]

In Appendix Figure B15, we also analyze changes in respondents' exposure to information about the firm across different potential learning channels. We see that per 25 percent consumer awareness of a firm's stance, there is on average a sharp 12.2 percentage point increase in the share of consumers who report having recently talked with someone about the brand (i.e., Word-of-Mouth exposure, Panel A) and a 5.6 percentage point increase in the share who report having recently seen advertising for the brand (Panel B). This suggests that consumers learn about firms' stances through multiple channels, with most of this learning coming from channels other than direct advertising by the firm.

We analyze stock price responses (sourced from Yahoo Finance) in Appendix Figure B16. We do not find clear evidence on average of immediate impacts from the firms' social stances on their stock prices among publicly-owned firms in the month of the firm's stance, and any such impacts are difficult to separate from typical trends and variation in the firms' stock prices.

7 Supply-Side Implications and Predictors of Firm Behavior

Having now estimated group-specific consumption responsiveness, we return to our stylized model and discuss when our estimates imply that taking social stances maximizes revenue as a function of a firm's baseline shares across consumer groups. We then analyze the extent to which the preferences of firms' different stakeholder groups and their ownership structure predict the direction of a firm's stance.

[37] Relative to our analysis of transaction data, our self-reported responses in consumer surveys have the benefit that they do not rely on imputed social preferences or no-stance counterfactuals, as we observe less seasonality in BrandIndex respondents' consideration or purchase intent and we directly observe self-reported demographics/social views in these consumer surveys. However, these self-reported BrandIndex responses could be "cheap talk" in that they are not reflected in consumers' actual purchase decisions. They also allow for only coarse measures of respondents' social view and only capture the extensive margin, which leaves us unable to fully quantify impacts on the firms' revenues (which depend on both extensive and intensive margin responses). This motivates our analysis of transaction data, through which we are able to address each of these limitations and to more fully quantify consumer responses and firms' revenue-maximizing behavior.

7.1 Net Revenue Impact by Counterfactual Baseline Shares

We analyze how the revenue impacts of the stances firms took would differ if they counterfactually faced different distributions of consumers. As described in Sections 2 and 6, the net revenue impact of a firm's stance depends on the direction of the firm's stance and on their baseline shares across consumer alignment groups, with more positive impacts when firm stances are better aligned with the firm's consumer base. Appendix Figure B17 estimates the cumulative revenue impact implied by our average stance effects, weighting these same responses by two sets of baseline shares: the actual τ_j-weighted averages among firms that took social stances, and the reversed baseline shares they would have faced had they taken the opposite For/Against stance. Because firm's social stances are better aligned with their actual existing consumers than are these opposite stance counterfactuals, we see that estimated cumulative net revenue impacts would be lower if firms faced consumer groups with these reversed baseline shares.

Appendix Figure B18, shows that baseline shares can vary even more dramatically in different contexts, such as across states or across industries. In Appendix Figure B19 Panels A and B, we map states by the estimated cumulative net sales impact after five months induced by taking a stance, using our average estimates of consumption effects (as shown in Figure 4 Panel B) and assuming that firms face the baseline shares of overall consumer expenditures within that state. We see that taking a stance aligned with the Against cluster is estimated to have positive revenue impacts in the South, Midwest, and Southwest and negative impacts in urban and coastal states, while the reverse is generally true if an average stance were instead taking in the opposite For direction. We thus estimate that a firm could on average benefit from or be hurt by taking either a For or Against stance on social issues given plausible distributions of consumer baseline shares, and that net revenue impacts will likely depend on the firm's geographic distribution of consumers and on its industry.

7.2 Predictors of Firm Behavior

We next analyze the extent to which the alignments of our social stance firms' different stakeholders predict stance directions (For vs. Against). We control for a firm's ownership structure as an indicator for whether the firm is publicly (rather than privately) owned. We measure consumer social alignments as the baseline-share-weighted average of $ShareAligned_{gj}$ (for alignment with For cluster positions) across groups. We measure employee, CEO, and board of director preferences as

the share of donations going to recipients generally aligned with the same direction, sourcing this data for employees from OpenSecrets and for the CEO and board members from Bonica (2016). We weight events in these regressions by τ_j.

Table 1 shows results from regressing an indicator for firms having taken a stance in this direction (rather than in the opposite direction) on these firm characteristics among our set of firm social stances. We see that firms take stances aligned with their employees and their consumers, but that the social preferences expressed in CEO and board donations are less strongly correlated with the direction of firms' social stances.

8 Conclusion

Summarizing our findings, we estimate consumer demand for social alignment with firms by analyzing consumption responses around events in which firms took salient stances on controversial social issues. We systematically identify 117 of these events based on Google Trends, news reports, contemporaneous consumer surveys, and queries to a large language model. We quantify consumer awareness of these events based on contemporaneous surveys. We then use transaction data from a major payment card company to partition cardholders' into large groups that are likely more aligned with or opposed to the firms' stance and to analyze how consumption responses depend on alignment with the firm's social stance.

We estimate that observed firm stances have positive revenue impacts on average, with considerable heterogeneity across stances and consumers. We use all other U.S. firms in the economy as potential control units in a synthetic difference-in-differences design to impute the counterfactual consumption that would have occurred at the firm had it not taken a stance, comparing observed consumption to this counterfactual to estimate treatment effects caused by the firm's stance.

Disaggregating consumption responses among consumers more aligned with vs. opposed to the firm's social stance, we find stark and diverging consumption responses in the month of the firm's social stance event, which provide clear evidence of consumer demand for socially-aligned firms. Donors to causes indicating social alignment with the firm's stance increase their consumption at the firm by 19 percent when one quarter of consumers were aware of the firm's stance, and donors similarly opposed to the firm's stance decrease their consumption by 11 percent in response to the firm's social stance. We observe similarly diverging but smaller responses consistent with demand for socially-aligned firms among non-donors depending on the predicted likelihood of

their alignment with the firm's stance.

Based on our consumer response estimates, we analyze when firm stances increase revenue for hypothetical firms facing arbitrary distributions of consumer social views. Firm stances have more positive revenue impacts on average when they are better aligned with the social views of their existing customer base, and we show how revenue-maximizing firm stances therefore vary across firms operating in different industries or geographic regions.

Combining our set of events with measures of the social alignments of a firm's employees, executives, and other stakeholders, we next analyze which stakeholders' preferences best predict the direction of a firm's stance. The direction of a firm's stance is best predicted by the social preferences of its consumers and employees, as well as by its public vs. private ownership status.

Together our results show that individuals' consumption decisions respond to their perceived social alignment with firms, and that this consumer response affects firms' incentives to take social stances and to otherwise engage with controversial social issues.

References

ANES Data Center (2016). ANES: American national election studies. Accessible from `https://electionstudies.org/data-center/`.

Arkhangelsky, D., S. Athey, D. A. Hirshberg, G. W. Imbens, and S. Wager (2021). Synthetic difference-in-differences. *American Economic Review 111*(12), 4088–4118.

Ashenfelter, O., S. Ciccarella, and H. J. Shatz (2007). French wine and the us boycott of 2003: Does politics really affect commerce? *Journal of Wine Economics 2*(1), 55–74.

Auger, P. and T. M. Devinney (2007). Do what consumers say matter? The misalignment of preferences with unconstrained ethical intentions. *Journal of Business Ethics 76*(4), 361–383.

Barari, S. (2023). Political speech from corporations is sparse, only recently liberal, and moderately representative: Evidence from social media. *Working Paper*.

Barboza, G., V. Pede, and S. Madero (2021). Shared social responsibility. dual role of consumers as stakeholders in firm strategy. *Social Responsibility Journal 17*(1), 48–68.

Baron, D. P. (2009). A positive theory of moral management, social pressure, and corporate social performance. *Journal of Economics & Management Strategy 18*(1), 7–43.

Barrage, L., E. Chyn, and J. Hastings (2020). Advertising and environmental stewardship: Evidence from the bp oil spill. *American Economic Journal: Economic Policy 12*(1), 33–61.

Bartling, B., R. A. Weber, and L. Yao (2014, 11). Do Markets Erode Social Responsibility? *The Quarterly Journal of Economics 130*(1), 219–266.

Barton, R., M. Ishikawa, B. Theofilou, and K. Quiring (2018). To affinity and beyond: From me to we, the rise of the purpose-led brand.

Becht, M., P. Bolton, and A. Röell (2003). Corporate governance and control. In *Handbook of the Economics of Finance*, Volume 1, pp. 1–109. Elsevier.

Bertrand, M., M. Bombardini, R. Fisman, and F. Trebbi (2020). Tax-exempt lobbying: Corporate philanthropy as a tool for political influence. *American Economic Review 110*(7), 2065–2102.

Bertrand, M. and E. Kamenica (2018). Coming apart? Cultural distances in the United States over time. Working Paper.

Bonica, A. (2016). Avenues of influence: On the political expenditures of corporations and their directors and executives. *Business and Politics 18*(4), 367–394.

Bonnefon, J.-F., A. Landier, P. R. Sastry, and D. Thesmar (2022). The moral preferences of investors: Experimental evidence. Technical report, National Bureau of Economic Research.

Broccardo, E., O. Hart, and L. Zingales (2022). Exit versus voice. *Journal of Political Economy 130*(12), 3101–3145.

Burbano, V. C. (2021). The demotivating effects of communicating a social-political stance: Field experimental evidence from an online labor market platform. *Management Science 67*(2), 1004–1025.

Business Roundtable (2019). Statement on the purpose of a corporation. *Retrieved on August*.

Cassidy, W. and E. Kempf (2022). The rise of partisan corporate speech.

Chavis, L. and P. Leslie (2009). Consumer boycotts: The impact of the iraq war on french wine sales in the us. *QME 7*(1), 37–67.

Chen, N. and Z. Z. Zhong (2019). Historical memory and brand preference. *Available at SSRN 3326975*.

Chen, T. and C. Guestrin (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pp. 785–794.

Csutora, M. (2012). One more awareness gap? the behaviour–impact gap problem. *Journal of consumer policy 35*(1), 145–163.

Davis, C. L. and S. Meunier (2011). Business as usual? economic responses to political tensions. *American Journal of Political Science 55*(3), 628–646.

DellaVigna, S. and M. Gentzkow (2010). Persuasion: Empirical evidence. *Annual Review of Economics 2*(1), 643–669.

Dimson, E., O. Karakaş, and X. Li (2015). Active ownership. *The Review of Financial Studies 28*(12), 3225–3268.

Elfenbein, D. W., R. Fisman, and B. Mcmanus (2012, 03). Charity as a Substitute for Reputation: Evidence from an Online Marketplace. *The Review of Economic Studies 79*(4), 1441–1468.

Fiorina, M. P. (2016). The political parties have sorted. *A Hoover Institution Essay on Contemporary American Politics*, 1–20.

Fos, V., E. Kempf, and M. Tsoutsoura (2021). The political polarization of us firms. *Available at SSRN 3784969*.

Fouka, V. and H.-J. Voth (2021). Collective remembrance and private choice: German-greek conflict and behavior in times of crisis.

Friedman, M. (1970). The social responsibility of business is to increase its profits. *The new york times magazine*.

Fuchs, A. and N.-H. Klann (2013). Paying a visit: The dalai lama effect on international trade. *Journal of International Economics 91*(1), 164–177.

Global Strategy Group (2018, January). A call to action in the age of Trump: 5th annual business & politics study. Technical report, Global Strategy Group.

Hart, O. and L. Zingales (2017). Companies should maximize shareholder welfare not market value. *ECGI-Finance Working Paper* (521).

Hart, O. D., D. Thesmar, and L. Zingales (2022). Private sanctions. Technical report, National Bureau of Economic Research.

Hedblom, D., B. R. Hickman, and J. A. List (2019). Toward an understanding of corporate social responsibility: Theory and field experimental evidence. Technical report, National Bureau of Economic Research.

Heilmann, K. (2016). Does political conflict hurt trade? evidence from consumer boycotts. *Journal of International Economics 99*, 179–191.

Hydock, C., N. Paharia, and S. Blair (2020). Should your brand pick a side? how market share determines the impact of corporate political advocacy. *Journal of Marketing Research 57*(6), 1135–1151.

Hydock, C., N. Paharia, and T. J. Weber (2019). The consumer response to corporate political advocacy: a review and future directions. *Customer needs and solutions 6*(3), 76–83.

Iyengar, S. and S. J. Westwood (2015). Fear and loathing across party lines: New evidence on group polarization. *American Journal of Political Science 59*(3), 690–707.

Kitzmueller, M. and J. Shimshack (2012). Economic perspectives on corporate social responsibility. *Journal of Economic Literature 50*(1), 51–84.

Klostermann, J., C. Hydock, and R. Decker (2021). The effect of corporate political advocacy on brand perception: an event study analysis. *Journal of Product & Brand Management*.

Larcker, D. F. and E. M. Watts (2019). Where's the greenium? *Rock Center for Corporate Governance at Stanford University Working Paper* (239), 19–14.

Liaukonyte, J., A. Tuchman, and X. Zhu (2022). Spilling the beans on political consumerism: Do social media boycotts and buycotts translate to real sales impact? *Available at SSRN 4006546*.

Lin, T. C. (2022). *The Capitalist and the Activist: Corporate Social Activism and the New Business of Change*. Berrett-Koehler Publishers.

Maks-Solomon, C. and J. M. Drewry (2021). Why do corporations engage in lgbt rights activism? lgbt employee groups as internal pressure groups. *Business and Politics 23*(1), 124–152.

McWilliams, A. and D. Siegel (2001). Corporate social responsibility: A theory of the firm perspective. *Academy of management review 26*(1), 117–127.

Meier, S. and L. Cassar (2018). Stop talking about how csr helps your bottom line. *Harvard Business Review*.

Milosevic, D., S. Andrei, and R. W. Vishny (2015). A survey of corporate governance. *The journal of finance 52*, 737–783.

Nielsen (2015, Oct). The sustainability imperative: New insights on consumer expectations.

Painter, M. (2021). Firm statements, consumer political beliefs, and the limits of stakeholder capitalism. *Available at SSRN 3557961*.

Panagopoulos, C., D. P. Green, J. Krasno, M. Schwam-Baird, and K. Endres (2020). Partisan consumerism: Experimental tests of consumer reactions to corporate political activity. *The Journal of Politics 82*(3), 996–1007.

Pandya, S. S. and R. Venkatesan (2016). French roast: consumer response to international conflict-evidence from supermarket scanner data. *Review of Economics and Statistics 98*(1), 42–56.

Rajan, R. G., P. Ramella, and L. Zingales (2022). What purpose do corporations purport? evidence from letters to shareholders. *Evidence from Letters to Shareholders (February 1, 2022)*.

Schoenmueller, V., O. Netzer, and F. Stahl (2023). Frontiers: Polarized america: From political polarization to preference polarization. *Marketing Science 42*(1), 48–60.

Sheffi, Y. (2021). Climate change: The real inconvenient truth. *Management and Business Review 1*(1).

Stein, J. C. (2003). Agency, information and corporate investment. *Handbook of the Economics of Finance 1*, 111–165.

The Economist (2021a). An eco-wakening. *The New York Times*.

The Economist (2021b). The political CEO.

Wicker, A. (2017). Conscious consumerism is a lie. here's a better way to help save the world. *Quartz*.

Figure 1: Consumer Awareness of Events, Based on BrandIndex Responses

Panel A: Unusual Awareness of News About Firm, Averaged Across Events

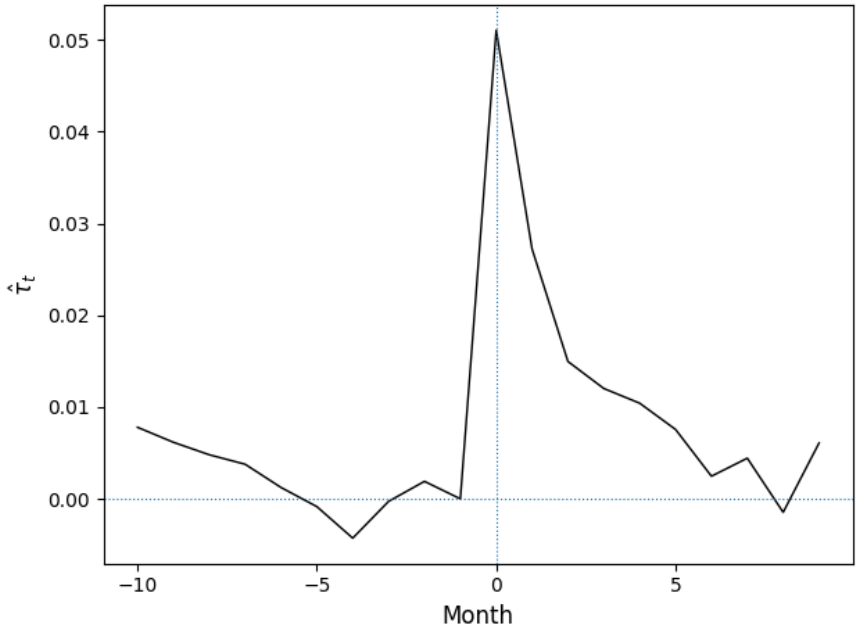

Panel B: Consumer Awareness Distribution Across Events (Histogram)

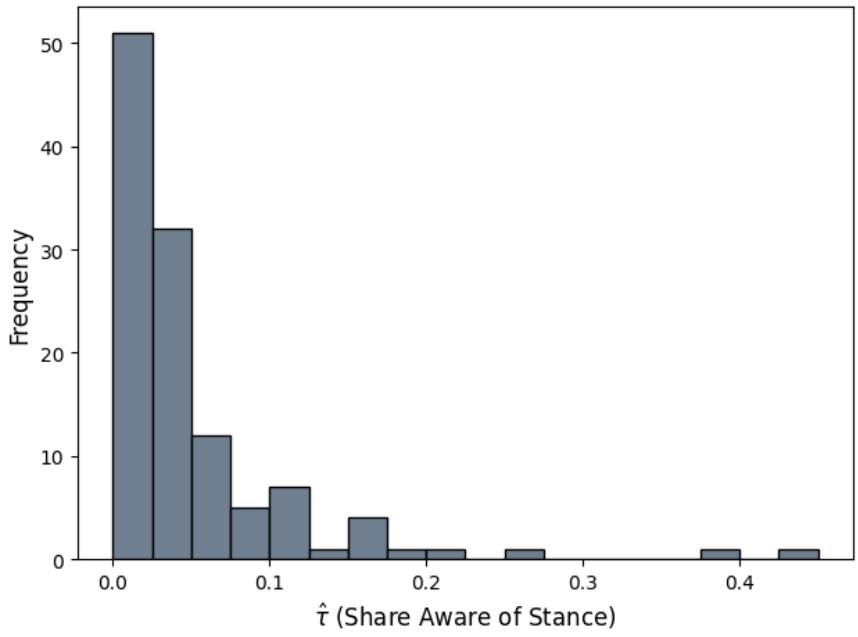

Note: Figure shows changes in consumer awareness of firms, based on responses in contemporaneous BrandIndex surveys to the following questions: *"Over the past two weeks, which of the following brands have you heard something POSITIVE/NEGATIVE about (whether in the news, through advertising, or talking to friends and family)?*". Defining a_t as the share who report having heard positive and/or negative news about the brand among respondents in month t, Panel A shows $\hat{\tau}_t := \frac{a_t - a_{-1}}{1 - a_{-1}}$, averaged by month across event-study firms. Months are defined as 4-week periods relative to the firm's event. Panel B shows a histogram summarizing across events our measure of consumer awareness ($\hat{\tau} := \frac{a_0 - a_{-1}}{1 - a_{-1}}$).

Figure 2: Donor Social Alignment Predictions

Panel A: Social Alignment Prediction Densities, among Donors by True Group

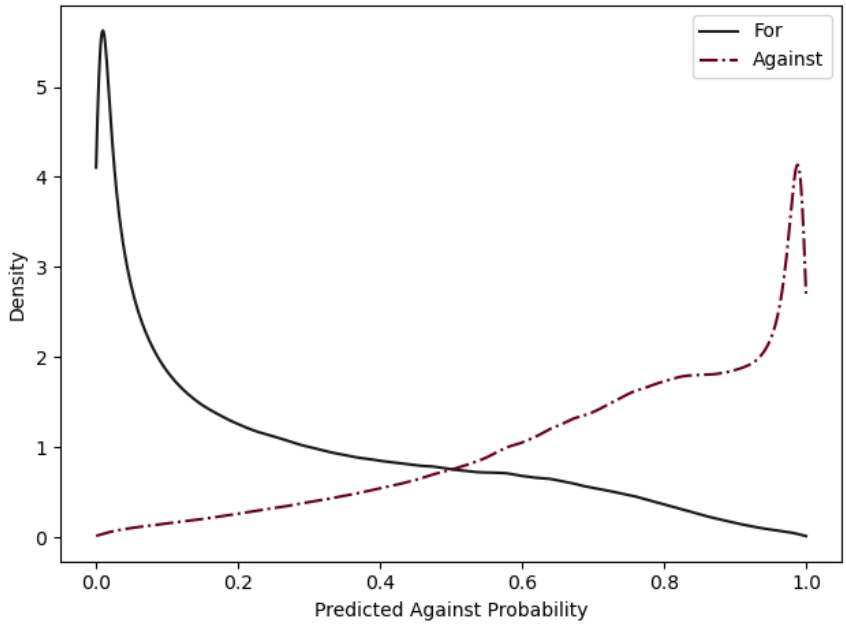

Panel B: Median Predicted Social Alignment Decile, among All Cards by State

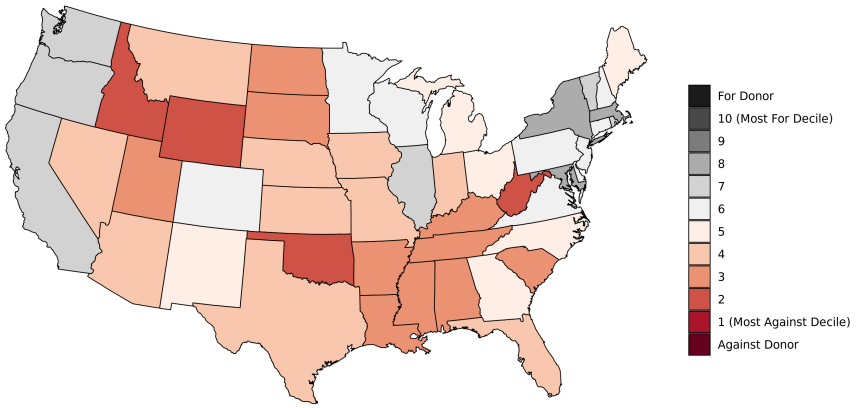

Note: Figure summarizes predicted social alignments based on transactions and demographics. Panel A shows predicted social alignment density distributions among the sample of cards with observed donations indicative of clear social alignments. Densities are plotted separately for donors to causes in the (arbitrarily labeled) "For" vs. "Against" donation clusters. Panel B maps (for each state) the median predicted probability of alignment with causes in the "For" cluster among all cards in that state, with deciles 10 and 1 denoting non-donors most likely to be aligned with vs. opposed to causes in this cluster, respectively.

Figure 3: Group Shares of Pre-Existing Consumption at Event Study Firms (Baseline Shares)

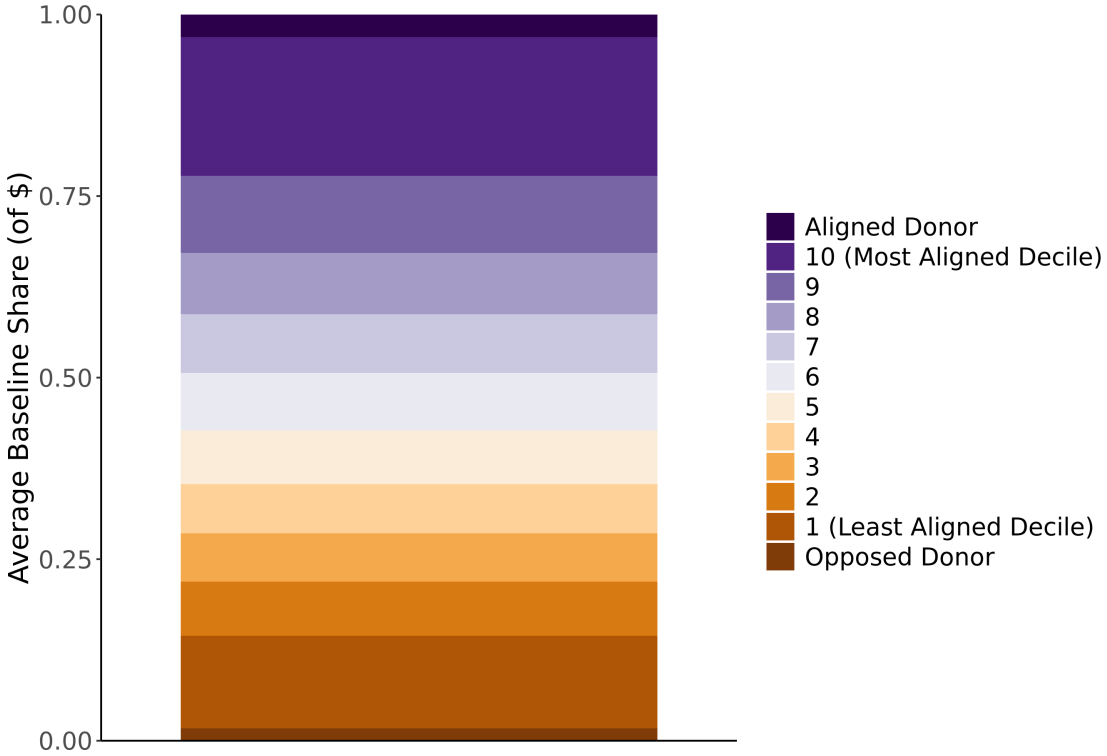

Note: Figure shows shares of consumption (in $) at social stance firms by alignment group in the year preceding these stances. Consumer groups are defined as described in Section 5, ordering consumers based on their predicted alignment with firm social stances. These baseline shares are averaged across social stance events, weighting events by consumer awareness of the firm's stance (τ_j, as defined in Section 4).

Figure 4: Estimated Causal Effects of Social Stances on Consumption at Firm

Panel A: On Overall Consumption

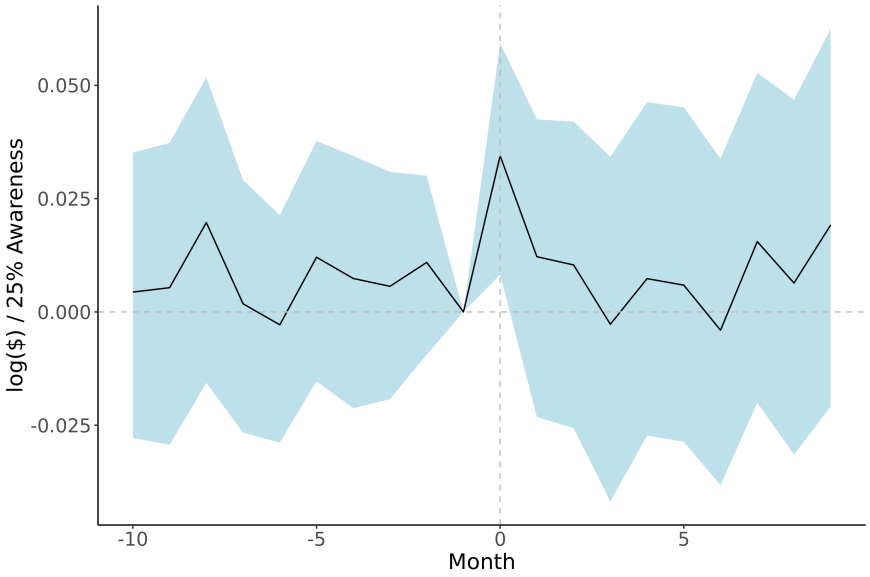

Panel B: On Consumption by Consumer Social View Group

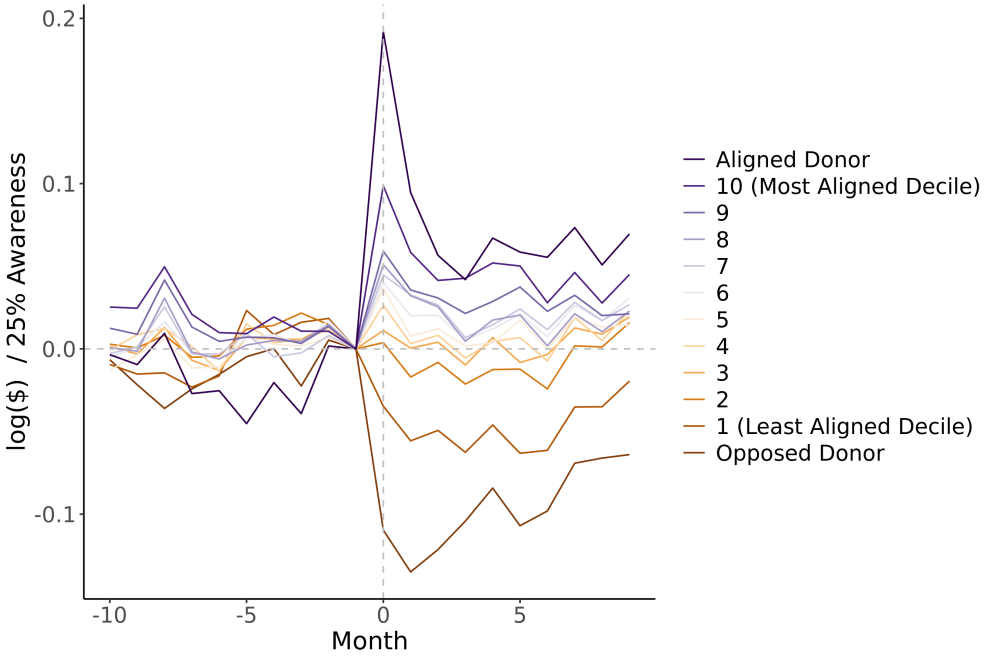

Note: Figure shows estimated causal effects of the firm's stance on log consumption in the months surrounding their social stances, overall and by consumer group. Panel A shows overall effects, calculated as the difference between observed consumption and a no-stance counterfactual predicted using a synthetic difference-in-differences design (see Section 6.2). Effects are scaled relative to consumer awareness and are averaged across firms using a precision-weighted average. 95% confidence intervals are constructed using a wild cluster bootstrap approach that accounts for uncertainty in our synthetic difference-in-differences counterfactuals. Panel B similarly provides causal estimates for the impact of the firm's social stance on consumption for each group (see Section 6).

Figure 5: Consumption Responses vs. Share Aligned, by Group

Panel A: Event-Month Consumption Response vs. Share Aligned, by Group

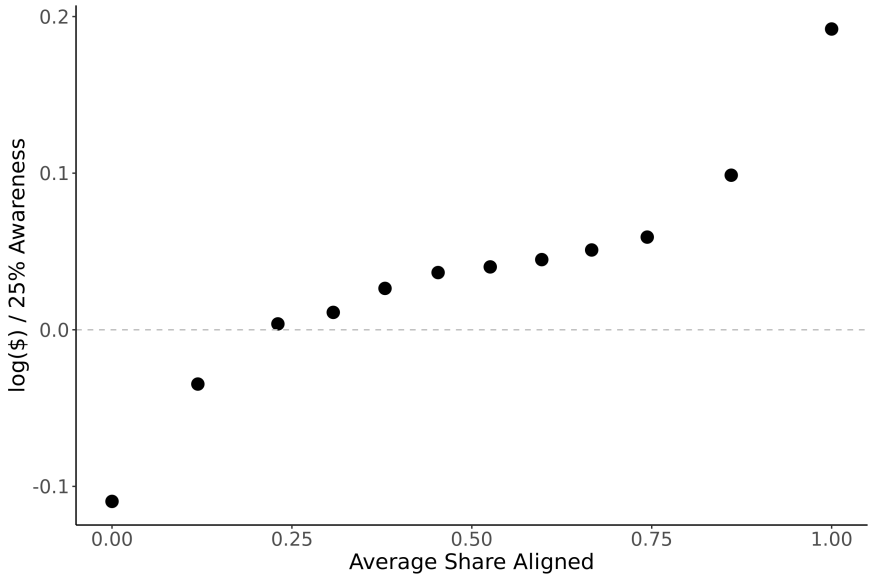

Panel B: Consumption Response Gradient vs. Share Aligned, Across Groups

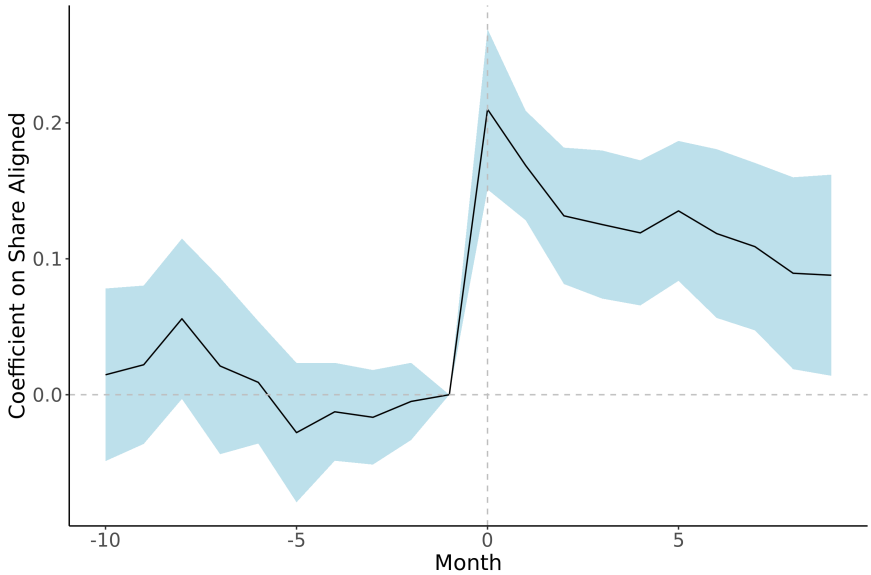

Note: Figure compares consumption responses (as shown in Figure 4 Panel B) by group to the mean share of individuals in that group predicted to be aligned with the firm's stance. Panel A shows for each alignment group the average estimated consumption response in the month of the firm's social stance event (y-axis, matching the t=0 value in Figure 4 Panel B) vs. the average share of consumers predicted to be aligned with the firm's stance (x-axis). Alignment shares for a given group are constructed as the mean probability of alignment among all non-donors in a given decile group (see Section 5.1 for detail), as 1 for the Aligned Donor group, and as 0 for the Opposed Donor group. Both averages are taken across firm-events and use the same precision weights described in Section 6.2. Panel B plots the coefficients β_t and 95% confidence intervals from the regression of consumption responses on share aligned described in Section 6.3, clustering standard errors by event.

Table 1: Predicting Stance Direction ($\mathbb{1}_{\{\text{Stance is For}\}}$), by Stakeholder Preferences

	(1)	(2)	(3)	(4)	(5)
Is Public	0.378**				
	(0.154)				
Consumer For %		0.012*			
		(0.007)			
Employees For Donate %			0.009***		
			(0.001)		
CEO For Donate %				0.003*	
				(0.002)	
Board For Donate %					-0.001
					(0.003)
Observations	117	117	101	37	42
R^2	0.157	0.060	0.449	0.128	0.012

Note: Signif. Codes: ***: 0.01, **: 0.05, *: 0.1. Table shows stance-level regressions in which the stance's direction is predicted by the preferences of different stakeholder groups. The dependent variable is an indicator for the stance being aligned with causes in the (arbitrarily labeled) "For" donation cluster. Predictors include the share of consumers estimated to be similarly aligned with positions in this donation cluster, the share of donations that went to similarly aligned groups from a firm's employees, CEO, and board of directors, as well as an indicator for being a publicly-owned (rather than privately held) firm. See Section 7 for detail on these measures. Events are weighted by estimated consumer awareness (τ_j). Heteroskedasticity-robust standard-errors in parentheses.

A Detail on Data and Measurement

A.1 Detail on Event Selection Procedures

Selecting Events Based on Google Trends:

We identify candidate events based on Google Trends searches for firm names in combination with particular keywords indicative of possible social stances. We generated a list of firm names over which to search by pulling the 10,000 largest U.S. firms from D&B Hoovers by 2021 revenues, excluding subsidiaries, public sector organizations, and non-profits. We pulled the name and tradestyle for each such firm from D&B Hoovers, and standardized these names (e..g, by removing common firm suffixes like Inc.). We then programmatically queried Google Trends searches for each firm name and keyword using the *pytrends* Python package. If no data was returned (due to insufficient searches) using the firm name, we repeated this pull using the company's tradestyle. We pulled search data for each firm and keyword at the month level (which indexed the largest value in a month to 100). We also separately pulled this series together with a common reference search across firms (so that searches for different firms could be compared in levels with a common reference index). We then created a composite index of these series by standardizing each series based on its mean and standard deviations over 6, 12, or 24 month backwards-looking rolling windows, and then took an unweighted average of these values across these three different windows as well as our self-indexed vs. commonly-indexed series. By averaging these self-indexed and commonly-indexed series, we valued events with unusual increases relative to typical patterns for a given firm+keyword combination, and which showed a significant increase in absolute (rather than just relative) searches through this inclusion of a common reference search. We then selected the 2,500 largest firm-months by this index as candidate months. For each candidate month, we then pulled daily Google trends searches for the firm+keyword combination in a 270-day window around the firm's event (which is the maximum window length over which Google Trends returns daily searches). We then returned candidate event firm-dates as those with the largest increase from a 28-day backward- to a 28-day forward-looking moving average for this daily series.

The precise keywords used to identify potential Google Trends searches were: *gay; transgender*; *immigration*; *political contributions*; *political*; *voting*; *controversy*; *buycott*; *boycott*; *gun control*; *abortion.* These keywords were selected as common keywords and topics present in discussions around a set of firm social stances that we first identified through a manual search. Our choice of keywords was also guided by the stance topics identified in Klostermann et al. (2021), and by querying OpenAI's ChatGPT for suggested keywords that could be used to identify social stances taken by firms based on Google Trends searches and news coverage.

Selecting Events Based on News Coverage:

We selected candidate events based on news coverage by initially pulling all articles from ProQuest's U.S. Newsstream that mention at least one of a union of social keywords anywhere in their text and which mention a firm as a subject (as identified by the ORG metadata field). In addition to the full list of social keywords used when analyzing Google Trends, we also include the following additional keywords when analyzing news articles: *racial*; *social issue*; *lesbian*; *queer*; *lgbtq*; *lgbtq+*; *lgbtqia*; *daca*; *gun*s; *second amendment*; *reproductive rights*. The smaller set of keywords used when analyzing Google Trends was motivated by computational constraints, as querying Google Trends searches for each additional keyword carried with it a higher computational cost than including additional keywords when filtering news articles from ProQuest's U.S. Newsstream.

For each firm mentioned as a subject in an article containing these keywords, we produced a daily time series counting the number of such articles mentioning the firm, and selected candidate event firm-dates as those with the largest difference between a 28-day forward- and a 28-day backward-looking moving average for this daily series. We also supplement this list of events based on new coverage by adding events identified in Klostermann et al. (2021), which identifies events by searching for any individual news articles that contain their own set of keywords indicative of corporate stances.

Selecting Events Based on BrandIndex Favorability Responses:

We selected candidate events based on BrandIndex surveys by first calculating for each firm the difference in net favorability towards the firm (calculated as the share with a positive impression less the share with a negative impression) between Democratic vs. Republican respondents. We similarly calculated this difference in the net share who reported having heard good less bad news about the firm in the last two weeks. We then selected candidate firm-dates as those with the largest difference in magnitude between a 28-day forward- and a 28-day backward-looking moving average for either of these daily series.

Selecting Events Based on Queries to GPT-4:

To generate a list of candidate social stance events from GPT-4, we provided the following prompt to OpenAI's ChatGPT on April 22nd, 2023 (at which point ChatGPT was on its "Mar 23" version): *"List 40 of the most notable and widely covered events in which individual companies took partisan stances on controversial social/political issues in the U.S. between 2010 and 2022, inclusive. For each event, provide the following variables:*
- company: company name,

- date: the start date of the company's stance event (MM/DD/YY format)
- ideology: the ideological direction (conservative or liberal) of the firm's stance
- description: a brief (2-6 word) description of the company's stance
Order events according to their notability (descending), and output this list in csv format."

Querying ChatGPT for more than 40 events at a time typically exceeded limits to its output length. As a result, we extended this list beyond 40 events through the following follow-up prompt:
"Extend this list by adding another 40 of the most notable and widely covered events in which individual companies took partisan stances on controversial social/political issues in the U.S. between 2010 and 2022, inclusive. For each event, provide the following variables:
- company: company name,
- date: the start date of the company's stance event (MM/DD/YY format)
- ideology: the ideological direction (conservative or liberal) of the firm's stance
- description: a brief (2-6 word) description of the company's stance
Order events according to their notability (descending), and output this list in csv format.
Avoid duplicates by not choosing any events which have the same company name and month as an event you've already suggested."
We then selected the first fifty suggestions from this combined list as candidate social stance events.

Combining and Filtering Candidate Events

For each of the candidate firm-event dates selected by the automated Google Trends, news coverage, BrandIndex, and GTP-4 based methods described above, we then manually check for the existence of a social stance taken by the firm around this date. We do so by searching the internet and querying news articles about the firm, and we filter out candidate events that were not associated with a social stance taken by the firm. This includes filtering out candidates for which we were unable to identify any salient event for the firm around this date. This also includes dropping events that were falsely flagged by our automated methods as a social stance, for example dropping candidates that were associated with a spike in social keyword activity because a shooting occurred at one of the firm's stores (without the firm taking a clear social stance in response) or because an executive used a racial slur. Having identified a clear social stance by the firm for each remaining candidate event, we also manually categorize the stances into topics and assign a For/Against direction based on discussions in the news articles and other online materials reviewed in our search. We also assign a tentative start date for each social stance event. We also filter out stances taken by firms that are not consistently identifiable as merchants in our transaction data, typically because they exclusively sell their products through to/other merchants or because they are particularly small firms.

After implementing this filtering procedure, we then take the union of candidates selected by our four different automated methods. We group together, as a single candidate event with multiple possible dates, candidates that occur within 28-days of each other for the same firm, using fuzzy string matching to identify different names that correspond to the same firm. We then finalize the single start date for each event as the date on which the firm initially took its social stance. In rare cases for which this date is not itself directly reported, we use the earliest publication date of news articles or other online materials that mention the firm's social stance.

Having identified a set of actual social stances taken by firms, we then use BrandIndex responses and other data to quantify consumer awareness of each firm's stance, as described in Appendix Section A.2. We then drop a small number of events by restricting our list to events which are the largest social stance events (in terms of consumer awareness) for that firm within a ± 2 year window. We also then drop three events for which we estimate that a non-positive share of consumers were aware of the firm's stance (because the share of consumers who reported hearing recent good or bad news about the firm decreased in the month following the firm's stance).

This procedure ultimately selects 117 social stance events, which were taken by 96 unique firms. We list each event's year, direction, estimated consumer awareness, and generic description in Appendix Table B1. Of these 117 social stance events, some are selected only by our Google Trends method (33.3%), by our news coverage method (26.5%), by our BrandIndex method (4.3%), and by our GPT-4 based method (3.4%). The remaining 32.5% of events are selected by multiple methods. The share of events selected by multiple methods increases to 60.6% when weighting events by their estimated consumer awareness, as events which were more salient to consumers are more frequently identified by multiple methods.

A.2 Detail on Event Size Measurement

A.2.1 Quantifying Consumer Awareness Based on BrandIndex Responses

As described in Section 4, for most events we use contemporaneous brand perception surveys from BrandIndex to quantify the share of consumers who were likely aware of the firm's stance. We refer to this share as "consumer awareness." We identify this share based on responses to the following question asked by BrandIndex: "*Over the past two weeks, which of the following brands have you heard something [positive/negative] about (whether in the news, through advertising, or talking to friends and family)?*" Defining a_{jt} as the share of BrandIndex respondents in event-time month t who report having heard something positive or negative about firm j in the past two weeks, we then define our BrandIndex-based estimate of consumer awareness as $\hat{\tau}_j := \frac{a_{j0} - a_{j,-1}}{1 - a_{j,-1}}$. This share represents the pre- vs. post-event-month change in the share of respondents who report having recently heard good or bad news about the firm, scaled by the share of respondents who were not

already reporting having heard recent news about the firm.

This denominator scaling accounts for the fact that the numerator will undercount awareness among respondents who hear about the firm's stance but who would already have reported hearing other news about the firm unrelated to its social stance. As one potential justification for this metric, suppose that a continuum of respondents become aware of the firm's social stance with i.i.d. probability τ_j and independently become aware of other news about the firm with i.i.d. probability γ_j, fixed over time. We can then see that our empirical measure provides an estimate of consumer awareness under these assumptions as follows:

$$
\begin{aligned}
a_{j,-1} &= \gamma_j \\
a_{j0} &= \mathbb{P}[AwareStance] + \mathbb{P}[AwareOther] - \mathbb{P}[AwareStance \cap AwareOther] = \tau_j + \gamma_j - \tau_j\gamma_j \\
\tau_j &= \frac{\tau_j(1-\gamma_j)}{1-\gamma_j} = \frac{a_{j0} - a_{j,-1}}{1-a_{j,-1}}
\end{aligned}
$$

The extent of this undercounting will in practice depend on the correlation between social stance awareness and awareness of other news about the firm. Here we assume independence between these two events, following calculations of persuasion rates in the literature (e.g., DellaVigna and Gentzkow, 2010) in using this $1-a_{j,-1}$ scaling as our adjustment factor.

In Appendix Figure B20, we plot $a_{j,t}$ (Panel A) and $\hat{\tau}_{jt} := \frac{a_{jt}-a_{j,-1}}{1-a_{j,-1}}$ (Panel B) by event. Panel A shows that firms vary substantially in the share of consumers reporting good/bad news about the firm in the pre-period (values range from near-zero to shares around 0.5), highlighting the need for this scaling adjustment in order to avoid differential undercounting across events. Panel A also shows reasonable consistency in a_{jt} within firms over pre-event months, consistent with γ_j being similar over time and suggesting that $a_{j,-1}$ is likely a reasonable proxy for the share of consumers that would hear about news other than the firm's social stance in month $t = 0$.

A.2.2 Quantifying Event Salience Based on News Coverage and Google Trends

As an alternative to our BrandIndex-based quantification of consumer awareness, we also proxy for the relative salience of each candidate event using data from ProQuest's U.S. Newsstream. To quantify the salience of each event, we define an event's size as the increase in the number of news articles mentioning the social stance firm (j) as a subject in the month following the event relative to the preceding month, i.e. $EventSize_j := \#articles_{j0} - \#articles_{j,-1}$.

We visualize the salience of these social stances to consumers in Figure B21. Panel A shows variation in news coverage of social stance firms by month around their social stance event, relative to the month preceding the firm's event and averaged across firm-events (i.e., showing

$[\sum_{j\in J}(\#articles_{jt} - \#articles_{j,-1})]/|J|$ for months $t \in [-10,9,\cdots,8,9]$). We see that on average across events, the firm taking a social stance is covered by 87 additional news articles in the month following relative to the month preceding its social stance event. This represents an unusual 54 percent increase in news coverage relative to the average number of articles covering the firm in the month preceding its stance (i.e., $[\sum_{j\in J}\#articles_{j,-1}]/|J|$=159). We have confirmed by looking at the text of these news articles that this increase in news coverage is primarily driven by the firm's social stance event itself, rather than by news covering some other aspect of the firm. News coverage is relatively constant during the months preceding the firm's event and is somewhat elevated in subsequent months following the firm's event month. This sharp spike in news coverage is consistent with the occurrence of an event (the firm's social stance) which is likely to be salient to consumers and to affect their perceptions of the social values associated with the firm, thus enabling our analysis of how individuals' consumption responds to these changes in perceptions.

In Panel B of this same figure, we use a histogram to show heterogeneity in event-month news coverage increases across events. We see that events vary in their induced news coverage. A handful of the largest social stance events see more than 1,000 article increases in news coverage in the month following the firm's event, suggesting that consumers are most likely to be aware of these events. Many of these events are much less salient to consumers, as the 75th percentile and median values are 64 and 15 news article increases, respectively.

We also show changes in log Google Trends searches for the firm around its social stance event in Appendix Figure B22, observing a sharp spike in Google searches in the month of the firm's social stance. Relative to our BrandIndex-based measure of awareness, we also observe greater month-to-month variation in both news coverage and Google searches in months far removed from the firm's social stance. This suggests that these alternative measures are likely noisier proxies for consumer awareness.

A.2.3 Imputing Consumer Awareness for Events Not Covered by BrandIndex

As mentioned in Section 4, 15 percent of firm-events in our sample are not covered by the BrandIndex dataset in the months around the firm's event. This means that we cannot quantify consumer awareness of these firm-events directly from BrandIndex responses, and must instead impute consumer awareness of these events using other data sources. We include the following as potential predictors in this imputation exercise: the (pre- vs. post-month) change in Google Trends searches for the firm (without specifying additional keywords); the change in the number of news articles mentioning the firm as a subject (without specifying additional keywords); and the change in the number of news articles mentioning the firm as a subject *and* that also include at least one of our news coverage social stance keywords. We also include as predictors changes in logs for each of the three metrics above. When any of these six variables are missing, we replace this missing

value with the average value of this predictor across events, and include as a potential predictor an indicator for whether the firm had a missing value for this variable.

To make these imputation predictions, we consider the following four methods: an elastic-net; random forest; XGBoost; and stepwise-selection. We select our preferred model by minimizing the five-fold cross-validated RMSE when predicting consumer awareness among the set of firms covered by BrandIndex data. This cross-validation procedure selects stepwise-selection with nine features as our preferred imputation method, including the following as controls in a linear regression (in addition to a constant term): the six changes in levels and in logs mentioned above, as well as missingness indicators for changes in Google Trends searches, for changes in log Google Trends searches, and for changes in the the number of news articles mentioning the firm in conjunction with social keywords. The cross-validated RMSE of these predictions was 0.04.

A.3 Detail on Construction of Other Variables

A.3.1 YouGov BrandIndex Survey Details and Question Text

To produce its BrandIndex dataset, YouGov owns and operates a syndicated global panel of more than 17 million respondents. More than four million of these respondents are located in the U.S., and we restrict our analysis only to this subset for consistency with the other data sources we analyze. Panel members sign up through a double opt-in process through which they register to join YouGov's panel, validate their email address, and start by sharing demographics about themselves (e.g., age, gender, and race). Panelists are then invited to complete brand preference surveys, in which they answer questions about multiple brands from a single product category (e.g., "Grocery Stores" or "Skin Care and Cosmetics") on a given day.[38] Panelists can only complete a particular survey on a given day if they receive an invitation to do so from YouGov, and YouGov employs a lock-out period following the completion of a survey to ensure that a given respondent does not complete multiple surveys within a short time window. Panelists are randomly assigned to product categories using a quota system to ensure that responses for each product-category$\times$day are in expectation nationally representative based on race, income, gender, and region (relative to U.S. Census data). YouGov also uses weights when aggregating responses to account for unexpected variation in completion rates, thereby ensuring that responses are also nationally representative ex-post. YouGov respondents receive points for their survey completion, which they can exchange for rewards like Amazon gift cards or movie tickets. YouGov collects responses from at least 5,000 U.S. respondents each day, collecting this data since June 3rd, 2007.

Within a given survey, respondents first select the brands that they are aware of within the

[38] In their U.S. subsample, YouGov elicits preferences regarding 2,000+ brands spread over 45+ product categories. Both numbers have varied over time.

product category (from a list of up to 40 brands). They then answer the remaining questions in the survey only for the brands of which they said they were aware. BrandIndex produces two kinds of metrics: 2-point metrics (e.g., Yes/No responses), and 3-point metrics (e.g., Positive, Negative, or Neutral). The exact wording of questions vary by product category to reflect the product category name, type of good, and typical purchase frequency. Here we provide the questions seen by YouGov BrandIndex respondents for each of the BrandIndex-based metrics used in our analysis, using as an example the exact question text for the "Dining: Fast Food" product category. For each question, we list the name given to this metric by YouGov, and specify whether the metric is on a 2-point or 3-point response scale.

- "Aided Brand Awareness" (2-point, initial question): *Which of the following restaurant chains have you *ever* heard of? Please select all that apply.*
- "Buzz" (3-point): *Over the PAST TWO WEEKS, which of the following restaurant chains have you heard something POSITIVE about (whether in the news, through advertising, or talking to friends and family)? / Now which of the following have you heard something NEGATIVE about over the PAST TWO WEEKS?*
- "Attention" (2-point): [Yes if respondent reported Positive and/or Negative "Buzz"]
- "Consideration" (2-point): *When you are in the market next to purchase food or drink, from which of the following would you consider purchasing?*
- "Purchase Intent" (2-point): *From which of these would you be most likely to purchase?* [Follow-up to "Consideration"]
- "Current Customer" (2-point): *Have you purchased food or drink from any of the following restaurant chains in the past 30 days?*
- "Former Customer" (2-point): *Have you ever purchased food or drink from any of the following restaurant chains?* [Excludes "Current Customers"]
- "Impression" (3-point): *Overall, of which of the following restaurant chains do you have a POSITIVE impression? / Now which of the following restaurant chains do you have an overall NEGATIVE impression?*
- "Word-of-Mouth Exposure" (2-point): *Which of the following restaurant chains have you talked about with friends and family in the PAST TWO WEEKS (whether in-person, online, or through social media)?*

- "Advertising Awareness" (2-point): *Which of the following restaurant chains have you seen an advertisement for in the PAST TWO WEEKS?*

We construct the variables used in our analysis from these metrics questions as follows. As described in Section 4, we define $a_{jt} = \frac{\sum_{i \in I_{jt}} w_i \mathbb{1}\{\text{Reported Positive and/or Negative Buzz}\}}{\sum_{i \in I_{jt}} w_i}$ as the share reporting having heard something positive and/or negative about the firm in the past two weeks, among all responses I_{jt} that asked about j's firm in event-month t.[39] Responses i are weighted by the survey weights w_i provided by YouGov to make responses nationally-representative for that product-category and day. Constructing $\hat{\tau}_{jt} = \frac{a_{jt} - a_{j,-1}}{1 - a_{j,-1}}$, Figure 1 Panel A then plots the average of $\hat{\tau}_{jt}$ across all event-firms in a given event-month, among events covered by BrandIndex. Constructing $\hat{\tau}_j$ as equal to $\hat{\tau}_{j0}$ for events covered by BrandIndex and imputing this value from Google Trends and news coverage when not covered by BrandIndex (as described in Appendix Section A.1), Figure 1 Panel B then plots a histogram of this consumer awareness measure across all 117 events.

Appendix Figure B1 shows averages of $a_{\tilde{g}jt}$ and $\hat{\tau}_{\tilde{g}jt}$ across firms, adding an alignment group dimension $\tilde{g}$ by calculating these metrics for each firm-event separately among respondents ($I_{\tilde{g}jt}$) split by their social alignment.[40] YouGov's 2-point metrics (with the exception of "Attention") and the splits by party affiliation are also only available starting in 11/13/2012. When producing any given time-series figure, we balance our panel by dropping events that aren't covered throughout the period for the metric shown in that figure (e.g., a hypothetical event with date 1/1/2013 would be dropped from any figure showing splits by party affiliation ten months prior).

When producing Appendix Figure B13 Panel A, we first calculate for each group, firm-event, and month $buzz_{\tilde{g}jt} = \frac{\sum_{I_{\tilde{g}jt}} w_i (\mathbb{1}\{\text{Reported Positive Buzz}\} - \mathbb{1}\{\text{Reported Negative Buzz}\})}{\sum_{I_{\tilde{g}jt}} w_i}$, i.e. the share of respondents in a given alignment group reporting positive news about the firm less the share reporting negative news. We then define our outcome metric for each group, firm-event, and month as $\frac{buzz_{\tilde{g}jt} - buzz_{\tilde{g}j,-1}}{\hat{\tau}_j}$ and calculate a precision-weighted average of this series across firm-events for a given month and group (i.e., weighting firm-events by $\hat{\tau}_j^2$). In Panel B, we similarly calculate favorability towards the firm using responses to "Impression" rather than "Buzz."

To produce Appendix Figure B14 Panel A, we first calculate for each group, firm-event, and month $considers_{\tilde{g}jt} = \sum_{I_{\tilde{g}jt}} \frac{w_i (\mathbb{1}\{\text{Included firm among answers to Consideration}\})}{\sum_{I_{\tilde{g}jt}} w_i}$, i.e., the share of respondents who report that they would consider purchasing from that firm when next in the market for its product category. We then similarly define our outcome metric as $\frac{considers_{\tilde{g}jt} - considers_{\tilde{g}j,-1}}{\hat{\tau}_j}$

[39] We do *not* exclude individuals who were unaware of the firm (in the "Aided Brand Awareness" question) from the denominator, although they did not answer subsequent questions about the brand.

[40] We define social alignment among BrandIndex survey respondents based on their self-reported party affiliation, oriented relative to our donation clusters and stances based on related donations in those clusters. This demographic question was answered previously and asks "*Generally speaking, do you think of yourself as a...? [Democrat, Republican, Independent, Other, Not Sure]*." We drop the small share of respondents answering "Other" or "Not Sure" to this question from all analyses of splits by alignment group in the BrandIndex data.

and calculate a precision-weighted average of this series across firm-events for a given month and group (i.e., weighting firm-events by $\hat{\tau}_j^2$). We similarly produce Panel B by calculating purchase intent (or more precisely that the respondent reports being most likely to purchase from the firm) using responses to "Purchase Intent" rather than "Consideration". We similarly produce Appendix Figure B15 Panels A and B using responses to "Word-of-Mouth Exposure" and "Ad Awareness" respectively.

B Appendix Exhibits

Appendix Figure B1: Consumer Awareness of Firm Social Stances, by Alignment

Panel A: Unusual Awareness of News About Firm ($\hat{\tau}_t := \frac{a_t - a_{-1}}{1 - a_{-1}}$)

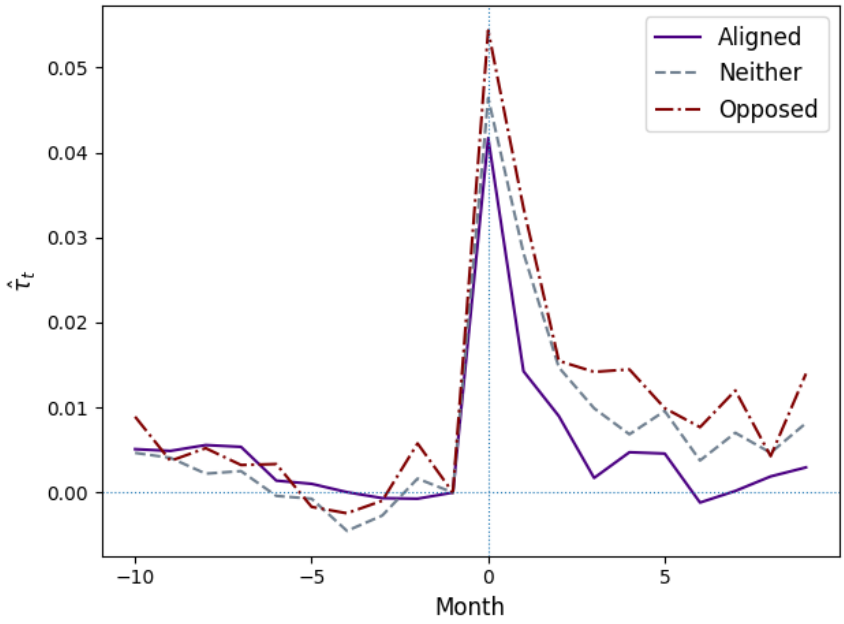

Panel B: Share Reporting Recent Good or Bad News (a_t)

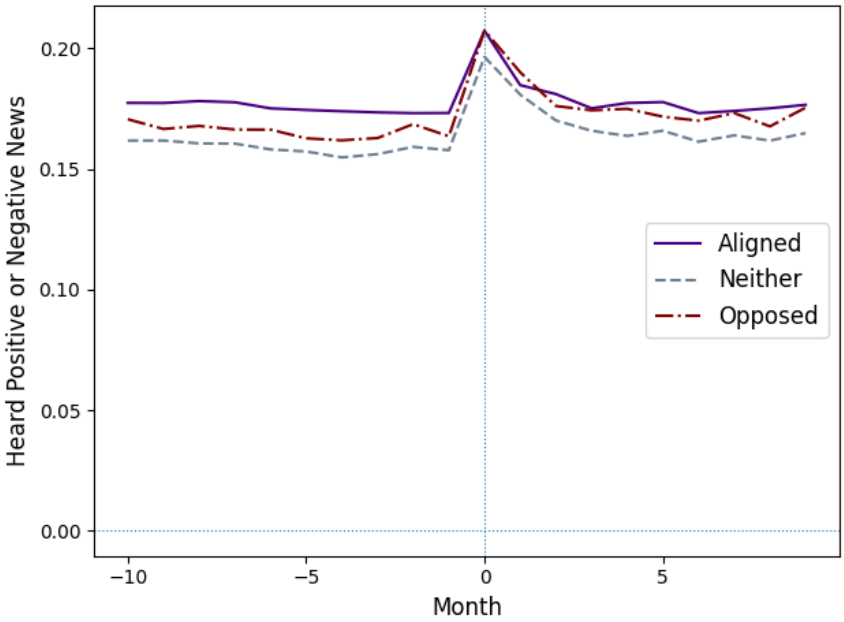

Note: Figure shows consumer awareness of firms' social stance events by social alignment, based on responses to BrandIndex surveys. Define a_t as the share who report having heard positive or negative news about the brand in the last two weeks among all respondents in month t. Panel A shows our consumer awareness measure $\hat{\tau}_t := \frac{a_t - a_{-1}}{1 - a_{-1}}$, averaged by month across event-study firms separately for respondents by alignment. Panel B similarly shows averages of a_t itself.

Appendix Figure B2: Median Predicted Social Alignment Decile, among All Cards by County

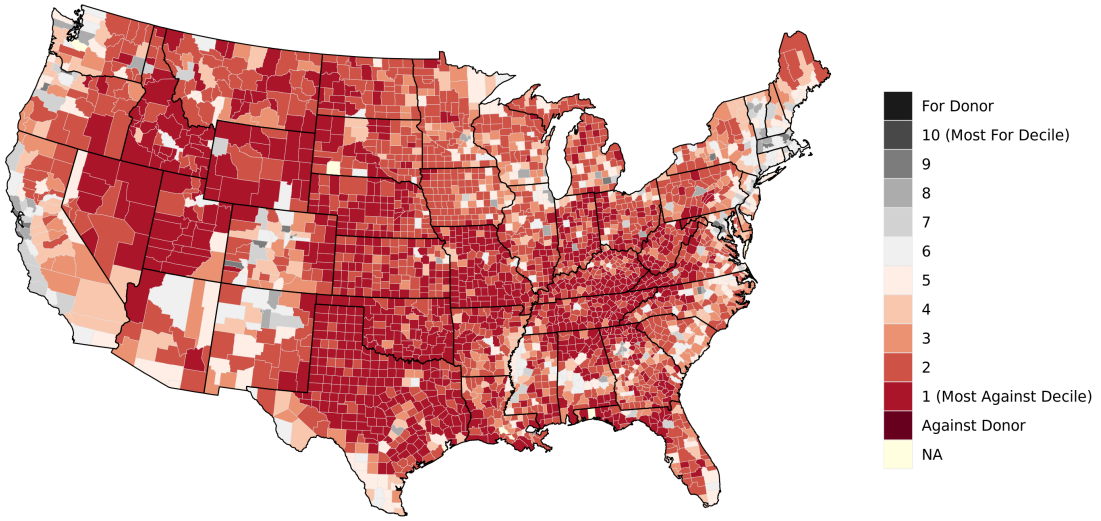

Note: Figure maps (for each county) the median predicted probability of alignment with causes in the (arbitrarily labeled) "For" donation cluster among all cards in that state, with deciles 10 and 1 denoting non-donors most likely to be aligned with vs. opposed to causes in this cluster, respectively.

Appendix Figure B3: Group Shares of Pre-Existing Consumption at Event Study Firms, by Position

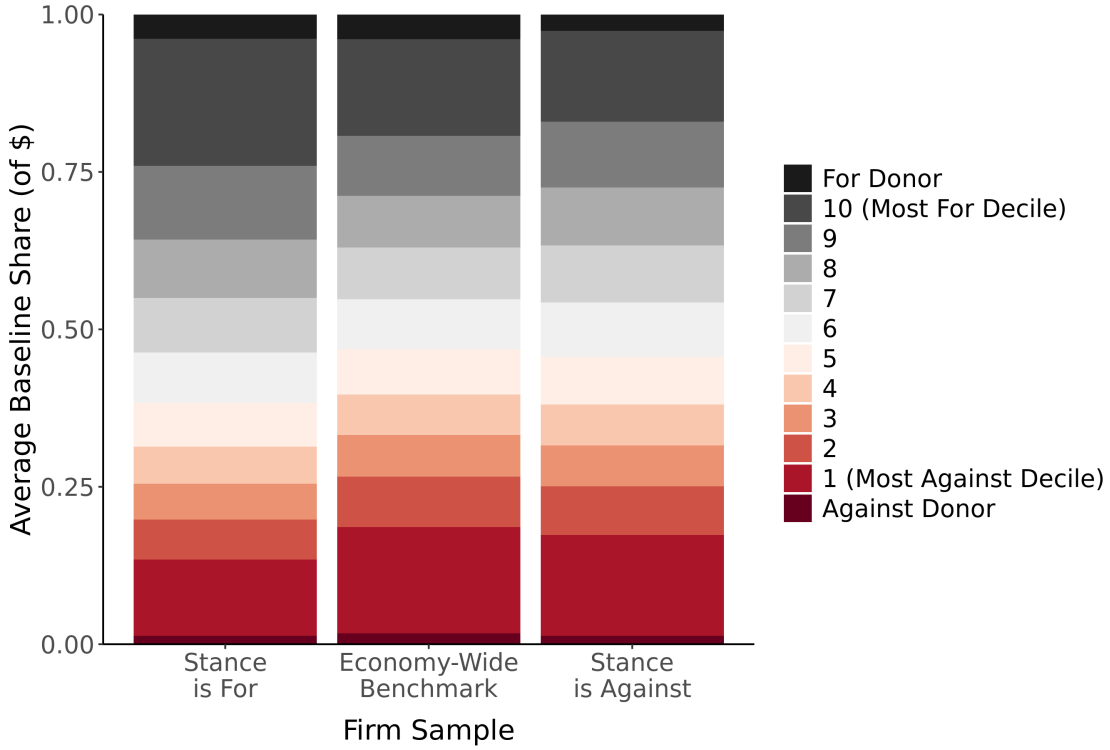

Note: Figure shows shares of consumption (in $) by group. Consumer groups are defined as described in Section 5, ordering consumers based on their predicted social preference alignment with the For donation cluster on social issues. The leftmost and rightmost columns show these baseline shares for consumption at firms taking social stances aligned with vs. opposed to donations in this cluster, respectively, in the year preceding these stances, which we refer to as baseline shares. Baseline shares are weighted by consumer awareness of the firm's stance (τ_j, as defined in Section 4) when averaging baseline shares across these firm events. In the middle bar, we show each group's share of consumption (in $) aggregating across all U.S. firms in the transaction data throughout the period studied (2008–2023Q1).

Appendix Figure B4: Consumption Responses by Group (vs. Group's Consumption at All Other Firms)

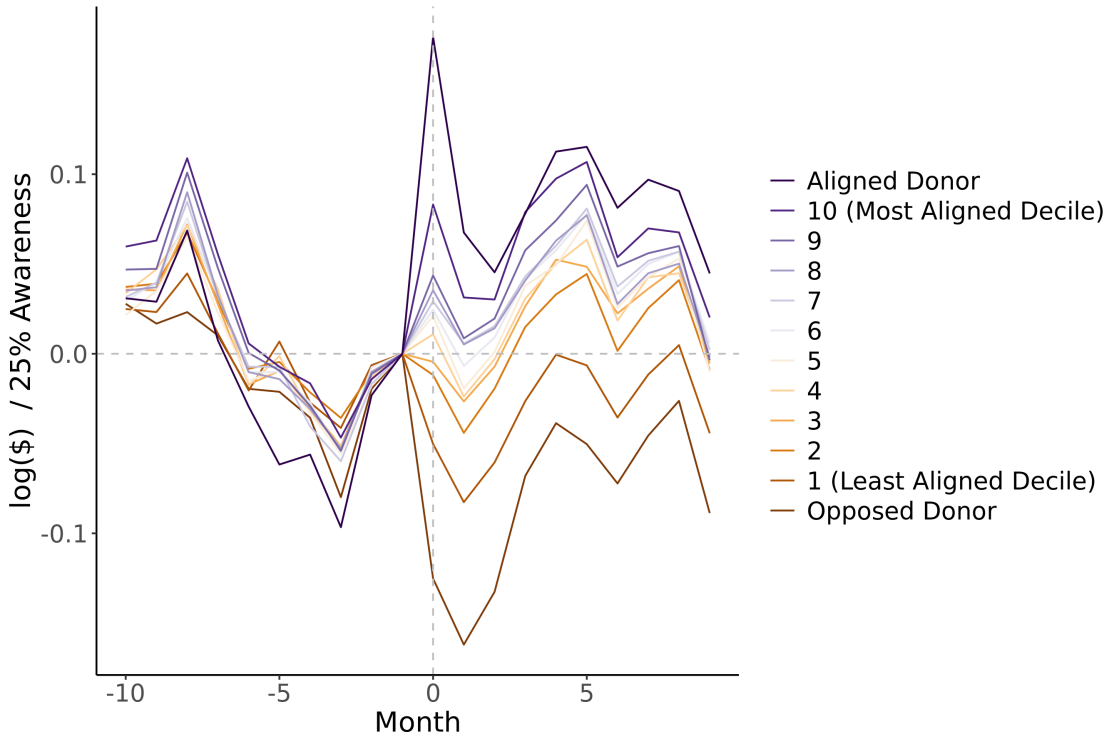

Note: Figure shows changes in log consumption at firms in the months surrounding their social stances, by consumer social alignment groups. Consumer social alignment groups are constructed as defined in Section 5.1. Changes in log consumption by group are normalized relative to the month before a firm's social stance and relative to changes in that group's consumption at all other firms in the economy. Changes are scaled relative to consumer awareness and are averaged across firms using a precision-weighted average, as described in Sections 4 and 6.2.

Appendix Figure B5: Predicting No-Event Counterfactual Consumption at Event-Study Firms

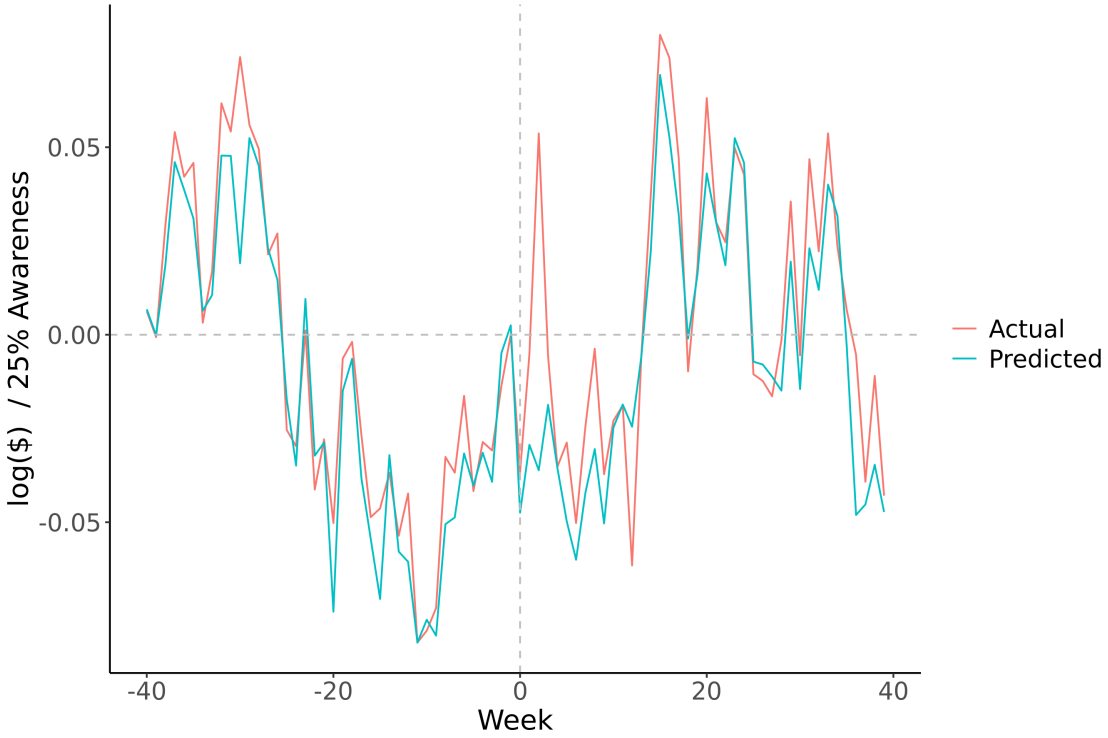

Note: Figure shows actual and predicted log consumption at firms taking social stances by event-week, normalized for visualization purposes relative to the month prior to the firm's social stance event and to changes in log consumption at all other firms in the economy. Log consumption in the absence of a firm's social stance is predicted using a synthetic difference-in-differences design as described in Section 6.2, using as predictors contemporaneous consumption at other firms and past consumption at the social stance firm. Changes are scaled relative to consumer awareness and are averaged across firms using a precision-weighted average, as described in Sections 4 and 6.2.

Appendix Figure B6: Differences in Consumption Response, Across Groups

Panel A: Aligned−Opposed Donors

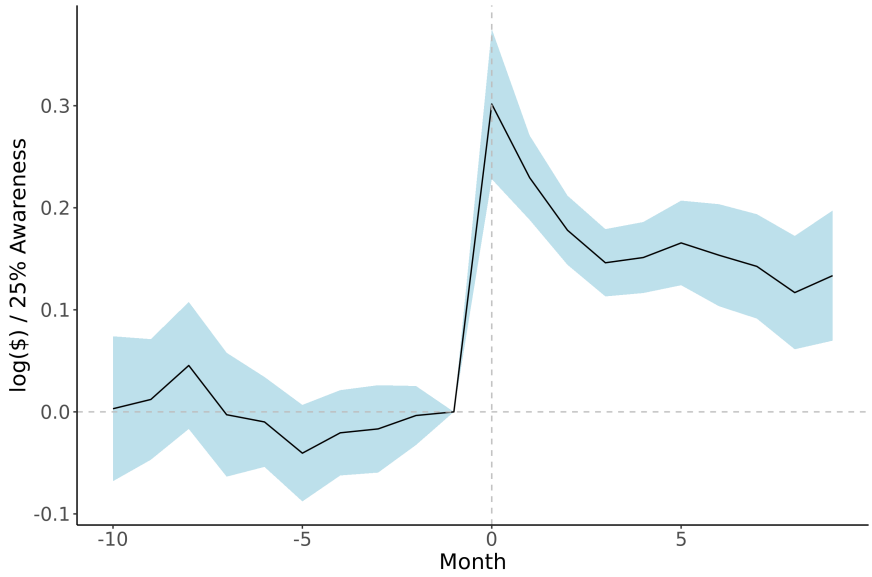

Panel B: Most Aligned−Most Opposed Non-Donor Deciles

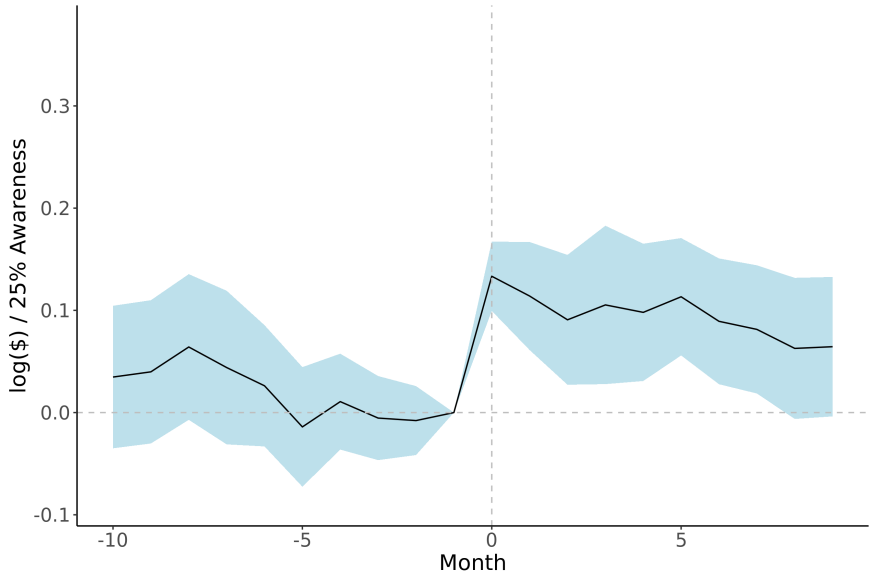

Note: Figure shows differences between the consumption responses of different groups (as shown in Figure 4), along with a 95% confidence interval for this difference. Panel A shows the consumption response difference among the Aligned vs. Opposed donor groups. Panel B shows the consumption response difference among the most aligned decile vs. most opposed non-donor decile groups. Responses are scaled relative to consumer awareness and are averaged across firms using a precision-weighted average, as described in Sections 4 and 6.2. Standard errors are clustered by event.

Appendix Figure B7: (One-Year-Prior Placebo) Changes in Consumption at Social Stance Firms, by Group

Panel A: Response Levels by Group (vs. Group's Consumption at All Other Firms)

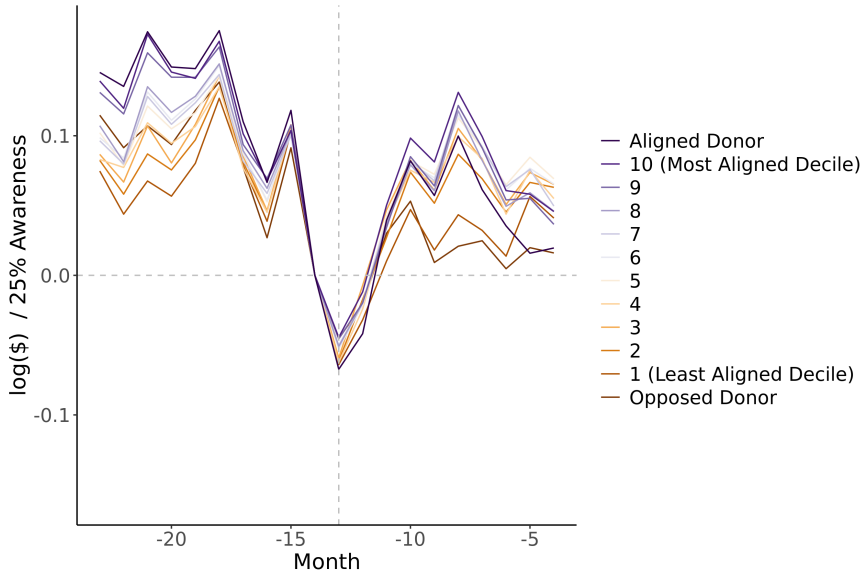

Panel B: Response Effects by Group (Shifting Levels to Match Estimated Overall Impact)

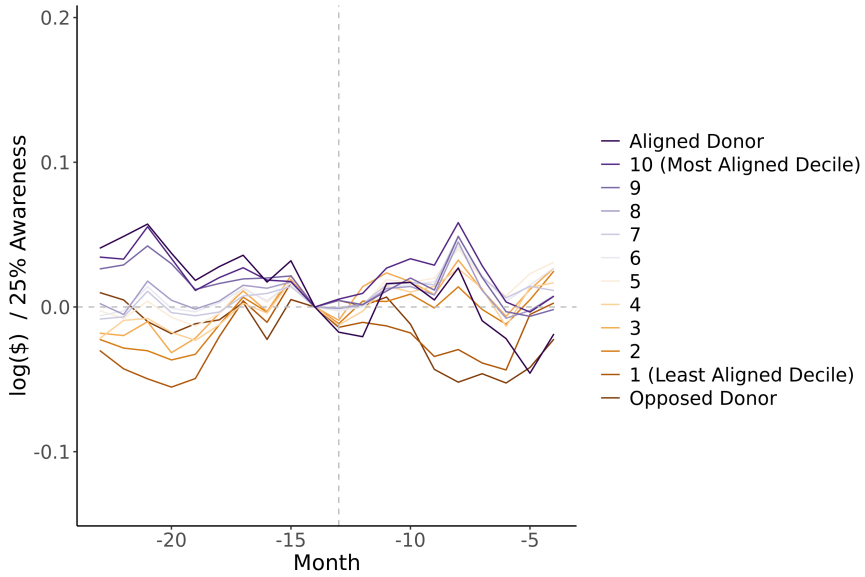

Note: Figure shows changes in log consumption at firms by consumer social alignment groups in a placebo exercise. In this placebo exercise, we rerun our analysis as if social stance events occurred one year prior to their actual date, including rerunning our synthetic DiD forecasts trained only on data prior to this placebo date. No actual social stance events occurred on these one-year leading dates. The y-axis range and all other specifications of Panels A and B follow Appendix Figure B4 and Figure 4 Panel B, respectively, except for the use here of placebo event dates.

Appendix Figure B8: Placebo Consumption Responses vs. Share Aligned, by Group

Panel A: Placebo Event-Month Consumption Response vs. Share Aligned, by Group

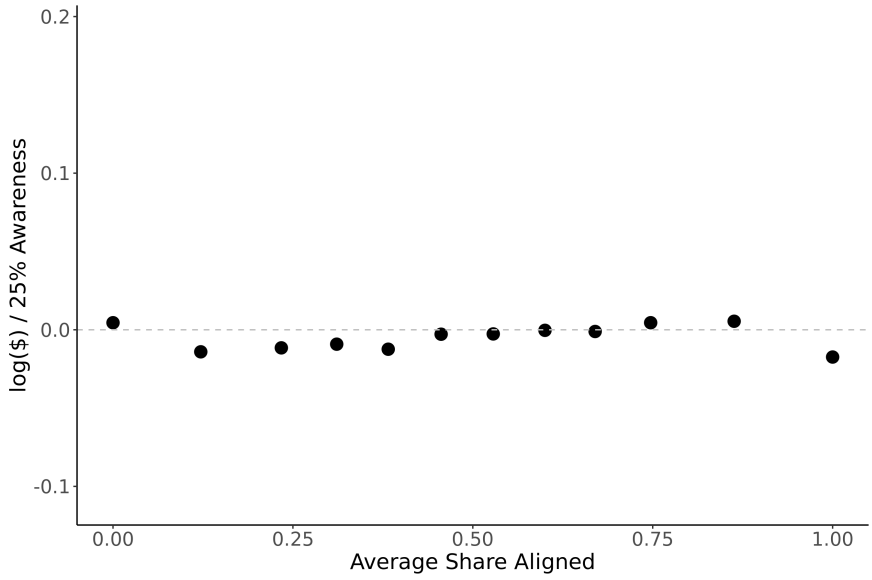

Panel B: Placebo Consumption Response Gradient vs. Share Aligned, Across Groups

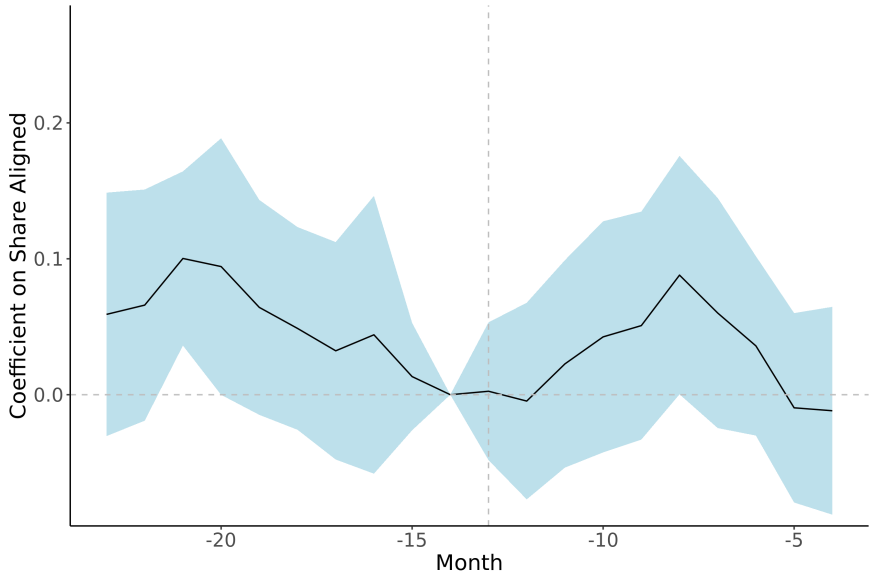

Note: Figure compares placebo consumption responses (as shown in Appendix Figure B7 Panel B) by group to the mean share of individuals in that group predicted to be aligned with the firm's stance. In this placebo exercise, we rerun our analysis as if social stance events occurred one year prior to their actual date, including rerunning our synthetic DiD forecasts trained only on data prior to this placebo date. Following Figure 5, Panel A shows for each alignment group the average estimated (placebo) consumption response in the month of the firm's social stance event (y-axis, matching the t=0 value in Appendix Figure B7 Panel B) vs. the average share of consumers predicted to be aligned with the firm's stance (x-axis). Panel B plots the coefficients β_t and 95% confidence intervals from the regression of consumption responses on share aligned described in Section 6.3, clustering standard errors by event.

Appendix Figure B9: Differences in Placebo Consumption Response, Across Groups

Panel A: Aligned−Opposed Donors

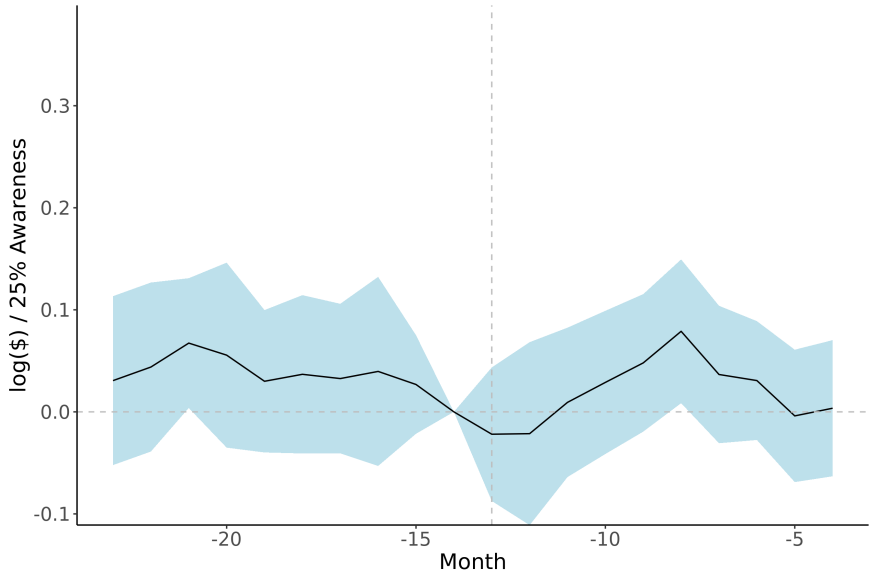

Panel B: Most Aligned−Most Opposed Non-Donor Deciles

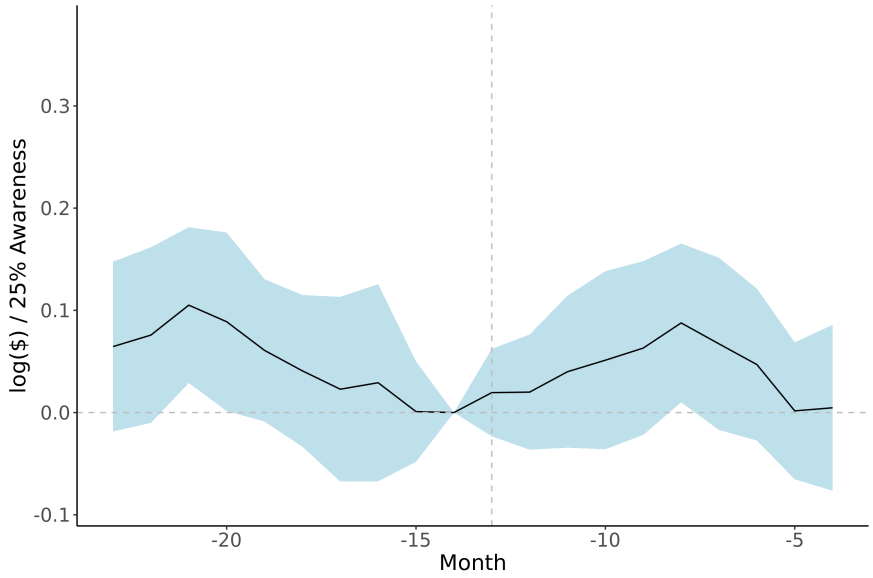

Note: Figure shows differences between the placebo consumption responses of different groups (as shown in Appendix Figure B7) , along with a 95% confidence interval for this difference. Panel A shows the consumption response difference among the Aligned vs. Opposed donor groups. Panel B shows the consumption response difference among the most aligned decile vs. most opposed non-donor decile groups. Responses are scaled relative to consumer awareness and are averaged across firms using a precision-weighted average, as described in Sections 4 and 6.2. Standard errors are clustered by event.

Appendix Figure B10: Changes in Consumption at Social Stance Firms, by Group (3-Year Pre-Period)

Panel A: Response Levels by Group (vs. Group's Consumption at All Other Firms)

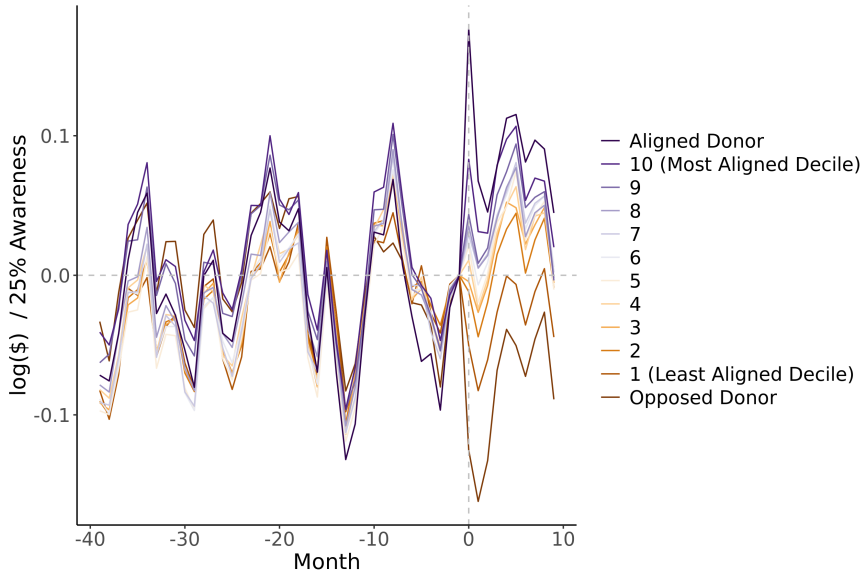

Panel B: Response Effects by Group (Shifting Levels to Match Estimated Overall Impact)

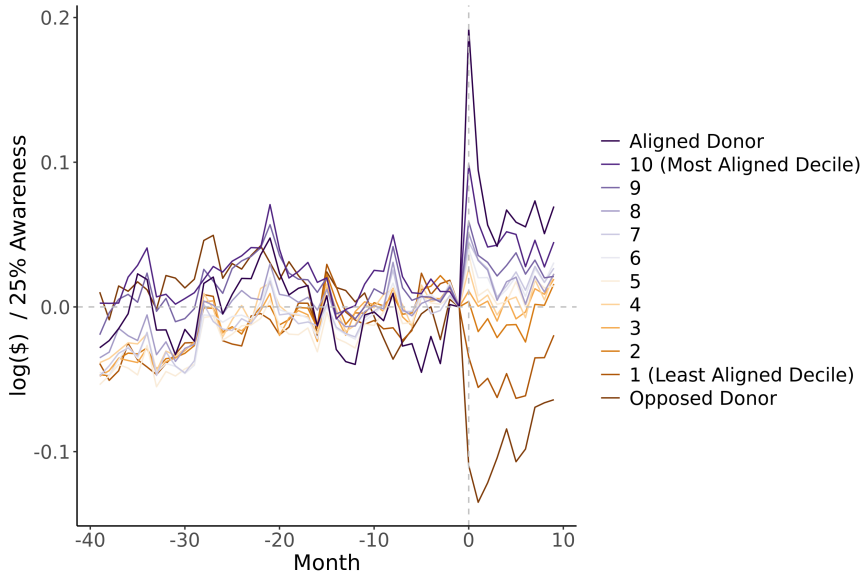

Note: Figure shows changes in log consumption at firms in the months surrounding their social stances by consumer social alignment groups. These panels have been modified to show 3-years of pre-event data (i.e., 39 pre-event 4-week "months"), with all other specifications for Panels A and B following Appendix Figure B4 and Figure 4 Panel B, respectively.

Appendix Figure B11: Response Levels by Group (2-Year Post-Period)

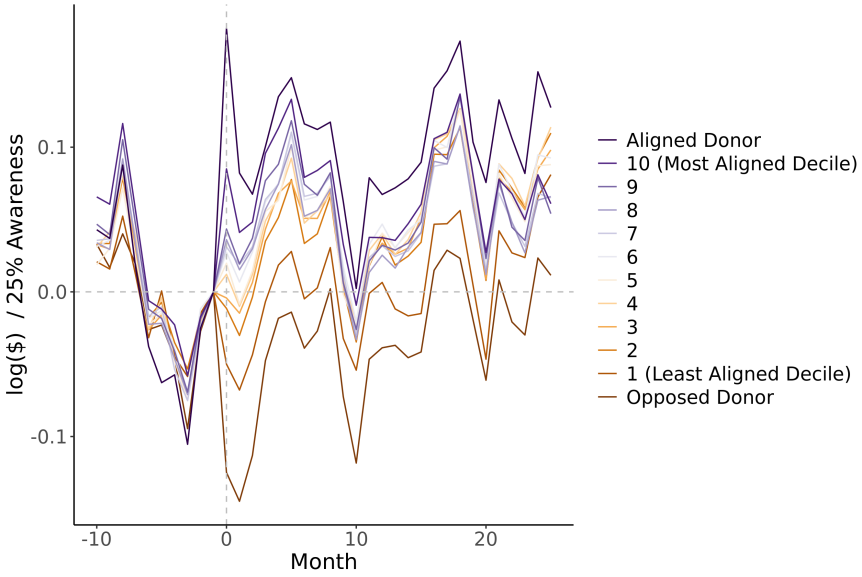

Note: Figure shows changes in log consumption at firms in the months surrounding their social stances by consumer social alignment groups. Figure modifies Appendix Figure B4 to show 2-years of post-event data (i.e., 26 post-event 4-week "months"), with all other specifications following this appendix figure.

Appendix Figure B12: Social Stance Consumption Impacts, by Cluster Alignment

Panel A: Overall Average Consumption Impacts

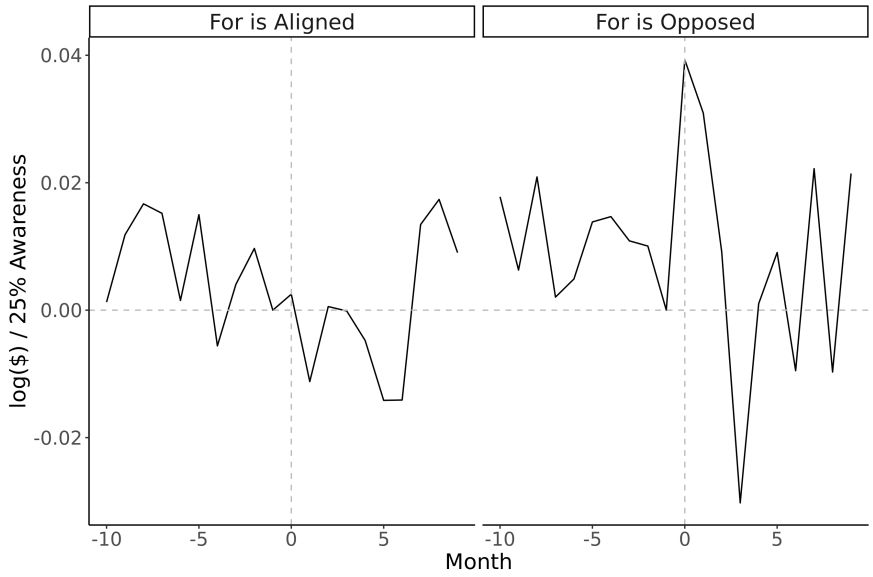

Panel B: Consumption Responses, by Group

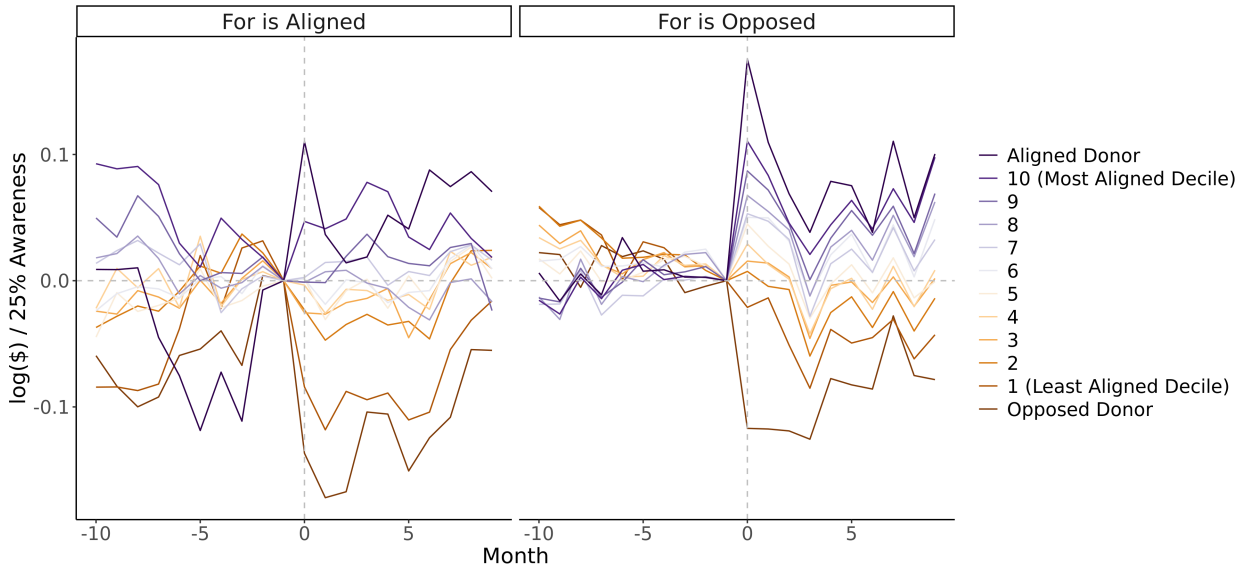

Note: Figure shows estimated overall and disaggregated consumption impacts separately among events where donors in the For cluster are likely aligned with vs. opposed to the firm's stance. Panel A extends Figure 4 Panel A in separately showing estimated overall consumption impacts. Panel B extends Figure 4 Panel B in separately showing estimated consumption responsiveness, by social alignment group. Groups in Panel B are colored according to their likely alignment with the firm's stance.

Appendix Figure B13: Interpretation of News by Alignment, in BrandIndex

Panel A: Net Favorability of News

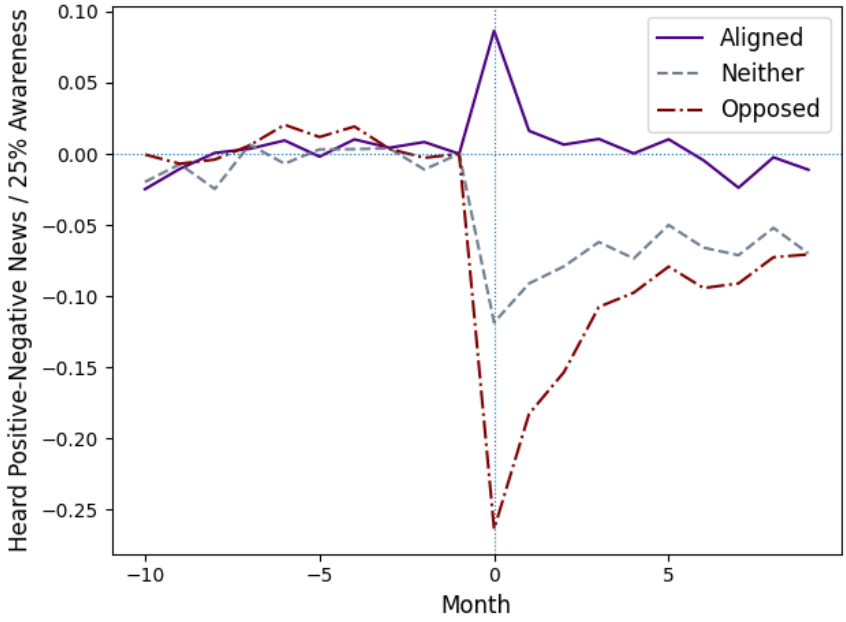

Panel B: Net Favorability towards Firm

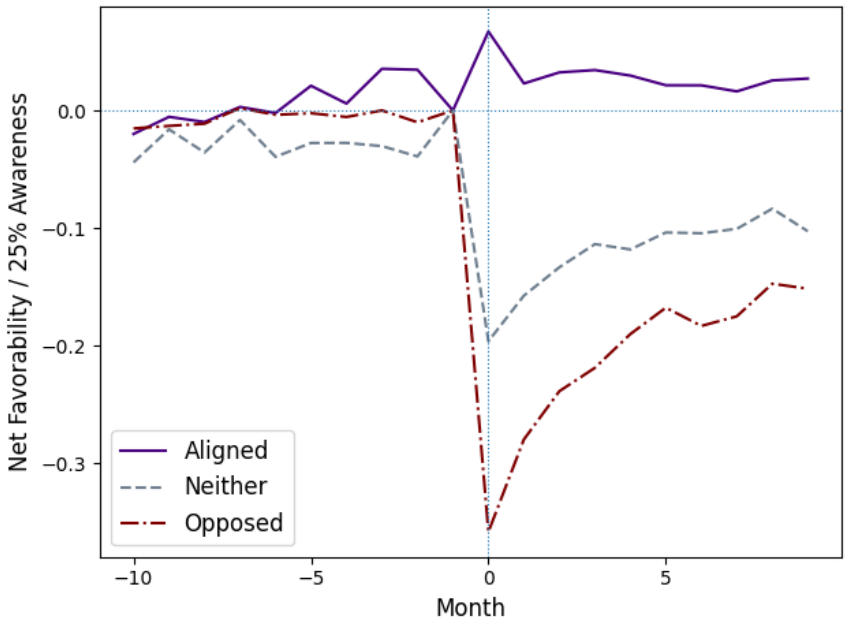

Note: Figure shows changes in favorability around firms' social stance events by social alignment, based on BrandIndex responses. Panel A shows favorability regarding news about the firm, coded as 1 if the respondent reported having heard "Positive" news about the firm in the last two weeks, -1 if reported having heard "Negative" news, and 0 otherwise. Panel B shows favorability towards the firm more generally by the respondents, again coded as +1 (Positive), -1 (Negative), or 0 (Neutral). Responses are scaled relative to consumer awareness, averaged across firms using a τ_j^2-weighted average, and normalized relative to the month before a firm's event ($t = -1$). Months are defined as 4-week periods relative to the firm's event.

Appendix Figure B14: Impact on Self-Reported Purchase Behavior by Alignment, in BrandIndex

Panel A: Would Consider Purchase at Firm

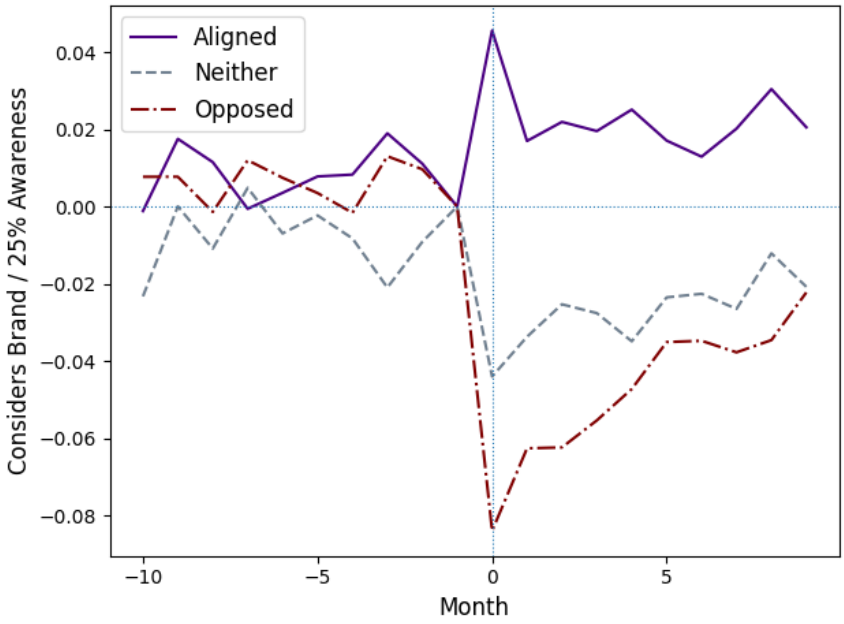

Panel B: Intend to Purchase from Firm

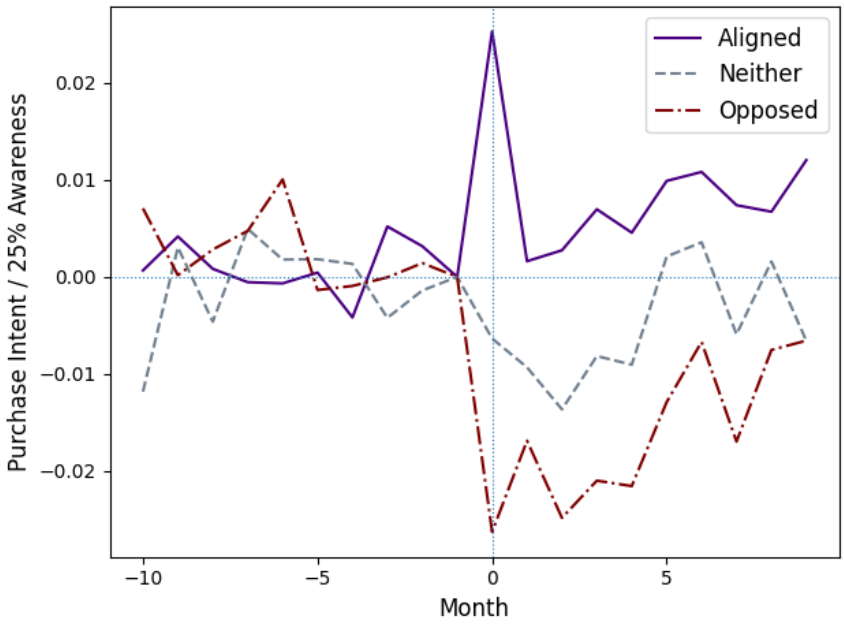

Note: Figure shows changes in self-reported purchase behavior at firms in the months surrounding their social stances, based on BrandIndex responses. Panel A shows an indicator for whether a respondent would consider purchasing from the firm when next shopping in that firm's market (1 if yes; 0 if no), and Panel B shows an indicator for whether a respondent would be most likely to purchase from that firm. Responses are scaled relative to consumer awareness, averaged across firms using a τ_j^2-weighted average, and normalized relative to the month before a firm's event ($t = -1$). Months are defined as 4-week periods relative to the firm's event.

Appendix Figure B15: Channels for Learning About Firm Stance by Alignment, in BrandIndex

Panel A: Word-of-Mouth Exposure

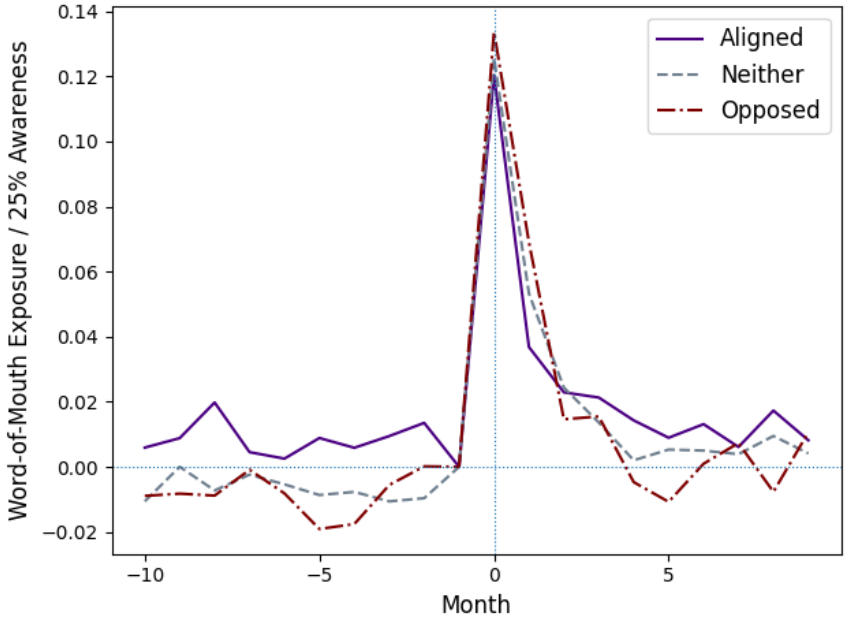

Panel B: Ad Awareness

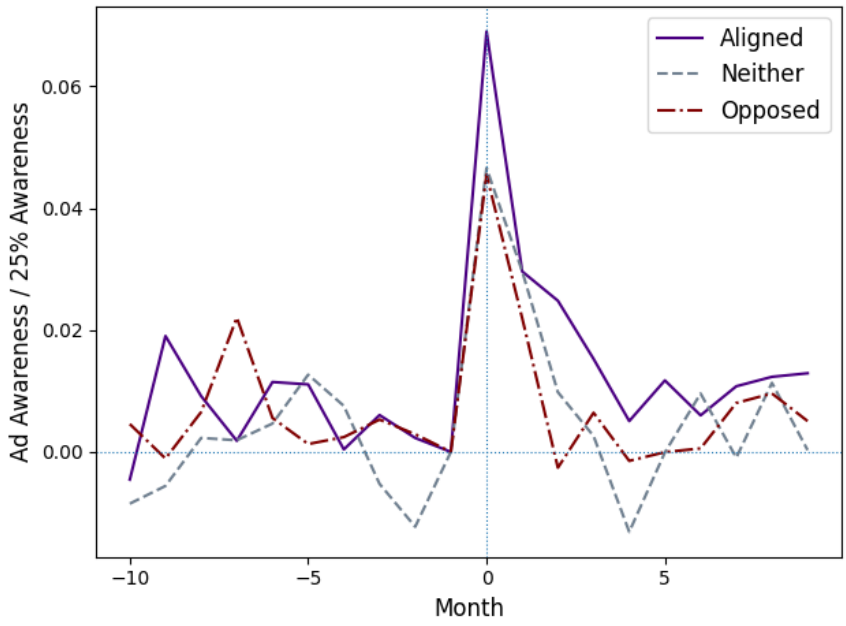

Note: Figure shows changes in exposure to information about firms, based on BrandIndex responses. Panel A shows an indicator for whether a respondent recently talked with someone about the brand (in-person, online, or through social media), and Panel B shows an indicator for whether a respondent recently saw an advertisement from that firm. Responses are scaled relative to consumer awareness, averaged across firms using a τ_j^2-weighted average, and normalized relative to the month before a firm's event ($t = -1$). Months are defined as 4-week periods relative to the firm's event.

Appendix Figure B16: Stock Returns of Event Study Firms vs. Event Size

Panel A: Stock Returns

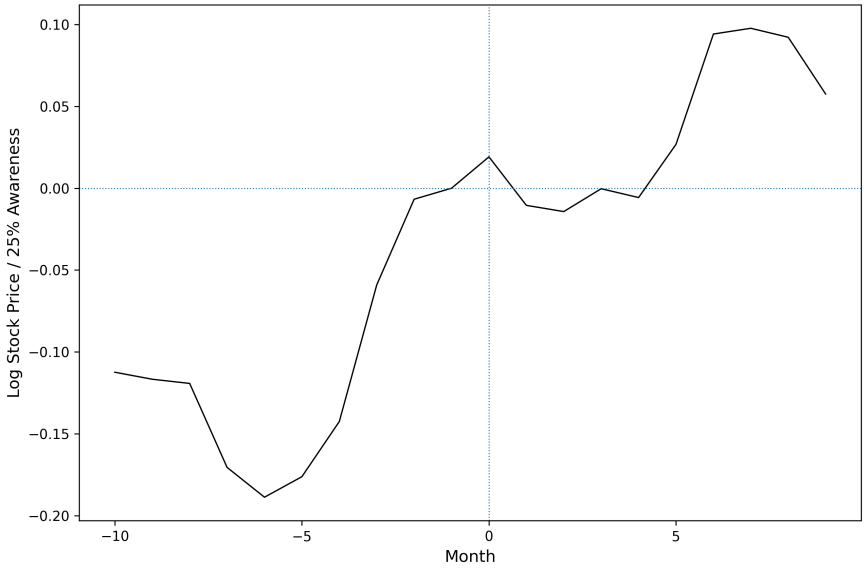

Panel B: Excess Stock Returns (vs. S&P 500 Benchmark)

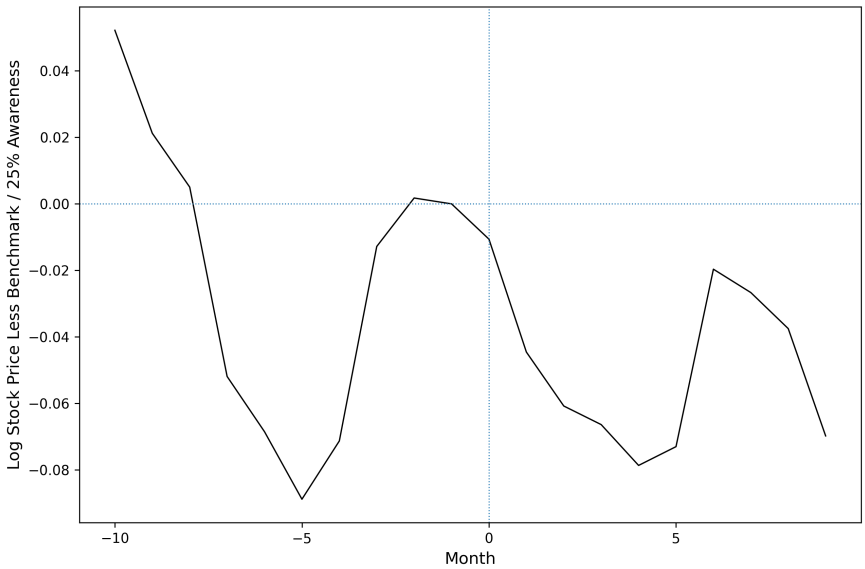

Note: Figure shows a τ_j^2 weighted-average of stock returns (Panel A) and excess stock returns (Panel B) per event size, averaged across firms taking a social stance around their event and normalized by week $t = -1$. Excess stock returns are defined using returns on the S&P 500 as a benchmark return, and consumer awareness (τ_j) of social stance j is defined based on contemporaneous brand perception surveys as described in Section 4.2.

Appendix Figure B17: Cumulative Net Sales Impact, by Alignment Direction

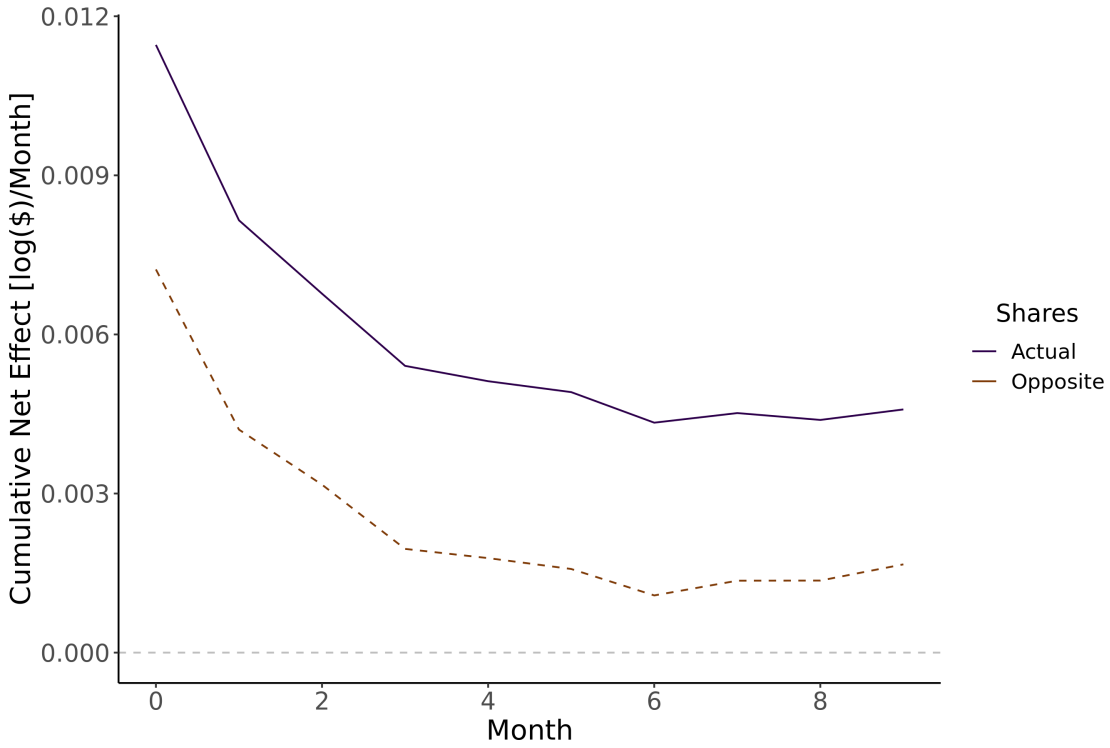

Note: Figure shows the cumulative monthly impact of aggregated group-specific consumption responses after a given number of months. These cumulative monthly impacts are estimated using the average stance consumption response estimates from Figure 4 Panel B. These effects are aggregated across groups using either the τ_j-weighted average baseline consumption shares shown in Figure 3 of firms' actual consumer base, or alternatively by the (reversed) baseline shares they would have faced had they taken counterfactual stances in the opposite For/Against direction on the same issue topic. See Section 7 for detail.

Appendix Figure B18: Variation in Baseline Group Consumption Shares, by Market

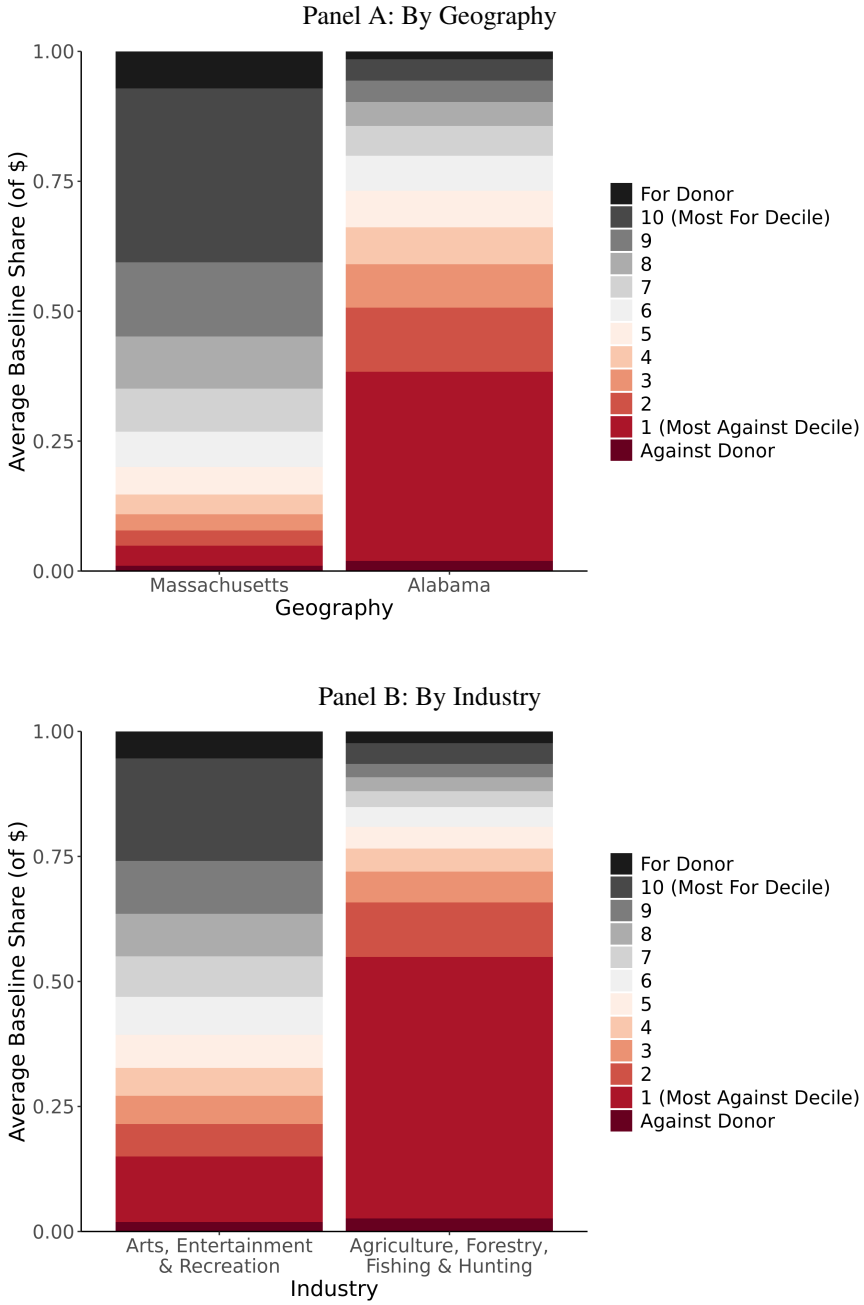

Note: Figure shows the share of consumption (in $) by group, which we refer to as baseline shares, in different markets. Consumer groups are defined as described in Section 5, ordering consumers based on their predicted social alignment with positions in the For donation cluster. Panel A shows each group's share of consumption (in $) by state for Massachusetts and Alabama separately, aggregating across all transactions throughout the period studied (2008-2023Q1) made by consumers based in the specified state. Panel B shows each group's share of consumption (in $) separately for two example industries ("Arts, Entertainment & Recreation" and "Agriculture, Forestry, Fishing & Hunting") aggregating all transactions made at firms in the specified industry throughout the period studied.

Appendix Figure B19: Cumulative Net Sales Impacts By State-Level Baseline Shares and Hypothetical Stance Direction

Panel A: Hypothetical Stance, Direction Aligned with For Cluster

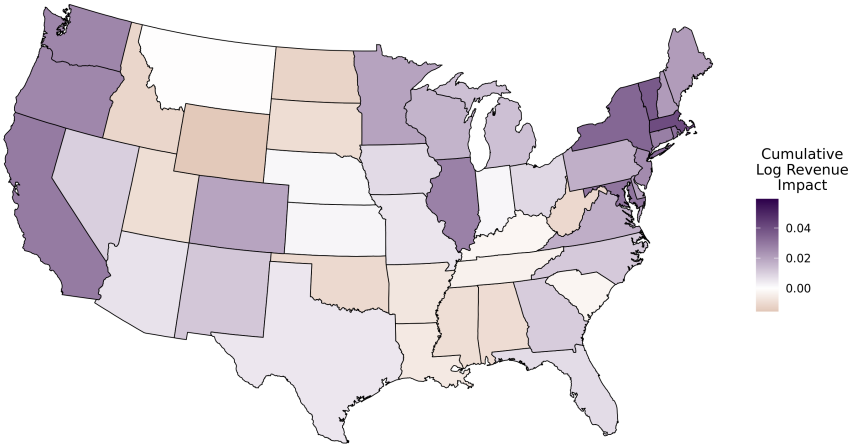

Panel B: Hypothetical Stance, Direction Aligned with Against Cluster

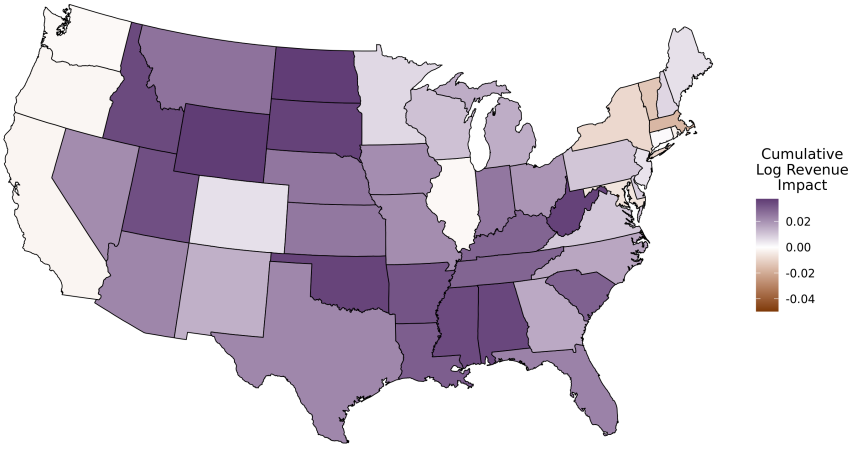

Note: Figure shows the average monthly log sales impact of a firm's social stance after five months, using average consumption response estimates identified in Figure 4 Panel B combined with state-specific baseline shares. These baseline shares are constructed using all consumption within a state, as shown in Appendix Figure B18 Panel A for Massachusetts and Alabama. Panels A and B, respectively, show the average monthly log sales impact (per 25% consumer awareness) induced by an average stance in the same direction vs. in the opposite direction as positions in the For donation cluster.

Appendix Figure B20: BrandIndex-Based Consumer Awareness, by Event

Panel A: Share Hearing News about Firm

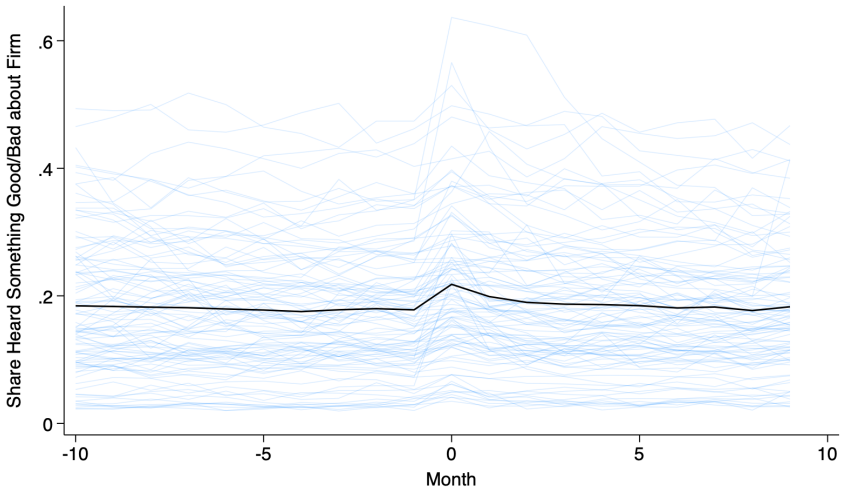

Panel B: Consumer Awareness of Firm Social Stances

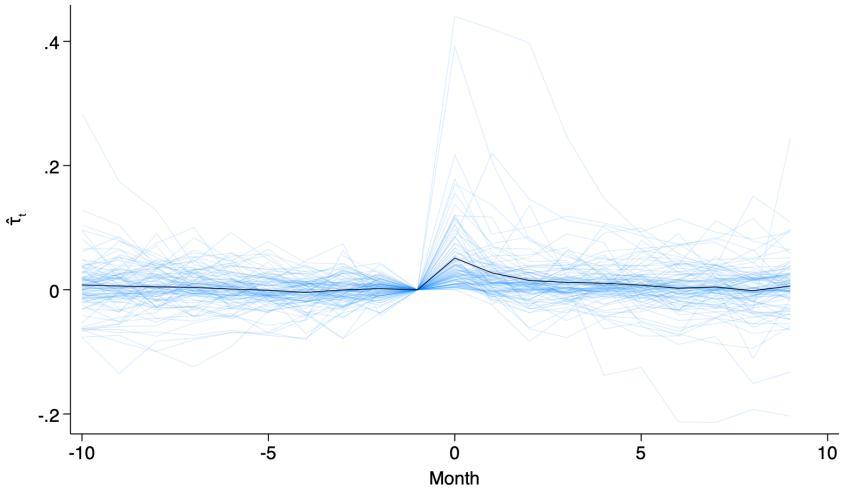

Note: Figure shows time trends related to our BrandIndex-based consumer awareness measure, plotting separate blue lines for each event covered by BrandIndex as well as a black line for the mean across events. Panel A shows monthly trends in the share of consumers who report having good or bad news about the firm in the last two weeks, as described in Appendix Section A.2. Denoting this share as a_{jt} for event-firm j in month t, Panel B then shows trends in $\hat{\tau}_{jt} := \frac{a_{jt}-a_{j,-1}}{1-a_{j,-1}}$. The value for a given line in month $t = 0$ gives our estimate of the share of consumers who were aware of that firm's social stance event (see Appendix Section A.2 for detail). Months are defined as 4-week periods relative to the firm's event.

Appendix Figure B21: News Coverage of Event-Study Firms

Panel A: Mean News Coverage of Social Stance Event-Study Firms Over Time

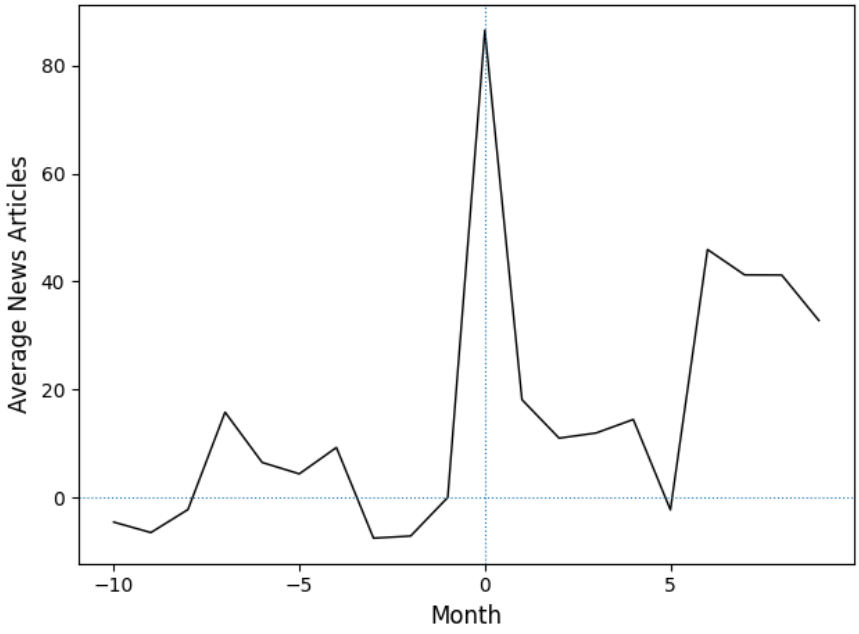

Panel B: News-Based Event Size Distribution (Histogram)

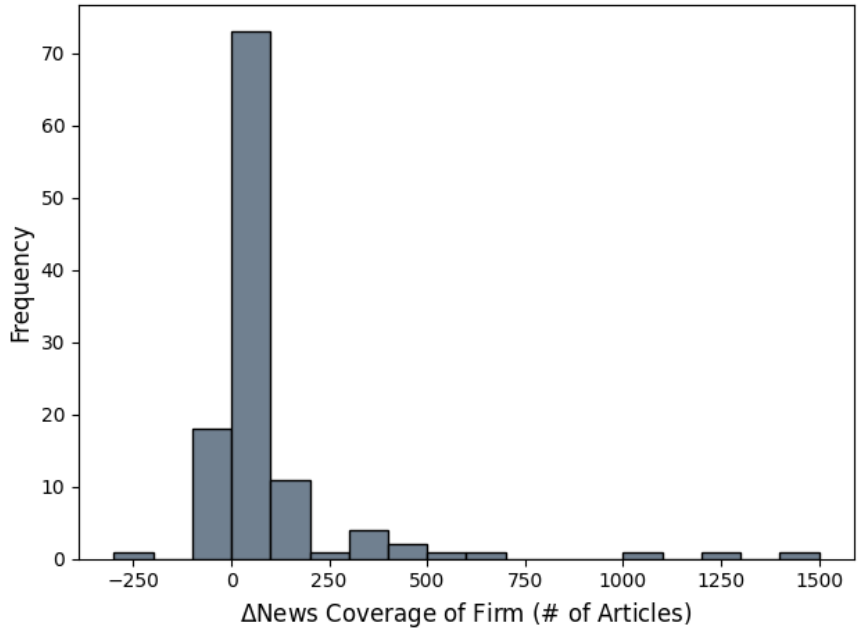

Note: Figure shows changes in news coverage of firms with social stance events around the date of this event. Panel A shows the number of TDM ProQuest U.S. Newsstream articles about the firm in each month (4-week window), averaged by month across event-study firms. Months are defined as 4-week periods relative to the firm's event. Social stance firms are the subject of 159 articles in month $t = -1$, and this value is normalized to zero in Panel A. Panel B shows a histogram summarizing across events the change in news coverage of the event-study firm between months $t = -1$ and $t = 0$, which is an alternative proxy for an event's size or salience as defined in Section 4.

Appendix Figure B22: Google Trends Searches for Event Study Firms

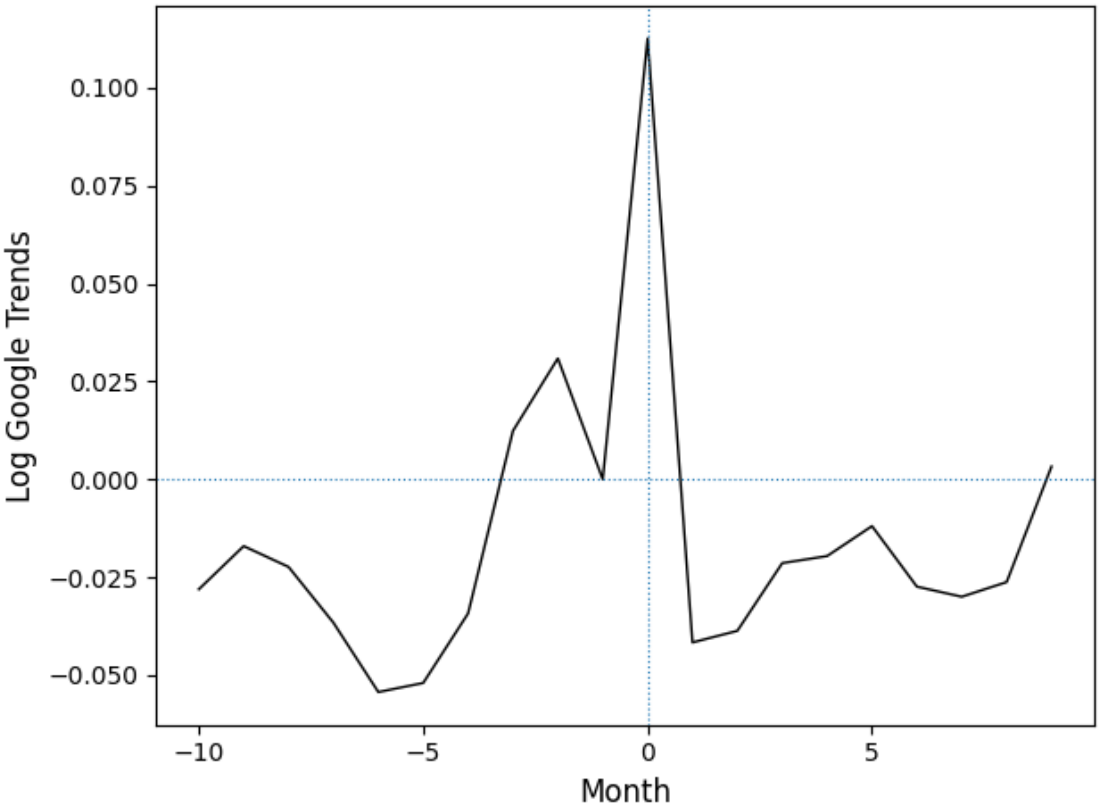

Note: Figure shows changes in log Google Trends searches for firms with social stance events around the date of this event. Changes are normalized relative to the month before a firm's event ($t = -1$). Months are defined as 4-week periods relative to the firm's event.

Appendix Table B1: Generic List of All Social Stance Events

Year	Direction	$\hat{\tau}_j$	Imputed $\hat{\tau}_j$	Description
2012	Against	0.440	No	Comments against same-sex marriage
2018	For	0.392	No	Endorsement of controversial racial justic protester
2014	Against	0.273	Yes	Contraceptive related corporate policy
2018	For	0.217	No	Gun control stance and policies
2017	For	0.179	No	Removed brand of controversial figure
2021	For	0.171	No	Banned account of controversial figure
2016	For	0.171	No	Transgender bathroom policy
2017	For	0.163	Yes	Pro-immigration stance and ad campaign
2020	For	0.156	No	Controversial diversity training
2017	Against	0.137	No	Criticism of protesters
2018	For	0.121	No	Severed ties with pro-gun group
2018	For	0.119	No	Severed ties with pro-gun group
2016	For	0.117	No	Prioritized privacy over law enforcement cooperation
2021	For	0.116	No	Stance on voting legislation
2015	For	0.113	No	Controversial diversity and inclusion campaign
2014	For	0.108	Yes	Pro-LGBTQ+ policy
2011	Against	0.106	No	Pulled ads from controversial program
2021	For	0.098	No	Support of controversial public policy
2018	For	0.091	No	Opposition to controversial figure
2013	For	0.086	No	Removed popular figure for anti-LGBTQ remarks
2021	For	0.086	No	Removed controversial video and app
2022	For	0.084	No	Opposition to anti-LGBTQ legislation
2018	Against	0.075	No	Donated to controversial cause
2020	For	0.074	No	Banned controversial group
2017	For	0.068	No	Opposed immigration restrictions
2012	For	0.067	No	Supported same-sex marriage
2015	For	0.062	No	Removed brand of controversial figure
2015	For	0.062	Yes	Supported same-sex marriage
2017	Against	0.060	No	Supported controversial figure
2021	Against	0.055	No	Donated to controversial causes
2021	For	0.055	No	Stance on voting legislation
2017	For	0.055	No	Employee fired over anti-diversity memo
2016	Against	0.052	No	Supported controversial figure
2018	Against	0.051	Yes	Continued ties with pro-gun group
2016	For	0.049	Yes	Supported Black Lives Matter
2018	For	0.049	No	Severed ties with pro-gun group
2014	For	0.048	No	Pro-LGBTQ+ campaign
2021	For	0.048	No	Supported abortion access
2021	Against	0.047	No	Produced program perceived as anti-transgender
2017	Against	0.045	No	Response to immigration restrictions
2012	For	0.044	No	Suspended sale of some types of firearms
2016	For	0.042	No	Opposition to controversial figure
2016	For	0.041	No	Opposition to anti-LGBTQ legislation
2018	For	0.040	No	Pro-LGBTQ+ stance
2016	For	0.040	No	Opposition to anti-LGBTQ legislation
2019	For	0.040	No	Asked customers not to bring their firearms in stores
2021	For	0.038	No	Supported abortion access
2021	Against	0.038	No	Donations to causes perceived as anti-abortion
2015	For	0.037	No	Pro-LGBTQ+ advertisement
2017	For	0.035	No	Criticized comment by controversial figure
2014	For	0.034	No	Asked customers not to bring their firearms in stores
2017	Against	0.034	Yes	Controversial donations by key stakeholder
2017	For	0.034	No	Stopped promoting brand of controversial figure
2017	For	0.032	No	Pulled ads from controversial program
2019	Against	0.032	No	Donations to causes perceived as anti-abortion
2017	For	0.031	No	Severed ties with controversial figure
2021	Against	0.030	No	Controversial donations by key stakeholder
2016	For	0.029	Yes	Opposed anti-LGBTQ legislation
2018	Against	0.029	No	Donations to causes perceived as pro-gun
2018	For	0.029	No	Severed ties with pro-gun group
2019	Against	0.027	No	Supported controversial figure
2016	For	0.027	No	Perceived opposition to controversial figure
2015	For	0.026	Yes	Opposed anti-LGBTQ legislation
2017	For	0.026	No	Opposed immigration restrictions
2020	For	0.026	No	Supported Black Lives Matter

Continued on next page

Appendix Table B1: Generic List of All Social Stance Events

Year	Direction	$\hat{\tau}_j$	Imputed $\hat{\tau}_j$	Description
2016	For	0.026	Yes	Supports pro-choice abortion activist
2017	Against	0.025	No	Continued selling brand of controversial figure
2021	Against	0.025	No	Opposition to controversial figure
2018	For	0.024	No	Supports gun control
2018	For	0.024	Yes	Pulled ads from controversial program
2018	For	0.023	No	Pulled ads from controversial program
2020	For	0.023	No	Imposed health mandates
2016	For	0.023	No	Supported Black Lives Matter
2022	For	0.022	Yes	Imposed health mandates
2021	For	0.022	No	Supported abortion access
2018	For	0.022	No	Severed ties with pro-gun group
2018	For	0.021	No	Confirmed severed ties with pro-gun group
2017	For	0.021	No	Removed brand of controversial figure
2020	For	0.020	No	Supported Black Lives Matter and racial justice initiatives
2014	For	0.020	No	Removed book from controversial figure
2016	Against	0.019	Yes	Supported controversial figure
2020	Against	0.018	Yes	Donated to controversial cause
2015	For	0.018	Yes	Transgender locker room policy
2012	Against	0.017	No	Criticism of healthcare reform
2021	For	0.017	No	Removed brand of controversial figure
2019	For	0.017	No	Asked customers not to openly carry firearms in stores
2017	For	0.016	No	Opposed immigration restrictions
2018	For	0.016	No	Severed ties with pro-gun group
2017	For	0.016	No	LGBTQ+ support/inclusion
2017	For	0.015	No	Opposed immigration restrictions
2011	For	0.015	No	Ad subverting gender stereotypes
2018	For	0.013	No	Pulled ads from controversial program
2016	For	0.013	No	Asked customers not to bring their firearms in stores
2017	For	0.010	No	Opposed controversial figure and changes to national monuments
2018	For	0.010	No	Severed ties with military-style weapons makers
2018	For	0.010	No	Pulled ads from controversial program
2018	For	0.010	No	Halted orders from supplier over gun sales
2021	Against	0.009	No	Donations to causes perceived as anti-abortion
2017	For	0.009	No	Removed brand of controversial figure
2017	For	0.009	No	Stopped promoting brand of controversial figure
2018	For	0.009	No	Severed ties with pro-gun group
2019	For	0.009	No	Opposed anti-abortion legislation
2017	Against	0.008	No	Stated intent to continue flying confederate-era flag
2017	For	0.008	No	Severed ties with controversial figure
2018	For	0.006	No	Severed ties with pro-gun group
2017	For	0.006	Yes	Opposed controversial figure and his supporters
2018	For	0.006	No	Ad criticizing policy of controversial figure
2021	Against	0.005	No	Continues ties with controversial figure
2014	For	0.004	No	Severed ties with controversial figure
2018	For	0.004	No	Pulled ads from controversial program
2017	For	0.004	No	Stopped selling controversial family's brand
2018	For	0.004	No	Severed ties with pro-gun group
2021	Against	0.004	Yes	Key stakeholder supported controversial figure
2020	For	0.003	No	Affirmative action hiring initiative
2017	For	0.002	No	Pulled ads from controversial program
2016	For	0.001	No	Pulled ads from controversial platform
2017	For	0.000	No	Opposed immigration restrictions

Note: Tables shows the complete list of social stance events we analyze. For each event, we provide the year, whether the stance was aligned with positions in the For vs. Against donation cluster, our estimate of the share of consumers who were aware of this event ($\hat{\tau}_j$), an indication of whether $\hat{\tau}_j$ was imputed from Google Trends data and news reports because this event was not covered by BrandIndex data, and a brief description of each event. These descriptions are generic, as we are unable to identify the firms included in our analysis under the terms of the agreement with our data provider.

Appendix Table B2: Most Influential Social Alignment Predictors

Direction	Type	Name
Against	Demographic (County)	Republican Vote Share
For	Transaction	*[National Online+Print Newspaper, Based in NY]*
Against	Transaction	*[Pro-Gun Membership Org]*
Against	Transaction	Other Political Organizations
Against	Transaction	Other Religious Organizations
For	Transaction	Other Restaurants
For	Transaction	*[Non-Profit Reproductive Healthcare Org]*
Against	Demographic (County)	Share Commutes by Auto
Against	Transaction	Other Sporting Goods Stores
For	Transaction	Other Theatrical Producers
For	Transaction	*[Non-Profit Human Rights Org]*
Against	Transaction	*[Christian Humanitarian Aid Org]*
For	Transaction	Other Parking Lots,Meters,Garag
For	Transaction	*[National Grocery Chain, Based in CA]*
For	Transaction	Other Book Stores
For	Transaction	Other Local Commuter Transport
For	Transaction	Other Colleges/Univ/Jc/Professi
For	Transaction	Other Motion Picture Theatres
For	Transaction	Other Eating Places And Restaur
For	Transaction	*[Online Marketplace for Handmade Goods]*
Against	Transaction	*[Cable+Media Company]*
Against	Transaction	*[Farm Supplies and Home Improvement Retailer]*
For	Transaction	*[Audio Streaming and Media Company]*
Against	Transaction	*[E-Commerce Platform]*
For	Transaction	Other Bars/Taverns/Lounges/Disc
For	Transaction	*[Food Delivery Platform]*
Against	Transaction	Other Automotive Parts Stores
Against	Transaction	*[Tool and Equipment Retailer]*
Against	Transaction	*[Identity Theft Software Provider]*
For	Transaction	*[National Railroad Company]*
Against	Transaction	*[Hunting and Outdoor-Related Retailer]*
Against	Transaction	*[Christian Catalog and Internet Retailer]*
For	Transaction	*[Newspaper, Magazine, and Media Company]*
Against	Transaction	*[Hunting, Fishing, and Outdoor-Related Retailer]*
Against	Transaction	*[Steakhouse Restaurant Chain]*
For	Transaction	*[Furniture Retailer]*
Against	Transaction	Other Religious Goods Stores
Against	Transaction	*[National Online+Print Newspaper, Based in NY]*
For	Transaction	*[Child-Focused Humanitarian Aid Org]*
Against	Transaction	*[Veterans Service Org]*
Against	Transaction	Other Public Golf Courses
For	Transaction	*[Online Encyclopedia]*
For	Demographic (County)	Share Commutes by Walk
Against	Demographic (Block)	Republican Vote Share
For	Transaction	*[International Humanitarian Medical Care Org]*

Note: Table shows a list of the 45 most influential transaction and demographic predictors of social alignment, as defined by the gain in predictive accuracy from including this predictor in our XGBoost prediction model. Predictors are in decreasing order of influence. For each predictor, we list the univariate direction it would predict for alignment with the For vs. Against donation clusters, its type, and the name of this predictor. We replace the names of individual merchants with generic descriptions (in brackets), as we are unable to identify individual merchants under our data agreement.