

Public Discourse and Socially Responsible Market Behavior*

Björn Bartling

Vanessa Valero

Roberto A. Weber

Lan Yao

July 8, 2024

We investigate the causal impact of public discourse on socially responsible market behavior. Across three laboratory experiments, having market participants engage in public discourse generally increases market social responsibility. These positive impacts are robust to variation in several characteristics of the discourse. We provide evidence that discourse strengthens beliefs that others support socially responsible exchange. However, relaxing requirements to engage in discourse sharply reduces its effectiveness. Our findings suggest that campaigns encouraging discussion of appropriate market behavior can have powerful sizable impacts on addressing inefficiencies due to market failures, but that policies encouraging broad public engagement may be important.

JEL Classification: C92, D62, D83, M14

Keywords: Public discourse, market failure, externalities, social responsibility, experiment, communication

* Bartling: Department of Economics, University of Zurich, bjoern.bartling@econ.uzh.ch; Valero: School of Business, University of Leicester, vanessa.valero@leicester.ac.uk; Weber: Department of Economics, University of Zurich, roberto.weber@econ.uzh.ch; Yao: School of Economics, Shanghai University of Finance and Economics, yao.lan@mail.shufe.edu.cn. We are grateful for support from the Swiss National Science Foundation (grant 100018_165943), the National Natural Science Foundation of China (grant 72394394, 72273077), Major Program of the National Social Science Foundation of China (grant 21&ZD082) and for helpful comments from Ernst Fehr, Florian Schneider, Paul Smeets, and participants at several seminars and conferences. We are also thankful for thoughtful advice from the editors and three anonymous reviewers, whose guidance was instrumental in designing our third study. We thank Xinquan Hu, Flavia Hug, Bartłomiej Kudrzycki and Xun Zhu for valuable research assistance. All studies in this paper comply with the formal requirements for exempt approval through the IRB review procedures of the Ethics Committee of the Faculty of Economics, Business and Informatics of the University of Zurich.

1. Introduction

The production, exchange and consumption of goods often produces negative impacts for the environment and for parties uninvolved in market transactions. Standard prescriptions for such problems are frequently infeasible due to the complexity of determining optimal policy and political inertia. As an alternative remedy, market actors may exhibit social responsibility, voluntarily internalizing externalities generated by their market activities (Bénabou and Tirole, 2010; Kitze and Shimshack, 2012). Public discourse campaigns often promote social responsibility, encouraging market actors to internalize harmful impacts. For example, the Fridays for Future movement and World Economic Forum both support discussions of how consumers and firms can take actions to mitigate climate change.¹ Other large and highly visible public campaigns advocate voluntary avoidance of fur and other animal products, single-use plastics and pesticides.

Whether such discourse impacts market behavior and societal outcomes remains an open question. Identifying the relationship between public discourse campaigns and socially responsible behavior faces the challenge that prominent campaigns are endogenous to the preferences and motivations of members in a society, making a causal interpretation of a campaign's impact on market behavior and outcomes challenging.² Moreover, while several studies in marketing demonstrate that nudges and primes can influence individuals' socially and environmentally responsible behavior (Goldstein, et al., 2008; White, Habib and Hardisty, 2019), it is unclear the extent to which such interventions are likely to arise endogenously through public discourse—where countervailing arguments may also arise—and, thus, the degree to which they influence equilibrium outcomes in competitive markets. Hence, a better understanding of how public discourse can affect socially responsible market behavior and outcomes remains necessary.

We report three laboratory studies, involving 2,457 participants and 187 independent markets, that explore the causal effect of public discourse on socially responsible market behavior.³ Our studies involve stylized product markets in which participants in the roles of buyers and sellers trade goods that vary in social impact. In all studies, sellers and buyers can exchange either a

¹ See <https://fridaysforfuture.org/> and <https://www.weforum.org/agenda/archive/climate-change> (accessed 05/18/23).

² For example, Levy and Mattsson (2023) study the impact of the MeToo movement on sex crime reporting, relying on the assumption that the strength of the movement's adoption in a society is independent of other factors that may impact the reporting of sex crimes. Madestam et al. (2013) show that rainfall-induced variation in participation in political rallies impacts subsequent voting behavior and policy outcomes. For a review of research demonstrating correlational impacts of social movements on political outcomes, see Amenta et al. (2010).

³ All our data are available in the repository Bartling et al. (2024).

harmful product that costs less to produce but creates external harm or a *responsible product* with higher production cost but no harm. The two products are otherwise identical. We measure market social responsibility by the market share of the responsible product type.

Our experimental variations focus on the impact of different forms of public discourse. In all cases, discourse consists of a one-time 8-minute free-form discussion via electronic chat. While existing experimental research—reviewed in the next section—documents that communication generally increases pro-social behavior and efficiency, this earlier work focuses on cases where agreements benefit all communicating parties and yield Pareto improvements. An important contribution of our work is to investigate whether communication yields positive impacts when agreements involve market actors *lowering their own earnings* to reduce external costs *for other parties*, an important feature of much real-world discussion concerning externality mitigation.

In Study 1, in which the external impact falls on passive participants in the role of third parties, our treatments vary who participates in discourse and what the participants know about their roles in the market when they communicate. Our first condition, *Veil*, represents an idealized form of discourse in which everyone—buyers, sellers and third parties—participates without yet knowing their market roles. As all participants are equally likely *ex ante* to be third parties, discourse might be most likely to produce “responsibilizing narratives” that advocate for socially responsible market behavior (Bénabou et al., 2020). This idealized discourse behind a Rawlsian “veil of ignorance” thus represents a potential upper bound for the effects of discourse.

Our second condition, *No Veil*, implements one change by informing participants of their role in the market prior to engaging in discourse. This reflects typical circumstances under which people engage in discourse. We expect more prevalent self-interested perspectives in discourse to result in lower subsequent levels of socially responsible behavior, consistent with other research showing that knowledge of one’s personal interests facilitates egoistic judgments and justifications for self-interested conduct (Babcock et al., 1995; Saccardo and Serra-Garcia, 2021).

Our third condition, *Exclusive*, restricts discourse to only buyers and sellers, who are aware of their roles as in *No Veil*, omitting those harmed by the externality. Public discourse often occurs among members of high-income countries, including firms and consumers whose conduct creates negative externalities on third parties absent from discourse. This condition allows us to study the importance of participation in discourse by impacted stakeholders, whose exclusion may decrease empathy toward the impacts they experience (Andreoni and Rao, 2011) and more easily yield

“absolving narratives” or self-serving justifications (Bénabou et al., 2020). We thus expect the exclusion from discourse of those harmed by the externality to further reduce social responsibility.

Finally, to explore how public discourse interacts with existing concerns for social impact, we study two populations in which earlier work found different levels of baseline market social responsibility. In Bartling et al. (2015), market experiments in Switzerland yielded substantially higher market shares for the responsible product than identical experiments in China. Public discourse may *complement* baseline social concern if a high share of concerned individuals is needed to promote persuasive arguments for further reducing external social harm. Alternatively, public discourse may act as a *substitute* for baseline levels of social responsibility if existing high levels limit opportunities for further increasing concerns for social impact.

We generally observe large, positive treatment effects in Study 1. First, public discourse in the *Veil* condition has sustained positive impact on market social responsibility. In Switzerland, discourse yields almost universal exchange of the responsible product, compared to a *Baseline* market share of about 50 percent without discourse. Contrary to our expectations, discourse generally retains its overall positive impact when eliminating the veil of ignorance and excluding third parties from discourse. Thus, we find that public discourse regarding appropriate market behavior can have profound and persistent positive impacts on market social responsibility.

Turning to the interaction between pre-existing levels of social concern and public discourse, we first replicate the observation of a lower baseline market share for the responsible product in China. However, introducing public discourse has comparable (in absolute terms) positive effects on market social responsibility in the two populations, suggesting that the effects of public discourse are largely independent of baseline levels of concern for social responsibility.

Study 2 introduces two additional potentially important features. First, we investigate whether discourse retains positive effects when the external impacts of market exchange are not borne by other laboratory participants. In laboratory settings, high levels of social proximity and norms of equality may make arguments in favor of mitigating impacts on other participants quite strong. Outside the laboratory, negative externalities typically involve more complex and distant impacts. We thus change the target of the externality to a charity working to mitigate climate change and economic inequality. Using this kind of external impact potentially facilitates arguments during discourse for why the charity may be ineffective or undeserving (Exley, 2020).

Our first condition, *Discourse*—similar to *Exclusive* from Study 1 except for this change in the design—replicates the positive impacts of discourse on market social responsibility.

The second change in Study 2 deals with the timing of discourse. In our other conditions, discourse occurs at the beginning of the experiment, before participants engage in market exchange. However, many situations of interest, like combating climate change, involve changing pre-existing market behaviors. It is plausible that increasing social responsibility may be more difficult in such contexts, where history-dependence may create inertia and a desire to advance arguments supporting the status quo or “business as usual.” Thus, in our *Experienced* condition, discourse occurs only after several rounds of market activity. However, we again find a positive impact of discourse regardless of whether it occurs after market experience.

Study 3 addresses several remaining open questions. First, in Studies 1 and 2 we prompt participants to discuss “appropriate” market conduct, reflecting real-world forums and campaigns focused on topics of social responsibility. But this potentially creates an unintended experimenter demand effect regarding what impact discourse should have. Study 3 eliminates such guidance and simply asks participants to “discuss the upcoming market activity.” We find that this neutrally framed *Discourse (Neutral)* condition again yields higher social responsibility than a *Baseline* condition without discourse.

Second, we investigate the impacts of varying forms of engagement in discourse. In most real-world public discourse, individuals choose whether to participate or can consume discourse generated by others as passive observers. We conduct an *Optional* condition, in which participants decide whether to engage in discourse, and a *Passive* condition, in which market actors only observe discourse generated by another group of subjects. The *Passive* condition produces positive impacts on market social responsibility, but slightly weaker than in *Discourse (Neutral)*. Most strikingly, the *Optional* condition yields no increase relative to the *Baseline*, despite the fact that the vast majority of participants enter the discourse forum.

Study 3 also includes questionnaire-based measures of values and expectations to help shed light on potential drivers of the impact of discourse. We find that public discourse strongly impacts the belief that others value and will act on concerns for social responsibility, though this impact is weakest in *Optional*, where we also observe no treatment effect on the share of socially responsible products. We interpret this as suggestive evidence that discourse may impact behavior partly by changing expectations about others’ values and behavior. This interpretation is further supported

by exploratory analysis of the content of discourse, which suggests that exposure to *others* advocating for responsible conduct influences participants' expectations of how others will behave and leads to more responsible market conduct. We also document that the single condition that yields no positive impact of discourse—*Optional* in Study 3—is the one in which participants receive the lowest exposure to others advocating for social responsibility.

Broadly, our work provides insights regarding the potential impact of public discourse on market behavior. We find that a one-time experience with discourse generally yields persistent positive impacts on socially responsible product market shares, even as we vary several features of our design to make such positive impacts less likely. We also provide evidence that discourse impacts beliefs about others' values and behaviors, as well as social norms governing market conduct. However, making discourse optional sharply reduces its positive impact, suggesting that effective discourse campaigns may be those that make widespread engagement mandatory.

2. Related Literature

Our research closely relates to a growing body of work that investigates the conditions under which individuals choose to *voluntarily* internalize the external impacts of their market activity.⁴ Much of this research uses laboratory experiments, allowing tight control of the market environment and the establishment of causal relationships. While this approach naturally raises questions about external validity, Engelmann et al. (2018) show a relationship between socially responsible laboratory behavior and socially responsible behavior in real-world product markets.

Related to our work, numerous studies demonstrate that communication can yield more pro-social and cooperative behavior, increasing efficiency, in other domains of social behavior, including social dilemmas (Dawes et al., 1977; Isaac and Walker, 1988; Ostrom et al., 1992; Bochet et al., 2006), coordination games (Cooper et al., 1992; Blume and Ortmann, 2007; Brandts and Cooper, 2007; Kriss et al., 2016) and incomplete contracting (Charness and Dufwenberg, 2006; Kessler and Leider, 2012; Krupka et al., 2017). Relatedly, appeals to act pro-socially can enhance cooperation in public good games (Dal Bo and Dal Bo, 2014; Antonakis et al., 2021).⁵

⁴ See, e.g., Rode et al. (2008), Bénabou and Tirole (2010), Dufwenberg et al. (2011, 2022), Bartling et al. (2015, 2019, 2022), Hainmueller et al. (2015), Kirchler et al. (2015), Pigors and Rockenbach (2016), Irlenbusch and Saxler (2019), Ockenfels et al. (2020), Sutter et al. (2020) and Danz et al. (2022).

⁵ Communication may not produce Pareto improvements in contexts when parties have strategic incentives to manipulate outcomes (see, e.g., Sutter and Strassmair, 2009; Hargreaves-Heap and Zizzo, 2011; Biais, Bisière and

This work suggests that discourse may similarly yield more efficient outcomes, with fewer externality-producing products, in our market experiments. However, an important distinction is that in our study discourse only improves efficiency and social responsibility if it convinces buyers and sellers *to incur greater personal monetary costs to benefit others*. Thus, our study is novel in investigating whether communication leads to more pro-social, but personally costly, actions.⁶

A separate line of research studies how communication influences distributional outcomes in dictator and bargaining games. For example, in a Baron-Ferejohn bargaining context, where one individual proposes how to divide a fixed pie and the proposal is implemented if approved by a majority of committee members, communication produces proposals closer to the theoretical prediction of minimum-winning coalitions that only reward coalition members (Agranov and Tergiman, 2014; Baranski and Morton, 2022). This suggests that communication can be employed to produce more favorable outcomes for a few, at the expense of others with less strategic power, which contrasts with our finding that discourse advantages third parties.

In dictator games, the presence and nature of communication can impact one-sided sharing. Allowing recipients to communicate with dictators increases giving; but allowing only dictators to send messages yields less sharing (Bohnet and Frey, 1999; Andreoni and Rao, 2011). In our experiments, this suggests that excluding those impacted by externalities from discourse may decrease social responsibility, perhaps even to levels below those in baseline conditions without discourse. While we find that the positive impacts of discourse sometimes decrease when excluding impacted third parties, socially responsible behavior is more prevalent, relative to the no-discourse baseline, even in such cases.

3. Study 1

We implement an experimental product market in which subjects in the roles of sellers and buyers can trade either a low-cost product that generates a negative externality for subjects in the role of

Pouget, 2014). Communication may facilitate dishonesty and less concern for external impacts by groups (see, e.g., d'Adda et al., 2017; Kocher, Schudy and Spantig, 2018).

⁶ Ellman and Pezanis-Christou (2010) also examined if communication promotes responsible behavior, even when it costs the communicating parties. They focus on two-person laboratory firms, where profitable production harms a third party. Communication has a positive impact on responsible behavior only in a vertical hierarchical structure (this treatment effect is statistically significant at $p=0.066$ in a one-tailed non-parametric test), but it has no impact in a horizontal structure. We interpret this as suggestive evidence that communication can increase pro-social concerns toward outsiders, but it does not address whether such effects extend to market contexts, as Bartling et al. (2015) find that pro-social behavior in market contexts differs from substantively similar non-market conduct.

third parties or a high-cost product that imposes no social harm. Our main interest is the market share of the latter product and how this is affected by the presence and nature of public discourse.

3.1 Experimental Design

3.1.1 *The Market*

A market comprises 16 participants: six sellers, five buyers and five third parties. Roles are randomly assigned and remain fixed throughout the experiment.

The market activity repeats for 24 periods. At the beginning of a period, each subject receives an endowment of 100 points. Sellers and buyers can earn additional points from their market transactions. There are two product types that differ only in their negative impact on the third parties and their production cost. The *responsible* product does not produce a negative externality but costs 10 to produce, while the *harmful* product creates a negative externality of 60 for third parties but costs nothing to produce. Both products have a value of 50 to a buyer.

In each period, sellers simultaneously select product types and prices (between 0 and 50) in a posted-offer market. Buyers then enter the market sequentially in a randomly determined order and observe a menu of up to six product offers, each consisting of a price and a product type for all product offers remaining unpurchased. We eliminate cross-period reputation by not showing identification numbers for other market participants and randomly ordering product offers in each period. Each buyer decides whether to purchase at most one product. Each seller whose product is sold earns the difference between the posted price and the production cost (a seller can only sell a single unit). Each buyer who decides to buy a product earns the difference between the product's value of 50 and the price paid. Sellers observe all product types and prices offered and sold.

Third parties are passive. They can neither sell nor buy but can incur losses depending on the types of products exchanged. In every period, third parties and buyers are randomly matched into pairs. If a buyer purchases a harmful product, then the randomly matched third party incurs a loss of 60. If the buyer purchases a responsible product or does not purchase any product, then the randomly matched third party incurs no loss.

3.1.2 *Treatments*

In our *Baseline* condition there is no discourse. Subjects directly proceed to the market game after learning their roles as sellers, buyers or third parties.

We implement three treatment conditions that introduce public discourse, a one-time 8-minute interval during which subjects can exchange messages via an electronic chat. The chat interface provides guidance for the public discourse by asking subjects to discuss how “socially appropriate” it is to trade the harmful product.⁷ The messages are not restricted, other than proscribing personally identifying, obscene or insulting statements.

Table 1: Overview of Experimental Conditions in Study 1

	<i>Baseline</i>	<i>Veil</i>	<i>No Veil</i>	<i>Exclusive</i>
$t=1$	–	Public discourse (all subjects)	–	–
$t=2$	All subjects learn their roles: seller, buyer or third party			
$t=3$	–	–	Public discourse (all subjects)	Public discourse (excl. third parties)
$t=4$	Subjects participate in the market game for 24 periods			

Table 1 presents our treatments. In the *Veil* condition, subjects engage in the one-time discourse period prior to learning their roles as sellers, buyers or third parties. That is, discussion in this condition takes place behind a Rawlsian “veil of ignorance,” though subjects learn their roles prior to the market activity. In *No Veil*, subjects learn their roles prior to the discourse period. Apart from the timing of information about subjects’ roles, *Veil* and *No Veil* are identical.⁸ Finally, in the *Exclusive* condition, subjects learn their roles prior to engaging in discourse, as in *No Veil*, but participation in discourse is limited to only sellers and buyers. In *Exclusive*, we also give third parties the possibility to discuss among themselves, to engage all subjects during the discourse period, though such discussion cannot directly affect market outcomes.

3.1.3 Social Norm Elicitation

In all conditions, we elicit social norms after the final market period, using the elicitation method of Krupka and Weber (2013). We elicit participants’ ratings of the social appropriateness of trading the harmful product. Subjects can choose from four possible responses: “very socially appropriate,” “somewhat socially appropriate,” “somewhat socially inappropriate” and “very

⁷ This statement aims to focus discourse on the relevant topic and reflects similar focusing in real-world campaigns. Study 3 eliminates the statement to address concerns that it may produce demand effects. We address the role of demand effects broadly in Section 7.

⁸ Subjects can refer to each other during the chat through fixed ID numbers. Messages are preceded by identifiers ranging from 1-16 in *Veil* and by letter-number combinations (e.g., S1 through S6 for sellers) that also identify roles in *No Veil*. However, subjects cannot subsequently match messages to individual market behavior.

socially inappropriate,” to describe such behavior. Responses are incentivized: each participant earns additional money (CHF 10 in Switzerland or CNY 10 in China) if that subject’s response corresponds to the most frequently chosen answer provided by the other subjects in a session.

3.1.4 Procedures

We conducted the study at the University of Zurich and the Shanghai University of Finance and Economics (between October 2015 and March 2017). We collected eight markets per treatment, both in Switzerland and in China; thus, a total of 1,024 subjects participated in the above four experimental conditions, half of them in Switzerland and half of them in China. We implemented the experiment using z-Tree (Fischbacher, 2007). Subjects received written instructions and had to answer control questions to ensure understanding of the market activity. We used German and Mandarin instructions based on those used in Switzerland and China for Bartling, et al. (2015). An English version of the instructions for Study 1 is in Appendix F.

Experimental points were converted into money at a rate of 10 points to CHF 2.50 in Switzerland and CNY 4 in China. The conversion rate aimed to match purchasing power across countries. We selected one period at random for payment at the end of a session. Subjects in Switzerland earned about CHF 49, on average, including a show-up fee of CHF 15; in China they earned about CNY 62, including a show-up fee of CNY 15. Sessions lasted about two hours.

3.2 Hypotheses

We assume that the market share of the responsible product (our measure of social responsibility), denoted by R , depends on the pre-existing level of social concern in a society, θ_j (with $j = \{\text{Switzerland, China}\}$ in our study), and on public discourse, where d_t captures the type of public discourse ($t = \{\emptyset, \text{Veil}, \text{No Veil}, \text{Exclusive}\}$). $R(\theta_j, d_\emptyset)$ thus denotes the market share of the responsible product in society j , in the absence of public discourse, as in our *Baseline* condition.

We first consider the impact of *Veil*, where we expect an idealized form of discourse—conducted behind a Rawlsian veil of ignorance—to have the best chance of increasing R . Placing all participants in a position where they might be harmed by negative externalities is likely to promote arguments against imposing such externalities, thereby increasing the level of social concern. We thus predict a positive effect of discourse on R in *Veil*, relative to *Baseline*.

$$\text{H1.1: } R(\theta_j, d_{\text{veil}}) > R(\theta_j, d_\emptyset)$$

Our next hypotheses deal with the impact of discourse on R as we add realistic features. First, when buyers and sellers know their roles—as beneficiaries of low-cost production—we expect them to be more inclined to generate self-serving arguments exculpating exchanging the harmful product and less easily swayed by appeals to social responsibility (Babcock, et al., 1995).

$$H1.2: R(\theta_j, d_{Veil}) > R(\theta_j, d_{No\ Veil})$$

Second, we expect that excluding third parties will further diminish discourse’s impacts on concern for third parties’ welfare (Andreoni and Rao, 2011).

$$H1.3: R(\theta_j, d_{No\ Veil}) > R(\theta_j, d_{Exclusive})$$

Note that we do not provide hypotheses for comparisons of R between *Baseline* and either *No Veil* or *Exclusive*, because the results could go either way. Even in *No Veil*, if enough participants care about efficiency or the welfare of third parties, they may promote socially responsible behavior through discourse and increase the market share of the responsible product. Alternatively, buyers and sellers may instead generate exculpating arguments, potentially reducing responsible behavior even relative to the *Baseline*. Moreover, even if the market share of the responsible product is higher in *No Veil* than in *Baseline*, whether it remains higher in *Exclusive*, once third parties are excluded from discourse, is not clear *a priori*. While we refrain from stating directional hypotheses, how the market share of the responsible product in *Baseline* compares with those in both *No Veil* and *Exclusive* are critical exploratory research questions.

Finally, a comparison of treatment effects in Switzerland and China allows us to investigate how pre-existing levels of social concern interact with public discourse. Bartling et al. (2015) found substantially lower market shares of the responsible product in the *Baseline* in China than in Switzerland, which we expect to replicate. We can then investigate, as an exploratory question, whether pre-existing social concern is a *complement* to or *substitute* for public discourse.

3.3 Results


3.3.1 Public Discourse Behind the Veil of Ignorance

Figure 1 shows the market shares of the responsible product, separately for Switzerland and China.⁹ Throughout the paper, we exclude the cases in which buyers made no purchase, thereby

⁹ Appendix Figure A.1 shows the market shares of the responsible product across periods. Figure A.2 shows the cumulative distributions of responsible product market shares by treatment in both countries.

imposing no externality. There are generally few such cases (1.8 percent in Switzerland and 3.5 percent in China in Study 1) and including them as responsible behavior does not substantively change our results.

Figure 1: Market Shares of the Responsible Product in Study 1



Notes: The figure considers completed transactions, ignoring the small number of cases in which a buyer did not purchase a product. The bars indicate 95-percent confidence intervals, calculated at the market level.

Comparing the two leftmost bars—*Baseline* and *Veil*—for each population reveals that social responsibility increases substantially in both countries when market actors learn their roles after engaging in discourse. The market share of the responsible product in Switzerland is about 50 percent in the *Baseline* condition but rises to almost 100 percent in *Veil*. In China, the *Baseline* market share of the responsible product is only 15 percent, much lower than in Switzerland,¹⁰ but the market share almost quadruples to about 60 percent in *Veil*. Wilcoxon rank-sum tests at the market (buyer) level indicate that the difference between *Baseline* and *Veil* is statistically significant in both countries; $p=0.001$ ($p<0.001$) for Switzerland and $p=0.005$ ($p<0.001$) for China.¹¹ The data thus support Hypothesis 1.1 in both countries.

¹⁰ The *Baseline* condition in this paper is identical to the *Baseline* in Bartling et al. (2015), which also studied the same two populations. The market shares of the responsible product in our *Baseline* conditions closely replicate those reported in Bartling et al. (2015), both in Switzerland (44 and 48 percent across two studies) and China (16 percent).

¹¹ Appendix Table A.1 provides p-values for all pairwise treatment comparisons. All Wilcoxon rank-sum tests reported in this paper are two-sided and use either the subject or market as the unit of observation.

Result 1.1: *Public discourse behind the veil of ignorance increases socially responsible market behavior.*

3.3.2 Alternative Forms of Public Discourse

Figure 1 shows that removing the veil of ignorance prior to discourse reduces the market share of the responsible product, from 96 to 87 percent in Switzerland and from 59 to 49 percent in China. Wilcoxon rank-sum tests indicate that the differences between *Veil* and *No Veil* are statistically significant at the market (buyer) level in Switzerland, $p=0.014$ ($p=0.001$), but not in China, $p=0.248$ ($p=0.145$). Hypothesis 1.2 is thus supported in Switzerland but not in China.

Result 1.2: *The positive impact of public discourse on socially responsible market behavior is slightly weaker when individuals discuss in front of rather than behind a veil of ignorance.*

Our third discourse condition, *Exclusive*, restricts participation in discourse to buyers and sellers. Figure 1 shows that the market share of the responsible product slightly increases from *No Veil* to *Exclusive* in Switzerland (from 87 to 92 percent) but decreases in China (from 49 to 37 percent). Wilcoxon rank-sum tests at the market (buyer) level indicate that the differences between *No Veil* and *Exclusive* are not consistently statistically significant; $p=0.140$ ($p=0.049$) for Switzerland and $p=0.293$ ($p=0.055$) for China. Thus, Hypothesis 1.3 is not supported in Switzerland and is directionally, but not statistically, supported in China.

Result 1.3: *Excluding those harmed by externalities from public discourse does not substantially weaken its positive impact on socially responsible market behavior.*

While we refrained from stating hypotheses regarding the market share of the responsible product in *No Veil* and *Exclusive* relative to the *Baseline*, Wilcoxon rank-sum tests indicate that it is significantly higher in *No Veil* and *Exclusive* than in *Baseline*, both in Switzerland and China, irrespective of whether we test at the market or buyer level ($p \leq 0.001$ in all tests in Switzerland; $p \leq 0.027$ in all tests in China; see Appendix Table A.1).¹²

¹² A standard interpretation of treatment effects from discourse opportunities is that such effects are due to the exchange of messages. However, such treatments also give subjects time to *reflect* on the relevant topic. In our experiment, even if subjects did not actively communicate, the 8-minute discourse interval might also affect behavior by prompting them to *reflect* about appropriate behavior (cf. Krupka and Weber, 2009). To shed light on the distinction between communication per se and reflection, we conducted a *post hoc* treatment in which subjects compose private statements regarding appropriate market behavior, not shared with others, for 8 minutes. This *Reflection* condition increases market social responsibility relative to *Baseline*. However, the effect of actual discourse extends beyond those of reflection in Switzerland, but not in China. This suggests a role for public campaigns that ask people to reflect on their behavior. Given its *post hoc* nature, we report the details of this *Reflection* condition in Appendix B.

Observation 1.1: *Public discourse increases socially responsible market behavior even when individuals discuss in front of the veil of ignorance and even when negatively affected third parties are excluded from the discourse.*

As a complement to the above non-parametric tests, Table 2 reports random-effects GLS regressions with buyers' product choices—i.e., whether a buyer purchases a responsible product in a period—as the dependent variable.¹³ We include binary variables, *Veil*, *No Veil* and *Exclusive*, which take on value 1 in the respective condition and 0 otherwise, measuring these treatment effects relative to the omitted *Baseline*. In models 2 and 4 we control for time effects by including the variable *Period*, taking on integer values between 1 and 24, and its interactions with the treatment variables. Models 1 and 2 present results for Switzerland and models 3 and 4 for China.

Table 2: GLS (random-effects) regressions of responsible buyer product choice

	Switzerland		China	
	(1)	(2)	(3)	(4)
<i>Veil</i>	0.457*** (0.057)	0.416*** (0.067)	0.446*** (0.100)	0.450*** (0.111)
<i>No Veil</i>	0.371*** (0.055)	0.436*** (0.067)	0.342*** (0.079)	0.419*** (0.102)
<i>Exclusive</i>	0.424*** (0.061)	0.382*** (0.077)	0.244** (0.093)	0.151 (0.121)
<i>Period</i>		-0.003 (0.002)		-0.004** (0.002)
<i>Period × Veil</i>		0.003 (0.003)		-0.000 (0.005)
<i>Period × No Veil</i>		-0.005 (0.004)		-0.006 (0.005)
<i>Period × Exclusive</i>		0.003 (0.003)		0.007** (0.004)
<i>Constant</i>	0.494*** (0.051)	0.533*** (0.061)	0.149** (0.067)	0.203** (0.084)
Observations	3770	3770	3705	3705
Subjects	160	160	160	160
R ²	0.371	0.371	0.247	0.249

Notes. The dependent variable in all models takes on value 1(0) if a buyer purchased a responsible (harmful) product. We omit 70 cases in Switzerland and 135 cases in China in which a buyer purchased no product. *Baseline* serves as the omitted category. *Period* takes on integer values between 1 and 24. Standard errors (in parentheses) are clustered at the market level, *** p<0.01, ** p<0.05, * p<0.1.

¹³ Probit and OLS regressions, reported in Appendix Tables A.2 and A.3, confirm the results shown in Table 2, including when we treat each market participant as a single observation.

Looking first at Switzerland, the coefficients for *Veil*, *No Veil* and *Exclusive* are large, positive, and statistically significant at the one-percent level in both models, indicating positive impacts of all discourse treatments, consistent with Observation 1.1. For China, model 3 similarly finds the effects of discourse to be positive and statistically significant, at least at the five-percent level. In model 4, the interaction *Period* \times *Exclusive* is positive and significant at the five-percent level, while the coefficient for *Exclusive* is not statistically significant, indicating that the effect of discourse develops over time in this condition in China. The results in Table 2 confirm our earlier observations from non-parametric Wilcoxon rank-sum tests.¹⁴

3.3.3 *Public Discourse and Pre-Existing Social Concern*

The substantially different *Baseline* levels of social responsibility in Switzerland and China allow us to investigate the relationship between public discourse and pre-existing levels of social concern. To this aim, we run regression models as in Table 2, but with the data pooled from both countries and with a country identifier for China and interaction terms of this variable with the separate treatments. None of the interaction terms is statistically significant, suggesting that the absolute increase in the market share of the responsible product does not differ between Switzerland and China. Similar results obtain in a specification where we interact the country identifier with a binary variable that takes on value 1 in any of the three discourse conditions. We report the regressions in Appendix Table A.5 and summarize these observations as follows:

Observation 1.2: *The positive effect of public discourse on socially responsible market behavior is independent of the pre-existing level of social concern.*

3.3.4 *Prices and profits*

We find that the higher production cost of the responsible product translates into higher market prices for these products. Across conditions, responsible products trade at a price of about 28 and harmful products at a price of about 22.¹⁵ Buyers and sellers who trade the responsible product thus share the additional production cost of 10. This implies that buyers forgo monetary payoffs when buying a responsible product and sellers obtain lower profits when offering these products

¹⁴ The coefficients for *Veil* confirm Result 1.1 for both countries. Post-estimation tests of equality of the coefficients for *Veil* and *No Veil* in models 1 and 3 fail to reject equality in China but not in Switzerland (see p-values from Wald tests in Appendix Table A.4), reflecting the mixed finding in Result 1.2. Post-estimation tests of equality of the coefficients for *No Veil* and *Exclusive* fail to reject equality in both countries, which confirms Result 1.3.

¹⁵ Appendix Figures A.3 and A.4 show average prices and sellers' average profits, respectively, in all conditions and both countries, separately for both types of products.


(sellers' sales probabilities do not vary substantially with the type of product offered).

While buyers and sellers are willing to pay higher prices and forgo profits, respectively, we also observe that market participants react to prices and expected profits. Appendix Table A.6 shows that buyers are more likely to purchase the responsible product if the price of the cheapest available responsible (harmful) product is lower (higher). Likewise, Table A.7 shows that sellers are more likely to offer the responsible product as the expected profit from offering the responsible product rather than the harmful product, based on prior rounds' profits, increases.

3.3.5 Social norms

At the end of each session, subjects rated the perceived social appropriateness of trading harmful products. Figure 2 shows mean appropriateness ratings in Switzerland and China. Harmful market behavior is generally perceived as more appropriate in China than in Switzerland. In addition, ratings of the appropriateness of harmful market behavior are lower in all discourse conditions than in *Baseline*, both in Switzerland and China, indicating that public discourse and subsequent market experience yield stronger norms against socially harmful market behavior.¹⁶

Figure 2: Effect of Public Discourse on Social Norms in Study 1



Notes. The figure shows the average rating of the appropriateness of exchanging the harmful product. “Very socially appropriate = 1,” “Somewhat socially appropriate = 1/3,” “Somewhat socially inappropriate = -1/3,” “Very socially inappropriate = -1.” The numerical rating values follow Krupka and Weber (2013). The bars indicate 95-percent confidence intervals, calculated at the market level.

¹⁶ Regression analysis in Appendix Table A.8 confirms the statistical significance of the treatment effects in Figure 2.

4. Study 2

Study 1 finds that various forms of public discourse increase socially responsible market behavior. However, it is important to consider whether particular features of our experiment may exaggerate this positive impact. For instance, an external effect that falls on another experimental subject, who differs from market actors only in a random draw of roles, makes earnings inequalities salient and clearly resolvable in the laboratory. Outside the laboratory, market externalities typically involve more distant impacts, not as easily quantifiable or remediable. In such contexts, discourse about the “right” thing to do might be more nuanced, facilitating exculpatory arguments. Study 2 thus changes the external impact to affect, rather than another experimental subject, a charity committed to fighting climate change and poverty. We additionally increased the production cost, to test the robustness of the effects of discourse to higher costs of producing the responsible product.

Moreover, public discourse about appropriate market behavior outside the laboratory typically occurs after people have experience engaging in externality-producing behaviors. Changing such established behavior might be particularly challenging, thus mitigating the impact that discourse can have on social responsibility, or perhaps even leading to rationalizations that further depress social responsibility. In Study 2, we thus also examine whether public discourse has a positive effect when it is introduced only after market participants trade for several periods.

4.1 Experimental Design and Procedures

A market in Study 2 comprises 11 subjects: six sellers and five buyers. At the beginning of every period, each participant receives an endowment of 100 points. In addition, we allocate to each market a 500-point donation (corresponding to 100 points for each buyer) to the charitable organization *Carbon Offsets to Alleviate Poverty* (COTAP), which funds programs fighting climate change and poverty in low-income countries (see <https://cotap.org/>).

Subjects can exchange two types of products, which differ in their impact on the charity. There is no reduction of the donation if a seller and a buyer exchange the responsible product (or no product at all), but exchanging a product with a negative external impact reduces the donation by 60. The harmful product costs nothing to produce, but the responsible product’s production cost is 20. Both types of products have identical value of 50 to a buyer. Except for the above differences, the design is otherwise the same as in Study 1 (see Section 3.1).

We implement three treatment conditions. In *Baseline*, as in Study 1, subjects proceed to

24 periods of the market game after learning their role as sellers or buyers. In *Discourse*, the market is preceded by eight minutes of public discourse involving sellers and buyers (third parties are not present in the lab), as in the *Exclusive* condition of Study 1. The third condition, *Experienced*, is identical to *Discourse*, except that sellers and buyers enter discourse only after eight periods of the market game (“Part I”). After discourse, market interaction continues for 16 additional periods (“Part II”). Table 3 provides an overview of the sequence of events in all conditions in Study 2. English translations of the original German instructions are provided in Appendix G.

Table 3: Overview of Experimental Conditions in Study 2

	<i>Baseline</i>	<i>Discourse</i>	<i>Experienced</i>
$t=1$	–	Public discourse	–
$t=2$			8 periods of market game (“Part I”)
$t=3$	24 periods of market game		Public discourse
$t=4$			16 periods of market game (“Part II”)

We conducted the study at the University of Zurich between May 2019 and June 2019, following the same procedures as in Study 1. We collected 16 markets per treatment, involving a total of 528 subjects. In all conditions, we elicit social norms at the end of the market, using the same elicitation method as in Study 1. On average, subjects earned about CHF 51, including a show-up fee of CHF 15. Sessions lasted about two hours.

4.2 Hypotheses

Based on the results of Study 1, we expect public discourse to have a positive effect on the market share of the responsible product, even when those impacted by the externality do not participate.

$$H2.1: R(\theta_j, d_{Discourse}) > R(\theta_j, d_{Baseline})$$

However, the *Discourse* versus *Baseline* comparison in Study 2 is not a direct replication of Study 1 because the distinct impact of the externality on a charity, rather than on other experimental subjects, may facilitate arguments against socially responsible market behavior, thus possibly yielding no positive impact, or even a negative impact, of public discourse on R .

As we discuss earlier, we expect any positive impact of discourse on R to be smaller when participants have a history of engaging in exchange that yields negative externalities.

$$H2.2: R(\theta_j, d_{Discourse}) > R(\theta_j, d_{Experienced})$$

Whether *Experienced* ends up at a higher, equal or lower market share of the responsible product than *Baseline* is a key research question for which we have no *a priori* predictions.


4.3 Results

4.3.1. Market Shares of the Responsible Product

Figure 3 shows the market shares of the responsible product over time.¹⁷ Across all periods, the market share is 50 percent in *Baseline*. The share of responsible products rises to 79 percent in *Discourse*, when discourse occurs before market interaction. Wilcoxon rank-sum tests at the market (buyer) level indicate that the difference in market shares between *Baseline* and *Discourse* is significant ($p=0.002$ ($p<0.001$)). This supports Hypothesis 2.1 and qualitatively replicates the impact of the *Exclusive* condition in Study 1, despite varying the nature of the external impact and increasing the cost of avoiding it.

Result 2.1: *Public discourse increases socially responsible market behavior even when the negative impact of the externality falls on a charity absent from discourse.*

Figure 3: Market Shares of the Responsible Product in Study 2



Notes: The figure considers completed transactions, ignoring the small number of cases in which a buyer did not purchase a product. Shaded areas indicate 95-percent confidence intervals, calculated at the market level.

¹⁷ Appendix Figure C.1 shows the cumulative distributions of responsible product market shares by treatment.

We next test whether public discourse has less positive impact when participants have prior market experience. Figure 3 shows that the market shares of the responsible product in Part II in *Discourse* and *Experienced*, which differ only in whether discourse took place either prior to or after Part I, are almost identical (76 vs. 79 percent; not significantly different in Wilcoxon rank-sum tests at the market (buyer) level; $p=0.568$ (0.357)). Hypothesis 2.2 is thus not supported.

Result 2.2: *Public discourse has equally strong positive impacts for experienced and inexperienced market participants on socially responsible market behavior.*

Moreover, while we posited no hypothesis regarding this comparison, the market share of the responsible product in Part II is higher in *Experienced* than in *Baseline* ($p \leq 0.001$ in Wilcoxon rank-sum tests at the market or buyer level).¹⁸

Observation 2.1: *Public discourse increases socially responsible market behavior even with experienced market participants.*

Random-effects GLS regressions in Table 4 confirm the above results.¹⁹ The dependent variable is a buyers' choice of the responsible product, and the binary variables *Discourse* and *Experienced* take on values of 1 in the corresponding conditions. In model 2 we control for time trends. The large and positive coefficients for *Discourse* indicate that, relative to *Baseline*, socially responsible behavior is higher in Part I following public discourse (Result 2.1). The coefficient for *Experienced* captures the higher frequency of responsible products exchanged in Part I in this condition, relative to the *Baseline*, that we observe despite random assignment to treatment. Taking this into account, the positive and highly significant coefficient for *Part II* \times *Experienced* indicates that the market share of the responsible product increases substantially following discourse, even after market experience (Observation 2.1). To evaluate the relative impacts of discourse with and without experience, we compare the impact of *Discourse* in Part II (sum of *Discourse* and *Part II* \times *Discourse*, which is 0.255 in model 1) with *Part II* \times *Experienced*. While the relative sizes of these effects suggests that discourse has slightly lower impact following market experience, this difference is not statistically significant (Wald test: $p=0.271$), consistent with Result 2.2.

¹⁸ Despite random assignment to treatment and identical instructions and procedures through Part I, the responsible product market share is higher in *Experienced* (63 percent) than in *Baseline* (50 percent) in Part I. Wilcoxon rank-sum tests at the market (buyer) level show that this unexpected difference is at least marginally significant; $p=0.083$ (0.036). Nevertheless, comparing the market shares of the responsible product in *Experienced* between Parts I and II confirms the positive impact of discourse; Wilcoxon rank-sum tests at the market (buyer) level, $p=0.041$ (0.001).

¹⁹ Probit and OLS regressions, reported in Appendix Tables C.1 and C.2, corroborate the results shown in Table 4.

Table 4: GLS (random-effects) regressions of responsible buyer product choice

	(1)	(2)
<i>Discourse</i>	0.326*** (0.073)	0.362*** (0.073)
<i>Experienced</i>	0.126* (0.064)	0.176*** (0.063)
<i>Part II</i>	-0.004 (0.025)	-0.038** (0.017)
<i>Part II</i> \times <i>Discourse</i>	-0.071* (0.039)	0.025 (0.037)
<i>Part II</i> \times <i>Experienced</i>	0.159*** (0.055)	0.291*** (0.053)
<i>Period</i>		0.003 (0.002)
<i>Period</i> \times <i>Discourse</i>		-0.008*** (0.002)
<i>Period</i> \times <i>Experienced</i>		-0.011*** (0.003)
<i>Constant</i>	0.509*** (0.051)	0.497*** (0.052)
Observations	5,619	5,619
Number of subjects	240	240
R ²	0.121	0.121

Notes: The dependent variable in all models takes on value 1 (0) if a buyer purchased a responsible (harmful) product. We omit 141 cases in which a buyer did not purchase a product. *Baseline* and *Part I* (periods 1 to 8) serve as omitted categories. *Part II* is a binary variable taking on value 1 for data from period 9 to 24 and 0 otherwise. *Period* takes on integer values between 1 and 24. Standard errors (in parentheses) are clustered at the market level; *** p<0.01, ** p<0.05, * p<0.1

4.3.2. Prices and profits

On average, the responsible product trades at higher prices (33) than the harmful one (23), yielding a price premium smaller than the marginal production cost of 20 for responsible products. Thus, as in Study 1, buyers and sellers share the burden of exchanging the responsible product.²⁰

²⁰ Appendix Figures C.2 and C.3 show average prices and sellers' average profits, respectively, in all conditions, separately for Parts I and II and for both product types. Tables C.3 and C.4 show that buyers respond to relative prices and sellers to expected profits, mirroring the analysis reported in Tables A.6 and A.7 for Study 1.

4.3.3. *Social norms*

Finally, we observe an effect of discourse and resulting market experience on social norms. Using a scale from “very socially inappropriate” (-1) to “very socially appropriate” (1), exchanging the harmful product is perceived to be substantially less socially appropriate in both the *Discourse* (-0.43) and *Experienced* (-0.53) conditions than in the *Baseline* (-0.14). This confirms similar impacts of discourse and subsequent market experience on social norms as in Study 1.²¹

5. Study 3

Public discourse substantially and consistently increases socially responsible market behavior in Studies 1 and 2. However, important questions remain regarding the robustness of this positive impact and the underlying mechanisms through which it operates. Our third study tests whether additional features, involving more limited forms of participation present in many real-world contexts, influence the impacts of discourse. We also introduce additional measures, elicited immediately after discourse, to study how discourse affects values and expectations. Finally, we modify the instructions to limit potential concerns that the framing of the communication interface in Studies 1 and 2 might produce experimenter demand effects.

5.1. Experimental Design

5.1.1. *Treatments*

Study 3 comprises four treatments. The *Baseline* condition without discourse is identical to the *Baseline* in Study 2. *Discourse (Neutral)* differs from *Discourse* in Study 2 only in that we omit the statement in the chat interface that participants discuss how “socially appropriate” or “socially inappropriate” it is to trade the harmful product. While non-laboratory public discourse often guides individuals to focus on issues of responsibility and appropriate behavior, in an experiment the prompt potentially creates a demand effect of the experimenter’s expectations. In *Discourse (Neutral)* we only instruct participants that they “have the opportunity to communicate” with others in a forum to “discuss the upcoming market activity.”

We conducted two additional conditions to investigate the degree to which engagement in public discourse is necessary for it to have positive effects. First, in *Optional*, discourse works

²¹ Appendix Figure C.4 presents means, as in Figure 2 for Study 1. Table C.5 tests treatment effects on elicited social norms; the coefficients for *Discourse* and *Experienced* are both negative and statistically significant ($p < 0.01$).

identically to *Discourse (Neutral)*, including the neutral framing, except that participation is voluntary. In this condition, each participant initially decides whether to engage in discourse. If at least two participants opt to engage, the chat interface opens. Participants who enter can leave at any time but cannot re-enter. The discourse period lasts up to eight minutes or until less than two participants remain. Subjects who do not enter the chat room neither produce messages nor read messages produced by others. Subjects who participate can see at any time how many other participants are engaged in the discourse and the market roles of those participating. The *Optional* condition investigates the impact of relaxing requirements to engage in discourse.

Second, in the *Passive* condition, participants cannot actively contribute to discourse but are instead exposed to discourse generated by a set of participants in the *Discourse (Neutral)* condition. Participants are told how the messages were generated—including the instructions shown to participants in *Discourse (Neutral)*—and know that all participants in their market observe the same discourse transcript. They then observe the same chat window as in one market from *Discourse (Neutral)*, with messages appearing with identical order and timing as in the original discourse. Participants in *Passive* receive no information about market outcomes for the group whose messages they observe. If the production of arguments by those involved in market exchange is necessary for discourse to be effective, *Passive* discourse may yield limited benefits. The *Passive* condition also provides insights into whether externally generated campaigns can shape the market behavior of individuals who play no role in their development.

5.1.2. *Elicitation of Values, Expectations and Social Norms*

To better understand why discourse impacts socially responsible market behavior, we measured participants' values and expectations through a questionnaire administered immediately after the discourse period—or, in the *Baseline*, after the instructions. This included several measures dealing with personal values (e.g., “I believe that it is important to trade the product that does not reduce the donation”) and expectations (e.g., “I am confident that other participants in my group will exchange the product that does not reduce the donation”). The list of all elicited variables and summary statistics are provided in Appendix Tables D.6 and D.7.

In Studies 1 and 2, we elicited social norms after the market experience, which prevents us from identifying whether discourse directly changes social norms, or whether such impact occurs only indirectly, through market experience. In Study 3, we measure social norms twice: once after the questionnaire but before the 24 periods of market interaction, and again after the market.

5.1.3. Procedures

The study took place at the University of Zurich between October 2021 and February 2022. We followed the same procedures as in Studies 1 and 2, except that instructions were in English, while sessions in Zurich for Studies 1 and 2 took place in German. The instructions for Study 3 are provided in Appendix H. We collected 14 markets in *Baseline*, 13 in *Discourse (Neutral)* and 19 in *Optional*; we oversampled *Optional* due to the expectation that analyzing variation in participation would require more observations. We also collected 13 markets in *Passive*, with each market observing the discourse transcript from one of the 13 markets in *Discourse (Neutral)*. A total of 649 subjects participated in Study 3. On average, subjects earned about CHF 51, including a CHF 15 show-up fee. Sessions lasted about two hours.

5.2. Hypotheses

Based on our earlier findings, we expect public discourse to have positive effects on the market share of the responsible product, R , even when participants are not instructed to focus their discussion on the social appropriateness of trading the product with external harm.

$$H3.1: R(\theta_j, d_{Discourse(Neutral)}) > R(\theta_j, d_{Baseline})$$

Conditional on finding a positive impact on R in *Discourse (Neutral)*, we expect that this positive impact will be smaller when market actors' engagement in discourse is weaker, either because they can avoid discourse altogether or they consume it passively.

$$H3.2: R(\theta_j, d_{Discourse(Neutral)}) > R(\theta_j, d_{Optional})$$

$$H3.3: R(\theta_j, d_{Discourse(Neutral)}) > R(\theta_j, d_{Passive})$$

We refrain from stating *a priori* predictions whether the market share of the responsible product will be greater in *Passive* and *Optional* than in the *Baseline*, though this is a key research question.


5.3. Results

5.3.1. Market Shares of the Responsible Product

Figure 4 shows the market shares of the responsible product across treatments. The market share is 59 percent in *Baseline*, which is close to the market shares in *Baseline* (50 percent) and in Part I of *Experienced* (63 percent) in Study 2.²²

²² Appendix Figure D.1 shows the market shares of the responsible product across periods.

Figure 4: Market Shares of the Responsible Product in Study 3



Notes: The figure considers completed transactions, ignoring the small number of cases in which a buyer did not purchase a product. The bars indicate 95-percent confidence intervals, calculated at the market level.

Discourse (Neutral) yields a substantially higher market share for the responsible product (78 percent) than *Baseline*. Wilcoxon rank-sum tests at the market (buyer) level indicate that this difference is statistically significant; $p=0.006$ ($p=0.001$). This supports Hypothesis 3.1, replicating the key finding from Studies 1 and 2 and indicating that our earlier results are not primarily driven by instruction wording encouraging discussion of “appropriate” behavior.

Result 3.1: *Public discourse with neutrally framed instructions increases socially responsible market behavior.*

We next examine whether discourse has weaker effects when participants can avoid engaging in discourse. Figure 4 shows that the market share of the responsible product in *Optional* is lower than in *Discourse (Neutral)* (62 vs. 78 percent), providing support for Hypothesis 3.2 (Wilcoxon rank-sum tests at the market (buyer) level; $p=0.081$ ($p=0.019$)).

Result 3.2: *The effect of public discourse on socially responsible market behavior is weaker when participation is optional.*

The market shares of the responsible product in *Optional* and *Baseline* are very similar and not significantly different (Wilcoxon rank-sum tests at the market (buyer) level; $p=0.743$ ($p=0.337$)).²³

Observation 3.1: *Optional public discourse does not increase socially responsible market behavior.*

Despite the potential opportunity cost of time associated with engaging in discourse, almost all participants did so: 91 percent (87 percent of buyers, 94 percent of sellers). Engagement does not seem to substantially correlate with socially responsible behavior. For example, buyers who entered the chat purchased socially responsible products 63 percent of the time, which is higher than the share for those who did not enter (53 percent), but this difference is not statistically significant (Wilcoxon rank-sum: $p=0.500$). At the market level, markets with participation rates above (8 markets) and below (11) the median participation rate both had market shares of 62 percent. Thus, it appears that actual engagement with discourse has at best a weak relationship with socially responsible market behavior, creating a puzzle regarding why discourse has limited positive impacts when it is optional. We return to this question later, in Sections 5.3.4, 6 and 7.

Turning to the *Passive* condition, the market share of the responsible product (71 percent) is slightly lower than the corresponding share in *Discourse (Neutral)* (78 percent). The difference is not statistically significant at the market level ($p=0.144$) but is significant ($p=0.040$) at the buyer level in Wilcoxon rank-sum tests.²⁴ We thus find modest support for Hypothesis 3.3, in that eliminating active participation in discourse somewhat lowers its effectiveness.

Result 3.3: *The effect of public discourse on socially responsible market behavior is smaller, though not consistently statistically significantly so, when participants passively consume public discourse.*

We also find that the market share in *Passive* is somewhat higher than in *Baseline* (71 vs. 59 percent). Wilcoxon rank-sum tests show that the difference is not statistically significant at the market level ($p=0.115$) but is marginally significant ($p=0.074$) at the buyer level.

²³ Interestingly, the cumulative distributions in Appendix Figure D.2 show that, relative to *Baseline*, *Optional* yields a wider dispersion of responsible product shares across individual markets—markets at the lower (upper) end of the distribution tend to have lower (higher) market shares in *Optional* than *Baseline*.

²⁴ Appendix Figure D.2 shows that the cumulative distribution of responsible product shares in individual markets for *Passive* lies between that of *Baseline* and *Discourse (Neutral)*. Since each of the 13 individual markets in *Passive* viewed the discourse from one *Discourse (Neutral)* market, we can investigate the relationship between market shares for “paired” markets. The Spearman correlation coefficient is positive (0.23), but not statistically significant.

Observation 3.2: *Public discourse increases socially responsible market behavior when participants passively consume it, but the effect of passive participation in discourse is not consistently statistically significant.*

The random-effects GLS regressions in Table 5 complement the non-parametric tests. The dependent variable is a buyer's purchase of a responsible product. The binary treatment variables *Discourse (Neutral)*, *Optional* and *Passive* take on values of 1 in the corresponding conditions.²⁵

Table 5: GLS (random-effects) regressions of responsible buyer product choice

	(1)	(2)
<i>Discourse (Neutral)</i>	0.201*** (0.074)	0.228*** (0.074)
<i>Optional</i>	0.038 (0.071)	0.103 (0.079)
<i>Passive</i>	0.126** (0.060)	0.148*** (0.056)
<i>Period</i>		-0.002 (0.002)
<i>Period × Discourse (Neutral)</i>		-0.002 (0.003)
<i>Period × Optional</i>		-0.005* (0.003)
<i>Period × Passive</i>		-0.002 (0.002)
<i>Constant</i>	0.585*** (0.040)	0.608*** (0.040)
Observations	6,933	6,933
Number of subjects	295	295
R ²	0.043	0.042

Notes: The dependent variable in all models takes on value 1 if a buyer purchased a responsible product and 0 if the buyer purchased a harmful product. We omit the 147 cases in which a buyer did not purchase a product. *Baseline* serves as omitted categories. *Period* takes on integer values between 1 and 24. Standard errors (in parentheses) are clustered at the market level. * significant at 10%, ** significant at 5%, *** significant at 1%.

Model 1 identifies treatment effects relative to the *Baseline*. The positive coefficient for *Discourse (Neutral)* indicates that neutrally framed discourse significantly increases socially responsible market behavior (Result 3.1). In contrast, *Optional* discourse yields little impact (Observation 3.1). Comparing the coefficients for *Optional* and *Discourse (Neutral)* yields a marginally statistically significant difference (Wald test: $p=0.056$), consistent with Result 3.2.

²⁵ We report coefficient estimates of probit and OLS regressions in Appendix Tables D.1 and D.2. In the probit models, the treatment effect for *Passive* is not statistically significant in model 1 and is marginally significant in model 2. In the OLS models, the coefficient for *Passive* is statistically significant in all models.

Finally, the statistically significant coefficient for *Passive* indicates a positive effect of being exposed to messages generated by others (Observation 3.2). This coefficient is about 60 percent the size of the one for *Discourse (Neutral)*, suggesting a weakened impact on market behavior when individuals do not participate in generating messages—though consistent with Result 3.3, the difference between these two coefficients is not statistically significant (Wald test: $p=0.325$).

Model 2 adds time trends, allowing these to differ across the treatment conditions. The positive coefficients for the treatments are, if anything, slightly larger. There is a general negative trend of the frequency with which buyers purchase socially responsible products over the course of the experiment. This negative trend is slightly larger in the three treatments with discourse, and is largest in *Optional*, suggesting that this treatment not only produces weaker initial effects, but also effects that lose their impact over time (see Appendix Figure D.1).

5.3.2. Prices and Profits

As in Studies 1 and 2, responsible products trade at higher prices (33) than harmful products (21). We thus find again that buyers and sellers share the additional cost of avoiding external impacts.²⁶

5.3.3. Social Norms

Study 3 measures norms of market conduct both before and after market interaction. This allows us to separately identify the direct effects of discourse from the joint effects of discourse and market interaction. Figure 5 shows mean ratings of social appropriateness, elicited both after discourse (after reading instructions in the *Baseline*) and following market interaction. Immediately after discourse, exchange of the harmful product is perceived as less socially appropriate in all three discourse conditions than in the *Baseline*, though the effect is strongest for *Discourse (Neutral)* and *Passive*, where we also observe stronger treatment effects on behavior.


Norms elicited after the market (dark bars) show that market experience exerts little change on norms in *Discourse (Neutral)* and *Passive*. However, in *Baseline* and *Optional*, social norms change to indicate weaker proscriptions against exchanging harmful products.²⁷ Thus, discourse

²⁶ Appendix Figures D.3 and D.4 show average prices and sellers' average profits, respectively, in all conditions separately for both types of products. The regression analyses reported in Tables D.3 and D.4 replicate the findings from Studies 1 and 2 that buyers react to prices and sellers react to expected profits.

²⁷ Regression analysis in Appendix Table D.5 indicates that *Discourse (Neutral)* and *Passive* both yield more negative social perceptions of exchanging the harmful product, relative to the *Baseline*, both before and after the market. These comparisons are at least marginally statistically significant (e.g., $p<0.086$ for all four coefficients in ordered probit models). The coefficients for *Optional* are also negative, but smaller in magnitude and not statistically significant, particularly after market experience.

appears to directly influence social norms, but we also find some evidence that subsequent market experience can further impact norms.

Figure 5: Effect of Public Discourse and Market Experience on Social Norms



Notes. The figure shows the average rating of the appropriateness of exchanging the harmful product. *Norms before market activity* refers to social norms measured prior to interacting in the market; *Norms after market activity* refers to social norms measured after market interaction. “Very socially appropriate = 1,” “Somewhat socially appropriate = 1/3,” “Somewhat socially inappropriate = -1/3,” “Very socially inappropriate = -1.” The numerical rating values follow Krupa and Weber (2013). The bars indicate 95-percent confidence intervals, calculated at the market level.

5.3.4. Impact of discourse on values and expectations

We next study the questionnaire responses collected directly after discourse (after instructions in the *Baseline*), to better understand how discourse changes factors potentially relevant for behavior. To identify common dimensions across the questionnaire items, we first conducted exploratory factor analysis on the nine items common to all treatments. We obtain three factors (with eigenvalues larger one) that jointly explain 70 percent of the variation across all items (see Appendix Table D.8). The loadings of specific items on factors (Table D.9) yield straightforward interpretations. First, *Beliefs about others* includes expectations that others value and expect to trade the responsible product. Second, *Personal values* includes personal support for exchanging different types of products. Finally, *Coordination* indicates the belief that everyone knows what to do when it comes to product types and prices.

Table 6 reports regressions that explore how discourse influences these values and beliefs. All three treatments have positive impacts on *Beliefs about others*, with the magnitudes of the coefficients suggesting stronger effects for the *Discourse (Neutral)* and *Passive* conditions, and weaker effects for *Optional*.²⁸ The impacts of the discourse treatments on this factor generally mirror the treatment effects on the share of responsible products. Model 2 reveals that the treatments have little effect on *Personal values*.²⁹ Finally, the third model indicates that both *Discourse (Neutral)* and *Optional* have strong impacts on *Coordination*, that is, on perceived agreement about prices and products, but there is little effect for *Passive*.³⁰

Table 6: Treatment effects on values and beliefs (OLS)

	(1) <i>Beliefs about others</i>	(2) <i>Personal values</i>	(3) <i>Coordination</i>
<i>Discourse (Neutral)</i>	0.498*** (0.180)	0.117 (0.118)	0.567*** (0.120)
<i>Optional</i>	0.283* (0.148)	-0.079 (0.086)	0.487*** (0.133)
<i>Passive</i>	0.394*** (0.091)	0.043 (0.101)	-0.073 (0.134)
<i>Constant</i>	-0.288*** (0.048)	-0.010 (0.069)	-0.266*** (0.091)
Observations	649	649	649
R ²	0.032	0.005	0.078

Notes: The dependent variable in each model is one of the factors constructed from the nine questionnaire items (see Appendix Table D.9). In all models, *Baseline* serves as the omitted category. Standard errors (in parentheses) are clustered at the market level; *** p<0.01, ** p<0.05, * p<0.1.

We also investigate the relationship between variation in the above factors and market behavior. Appendix Table D.12 reports regressions of a buyer's decision to purchase a socially responsible product or a seller's decision to offer a socially responsible product on *Beliefs about*

²⁸ However, Wald tests of the equality of *Optional* with *Discourse (Neutral)* and with *Passive* are not statistically significant (p=0.340 and p=0.490, respectively).

²⁹ Appendix Table D.10 reports analyses using each questionnaire item separately. The results are generally consistent with the aggregated analysis in Table 6. The largest treatment effects of *Discourse (Neutral)* are on items that load heavily on *Beliefs about others* and *Coordination*. *Discourse (Neutral)* has some impact on two of the three items related to *Personal values*, though these effects are smaller. *Passive* mainly impacts items related to *Beliefs about others*. The impacts of *Optional* are strongest for items related to *Coordination*.

³⁰ Since the focus here is on individual beliefs and behavior, rather than on product market shares, we also consider buyers and sellers, separately. Appendix Table D.11 repeats the analysis in Table 6 for buyers and sellers, generally finding similar patterns. For both roles, the coefficients measuring treatment effects on *Beliefs about others* are largest for *Discourse (Neutral)* and *Passive* and there are no treatment effects on *Personal values* for either role. Treatment effects on *Beliefs about others* and *Coordination* tend to be stronger for sellers than for buyers.

others, *Personal values* and *Coordination*.³¹ *Personal values* are strongly related to responsible market behavior, indicating that personal convictions regarding social responsibility translate into actions buyers and sellers take in the market, though Table 6 indicates that this factor is not influenced by discourse. The factor *Beliefs about others*—which does increase following discourse, particularly in the *Discourse (Neutral)* and *Passive* conditions—is also positively and statistically significantly related to socially responsible market behavior for both buyers and sellers. Finally, *Coordination* is positively related to market behavior for sellers, but not buyers, and the relationships are generally weaker than for *Beliefs about others*.

We interpret the above analysis as providing suggestive evidence that discourse may impact behavior, at least in part, by strengthening beliefs that others support exchanging socially responsible products. Of course, the analysis here is exploratory and correlational and should be interpreted cautiously.

6. Content of the Public Discourse

In this section, we briefly report exploratory analysis of the content of the discourse from our three studies. Most subjects were actively engaged: the proportion sending at least one message was 98, 95 and 88 percent in Switzerland in Studies 1, 2 and 3, respectively, and 98 percent in China.

Independent coders classified the content of the messages into categories.³² Our analysis focuses on two main categories: messages classified as advocating either *No Impact* (“responsibilizing” messages advocating for the responsible product or against the harmful one) or *Impact* (“exculpatory” messages advocating the harmful product). Using these two categories, we construct a measure of the degree to which a participant argued for responsible market conduct: for each participant, we summed the number of *No Impact* messages sent by that participant and subtracted the number of *Impact* messages sent, to construct the variable *Prosocial*.

Table 7 reports the distributions of participants’ *Prosocial* classifications across studies and conditions.³³ A few observations are apparent from the table. First, across most conditions in

³¹ Appendix Table D.13 reports analogous analysis using the individual questionnaire items.

³² Four coders evaluated each statement; we classify a statement as belonging to a category if at least three coders assigned it to that category. Appendix Section E.1 describes all categories and details of the coding procedure.

³³ Appendix Table E.8 reports ordered probit regressions testing whether the categorization frequencies in Table 7 differ across discourse conditions. Only two conditions, *Exclusive* in China (Study 1) and *Optional* (Study 3) yield significantly less *Prosocial* types than the other conditions in the respective studies. To corroborate that *Prosocial* classifications in Table 7 correspond sensibly to participants’ communication strategies, we study how often the three

Switzerland, the modal type is *Prosocial* > 0, where a participant sends more messages encouraging *No Impact* than *Impact*. The one exception is the *Optional* condition of Study 3, in which only 33 percent of market actors are classified as *Prosocial* > 0. Second, the frequencies of *Prosocial* > 0 types are lower in all conditions in China than in Switzerland; in China, the modal classification is always the neutral type (*Prosocial* = 0). Third, in China, we observe fewer participants classified as *Prosocial* > 0 in the *Exclusive* condition than in the other two conditions. Importantly, all three of these observations track the variation in the market shares of the responsible product across our studies.

Table 7: Proportions of subjects classified according to prosocial orientation in discourse

	Study 1						Study 2		Study 3	
	Switzerland			China						
	<i>Veil</i>	<i>No Veil</i>	<i>Exclusive</i>	<i>Veil</i>	<i>No Veil</i>	<i>Exclusive</i>	<i>Discourse</i>	<i>Experienced</i>	<i>Discourse (Neutral)</i>	<i>Optional</i>
<i>Prosocial</i> > 0	57%	65%	61%	36%	33%	18%	49%	50%	49%	33% (36%)
<i>Prosocial</i> = 0	36%	34%	34%	59%	63%	66%	43%	40%	46%	56% (51%)
<i>Prosocial</i> < 0	7%	1%	5%	5%	4%	16%	8%	10%	6%	12% (13%)

Notes. The modal type in each column is shaded. *Prosocial*>0 (*Prosocial*<0) corresponds to participants who sent strictly more (fewer) messages advocating for the responsible product (*No Impact*) than for the harmful product (*Impact*). *Prosocial*=0 corresponds to participants who sent equal numbers of messages (possibly zero) of both types. For Study 1, third parties are excluded from analysis for the *Exclusive* condition. For Study 3 *Optional*, participants who did not enter the discussion forum are classified as *Prosocial* = 0 and the numbers in parentheses indicate the frequencies when we exclude participants who did not enter the discussion forum.

Table 7 suggests that communication patterns—particularly, the frequency with which subjects advocate for socially responsible market behavior—may play a role in shaping market behavior. In Appendix Section E.3, we describe additional analysis of the relationship between experiences with prosocial discourse and the subsequent beliefs and behavior of market participants. While this exploratory analysis should be interpreted cautiously, it suggests that the degree to which an individual is exposed to prosocial arguments from others is strongly correlated

types sent messages appealing to *Fairness* and *Self-interest*, two other categories in our content analysis (see Tables E.1 to E.3). Appendix Table E.9 shows that, across all studies, messages appealing to *Fairness* (*Self-interest*) are sent more frequently by individuals classified as *Prosocial* > 0 (*Prosocial* < 0) than by either of the other types.

with subsequent expectations about others' market behavior (the variable *Beliefs about others*, discussed in Section 5.3.4), social norms of market conduct elicited before market experience and socially responsible market behavior.

7. Discussion

Across three studies, we find significant positive impacts of discourse on socially responsible market behavior in 10 out of 11 independent treatments. These effects range from a 12-percentage-point increase (*Passive* condition in Study 3) to a 47-percentage-point increase (*Veil* in Switzerland in Study 1). This is despite most of our treatments, after *Veil*, having been designed to introduce features expected to diminish or even reverse the positive impacts of discourse. Moreover, the large positive impacts of discourse in most of our treatments are striking, given the general stability of the baseline level of market social responsibility across our studies and in earlier work—typically robust around 50 percent in experiments conducted in Switzerland—and the fact that other interventions not involving discourse generally find modest or no effects on market social responsibility (Bartling, et al. 2015, 2019, 2022). It is also noteworthy that we observe an exception to the positive impacts of discourse. Making participation in discourse optional reduces its effectiveness, even though a large majority of participants opt to enter the discourse forum.

In this section, we draw together evidence collected across our studies, to obtain insights about when discourse works and when it fails, what makes discourse effective, and what these observations imply for the generalizability of our results. This discussion is *post hoc* and speculative, and we provide several conjectures that can be tested in future research, rather than definitive conclusions.

Communication and social behavior. Numerous studies document positive impacts on pro-social behavior of providing individuals with opportunities to communicate. However, in this work communication facilitates more favorable outcomes to those communicating (e.g., Ostrom, et al., 1992). Our study documents that communication between individuals can also yield substantial increases in pro-sociality when this impact is financially harmful to the communicating parties. Even in treatments in which *no* parties engaging in discourse benefit financially from socially responsible conduct (e.g., the *Exclusive* condition in Study 1 and all conditions in Studies 2 and 3) we generally observe that discourse increases prosocial conduct.

As with other work on the positive impacts of communication (e.g., Charness and Dufwenberg, 2006; Kessler and Leider, 2012), we find some evidence that an important aspect of how discourse influences behavior is through its impact on individuals' expectations of others' values and behavior. Discourse often yields a high frequency of participants advocating for socially responsible market behavior, particularly in Switzerland (Table 7), and the prevalence of such prosocial arguments correlates with increased expectations that others support socially responsible exchange and with higher market shares for the socially responsible product (see Appendix Section E.3). This suggests that public discourse may be an instrument for addressing individuals' misperceptions of the degree to which others support costly actions to mitigate negative externalities, or “pluralistic ignorance” (Sparkman, et al., 2022).

Impacts of discourse across populations. Study 1 investigates the impact of discourse in two student populations with different levels of baseline social concern (Bartling et al., 2015). Rather than finding public discourse to complement or substitute these baseline levels, we observe impacts similar in absolute magnitude in the two populations. This provides an initial indication that the impacts of discourse may generalize across different populations, suggesting positive impacts of discourse campaigns even in societies with varying levels of baseline social concern.

The production and consumption of discourse. Our studies also provide evidence on the separate impacts of individuals producing arguments in discourse and consuming arguments produced by others. In the *Passive* condition in Study 3, in which individuals observe discourse produced by another set of individuals, the impact on the market share of the responsible product is 63 percent of the impact of *Discourse (Neutral)*. In the supplemental *Reflection* conditions in Study 1 (see Appendix B), in which individuals privately produce arguments, we observe increases in socially responsible product market shares that are 47 and 83 percent, respectively, of the impacts in the corresponding discourse conditions in Switzerland and China. Similarly, our analysis of the impacts of what participants write and what they observe others writing (Appendix Section E.3) supports roles for both one's own communication and the communication generated by others. Thus, both the production and consumption of arguments appear to independently influence concern for social impact, though their effects appear smaller than the joint effect of discourse involving the simultaneous production and consumption of arguments. This observation suggests some role for public campaigns that either ask individuals to reflect on their behavior or

that present them with arguments supporting acting responsibly (Goldstein, et al., 2008; Krupka and Weber, 2009; Dal Bo and Dal Bo, 2014), though richer discourse may yield additional benefits.

The role of demand effects. An important consideration in our research involves the degree to which behavior is influenced by a “demand effect”—the expectation that an authority figure like the experimenter wants participants to act a certain way (Zizzo, 2010). For example, by asking participants to discuss the “appropriate” thing to do, some of our conditions may have implicitly encouraged them to act more responsibly.³⁴ Our most direct test of demand effects is a comparison between the *Discourse* and *Discourse (Neutral)* conditions in Studies 2 and 3, with the former encouraging discussion of “appropriate” conduct and the latter omitting such a reference. The fact that these conditions yield virtually identical market shares for the socially responsible product (79 versus 78 percent) provides some reassurance that our treatment effects are unlikely to be entirely driven by demand effects. However, the fact that the increase in market share, relative to the *Baseline*, is lower for *Discourse (Neutral)* in Study 3 than for *Discourse* in Study 2 means that we cannot entirely rule out that demand effects might have played some role in our experiment.³⁵

Another possible interpretation of the role of demand effects involves the possibility that different treatments, by their nature, may have provided participants with cues about the experimenter’s intentions or preferences. Thus, for example, the *Optional* condition in Study 3, the only discourse condition yielding no increase in market social responsibility, might have been interpreted as a condition where the possibility of avoiding discourse signaled the experimenter’s expectation it would not be effective.³⁶ While we cannot rule out this possibility, this explanation would also have to account for why other treatments in which the experimenter gave participants potential excuses to avoid socially responsible behavior did not create similar demand effects. For example, the *Exclusive* condition in Study 1 (which excluded third parties impacted by market exchange from discourse) or the *Experienced* treatment in Study 2 (in which participants’ own prior behavior could justify “business as usual”) could have similarly produced cues indicating the

³⁴ It is worth noting, however, that we only asked subjects in Studies 1 and 2 to discuss *whether* it is socially appropriate or inappropriate to trade the product with impact. Thus, even the wording used in our first two studies is likely to create only weak demand effects, if any, in contrast to stronger forms studied by, e.g., De Quidt et al. (2018).

³⁵ The difference in the size of the treatment effects could be caused by random variation in the *Baseline* market shares between Studies 2 and 3. To provide evidence on whether the difference may be more substantive, we conducted regressions using pooled data from Studies 2 and 3 comparing the treatment effect of *Discourse* (relative to *Baseline*) in Study 2 to that of *Discourse (Neutral)* in Study 3 (see Appendix Table D.14). The interaction term measuring whether the impact of discourse differs across studies is negative but not statistically significant, suggesting no substantive impact of the instruction wording on the effect of public discourse on market social responsibility.

³⁶ We thank an anonymous reviewer for pointing out this possible interpretation of the *Optional* condition.

experimenter's expectation that discourse would have no positive impacts. Indeed, it is worth restating that our objective in designing *all* treatments after the *Veil* condition in Study 1 was to introduce features likely to make discourse less effective. However, discourse nevertheless had large positive impacts in all conditions but *Optional*.³⁷ Moreover, if the *Optional* condition created an experimenter demand effect, it would likely operate through a perceived encouragement for individuals to avoid discourse, which is inconsistent with the observed high participation rate.

Finally, we believe it is important to note that mechanisms similar to demand effects may play an important role in promoting socially responsible behavior outside the laboratory. Requests from politicians and social leaders to engage in more responsible market conduct may play roles in achieving behavior change through processes similar to “demand effects.” We believe that such channels are an important avenue for study (Dal Bo and Dal Bo, 2014; Antonakis, et al., 2022).

The importance of mandatory engagement. Our study highlights an instance where discourse is ineffective, specifically when engagement in discourse is optional. Our expectation was that if discourse proved less effective in *Optional*, this would be through reduced participation. However, this seems to not be a principal driver—*Optional* yields relatively low market shares for the responsible product, despite high levels of participation. Here, we draw on a few observations that provide suggestive indications for why *Optional* discourse may be ineffective.

First, Table 7 documents that participants are less likely to advocate for socially responsible market behavior in *Optional* than in the other discourse conditions in Switzerland. This does not appear to be driven primarily by low participation, as the vast majority of participants enters the discourse. However, our data do not provide a clear explanation for why communication differs in *Optional*. Speculative reasons include, for example, that the act of initially voluntarily engaging in discourse may create a form of “moral licensing” that diminishes the perceived obligation to subsequently advocate for socially responsible exchange (Bénabou and Tirole, 2011) or that those who forgo discourse are those most likely to advocate for social responsibility, perhaps from a sense of obligation (Lazear, et al., 2012).

³⁷ As an additional test, we examined the content of the discourse for instances in which subjects referred to the expectations or desires of the experimenters. Across all studies, we identified only four such statements (and only five additional ones in which the experimenter was mentioned in any manner): two in the *Veil* condition in China, one in the *Veil* condition in Switzerland and one in the *Optional* condition, which suggests that experimenter demand effects (and the experimenter in general) did not play a prominent role, at least not in discourse between subjects.

Second, market actors in *Optional* leave the discourse period with lower confidence that others support exchanging the socially responsible product (Table 6) and with slightly weaker social norms against exchanging the harmful product (Figure 5) than in either *Discourse (Neutral)* or *Passive*. They also start the market activity with lower market shares for the socially responsible product (Appendix Figure D.1). Across all these measures of the initial strength of concern for mitigating negative social impacts, *Optional* lies roughly halfway between the *Baseline* and *Discourse (Neutral)*, though our study lacks power to precisely test the strength of these initial differences. Nevertheless, these differences are consistent with the possibility that exposure to others' prosocial arguments—which Table 7 shows occurs less frequently in *Optional*—plays an important role in establishing expectations and norms that subsequently support socially responsible market behavior. In Appendix Section E.3 we provide evidence that exposure to others' arguments advocating for social responsibility is positively correlated with the expectation that others support socially responsible exchange, with social norms proscribing such exchange and with subsequent behavior in the market.

Third, the market share of the responsible product in *Optional* decreases over time, ending up at levels indistinguishable from those in the *Baseline* (Appendix Figure D.1). This is also consistent with the decline in *Optional* for social norms against exchanging the harmful product following market experience (Figure 5). Interestingly, the declining market share of the responsible product seems to be stronger for those groups that did not have full participation in discourse (see Appendix Figure D.5 and Table D.15), suggesting a potential indirect role for full engagement, even when voluntary, in supporting socially responsible exchange over time.

Thus, rather than a single factor being responsible for the limited efficacy of *Optional* discourse, it seems that this experimental condition is less effective for promoting socially responsible market exchange due to a combination of differences that arise during discourse, in initial beliefs and behavior following discourse and through market experience. It is important to note, however, that this interpretation is *post hoc* and speculative, and should be taken with caution.

Finally, one important implication of our finding that *Optional* engagement with discourse limits its positive impact is that, where possible, policymakers should encourage mandatory engagement in discussions of what it means to act responsibly in markets. This might include, for instance, educational interventions such as civics classes that promote discussions of social values (Winthrop, 2020), perhaps particularly at young ages before people learn of their roles in society

(as in our *Veil* conditions), or public campaigns that are highly visible to all members of a society (as in our *Passive* condition).

8. Conclusions

We investigate the impact of public discourse on the exchange of responsible products. Across three laboratory studies, our main finding is that all but one of 11 comparisons between no-discourse baselines and conditions that expose subjects to some form of discourse increase the share of responsible products subsequently traded in a market. The lone discourse condition that appears to produce no substantial increase in socially responsible market behavior occurs when participants have the option to forgo discourse—even though most participants voluntarily enter discourse. Our results thus provide evidence that it is possible for collections of anonymous individuals to use discourse to convince one another to shift their market behavior to incur greater personal costs in return for reducing externalities imposed on others uninvolved in market exchange.

Naturally, our work requires several caveats. While our treatment comparisons vary the nature of public discourse in many ways, there are potentially important factors that we omit. This includes sources of heterogeneity among market actors, such as nationality and income, which may make it difficult for discourse to yield sustainable agreements. Furthermore, the specific nature of discourse—e.g., through policy institutions, non-profit public campaigns and the media—can vary in many ways that extend far beyond the simple kind of discourse implemented in our study. Nevertheless, our design can be easily extended to incorporate additional features worth studying.

Additionally, the degree to which laboratory findings generalize to non-laboratory settings is an important concern with all laboratory studies. Our study is motivated by the observation of large public discourse campaigns intended to promote more responsible market conduct, and the recognition that studying the causal impacts of such campaigns is challenging to address with non-laboratory data. Our laboratory data provide an indication that the kind of discourse encouraged by such campaigns *can* facilitate socially responsible market behavior. We make no claims that it will always, or even necessarily frequently, do so. Indeed, we document that making participation optional—a feature of most natural settings—may drastically reduce the positive impacts of discourse. Nevertheless, our study documents instances in which discourse can have positive and large impacts and provides preliminary insights into the mechanisms that might be at work.

References

- Amenta, E., Caren N., Chiarello, E., and Su, Y. (2010). The Political Consequences of Social Movements. *Annual Review of Sociology*, 36, 287-307.
- Andreoni, J., and Rao, J. M. (2011). The Power of Asking: How Communication Affects Selfishness, Empathy, and Altruism. *Journal of Public Economics*, 95(7-8), 513-520.
- Antonakis, J., d'Adda, G., Weber, R. A., and Zehnder, C. (2021). "Just words? Just speeches?" On the economic value of charismatic leadership. *Management Science*.
- Agranov, M. and Tergiman, C. (2014). Communication in Multilateral Bargaining. *Journal of Public Economics*, 118, 75-85.
- Babcock, L., Loewenstein, G., Issacharoff, S., and Camerer, C. (1995). Biased Judgments of Fairness in Bargaining. *American Economic Review*, 85(5), 1337-1343.
- Baranski, A. and Morton, R. (2022). The determinants of multilateral bargaining: A comprehensive analysis of Baron and Ferejohn majoritarian bargaining experiments. *Experimental Economics*, 25(4), 1079-1108.
- Baron, R. M., and Kenny, D. A. (1986). The moderator–mediator variable distinction in social psychological research: Conceptual, strategic, and statistical considerations. *Journal of Personality and Social Psychology*, 51(6), 1173.
- Bartling, B., Valero, V., and Weber, R. (2019). On the Scope of Externalities in Experimental Markets. *Experimental Economics*, 22 (3), 610-624.
- Bartling, B., Valero, V., and Weber, R. (2022). The Causal Effect of Income Growth on Consumer Social Responsibility. *SSRN Discussion Paper No. 3249788*.
- Bartling, B., Weber, R., and Yao, L. (2015). Do Markets Erode Social Responsibility? *Quarterly Journal of Economics*, 130(1), 219-266.
- Bartling, B., Valero, V., Weber, R. and Yao, L. (2024). Replication data for: Public Discourse and Socially Responsible Market Behavior. American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor].
- Bénabou, R., Falk, A., and Tirole, J. (2020). Narratives, Imperatives, and Moral Persuasion. *NBER working paper*.
- Bénabou, R., and Tirole, J. (2010). Individual and Corporate Social Responsibility. *Economica*, 77(305), 1-19.

- Bénabou, R., & Tirole, J. (2011). Identity, morals, and taboos: Beliefs as assets. *The Quarterly Journal of Economics*, 126(2), 805-855.
- Biais, B., Bisière, C., and Pouget, S. (2014). Equilibrium Discovery and Preopening Mechanisms in an Experimental Market. *Management Science*, 60(3), 753-769.
- Blume, A., and Ortmann, A. (2007). The Effects of Costless Pre-play Communication: Experimental Evidence from Games with Pareto-ranked Equilibria. *Journal of Economic Theory*, 132(1), 274-290.
- Bochet, O., Page, T., and Putterman, L. (2006). Communication and Punishment in Voluntary Contribution Experiments. *Journal of Economic Behavior and Organization*, 60(1), 11-26.
- Bohnet, I., and Frey, B. S. (1999). Social distance and other-regarding behavior in dictator games: Comment. *American Economic Review*, 89(1), 335-339.
- Brandts, J., and Cooper, D. J. (2007). It's what you say, not what you pay: an experimental study of manager employee relationships in overcoming coordination failure. *Journal of European Economic Association*, 5 (6):1223-1268.
- Charness, G., and Dufwenberg, M. (2006). Promises and Partnership. *Econometrica*, 74(6), 1579-1601.
- Cooper, R., DeJong, D. V., Forsythe, R., and Ross, T. W. (1992). Communication in Coordination Games. *Quarterly Journal of Economics*, 107(2), 739-771.
- d'Adda, G., Darai, D., Pavanini, N., and Weber, R. (2017). Do Leaders Affect Ethical Conduct? *Journal of the European Economic Association*, 15(6), 1177-1213.
- Dal Bo, E., and Dal Bo, P. (2014). "Do the right thing:" The effects of moral suasion on cooperation. *Journal of Public Economics*, 117, 28-38.
- Danz, D., Engelmann, D., and Kübler, D. (2022). Do Legal Standards Affect Ethical Concerns of Consumers? *European Economic Review*, 144, Article 104044.
- Dawes, R. M., McTavish, J., and Shaklee, H. (1977). Behavior, Communication, and Assumptions about other People's Behavior in a Commons Dilemma Situation. *Journal of Personality and Social Psychology*, 35(1), 1.
- De Quidt, J., Haushofer, J., and Roth, C. (2018). Measuring and Bounding Experimenter Demand. *American Economic Review*, 108(11), p. 3266-3302.
- Dufwenberg, M., Heidhues, P., Kirchsteiger, G., Riedel, F., and Sobel, J. (2011). Other-regarding Preferences in General Equilibrium. *Review of Economic Studies*, 78(2), 613-639.

- Dufwenberg, M., Johansson-Stenman, O., Kirchler, M., Lindner, F., and Schwaiger, R. (2022). Mean Markets or Kind Commerce. *Journal of Public Economics*, 209, 104648.
- Ellman, M., and Pezanis-Christou, P. (2010). Organizational Structure, Communication, and Group Ethics. *American Economic Review*, 100(5), 2478-91.
- Engelmann, D., Friedrichsen, J., and Kübler, D. (2018). Fairness in Markets and Market Experiments (No. SP II 2018-203). WZB Discussion Paper.
- Exley, C. (2020). Using Charity Performance Metrics as an Excuse Not to Give. *Management Science*, 66(2), 553-563.
- Goldstein, N., Cialdini, R., and Griskevicius, V. (2008). A Room with a Viewpoint: Using Social Norms to Motivate Environmental Conservation in Hotels. *Journal of Consumer Research*, 35(3), 472–482.
- Hainmueller, J., Hiscox, M. J., and Sequeira, S. (2015). Consumer Demand for Fair Trade: Evidence from a Multistore Field Experiment. *Review of Economics and Statistics*, 97(2), 242-256.
- Hargreaves Heap, S. and D. Zizzo (2011). Emotions and Chat in a Financial Markets Experiment. *SSRN Working Paper 1783462*.
- Irlenbusch, B., and Saxler, D. (2019). The Role of Social Information, Market Framing, and Diffusion of Responsibility as Determinants of Socially Responsible Behavior. *Journal of Behavioral and Experimental Economics*, 80, 141-161.
- Isaac, R. M., and Walker, J. M. (1988). Communication and free-riding behavior: The voluntary contribution mechanism. *Economic Inquiry*, 26(4), 585-608.
- Kessler, J. B. (2017). Announcements of support and public good provision. *American Economic Review*, 107(12), 3760-3787.
- Kessler, J. B., and Leider, S. (2012). Norms and Contracting. *Management Science*, 58(1), 62-77.
- Kirchler, M., Huber, J., Stefan, M., and Sutter, M. (2015). Market Design and Moral Behavior. *Management Science*, 62(9), 2615-2625.
- Kitzmueller, M., & Shimshack, J. (2012). Economic Perspectives on Corporate Social Responsibility. *Journal of Economic Literature*, 50(1), 51-84.
- Kocher, M., Schudy, S., and Spantig, L. (2018). I Lie? We Lie! Why? Experimental Evidence on a Dishonesty Shift in Groups. *Management Science*, 64(9), 3971-4470.

- Kriss, P. H., Blume, A., and Weber, R. A. (2016). Coordination with Decentralized Costly Communication. *Journal of Economic Behavior and Organization*, 130, 225-241.
- Krupka, E. L., Leider, S., and Jiang, M. (2017). A Meeting of the Minds: Informal Agreements and Social Norms. *Management Science*, 63(6), 1708-1729.
- Krupka, E., and Weber, R. A. (2009). The focusing and informational effects of norms on prosocial behavior. *Journal of Economic Psychology*, 30(3), 307-320.
- Krupka, E. L., and Weber, R. A. (2013). Identifying social norms using coordination games: Why does dictator game sharing vary?. *Journal of the European Economic Association*, 11(3), 495-524.
- Lazear, E. P., Malmendier, U., and Weber, R. A. (2012). Sorting in experiments with application to social preferences. *American Economic Journal: Applied Economics*, 4(1), 136-163.
- Levy, R., and Mattsson, M. (2023). The Effects of Social Movements: Evidence from #MeToo. *SSRN Working Paper 3496903*.
- MacKinnon, D. P., Fairchild, A. J., and Fritz, M. S. (2007). Mediation analysis. *Annual Review of Psychology*, 58, 593-614.
- Madestam, A., Shoag, D., Veuger, S., and Yanagizawa-Drott, D. (2013). Do political protests matter? Evidence from the tea party movement. *Quarterly Journal of Economics*, 128(4), 1633–1685.
- Ockenfels, A., Werner, P., and Edenhofer, O. (2020). Pricing Externalities and Moral Behaviour. *Nature Sustainability*, <https://doi.org/10.1038/s41893-020-0554-1>.
- Ostrom, E., Walker, J., and Gardner, R. (1992). Covenants with and without a sword: Self-governance is possible. *American Political Science Review*, 86(2), 404-417.
- Pigors, M., and Rockenbach, B. (2016). Consumer Social Responsibility. *Management Science*, 62(11), 3123-3137.
- Rode, J., Le Menestrel, M., and Hogarth, R. (2008). Ethical Differentiation and Market Behaviour: An experimental approach. *Journal of Economic Behavior and Organization*, 66, 265-280.
- Saccardo, S., and Serra-Garcia, M. (2021). Cognitive Flexibility or Moral Commitment? Evidence of Anticipated Belief Distortion, *CESifo Working Paper No. 8529*.
- Sparkman, G., Geiger, N., and Weber, E. U. (2022). Americans experience a false social reality by underestimating popular climate policy support by nearly half. *Nature Communications*, 13(1), 4779.

- Sutter, M., Huber, J., Kirchler, M., Stefan, M., and Walzl, M. (2020). Where to Look for the Morals in Markets? *Experimental Economics*, 23, 30-52.
- Sutter, M., and Strassmair, C. (2009). Communication, Cooperation and Collusion in Team Tournaments - An Experimental Study. *Games and Economic Behavior*, 66(1), 506-525.
- White, K., Habib, R., and Hardisty, D. (2019). How to SHIFT Consumer Behaviors to be More Sustainable: A Literature Review and Guiding Framework. *Journal of Marketing*, 83(3), 22-49.
- Winthrop, R. (2020). The need for civic education in 21st-century schools. Brookings Institution Policy 2020 Report, June 4 available at https://www.brookings.edu/wp-content/uploads/2020/04/BrookingsPolicy2020_BigIdeas_Winthrop_CivicEducation.pdf.
- Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics*, 13, 75-98.