

Correcting Consumer Misperceptions about CO₂ Emissions *

Taisuke Imai¹ Davide D. Pace¹
Peter Schwardmann² Joël J. van der Weele³

¹Ludwig-Maximilians-Universität München

²Carnegie Mellon University

³Amsterdam School of Economics, Tinbergen Institute

August 24, 2023

Abstract

Policy makers champion information provision about carbon impact, on the premise that consumers are willing to mitigate but are poorly informed about how to do so. We empirically test this argument and reject it. We collect an extensive new dataset and find both large misperceptions of carbon impact and clear preferences for mitigation. Yet, in two separate experiments, we show that large beliefs corrections have no effect on consumption in a large representative sample. Our null result is well-powered and highly informative, as we target information for maximal impact. It questions the potential of information policies to fight climate change.

JEL Classification codes: C81, C93, D84, Q54

Keywords: climate change, carbon emissions, information provision, consumer behavior

*We thank Sára Khayouti, Antonia Kurz, and Tom Zegers for excellent research assistance. We thank David Bilén, Raphael Epperson, Henning Hermes, Thao Hanh Le, Yves Le Yaouanq, Florian Schneider, and seminar participants at LMU Munich, University of Pittsburgh, BAAdW, 6th Workshop on Experimental Economics for the Environment, ESA 2021 Global Meetings, Osaka University ISER, CRC Workshop “Memory, Narratives and Belief Formation”, Düsseldorf Institute for Competition Economics, French Experimental Talks, ASFEE 2022, CESifo Area Conference on Behavioral Economics 2022, Berlin Behavioral Economics Seminar, BGSE/briq Applied Microeconomics Workshop, EAERE 2023, and EEA-ESEM 2023, for helpful comments. Financial support from Deutsche Forschungsgemeinschaft through CRC TRR 190 (Imai, Pace, Schwardmann), the Bavarian Academy of Sciences and Humanities (Schwardmann), the NWO in the context of VIDI grant 452-17-004 (van der Weele), the Amsterdam Center for Behavioral Change of the University of Amsterdam, and the initiative, “A Sustainable Future,” at the University of Amsterdam (van der Weele), are gratefully acknowledged.

1 Introduction

Reducing the emission of greenhouse gases is one of the most pressing challenges of our time. Carbon pricing, a potential remedy, is politically contentious. Thus, policy makers frequently stress the role of information about CO₂ emissions to consumers and producers. For instance, the European Commission’s “Farm to Fork Strategy” proposes an extensive carbon labeling strategy, while its “New Consumer Agenda” argues for “more reliable information on sustainability” (European Commission, 2020). In the US, the proposed, but ill-fated, American Clean Energy and Security Act of 2009 contained provisions to study and implement carbon information aimed at consumers (Waxman and Markey, 2009), while the Department of Agriculture and the EPA implement greenhouse gas labels for various products. Corporations are also interested in carbon labeling, as evidenced by initiatives from large European retail chains like TESCO, Casino, and E.Leclerc (Taufique et al., 2022).

Underlying all these information initiatives is an implicit argument: It presumes that people care about reducing emissions but that they do not have sufficient awareness of climate impact and may underestimate the impact of their actions. From these premises, it follows that when misperceptions about emission sizes are corrected, consumers adjust their behavior and reduce emissions. Indeed, the Commission states in relation to the green transition that it “aims to ensure [...] that consumers have better information to be able to make an informed choice.”¹

In this paper, we test this argument empirically and show that it is flawed. To do so, we proceed in several steps. First, to test the argument’s premises, we survey a representative sample of US consumers ($N = 1,022$). We collect point estimates and belief distributions about the carbon impact of several products and actions. We then measure valuations of carbon emissions for the same consumers, using a willingness to pay for different amounts of carbon offsets. We find that consumers generally underestimate carbon impact; the largest underestimates exist for high-carbon-impact food categories such as beef and coffee. Valuations of carbon emission reductions are relatively high, but the marginal willingness to mitigate declines strongly with emission size. These findings partially replicate previous results, but with incentivized methods and in a representative sample, and confirm the two premises of the policy makers’ argument.

We then leverage these data to develop the strongest possible test of the effect of information on consumer behavior. We use a structural model where consumers derive disutility from the (expected) emissions linked to their actions. We compare the individual disutility of consumption given a person’s subjective beliefs about emissions, with a counterfactual where the belief distribution is replaced by the true value of carbon emissions, as measured by the latest scientific estimates. The predicted information

¹See press release on the New Consumer Agenda, https://ec.europa.eu/commission/presscorner/detail/en/IP_20_2069, accessed 23rd August 2023.

impact is thus based on consumers who are both uninformed and willing to mitigate, and it takes into account the degree of uncertainty among consumers and the (diminishing) responsiveness to emission size. The model allows us to predict the products for which information provision generates a maximal behavioral response.

Finally, we test our predictions in two experiments, each in a large representative sample. Our experiments focus on the demand for beef and poultry; while these products are part of the same food category (meat), beef has almost 10 times the carbon impact of poultry in CO₂ equivalents. Our participants understand that beef is more polluting than poultry, but they think that the difference between them is much smaller than it actually is. In line with this, our structural model, applied to the representative survey data, predicts that information on beef should have a large impact on demand. Instead, the impact on demand by providing information on chicken should be small or non-existent.

Our first experiment gives some of the participants in our initial survey precise information about the carbon impact of these (and other) products. We contacted the same subjects two weeks later and asked their recollection of the information and its impact on their consumption patterns. While subjects remembered some of the information and shifted their beliefs, they did not report changes in consumption, either in the interim period or in their intentions for the future.

The second “Butcher” experiment collects behavioral measures of consumption. We recruited $N = 2,081$ subjects of a representative sample of US consumers via an online platform, and we elicited their willingness to pay for a package of meat from a premium online butcher. We incentivized behavior by implementing the choices of randomly selected participants. In four between-subject treatments, we varied the type of meat (beef vs. poultry) and whether we provided information about the carbon emissions associated with the product in question. All conditions feature prominent mentions of the climate change impact of *some* products in order to keep the salience of climate change constant across settings and thereby isolate the effect of information on consumption that works through beliefs.

While our intervention successfully corrects misperceptions, as evidenced by an upward shift in beliefs about beef impact, we again find no change in the demand for either beef or chicken. This null result is true for all subgroups in our sample and robust among those whose beliefs responded to the intervention. Our design allows us to rule out that this result is driven by pessimism about substitute products, by meat-eaters being already well informed, by an overly noisy measure of demand, or by a non-replicable statistical fluke. We also rule out behavioral channels like an intention-action gap.

We conclude that even successful corrections of underestimations do not increase voluntary mitigation in our representative samples. This refutes the conclusions of the policy maker’s implicit argument, as well as the standard intuitions of economic decision-making embodied in our model. Our null result is very informative: our well-

powered experiments target a population that cares about avoiding emissions and a set of products for which misperceptions are large.

Our results indicate a disconnection between knowledge or beliefs and actual behavior in everyday consumption decisions. Unlike abstract elicitations of green preferences, like our measure of willingness to mitigate, the decision to purchase meat involves a large number of potentially relevant factors. Our experiments show that the size of CO₂ emissions, even when made salient, is not an important decision factor among our participants. The results suggest that policy makers should temper their enthusiasm about information provision.

These insights contribute to a literature on measuring beliefs about climate and to one about carbon labeling. Concerning the former literature, we show that addressing misperceptions or even underestimations about impact need not lead to behavioral change. Instead, the literature on carbon labeling finds small but discernible effects of labels on consumption. Our results imply that, contrary to the explanation commonly evoked in this literature, the effect of labels is unlikely to work through changing beliefs about CO₂ emissions. Instead, it may work through changing the salience of climate change or the perception of social norms, factors that were kept constant in our experiments. We now discuss these contributions in more detail.

The literature on measuring misperceptions features a number of papers that elicit broad knowledge of the climate change phenomenon and link it to measures of concern and policy support (Tobler, Visschers and Siegrist, 2012; Shi et al., 2016; Klenert et al., 2018; Dechezleprêtre et al., 2022; Fairbrother, 2022). Attari et al. (2010) find that people underestimate the energy use associated with different activities. Closest to our paper, Camilleri et al. (2019) elicit perceptions of greenhouse gas emissions associated with the production and transportation of food and the use of several electric appliances. Participants underestimate emissions for all products and activities, especially those in the food domain.

We go beyond eliciting unincentivized point estimates of carbon impact, by administering incentivized elicitations of belief distributions and combining them with revealed preferences over mitigation in a structural model.² Our approach has the potential to overturn conclusions about the optimal targeting of information that are derived solely on the basis of point estimates. Yet, we find that the predictions of our structural model are broadly in line with results in Camilleri et al. (2019) and with the results we would have obtained looking only at beliefs. Therefore, our representative survey and

²There is a large literature on willingness to pay to reduce climate impact, often using unincentivized surveys and contingent valuation methods (see Nemet and Johnson (2010) for a review) and the literature on willingness to pay to reduce emissions from specific sources like car transport (Hulshof and Mulder, 2020) or flights (Bernard, Tzamourani and Weber, 2022). Two recent studies use incentivized revealed preference techniques to elicit WTP for a single emissions amount. Löschel, Sturm and Vogt (2013) find an average WTP to buy emissions offsets for one ton of CO₂ of €12, whereas Diederich and Goeschl (2014) find a mean of €6.30. Andre et al. (2021) show that the willingness to donate to a charity to fight climate change is affected by perceived social norms.

structural model lend robustness to established results. However, our experiment also demonstrates that the presence of misperceptions does not imply that their correction yields behavioral change. More generally, our results reveal limits of survey evidence in guiding policies related to voluntary climate change mitigation. We find that economic primitives such as the valuation of carbon emissions and beliefs about their size, measured with state-of-the-art elicitation techniques, have little predictive power over consumer decisions in our experiments.

The literature on labeling studies the effect of climate labels that code high and low-impact consumption in easily digestible ways (see Taufique et al. (2022) for a summary). While most studies in this literature focus on hypothetical choices, several papers have looked at real consumption choices in the context of restaurants or university canteens, sometimes studying labels in combination with another information intervention, like posters (e.g., Spaargaren et al., 2013; Visschers and Siegrist, 2015; Brunner et al., 2018; Soregaroli et al., 2021; Lohmann et al., 2022). Other studies have provided shoppers in (online) supermarkets with informative labels about specific products or shopping baskets (Vlaeminck, Jiang and Vranken, 2014; Elofsson et al., 2016; Perino, Panzone and Swanson, 2014; Kanay et al., 2021; Bilén, 2022), or informed them via a cell phone app (Fosgaard, Pizzo and Sadoff, 2021). Most of these papers find a small and short-lived effect of labels on behavior and ultimate emissions.³ However, null results have been reported for specific products like detergents (Kortelainen, Raychaudhuri and Roussillon, 2016). Some studies look at the effect of labels on meat consumption specifically. Camilleri et al. (2019) conduct an experiment where participants were asked to purchase a can of soup. Participants were less likely to buy high-carbon-impact beef soup when a GHG impact label was present. Bilén (2022) finds suggestive evidence that when carbon labels are introduced in a supermarket, customers reduce their purchases of beef.⁴

It is instructive to compare these positive (though typically small) effects on green consumption with our null result. Our experiment is designed to study the effect of changing beliefs about CO₂ emissions on climate-friendly consumption while keeping the salience of climate change constant. Instead, the introduction of climate labels in the above studies may have yielded behavior change by increasing the salience of climate change or by changing the perceived social norms, channels we rule out. Indeed, there is recent evidence for salience channel by Schulze Tilling (2023), who shows that simply making carbon emissions salient through belief elicitation without giving additional

³Labeling has also been shown to affect energy-saving (Allcott and Taubinsky, 2015). However, it is unclear whether these results extend to CO₂ emissions since energy costs are paid by the consumer, but emissions are not. Indeed, in an experiment concerning a hypothetical choice for a water heater, Newell and Siikamäki (2014) show that CO₂ emission information is less effective in inducing sustainable choices than informing people about energy costs.

⁴Moreover, a review by Bianchi et al. (2018) finds that information can affect intentions to buy meat. Carlsson, Kataria and Lampi (2022) finds substantial resistance to switching away from meat among Swedish consumers. Jalil, Tasoff and Bustamante (2020) show that a 50-minute lecture on meat consumption reduces purchases of meat-based meals at the university canteen.

information has a similar effect to providing labels. Although she also finds a role for beliefs, her evidence points to “the direction of attention being at least as important of a channel as the correction of biased beliefs.”

Understanding the channel through which information can change behavior is more than a theoretical curiosity because it matters for policy. If beliefs are key, then education can play an important role. If the effect of information is driven by salience or perceived social norms, then policy makers will have to design interventions that effectively change the attentional, emotional, or social context at the point of purchase. We provide clean evidence that beliefs about CO₂ emissions are of second-order importance in driving green consumption. An immediate corollary of this result is that much remains to be understood about why and when information and labels actually affect consumption behavior.

2 Climate Survey

Our initial survey measures consumers’ existing beliefs about CO₂ emissions generated in the production of common consumer goods, as well as their willingness to pay (WTP) to avoid CO₂ emissions. These quantities subsequently serve as inputs for a structural model that allows us to make predictions about the provision of information, as we explain in Section 3. Figure 1 shows the four tasks that constitute the survey (see the first four tasks in “Session 1”; the other tasks will be discussed in Section 4.1 below). The first task asked general questions about climate change facts and the social cost of carbon. The next two tasks focused on eliciting beliefs, where we collect both point beliefs and belief distributions of CO₂ emissions from several common consumer products and activities. The last task elicited willingness to pay for mitigating CO₂ emissions.⁵ After participants completed all four tasks, we asked them about their demographics and revisited the products and activities from tasks two and three to ask them about their consumption frequency in these categories.

Our elicitation methods used incentive-compatible payment schemes developed in the experimental economics literature, while keeping the instructions and the interface as simple and participant-friendly as possible to allow for a representative sample to take part. Below we elaborate on each of the elicitation procedures in more detail. Online Appendix A.1.2 contains additional information about the steps we took to maximize the data quality.

⁵The survey had one additional part that we analyze in a separate paper. At the end of the survey, we provided subjects with information about the actual impact of a subset of the product list (three or six randomly selected products). We then re-invited the subjects two weeks later to test their recollection of this information.




Figure 1: Timeline of the climate survey.

Belief elicitation

At the start of the survey, we elicited participants’ beliefs about the CO₂ emissions generated by driving one mile by car. We then elicited beliefs about 12 common consumer products and activities listed in Table 1. We included food items, the use of household appliances, and transportation. We provided participants with information about the product specification and the type of emissions we considered. Table 1 presents the scientific estimates we used to incentivize the guesses together with their source.⁶ We took these estimates from top-tier academic journals or from the estimates the UK government uses for its environmental regulations. We disclose these scientific sources only at the end of the experiment.

To make the answers more meaningful to subjects, we did not elicit emissions in grams, but asked about the number of miles by car one needs to drive to emit as much CO₂ as the product in question, an approach in line with previous studies (Camilleri et al., 2019). Since we also elicited the conversion from a mile driven by a car to grams of CO₂, we can convert all measures to the perceived grams equivalent (see Table A.3 and Figure A.9 in the Online Appendix). Moreover, the model we describe in Section 3 further mitigates any concern that systematic misperceptions about the CO₂ emissions associated with driving bias our predictions, because these predictions will be independent of the denomination of CO₂ emissions.

We divided the belief elicitation into two parts. We first elicited a point estimate for the modal value of the emissions. Participants indicated how much CO₂ each of the 12 products in Table 1 emitted relative to driving one mile by car. Participants answered all 12 questions on one page, and the order of the products was randomized across participants (Figure 2A). In the rest of the experiment, the same order was used every time participants answered additional questions about these products. To help participants keep track of their guesses and the rankings of the products, we presented an interactive box summarizing their (current) answers at the bottom of the page, including the ranking of the products by estimated impact. We incentivized a correct point estimate with a \$5.36 (£4) bonus. We considered an estimate correct if it was

⁶Participants could learn the detail of what the scientific source took into account in calculating the size of CO₂ emissions. See Table A.1 in the Online Appendix.

Table 1: List of consumer products and actions.

	Quantity	Emission size		Source
		Estimate	Unit	
Beer	12 fl oz	1.46	mile	Poore and Nemecek (2018)
Phone call	1 hour	1.55	mile	Smith et al. (2013)
Microwave	1000W, 2 hour	1.76	mile	UK BEIS (2020)
Milk	1 cup	2.60	mile	Poore and Nemecek (2018)
Egg	6 eggs	4.81	mile	Poore and Nemecek (2018)
Poultry meat	7 oz	6.78	mile	Poore and Nemecek (2018)
Shower	Average usage	3.90	mile	Hackett and Gray (2009)
Dark chocolate	100g	16.03	mile	Poore and Nemecek (2018)
Coffee beans	1 lb	44.41	mile	Poore and Nemecek (2018)
Beef	7 oz	68.39	mile	Poore and Nemecek (2018)
Flight	SFO to LAX	304.60	mile	UK BEIS (2020)
Gas heating	One month	606.68	mile	Padgett et al. (2008)
Car	Drive 1 mile	291.00	gram	UK BEIS (2020)

within a 5% interval from the scientific estimate. This incentive scheme truthfully elicits the mode of the subjective probability distribution about the scientific estimate (Schlag, Tremewan and van der Weele, 2015).⁷

In order to understand the participants’ confidence in their answers, we then elicited the subjective probability distribution about the size of CO₂ emissions. For each product, we presented five “bins” around the point estimate participant reported in the first part and asked the participant to allocate 20 balls into these five bins. We told participants that each bin represents an interval that might contain the scientific estimate and that they should allocate the balls to represent their level of confidence that the estimate is in fact in that bin. Figure 2B provides an illustration. We incentivized the elicitation by randomly selecting one of the bins and scoring the answer according to a randomized quadratic scoring rule. This mechanism encourages participants to truthfully reveal their belief that the scientific estimate falls in a particular bin (Schlag and van der Weele, 2013). To keep things simple and avoid information overload, we did not provide participants with the exact details of the scoring rule, which were available with a mouse click, but told them that they would maximize their expected earnings by answering truthfully, an approach suggested by Danz, Vesterlund and Wilson (2022).

⁷We did not incentivize the questions about the CO₂ emissions and the social cost of driving one mile by a car as we realized that answers to these questions can be straightforwardly obtained on Google.

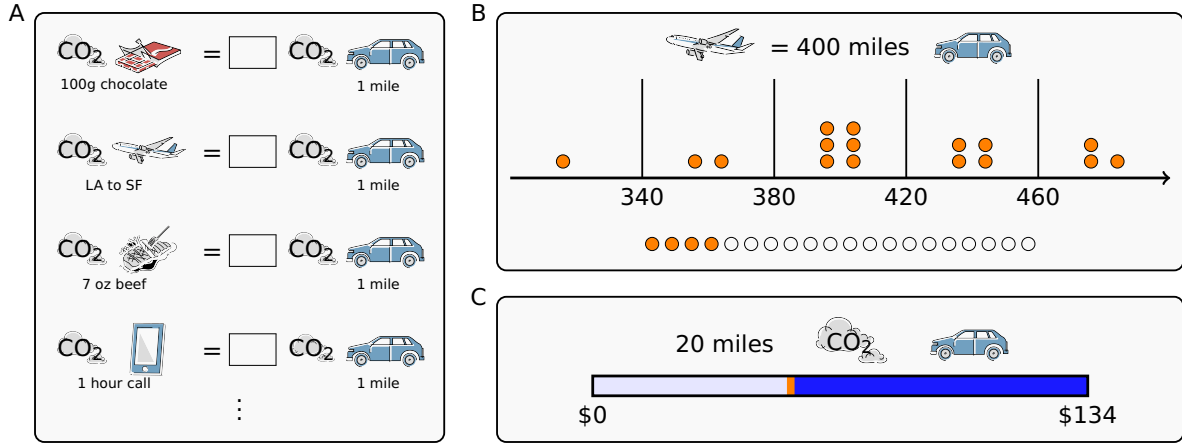


Figure 2: Illustration of the belief and WTM elicitation interface. (A) Point-belief elicitation task. (B) Bins-and-balls belief elicitation task. (C) WTM elicitation task. *Notes:* Panel B shows an example in which a participant stated 400 in the previous point belief elicitation task and is now asked to allocate 20 balls into five bins, centered around this number. See Online Appendix A.1 for screenshots of the interface.

Willingness to mitigate

After the belief tasks, we elicited the participants’ willingness to pay for mitigating CO₂ emissions of different sizes. We call this measure *willingness to mitigate (WTM)*. To introduce real consequences in the elicitation task, we offered participants trade-offs between monetary payments and carbon offset certificates. More precisely, we used donations to Carbonfund.org (<https://carbonfund.org/>), a charity that finances various projects to offset CO₂ emissions and offsets one ton of CO₂ for every \$10 donated.

To cover the amounts of the emissions generated by all the consumer products we asked in the survey, we elicited the WTM for eight levels of CO₂ emissions, corresponding to emissions generated by driving 1, 5, 20, 50, 100, 200, 450, and 700 miles by car. Participants expressed their WTM to offset these amounts of CO₂ using a slider between \$0 and \$134 (£100), see Figure 2C.⁸ The interface was designed to help participants make consistent choices and avoid anchoring. To this end, the sliders for each emission quantity were all displayed on the same screen, and the bottom of the page featured a graphical summary of reported WTMs by emission quantity (see Online Appendix A.1).

We incentivized the WTM with a Becker–DeGroot–Marschak (BDM) mechanism, which means that reporting the true WTM is in the best interest of the participant.⁹ To make sure our donations were credible to participants, we emphasized that our ethics committee does not allow misleading instructions, and promised to send them

⁸Participants could also express their WTMs either in GBP (between £0 and £100), the official currency of Prolific, or in USD (between \$0 and \$134).

⁹We randomly selected one number from a discrete set of values between 0 and 100. If the number was bigger than what the participant reported, we paid the participant a bonus equal to the randomly selected number. If, instead, the number was smaller than the participant’s report, we donated to Carbonfund.org as much money as needed to compensate for the CO₂ emissions stated in the question.

the carbon offset receipts from the experiment. The method above provides data that are censored at \$134. To mitigate this problem, we added a second, unincentivized set of questions. For every emission level for which a participant reported a WTM of \$134, we asked the participant to indicate for which amount of money he or she would have agreed to allow the emissions. The participant could either type in a number or check a box to signal that no monetary compensation would have been enough.

At the end of the session, we asked a series of questions about demographic background, consumption habits (about the 12 products), and attitudes toward climate change. See Online Appendix A.1 for the complete list of questions.

Implementation

We recruited 1,430 participants on Prolific (<https://www.prolific.co/>) between the 3rd and 6th December 2020, and 1,022 of those completed the whole survey.¹⁰ We restricted participation to US residents and we aimed to collect a sample representative for age, gender, and ethnicity.¹¹ Our sample is on average 42.7 years old (SD = 15.4) and 48.3% of the participants identified themselves as male. Table A.2 in the Online Appendix shows the demographic characteristics of the sample.

To make the instructions as accessible as possible, we used slides that displayed the instructions step by step with explanatory images complementing the written text. Besides, we divided the instructions into 5 blocks. After each block, we asked participants to answer several comprehension questions. We did not allow subjects to continue with the experiment until they answered all the questions of each block correctly. In total, participants had to answer 21 comprehension questions.

At the end of the experiment and for every participant, we randomly selected one question from the entire study. Depending on the participant’s answer to that question and luck, we paid them a bonus. This incentive mechanism elicits truthful answers in experiments with multiple tasks (Azrieli, Chambers and Healy, 2018).¹² Participants received \$10.05 for completing the study plus a variable bonus depending on their answers (mean = \$2.67, SD = 4.31).¹³ The median survey completion time was 55 minutes.

¹⁰We ran extra sessions on 21st and 22nd December 2020 to recover some participants’ demographic data. These data were not originally saved due to a failure in the survey code. We managed to retrieve the data of 67 of the 69 participants for which the failure was verified. Only the demographic questions were asked in these extra sessions.

¹¹We noticed that participants in the oldest age bracket (above 58 years old) particularly struggled with the comprehension questions about the WTM resulting in many dropouts on the page where those questions were asked. As subjects in this demographic category were hard to recruit, we opted to give them a second chance to complete the experiment. On 7th December 2020, we invited them to re-start the experiment from the WTM instructions and we gave them the solutions to two of the 7 related comprehension questions. Of the 41 subjects that were allowed to restart the experiment, 22 completed it.

¹²The probability with which a question was selected for payment was not uniform but depended on the part of the experiment that the question came from. In the instructions, we informed participants of the probability that the question was drawn from each of the different tasks of the experiment.

¹³Participants received the completion reward and the bonus only if they completed the second

Table 2: Summary statistics of elicited (point) beliefs about CO₂ emissions from 12 consumer products and activities.

Product	Emissions	Unit	Belief			
			Q1	Median	Q3	Under-est.
Beer	1.46	miles	0.50	1.20	6.00	0.516
Phone call	1.55	miles	0.40	1.00	5.00	0.549
Microwave	1.76	miles	0.80	2.15	10.00	0.406
Milk	2.60	miles	0.50	2.00	8.00	0.570
Shower	3.90	miles	0.50	1.50	5.00	0.689
Egg	4.81	miles	0.50	1.50	6.00	0.697
Poultry	6.78	miles	0.60	2.50	10.00	0.676
Chocolate	16.03	miles	0.40	1.20	8.00	0.831
Coffee	44.41	miles	0.50	2.00	10.00	0.885
Beef	68.39	miles	1.00	5.00	20.00	0.858
Flight	304.60	miles	10.00	150.00	600.00	0.586
Gas heating	606.68	miles	3.00	20.00	100.00	0.919
Car (drive 1 mile)	291.00	grams	5.03	85.00	403.00	0.677

Notes: The last column “Under-est.” shows the fraction of participants who underestimated the size of emissions.

2.1 Results

Beliefs. Participants estimated CO₂ emissions from 12 common consumer products and activities in terms of miles of driving by car. Table 2 shows summary statistics of reported (point) beliefs and Figure 3A plots them against scientific estimates of CO₂ emissions.^{14,15} Median beliefs lie below the identity line for all but one (microwave) products, indicating that participants underestimated the size of CO₂ emissions. This is in line with findings in Camilleri et al. (2019), despite differences in the sets of products, elicitation methods, and the reference items (lightbulb vs. car).

The fraction of participants who underestimated the size of emissions varies from 41% (microwave) to 92% (gas heating), with this fraction increasing in the true size of the emissions. Flying is a notable exception to this trend: it is a highly polluting activity but its emissions are underestimated only by 59% of participants. This could be due to the ample coverage of emissions from flying from media outlets, or because subjects simply took as an estimate the driving distance between San Francisco and Los Angeles (≈ 350 miles), which is close to the right answer.

part of the experiment. This second part of the experiment took place two weeks after the first. Participants that completed both parts of the experiment received a total completion reward of £10 and an average bonus of £2.20. Following the participants’ decisions in the experiment, we donated \$88 to Carbonfund.org, offsetting 8.8 tons of CO₂ emissions.

¹⁴We focus on median beliefs since there are several extreme outliers.

¹⁵Figure A.8 in the Online Appendix shows empirical CDFs of reported CO₂ emission sizes for each product.

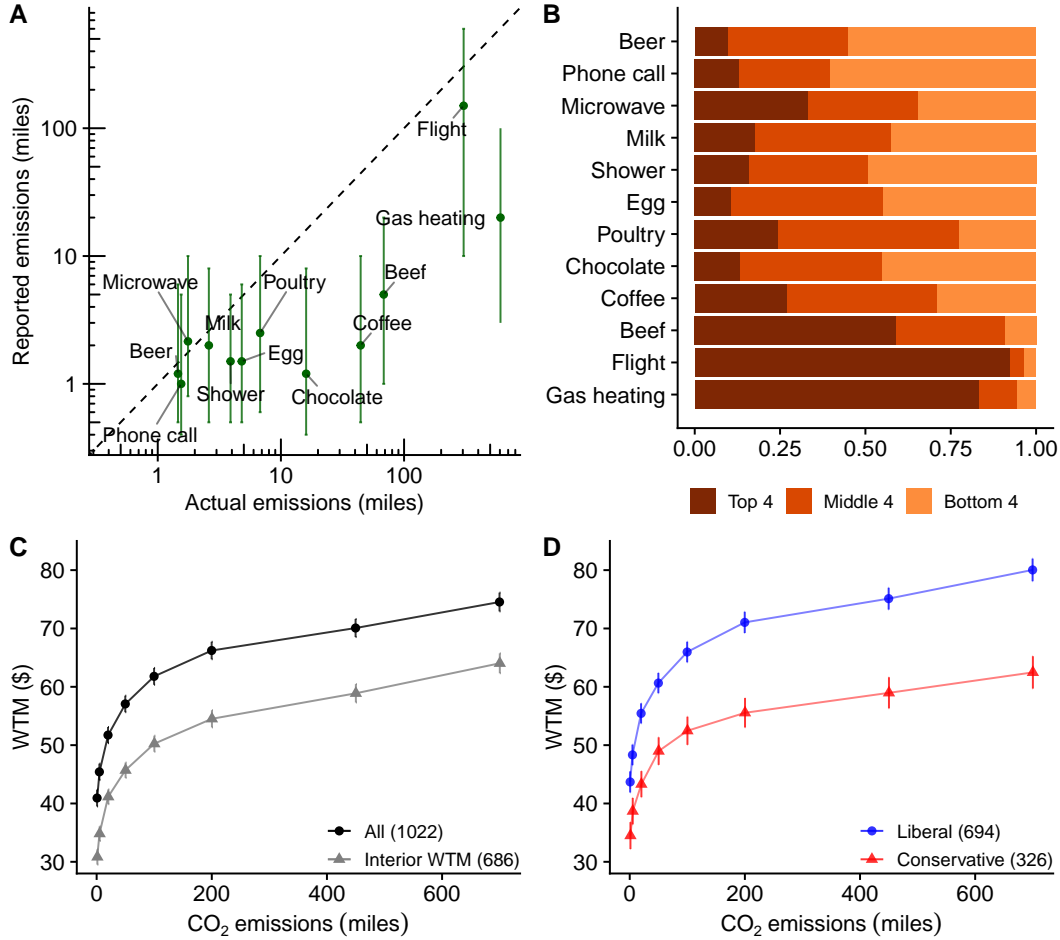


Figure 3: Beliefs and willingness to mitigate. (A) Summary statistics of reported CO₂ emissions (median and IQR). Axes are on a logarithmic scale. (B) Ranking of reported emission sizes. Products are sorted by the true emission size from low to high. (C) Concave WTM (mean and SEM). (D) WTM and political view (mean and SEM). *Notes:* In panels C and D, numbers in parentheses indicate the number of observations. In panel D, “somewhat liberal” and “somewhat conservative” are grouped into liberal and conservative, respectively.

Even though participants misperceived the size of CO₂ emissions from each product, they had a good understanding of which products emit more CO₂. As Figure 3B shows, the “true” ranking of emission sizes based on scientific estimates and the ranking “revealed” by each participant’s estimate are positively correlated.¹⁶

All the qualitative results of this section replicate if we express participants’ beliefs in terms of grams of CO₂ using their beliefs about the CO₂ emissions linked with driving one mile by car. Figure A.9 in the Online Appendix shows that, since participants underestimate the grams of CO₂ emitted when driving, the underestimation is more severe if we express the beliefs in grams.

¹⁶We also calculated Spearman’s rank-order correlation between the actual ranking of CO₂ emissions and “revealed” ranking of emissions for each participant. About 95% of the participants exhibited a positive correlation, and 45.6% of the participants exhibited a statistically significant positive correlation (two-sided, $p < 0.05$). The average correlation coefficient is $\rho = 0.559$.

Taken together, the belief elicitation tasks in the climate survey suggest that consumers significantly underestimate the size of CO₂ emissions associated with common consumer products and activities, but they have more accurate perceptions about the ordinal ranking of CO₂ emissions.

Willingness to mitigate. We now turn to participants’ willingness to mitigate CO₂ emissions. Note that we elicited WTM for eight levels of CO₂ emissions, that correspond to emissions generated by driving 1, 5, 20, 50, 100, 200, 450, and 700 miles by car. On average, participants have positive and sizable WTM for all levels of CO₂ emissions, and they exhibit a concave pattern (Figure 3C, dark line). Moving from emissions equivalent to driving 5 miles to 20 miles, a four-fold growth, increases the WTM by \$6.3 on average, while moving from 5 to 200 miles, a jump 10 times as large as the previous one, pushes the average WTM by only \$20.8. The marginal willingness to pay for mitigation decreases as the emission size increases, confirming findings in Pace and van der Weele (2020). This pattern is not due to top-censoring at \$134— the concave pattern is preserved even when we focus on 686 participants whose WTMs are all strictly between \$0 and \$134 (Figure 3C, light gray line). See Tables A.4 and A.5 in the Online Appendix for summary statistics of WTMs and the number of “corner” observations for each level of emissions.

As in elicited beliefs, we observe strong correlations between WTM and some of the demographic characteristics. Participants who identified themselves as liberal on the political spectrum have uniformly higher WTM than conservative participants (Figure 3D). Female participants have higher WTM than male participants, and participants in the age ranges of 18-37 and 58 and older have higher WTM than those between 38 and 57 years of age (Figures A.10 and A.11 in the Online Appendix).

Figure 3C shows a smooth and concave WTM curve at the aggregate level, but it masks substantial heterogeneity across participants. There are 52 participants who “do not care” about CO₂ emissions and request \$0 for all eight levels of emissions, and there are 77 participants who are “deontological” and request \$134 all the time. We can classify the shape of the WTM curve. We observe that 31% of individual-level WTM curves are concave, and 28% of WTM curves are non-monotonic. Less than 10 are convex. There are only 44 cases of decreasing WTM curve, an irrational pattern of WTM that is not captured by small mistakes. See Online Appendix A.2 for details.

In the next section, we describe how to combine these measures for the prediction of information provision.

3 Modeling the Impact of Information

In this section, we outline a simple formal framework to combine beliefs about the impact and willingness to mitigate and produce a prediction about the resulting con-

sumer decision. The key assumption is that consumers suffer a cost from the expected emissions produced by their actions and that they make utility-maximizing decisions about the quantities of emissions. Our approach is inspired by findings that subjects make rationalizable trade-offs about payoffs for themselves and others that allow for the construction of a utility function (Andreoni and Miller, 2002; Fisman, Kariv and Markovits, 2007).

Consider a consumer who gets material utility v from purchasing a good or activity. We assume that the good (or activity) is sold at a market price of p and is associated with a quantity of CO₂ emissions $c \geq 0$. The consumer’s utility from consuming the product is:

$$U(v, p, c) = v - p - w(c),$$

where $w : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ captures the psychological cost from CO₂ emissions. We assume w is strictly increasing and $w(0) = 0$.

In writing the preferences in this way, we are making two assumptions. First, for simplicity, we assume that the consumer’s overall utility is additively separable in v and in the psychological cost of emitting CO₂. Second, we assume that the psychological cost only depends on the emissions associated with the current purchase and not on the emissions linked to previous consumption of the same or other products. This last assumption finds support in our willingness to mitigate data. For us to observe the concavity of the function w , it must be the case that the consumers consider the emissions they can offset in the experiment separately from all the emissions they have generated so far. Without this “narrow bracketing” of emissions, participants with a concave WTM would report a flat WTM curve in the survey.¹⁷

We assume that the consumer may not have precise knowledge about emission sizes c , but has some beliefs about them. Let F denote her belief about c . With this subjective belief and following standard expected utility, the consumer’s preferences can be expressed as

$$U(v, p, c) = v - p - E_F[w(c)].$$

Two key ingredients in this framework are the function w capturing psychological cost and the subjective belief about CO₂ emissions F . The climate survey we discussed above is designed to measure these two quantities as precisely as possible. Remember that we used “miles driving a car” as the common unit of emission size in the belief and WTM elicitation tasks in the survey.

The WTMs stated by each participant provide information about w . Requesting a bonus of y_m to allow emitting CO₂ corresponding to emissions generated by driving m

¹⁷Narrow bracketing has also been documented in choices over monetary outcomes (Rabin and Weizsäcker, 2009; Ellis and Freeman, 2020) and in work choices (Fallucchi and Kaufmann, 2021). The concavity of the WTM function also implies that narrow bracketing is essential for an information campaign to have any effect on behavior. Given the beliefs and consumption levels of the average US consumer, broad bracketing implies that they will be on a flat part of w .

miles by a car, c_m , reveals

$$y_m = w(c_m),$$

assuming a linear utility for money. Using eight pairs of observed (c_m, y_m) and extrapolating (see Online Appendix A.3), we can recover w for each participant. Hereafter we will refer to w as the WTM function.

Similarly, we use the second part of the belief elicitation, the bins-and-balls task, to recover subjective belief *distribution* F_k for each product k . See Online Appendix A.3 for details.

Quantifying the effect of information

Given a WTM function w and a subjective belief distribution F about CO₂ emissions associated with a good or activity, we can calculate the *expected WTM*,

$$\bar{W}(w, F) = E_F[w(c)] = \int w(c) dF(c).$$

This quantity captures the extra amount of money a consumer is willing to pay in order to consume an imaginary, “carbon-neutral,” version of the good or activity, taking into account the lack of knowledge about the actual size of CO₂ emissions.

We model an *information policy* as a device that shifts consumer i ’s belief about CO₂ emissions associated with good k from F_{ik} to F_k^* , a degenerate distribution at the “true” size of CO₂ emissions.¹⁸ The difference in expected WTM before and after information for each consumer i and product k is given by

$$\Delta_{ik} = \bar{W}(w_i, F_k^*) - \bar{W}(w_i, F_{ik}).$$

If $\Delta_{ik} > 0$, information raises the psychological cost from consuming a unit of good k for consumer i through a change in her beliefs. If this increase is large enough, information may result in a change in consumer i ’s buying behavior.

Finally, we define the effect of information provision on the consumption of good k , Δ_k , as the sample average of Δ_{ik} with respect to a reference group of agents G :

$$\Delta_k = \frac{1}{|G|} \sum_{i \in G} (\bar{W}(w_i, F_k^*) - \bar{W}(w_i, F_{ik})).$$

Again, if $\Delta_k > 0$ and demand is downward sloping, then information is predicted to result in a decrease in buying behavior in target group G .

Several features of our structural model bear mentioning. First, the effect of an

¹⁸Note that we impose an assumption that the consumer trusts the information and fully updates her belief, but the framework can easily accommodate the possibility that the updated belief is not exactly F_k^* , reflecting the idea that the consumer has some doubt in the information or has difficulty in giving up her original belief.




Figure 4: Predicted effect of information provision Δ_k for each product. *Notes:* The reference group G is the entire sample of survey participants ($N = 1,022$). Bars indicate SEM.

information campaign Δ_k has a simple interpretation: providing accurate information on the CO₂ emissions of product k increases the average subjective cost of consuming product k by Δ_k dollars. Therefore, Δ_k can be thought of as the equivalent of a price increase. As with a price increase, the ultimate effect of information on consumption choices will be mediated by a product’s elasticity of demand, something we will address in the next section.

Second, because our model combines beliefs and willingness to mitigate CO₂ emissions that were both expressed as miles-driven-in-a-car equivalents, the unit of denomination of CO₂ emissions drops out of our prediction. This allows us to use an intuitive and common way of denominating CO₂ emissions while assuring that any systematic misperceptions about the climate impact of driving do not affect our predictions.

Prediction

We now calculate our measure of the effect of information provision using the data from the survey. Taking the entire sample of 1,022 participants as the reference group G , we obtain Δ_k for each product k as shown in Figure 4.

We observe a substantial variation in the effect of information provision. We expect a positive effect for five products (gas heating, beef, coffee, flight, chocolate), no effect for three products (shower, poultry, egg), and a negative effect for four products (phone call, milk, beer, microwave). Note that we expect a larger effect of information for products with larger CO₂ emissions: the ordering in Figure 4 is almost the mirror image of the ordering in Table 1. This is because the fraction of participants who underestimates the size of emissions is larger for these products, and our measure favors these participants as long as their WTM function responds to the size of emission (i.e., w is not constant on the relevant range). These predictions have received some support in the empirical literature. For instance, the negative effect for electrical appliances has




Figure 5: Predicted effect of information provision for each demographic group. *Notes:* (D) “Somewhat liberal” and “somewhat conservative” are grouped into liberal and conservative, respectively. (E) “Are you prepared to reduce your future consumption of beef/poultry in light of its CO₂ emission footprint?” (F) “How difficult would it be to reduce your current consumption of beef/poultry by half?” (G) “How many times do you eat beef/poultry per week?” Bars indicate SEM.

been documented in several empirical papers (Rodemeier and Lösschel, 2020; d’Adda, Gao and Tavoni, 2022). In a labeling intervention in an online Swedish supermarket, Bilén (2022) observes an effect for beef, but not poultry.

Taking different subgroups of participants as the reference group G , we can also quantify Δ_k depending on the target population. Figure 5 conducts such an exercise, focusing on two meat products, beef and poultry, that will be the subject of the experiment in the next section. While panel A shows the aggregate effect, panels B-G disaggregate the predictions across several subgroups. These panels illustrate the advantages of integrating preferences and beliefs over simpler approaches, like simply targeting populations with a high willingness to pay. For instance, the model predicts a larger effect for males than females (panel C), and for participants who have conservative political views than those with liberal views (panel D), despite the fact that in both cases, the former group has a lower WTM (see Figure A.10 in the Online Appendix). The reason is that these groups also have larger underestimations of climate impact, which more than offsets their lower WTM, resulting in a higher predicted impact of information.

Moreover, we can assess the robustness of our model’s prediction for beef consumption. The predicted effect of an information campaign may be interpreted as a “subjective price increase” of the product under investigation. Just like with a conventional

change in prices, a price increase will have little effect on demand if it is primarily experienced by individuals whose demand is inelastic or by individuals who do not consume the product, to begin with. Thus, one might ask whether the effect differs between groups that might have different elasticities of demand, based on self-reported consumption patterns in the survey.

Such an exercise is shown in panels E-G of Figure 5. The predicted effect of an information campaign is higher for those who are more prepared to reduce future meat consumption in light of its CO₂ emissions (panel E), those who find it “not difficult” to reduce beef consumption and hence should have more elastic demand (panel F), and those who consume beef below the median frequency (panel G). However, in each case, the effects of these splits are relatively small, illustrating that our predictions about interventions for beef are robust to prevailing demand levels and elasticities.¹⁹

4 Information Experiments

We now test the predictions we derive from our calibrated structural model in Section 3. To this end, we compare the effect of information between beef and poultry meat. There are three main reasons for choosing these two products. First, meat products are an important application, as meat (and especially beef) consumption makes a meaningful contribution to climate change and is one of the main sources of emissions that are under the direct control of consumers.²⁰ Second, these two products are comparable in many respects as they fall into the same food category and may be considered substitutes for certain purposes. Third, despite their similarity, these two products have very different predicted effects of information provision, as we show in Figure 4. While the predicted effect of information on beef consumption is among the very highest on our product list, it is approximately zero for poultry. This is mainly because beef production is about 10 times more carbon-intensive than poultry production, an effect that is not incorporated into the expectations of consumers, and hence subject to correction through information provision.²¹ Thus, the main hypothesis that we test in our experiments is that information provision about carbon impact will have a bigger impact on consumer

¹⁹The prediction is based on the consumption of 7 oz of beef and poultry, the size of meat products participants reported their beliefs about CO₂ emissions. Figure A.13 in the Online Appendix shows the prediction about 5 lb (80 oz) of beef and poultry, the size of meat products offered to participants in the Butcher Experiment, by “scaling up” their belief distributions by the factor of 80/7, which shifts Δ_k upward for both products. The overall prediction is different in absolute terms (e.g., the bottom panel of Figure A.13 shows a positive overall effect of information even for poultry), but qualitatively the results do not change: the model still predicts larger effects of information for beef.

²⁰Alexandre Koberle (Grantham Institute for Climate Change, Imperial College London) writes that “Next to flying less, it is probably right to say that, as individuals, reducing beef consumption is the most significant contribution directly under our control” (Vetter, 2020).

²¹This difference results mainly because beef involves the release of large amounts of methane, a greenhouse gas with about 30 times the warming equivalent of CO₂, and because beef requires large amounts of feed which spurs deforestation (Poore and Nemecek, 2018).

demand for beef products than for chicken products.

The first experiment takes part in the context of the climate survey, where we informed some subjects about the true impact, and measured changes in self-reported consumption two weeks later. The second experiment involves a new sample, where we study an incentivized choice to buy meat from an online butcher. We see our two experiments as complementary. The first tells us about the effect of the information on the same subject pool we used to generate our predictions, and is informative about belief recall and subsequent consumption. The second experiment tests our predictions out of sample, involves actual choices, and allows exploration of some of the mechanisms.

4.1 Survey Experiment

This experiment took place with the participants of the climate survey described in Section 2. The first part of the experiment (“Session 1”) took place at the end of the survey. The second part (“Session 2”) took place about two weeks later (between the 16th and 21st of December 2020), when we re-contacted the participants for another round of the survey. Figure 1 shows the timeline of tasks in Session 2, which we now discuss in more detail.

Session 1. After the WTM elicitation, participants received information about the emissions associated with a few products randomly selected from the 12 products in our survey.²² The information consisted of the latest scientific estimate for the carbon impact of the product, and featured a link to the source of the information. Our treatments are thus on the subject-product pair level: if a subject is informed about a particular product’s emissions, the pair is in the *Info* treatment, if not it is in the *No Info* treatment. Table A.6 in the Online Appendix shows information about the number of observations in each treatment, split by product.

To assure subjects engaged with the information, they first had to repeat it back to us (Figure A.5 in the Online Appendix). We then presented them with a surprise short-term memory task, where they had to repeat the information again and each correct answer earned additional £0.20 bonus (Figure A.6 in the Online Appendix). Finally, we asked several questions about the participants’ consumption of the products we used in the survey. Among these questions, we asked if the participants plan to reduce their consumption of each of the products due to their CO₂ emissions.

Session 2. We called back participants 13 to 18 days after the day of the first session. Of our original 1,022 participants, 946 (92.5%) completed the second session. At the

²²We implemented three different conditions: 322 subjects received information about 6 of the 12 products we used in the survey, 344 received only about 3 of the products, and the remaining 356 subjects received no information. We designed these different information treatments to understand questions related to recall and information overload, but we will not analyze those issues in this paper.

start of the survey, we elicited participants' beliefs about the emissions of the 12 products, some of which they may have been informed about in Session 1. We followed the exact same procedures as in the first survey, eliciting both point estimate and subjective probability distribution. We preserved the order in which the products were presented across the two sessions. At the end of the survey, every participant received an overview with information about the CO₂ emissions of every product used in the experiment. Finally, we asked subjects whether they changed their consumption of any of the products in light of the carbon impact.

Results

To understand if the information still affected beliefs two weeks later, and as a check on whether our treatment worked, Figure 6 shows CDFs of elicited beliefs in both sessions, separated by whether they received information about the products. As a randomization check, Panel A and B show that beliefs about beef and poultry do not differ with statistical significance in Session 1, before any subject received information (two-sample Kolmogorov-Smirnov test; beef, $p = 0.151$; poultry, $p = 0.810$). Panels C and D show beliefs have shifted and underestimation is reduced (two-sample Kolmogorov-Smirnov test; beef, $p < 0.001$; poultry, $p = 0.004$), although only a small minority remembers the actual value and a large majority still underestimates the impact of beef.

Next, we ask whether the information given in Session 1 affected intentions to reduce meat consumption, stated immediately after receiving the information, or the consumption in the interim period. Our hypothesis, following the predictions in Figure 4, was that we would see larger reductions for beef than for poultry. However, as Figure 7 shows, there is no discernible difference between the treatment and the control condition for beef. For reasons of space, equivalent graphs for poultry can be found in the Online Appendix (Figure A.14): although about fewer subjects report to have reduced their poultry consumption (10 percentage points compared to beef), there is again no difference between *Info* and *No Info*. Indeed, a logistic regression that regresses intentions or behavior on a dummy for the *Info* treatment finds no significant effects - as illustrated by the estimated marginal effects in Figure 8 (shaded grey). These results are robust if we focus on only those subjects who updated their beliefs in the right direction between the two sessions, and hence engaged demonstrably with the information (see Figure A.15 in the Online Appendix).

Thus, we conclude that belief changes did not translate into intentions to reduce the consumption of either meat product, nor into reductions in actual consumption. This evidence is all the more striking since behavior is self-reported, so it would have been cheap for subjects to report a socially desirable change.




Figure 6: Impact of information on beliefs about the impact of beef and poultry. *Notes:* Panels A and B show beliefs from the original survey in Session 1, before information was given. Panels C and D show beliefs in Session 2. Vertical lines correspond to the “true” size of CO₂ emissions (6.78 miles for poultry and 68.39 miles for beef). Numbers in parentheses indicate the number of observations.




Figure 7: Impact of information on intentions and consumption across treatments. *Notes:* (A) Intentions to reduce beef consumption reported in Session 1, after information provision in the *Info* treatment. (B) Actual consumption changes in Session 2.

Effect on non-meat products. Moving beyond meat, we ask if there is any effect of information on the consumption of the remaining products. Here, the predictions for each product differ quite sharply according to our model, with big effects predicted for gas heating, and negative effects for microwave (see Figure 4). To evaluate the effect of information for each product, we repeat the logistic regression of either the intentions to reduce consumption or actual consumption on a dummy for the *Info* treatment. Figure 8 ranks the product by the predicted impact of information (highest on the left), and shows marginal effects for the resulting coefficient estimates. There are no statistically significant effects for any of the products, with the exception of Beer, but




Figure 8: The effect of information on the intended and actual reduction of consumption. *Notes:* Logistic regression was run on each product separately. The dependent variable is a binary indicator “intend to reduce/actually reduced consumption of [product]” and the independent variable is the indicator “received information about CO₂ emissions from [product]”. Products are ordered (descendingly) by the predicted effect of information as in Figure 4. Bars indicate 95% CI.

this latter effect does not survive a correction for multiple hypothesis testing. We conclude that information does not have an effect on self-reported consumption patterns in our experiment.

4.2 Butcher Experiment

In this experiment, our goal is to provide evidence on the role of information in an actual consumption decision. We offered participants an opportunity to purchase a bundle of high-quality meat products, either 10 beef sirloin steaks or 10 skinless chicken breasts. We kept the features of bundles as close as possible: they were sold on a premium online butcher Porter Road (<https://porterroad.com/>); they weighed about 5 lb (≈ 2.3 kg); they cost \$100 (at the time of designing the experiment in 2021); they were pasture-raised in the US without hormones and antibiotics. We provided these descriptions in the relevant part of the instructions.

Across treatments, we varied between subjects whether the participants received information on the CO₂ emissions associated with beef and poultry meat (Info treatment) or not (NoInfo treatment). In keeping with our climate survey, we provided the information in terms of the number of miles by car one needs to drive to emit as much as 1 lb of the meat. We pinned down participants’ beliefs about the car CO₂ emissions by including a scientific estimate of these emissions (in ounces) in the instructions. In this way, we made sure that our information treatment could only impact the beliefs about the meat. The information about car emissions was available in all treatments.

As an additional manipulation, we varied whether the participants were first offered




Figure 9: Timeline of the butcher experiment.

the beef bundle (BeefFirst treatment) or the poultry bundle (PoultryFirst treatment). For these two products, subjects remained in the same information treatment. We test our main hypothesis about the differential impact of information for the two products using the first product offered in the experiment. The second part allows us to evaluate spill-over effects, whereby information about beef affects beliefs and WTP for poultry or vice versa.

The timeline of the experiment is illustrated in Figure 9. The experiment has two parts, one per each of the products we offer. The two parts followed the same structure. Each part of the experiment started with a description of the bundle the participants could purchase as well as its retail value (\$100). We then asked the participants to guess the average CO₂ emissions associated with the production and distribution of 1 lb of the type of meat that they were offered. As in the climate survey, participants expressed their guesses in terms of CO₂ emitted by driving one mile by car.²³

To help participants to get a sense of the magnitudes of emissions, just before they could express their guesses, we informed them of how the CO₂ emissions from driving one mile by car compared with the emissions generated by the production and distribution of 12 fl oz of beer and by taking a plane from Los Angeles to San Francisco. We provided this baseline information to all participants to keep the salience of emissions and possible norms around low-carbon consumption constant across treatments. To incentivize belief elicitation, we used the same sources of scientific estimates as in the climate survey and we rewarded accurate guesses (those within $\pm 5\%$ of the scientific estimate) with a \$0.5 bonus.

Next, we had our treatment manipulation. The participants in the Info treatments were informed about the average emissions associated with the meat product they could purchase. To make sure that the participants paid attention to the information, we asked them to identify the true size of the emissions among three possible options. The participants in the NoInfo treatments, instead, saw three random numbers and answered a similar question.²⁴

²³We did not elicit belief distributions to fit the survey in the time constraint of 15-20 minutes.

²⁴In both treatments, participants were allowed to proceed regardless of their answers. However, participants who answered incorrectly received an alert warning them of the mistake and repeating the correct answer.

We then elicited participants’ WTP using a two-stage multiple price list (MPL) with forced single switching.²⁵ On the first list, participants saw 11 choices between two options: the left option is the meat bundle and the right option is the monetary bonus ranging from \$0 to \$100 in \$10 increment. In the remainder, we refer to this bonus as the “price”, although it was not framed as such in the experiment. The second list “zoomed in” around the switching point and asked another nine questions. With this procedure, we measured WTP in the precision of \$1.²⁶ The instructions encouraged the participants to think about their own valuation of the meat bundle and to use this valuation to make the decisions.

After completing the MPL task, we asked participants to guess one more time the size of the emissions associated with the meat product they had the opportunity to purchase. This second guess was not incentivized.

The second part of the experiment followed the same structure as the first one, but it asked participants about their beliefs and WTP for the other meat bundle—the poultry bundle if the first part was about beef, and the beef bundle if the first part was about poultry. Thus, in the Info treatments, participants saw the information about the CO₂ emissions associated with the new meat bundle together with all the information previously provided. In the NoInfo treatment, instead, participants saw four randomly generated numbers.

At the end of the experiment, we asked the participants about their meat consumption patterns, attitudes toward climate change, and trust in the experimenters. We also asked for their contact information (both home address and email) to deliver the meat product or the monetary bonus, if any.

Implementation

We recruited participants on the platform Lucid between 31st March 2022 and 15th April 2022.²⁷ We focused on participants who consume meat and excluded those who lived outside contiguous US states due to shipment requirements by Porter Road.²⁸

²⁵We used an MPL instead of the slider interface from the climate survey since we elicited only two valuations in this experiment while we elicited eight in the survey. The small number of valuations makes an elicitation strategy that requires simpler instructions (MPL) preferable to a strategy that requires more complicated instructions but allows the participants to input their decisions more quickly.

²⁶We used a BDM procedure to make this two-stage MPL incentive compatible. We randomly selected a price (an integer) between 1 and 100 to determine whether the participant receives the monetary reward or the meat bundle. Each price has the same chance of being extracted *independently* of the participant’s choice in the first multiple price list. If the randomly selected price was not the one the participant had seen, we inferred his or her choice for this price from the choices for the other price levels. This strategy was feasible because we forced a single switching and hence we enforced consistency in choices.

²⁷Lucid was acquired by Cint (<https://www.cint.com>) in January 2022, but still operated under the old name at the time of our experiment.

²⁸To enhance data quality, we included five attention checks and three comprehension questions about the instructions. Participants were excluded if they failed any of the attention checks or if they needed more than five attempts to answer the comprehension questions correctly.

2,081 participants satisfied the pre-registered inclusion criteria: 1,047 were assigned to the NoInfo treatment and 1,034 were assigned to the Info treatment.²⁹ Participants are representative along gender and age. Table B.1 in the Online Appendix shows that demographic characteristics are balanced across treatments. Our sample is on average 46.8 years old ($SD = 17.1$) and 48.4% of the participants identified themselves as male. The median survey completion time was 17 minutes.

We implemented one of the two MPL decisions for one in every 20 participants and delivered the meat bundle (beef or poultry, depending on the selected MPL) or the monetary bonus, based on the participant’s choice for the randomly selected price level. Finally, one (lucky) participant received a \$500 completion reward. All bonus amounts were paid using Amazon gift cards. We preregistered our hypotheses and sample sizes on Aspredicted.org, the preregistration is available in the Online Appendix B.2.

Results

Following our preregistration, we focus on the belief and WTP data from the first part of the experiment for a clean analysis of the treatment effect. This means that belief and WTP data about the beef bundle come from BeefFirst treatments ($N = 1,048$) and the data about the poultry bundle come from PoultryFirst treatments ($N = 1,033$).

As in the climate survey discussed in Section 2, participants exhibited a significant underestimation of the size of CO₂ emissions from beef and poultry. Figure 10 shows that the magnitude and the prevalence of underestimation are more significant in the experiment as compared to the survey—median beliefs are much lower in the experiment (even though the quantity of meat products presented to the participants was more than twice as large as the quantity used in the survey) and the fraction of participants who underestimated the emission size was 92.7% for beef and 89.4% for poultry, respectively. Like in our survey, we see a large difference in the size of underestimation between the two products: the absolute level of underestimation for the median subject is 153 miles for beef and 14.4 miles for poultry, respectively.

Participants were initially equally uninformed about CO₂ emissions across treatments. The distributions of prior beliefs (asked before WTP) show no differences between Info and NoInfo treatments for both meat products (Figure 11AB). Providing information successfully shifted the beliefs of many participants in the treated groups, as evident in jumps in the distributions of posterior beliefs (asked after WTP), illustrated in Figure 11CD. In particular, 64.8% (337/520) of participants moved their beliefs to the correct value for beef, and 51.0% (262/514) did so for poultry.³⁰

Remember that our model in Section 3 predicts that information has a positive im-

²⁹Number of participants in each treatment is: 520 in the BeefFirst, Info treatment, 528 in the BeefFirst, NoInfo treatment, 514 in the PoultryFirst, Info treatment, 519 in the PoultryFirst, NoInfo treatment.

³⁰If we allow a margin of $\pm 10\%$, the number increases to 68.3% (351/514) for poultry.




Figure 10: Empirical CDFs of beliefs about CO₂ emissions from two samples. (A) Beef. (B) Poultry. *Notes:* The size of meat products for belief elicitation was 7 oz in the climate survey and 1 lb (= 16 oz) in the butcher experiment. For the data from the butcher experiment, we focus on belief data from the first elicitation in the first part of the experiment. Vertical dashed lines correspond to the “true” size of CO₂ emissions (A: 155 miles for the butcher experiment and 68.39 miles for the climate survey; B: 15.4 miles for the butcher experiment and 6.78 miles for the climate survey).

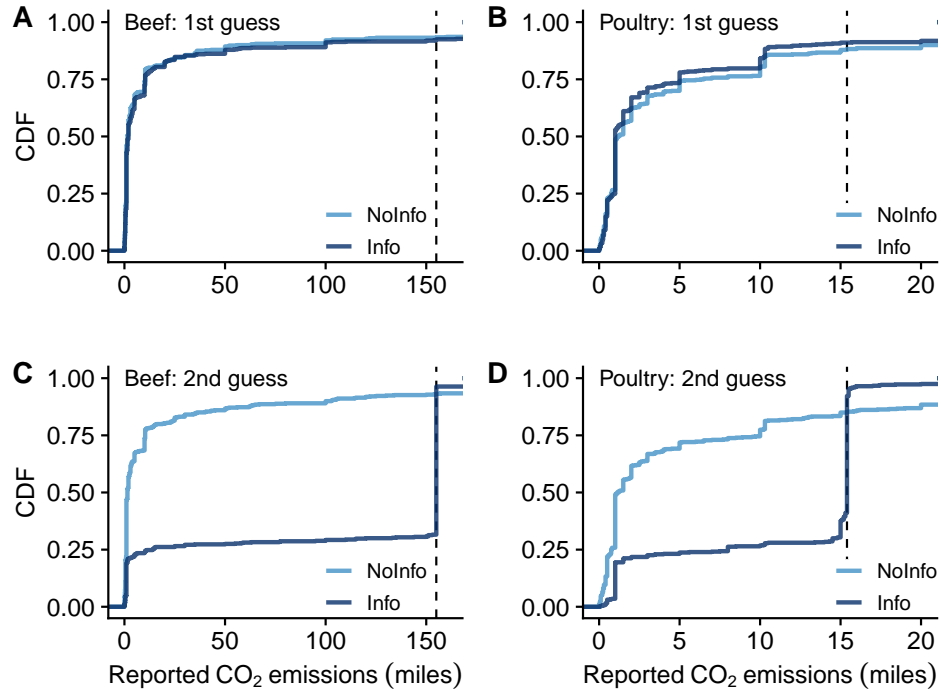


Figure 11: Beliefs about CO₂ emissions from two meat products. *Notes:* We focus on the data from the first part of the experiment (panels AC: BeefFirst treatments; panels BD: PoultryFirst treatments). Vertical lines correspond to the “true” size of CO₂ emissions (15.4 miles for poultry and 155 miles for beef).

pact in the direction of reducing the demand for beef but has no impact on the valuation of poultry. In the experiment, these predictions are translated into a *decrease* in average WTP for the beef bundle and no effect for the poultry bundle. These predictions are not supported in the data. Figure 12A shows the WTP for meat products by treatment. If anything, there is a small *upward* movement in the valuation of the beef package after




Figure 12: (A) Average willingness to pay for the first meat product. (BC) The proportion of participants buying the meat product at each price. *Notes:* We focus on the data from the first part of the experiment. In panel A, Bars indicate SEM. Figure B.4 in the Online Appendix shows the CDFs of WTPs.

information provision. Average WTPs are not significantly different between treatments for both products (beef: $t(1046) = -1.200$, $p = 0.230$; poultry: $t(1031) = 0.938$, $p = 0.349$). Panels B and C of Figure 12 give a more complete overview of demand and show the proportion of buyers for each price, confirming that there is no discernible difference between the treatments.

Table 3, column (1) shows the effect of information on beef valuation in a regression analysis. This “null” finding is robust to the inclusion of several control variables in the regression (Table B.2 and Figure B.6 in the Online Appendix). Several of those covariates have sensible signs: we find a higher WTP for beef for those subjects who report above-average beef consumption, or who report that it is difficult to reduce beef consumption. We also find a lower WTP for both beef and poultry amongst women and younger individuals. In the Online Appendix Figure B.6, we also conduct an analysis of the treatment effect by subgroup. For all subgroups, we cannot reject the null hypothesis that the information effect for beef is zero.

Finally, to replicate our result of the first survey experiment, we look at participants’ stated intentions about future consumption in the post-experimental questionnaire. At the end of the experiment, we asked participants “Do you intend to reduce your beef/poultry consumption in light of its CO₂ emissions?” and they answered on a Likert scale from 1 to 5. A chi-squared test of independence shows no differences in response distribution between Info and NoInfo treatments for intention to reduce beef (Figure 13; $\chi^2(4) = 0.964$, $p = 0.915$). For extra power, we pool the datasets from our this experiments and the survey experiment on consumption intentions and run a similar logistic regression with an information dummy and an experiment dummy. The results, reported in Table B.3 in the Online Appendix, show no effect of information on intentions.

In summary, all our results indicate that there is no connection between beliefs about




Figure 13: Distribution of responses to a survey question: “Do you intend to reduce your beef consumption in light of its CO₂ emissions?” *Notes:* We focus on 1,013 participants in the BeefFirst treatments. Participants responded on a 5-point Likert scale (1: “No.” 2: “Yes, I am prepared to reduce my current consumption by about 10%.” ... 5: “Yes, I am prepared to reduce my current consumption by more than 50%.”).

emissions size and consumption decisions.

4.3 Interpretation of the Null Effect in the Butcher Experiment

We now turn to investigate possible reasons for the observed null effect of information about CO₂ emissions on the demand for meat. We focus on beef, where we predicted that information should affect willingness to pay negatively and decisively, and on the Butcher experiment, where we elicited additional variables to interpret a potential null effect.

Were participants’ beliefs insensitive to the information treatment? In both the Info and the NoInfo treatments, we measure beliefs twice (Figure 9). In the Info treatment, the second belief, or posterior, is measured after information about beef consumption is provided. Participant’s posterior is affected by the information treatment and exhibits, on average, less optimism about CO₂ emissions (see Figure 11CD). This shows that participants’ beliefs were changed by the information they saw. However, these belief changes do not translate into differences in WTP. Column (2) of Table 3 shows regression results of WTP on a dummy for the Info treatment, including only participants in the latter treatment who responded to information by updating their beliefs upward. While the coefficient on the Info treatment declines relative to the full sample (column (1)), the null effect remains.

Did participants become more pessimistic about other meat products? It is possible that information about beef made participants more pessimistic about other meat products. This would limit the options for (low carbon) substitution, rendering demand for beef inelastic in information. We can address this point in several ways.

Table 3: Interpretation of the null effect of information on WTP for beef.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Info	2.743 (2.285)	0.528 (2.448)	2.914 (2.400)	1.673 (2.847)	5.277 (3.284)		4.835 (3.191)	-0.860 (2.292)
Belief (poultry)			-0.011 (0.012)					
Difficulty (beef)						3.074*** (1.004)		
Constant	32.225*** (1.590)	32.225*** (1.590)	33.018*** (1.660)	32.912*** (1.984)	34.574*** (2.263)	25.580*** (2.942)	35.062*** (2.162)	33.080*** (1.616)
First product	Beef	Beef	Beef	Beef	Beef	Beef	Beef	Poultry
Observations	1,048	901	1,032	672	529	991	579	1,013
R ²	0.001	0.0001	0.002	0.001	0.005	0.010	0.004	0.0001

Notes: The dependent variable is WTP for beef. Samples are as follows. (1) All participants in the BeefFirst treatments. (2) Participants in the NoInfo treatment, and those in the Info treatment who responded to information by updating their beliefs upward. (3) Participants who completed both parts of the experiment. One subject who reported an extremely large belief ($\approx 1.46 \times 10^9$) about CO₂ emissions from poultry is excluded. (4) Participants in the BeefFirst treatments who self-proclaimed to care about the environment (based on the response to the question "How severe do you consider the problem of climate change?"). (5) Participants in the BeefFirst treatments who self-reported consuming beef at least three times per week. (6) Participants in the BeefFirst treatments who self-reported consuming beef. (7) Participants in the BeefFirst treatments who expressed trust in us actually sending meat. (8) All participants in the PoultryFirst treatments. Robust standard errors are reported in parentheses. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

The first is to directly control for this spillover in beliefs. In the BeefFirst treatment, we measure participants' beliefs about the CO₂ emissions associated with poultry after the participants received information about and stated their willingness to pay for beef. We find that participants do indeed become much more pessimistic about poultry after receiving information about beef. About 63% of the participants in the BeefFirst, Info treatment (317/505) overestimated the size of CO₂ emissions from poultry (reported numbers above 15.4 miles) and 48 subjects reported 155 miles, which is exactly the size of CO₂ emissions from beef they learned about in the first part of the experiment (see Figure B.5 in the Online Appendix). However, this updating about a substitute product does not appear to be an important mediator of the information effect on beef demand: the null effect persists after controlling for the beliefs associated with poultry consumption (Table 3, column (3)).

In addition, we can look at the case where beef is the second product participants can buy. Here, by the time participants state their willingness to pay for beef in the Info treatment, they have received information on both poultry and beef. This group is therefore aware of a climate-friendly substitute. However, we find no treatment effect in the second product either (Table 3, column (8)).

Finally, this possible reason behind the null result is undermined by the analysis of the stated future consumption intentions in the *Info* treatment. By the time the participants stated these intentions, they had received information about both beef and poultry, and hence knew that poultry was a cleaner substitute. Yet information does not change their intended behavior.

Preaching to the choir. One reason that information may have little impact on CO₂ emissions is that prior optimism about CO₂ emissions is concentrated among individuals who have little willingness to mitigate. The info treatment would then correct the beliefs of only those who have no interest in mitigation, and speak to those who are already well-informed. Our structural model was explicitly designed to make predictions that take this mismatch into account, so our initial predictions, based on the representative survey, are not subject to this concern. To see whether these concerns could matter in the experiment, we can restrict our analysis to those participants who self-proclaim to care about the environment. The null effect persists in this restricted sample (Table 3, column (4)).

Are near-vegetarian driving the results? If only near-vegetarians are optimistic about the CO₂ emissions associated with meat consumption, then providing this information will do little to curb the demand for meat. Of course, this state of affairs is ex-ante implausible, but, for the sake of completeness, we can provide an explicit test for this hypothesis by restricting our dataset to participants who consume meat at least three times per week (i.e., above-median frequency). The null effect persists in this

restricted sample (Table 3, column (5)).

Do participants suffer from an intention-action gap? An intention-action gap would manifest itself as a stated intention to reduce meat consumption in the future, but a failure to do so in the present. The underlying reason could be a preference for immediate gratification or a self-control problem. Yet, as we reported above, and as Figure 13 shows, intentions to reduce beef consumption are not much affected by the information treatment. Thus, the null effect of information does not stem from a failure to implement virtuous plans, but from a failure to make such plans, to begin with.

Does the information cause participants to decrease their consumption of lower-quality meat outside of the experiment? A key challenge of our experimental setup is to sell participants a product that they find appealing. To this end, we used high-quality meat. But this may invite the concern that participants respond to information by demanding less, but better, meat. If this were the case, then the information treatment may decrease average meat consumption, but not the willingness to pay for the meat we sell to participants. Again, the fact that information does not impact participants' stated intention to consume beef rules out this conjecture.

What if people chose directly between beef and poultry? Many shopping environments offer a choice between variations of comparable products. Information about CO₂ emissions might be effective in such environments as it makes the low-emission alternative salient relative to the other products, and hence more attractive. While this logic depends on changes in the relative valuations of the two products that our experiment is designed to detect, we do not test substitution directly.

However, we can test the impact of information about substitute products, by looking at valuations of product 2 in the information treatment. In particular, when learning the emissions associated with product 2, subjects already know about the emissions of the potential substitute. This should cause an (additional) negative effect of information on beef demand (as a cleaner alternative is known to be available), and a positive effect of information on poultry demand (as the alternative is more polluting).³¹

To evaluate these hypotheses, we regress the WTP for both meat products separately on a dummy for Info, timing ("Second product"), and the interaction between these variables. Following our reasoning above, the coefficient for the interaction between *Second product* and *Info* should be positive for poultry and negative for beef. Table 4 shows the results of these regressions. The estimated coefficients are in the hypothesized direction, but not (close to) statistical significance. Thus, we cannot rule out that knowledge about substitute products induces a shift in valuation, but the effect, if

³¹Note that effect of the timing on beef valuations should be smaller than that for poultry, because subjects' beliefs about poultry are already relatively well calibrated in the absence of information.

Table 4: The effect of knowledge about the emissions of substitute products.

	(1)	(2)
	WTP (poultry)	WTP (beef)
Info	-2.028 (2.162)	2.743 (2.285)
Second product	-0.202 (2.189)	0.855 (2.267)
Info \times Second product	3.755 (3.100)	-3.603 (3.237)
Constant	29.517*** (1.544)	32.225*** (1.590)
Observations	2,054	2,061
R^2	0.001	0.001

Notes: The dependent variable is WTP for poultry (1) or beef (2). The “Second product” dummy takes 1 if the WTP for poultry (beef) is measured in the BeefFirst (PoultryFirst) treatment. Robust standard errors are reported in parentheses. *: $p < 0.1$; **: $p < 0.05$; ***: $p < 0.01$.

it exists, is relatively small. The role of substitution in information provision is an important question for future research.

Do participants react to information not by demanding less beef, but by off-setting the CO₂ emissions of their consumption outside of the experiment?

We deem this hypothesis unlikely. It requires individuals to care about mitigating CO₂, to take into account and feel the pain of their meat consumption emitting CO₂, but to be completely inelastic in their meat consumption. Empirically the price elasticity of demand for beef steaks in the US is between -0.42 and -0.52 , making beef demand far from inelastic (Dong, Davis and Stewart, 2015). So if learning about the CO₂ emissions increases the subjective cost of buying meat, it seems unlikely that participants do not use the rather elastic margin of adjustment that is a decrease in the WTP for meat, and instead adjust only buy purchasing offsets outside of the experiment.

Is reducing meat consumption particularly hard?

It could be that people find it impossible to reduce their consumption while they would be happy to reduce the consumption of other products. Two observations are inconsistent with this hypothesis. First, as just discussed, people are able to eat less meat when its price goes up. Second, the survey experiment presented in Section 4.1 shows no significant effect of information for a large and heterogeneous set of products.

Does our willingness to pay measure suffer from noise, misinterpretation, or lack of trust?

A possible reason for a null effect of the information treatment may be that our measure of demand is very noisy. If our WTP measure is a very poor

proxy for actual demand, then it would follow that this measure does not necessarily change with new information, even if this information would have had an impact on participants' actual demand for meat. To shed some light on this possible reason for a null effect, we ask whether our willingness to pay measure is correlated with other measures of preferences for meat. This would not be the case if WTP was very noisily measured. We find that WTP for beef is significantly correlated with participants' self-reported difficulty in reducing beef consumption if they had to (Table 3, column (6)).

A related worry may be that despite our elaborate efforts to be credible, (some) participants did not believe us that we would actually send them the meat they purchased with positive probability. Then, what they answered in the willingness to pay elicitation may not reflect their sincere demand for beef. To test this hypothesis we ask whether there was a treatment effect among those who expressed a lot of trust in us actually sending meat in the post-experimental survey.³² The null effect persists in this restricted sample (Table 3, column (7)).

A final, somewhat related, concern is that the participants misunderstood our WTP question and thought they had to indicate the (socially) fair price for the beef shipment. This misunderstanding could generate a null result if some participants in the Info treatment thought that the fair price should be higher due to the high emissions.

Several considerations assure that this misunderstanding is unlikely. First, the word "price" did not appear in the experiment: subjects made a sequence of binary buying decisions from which we infer a WTP. Second, we advised the participants explicitly to use their valuation of the meat to make their decisions. Third, the instructions did not contain any reference to CO₂ offsets or to other environmental actions associated with the product (and indeed there was no such offset), so there was no reason to pay more out of fairness concerns. Finally, if the information made participants think that the fair price is higher, we should find that information reduces the intention to consume beef. However, as we discussed above, we do not find evidence for this treatment effect.

Was the null effect a fluke? Even relatively well-powered studies may sometimes result in erroneous null effects. Three results speak against this hypothesis. First, we can ask whether there is any correlational evidence that beliefs about CO₂ are predictive of the willingness to pay for meat. While any such evidence is subject to the usual caveats and endogeneity concerns, a strong negative correlation between beliefs about CO₂ emissions and WTP in the NoInfo treatment should give us pause in interpreting the null effect of the info treatment. We find that prior beliefs in the NoInfo treatment do not correlate with meat consumption.

Second, we can use the comparison of the Info and NoInfo treatments when beef was offered in the second part as a replication experiment. Of course, because these data

³²Participants responded to the question "Do you trust that the researchers will indeed ship meat products as described in the instructions?" on a 5-point Likert scale (1: not at all; 5: completely).

stem from Part 2 of the experiment, the treatment comparison is less tightly controlled, with information about poultry possibly also bearing on participants' willingness to pay for beef. At the same time, it is hard to construct an explanation of how this additional information would lead to a null effect. We find that experiment 2 also features null effects of the information treatment.

Third, the lower bound of the 95% confidence interval for the effect of information on the willingness to pay for beef is $-\$1.74$. Hence, even if the information has an effect that we are not powered to detect, this effect is likely less than 2% of the market price of the meat.

Finally, and as we have already seen, the information does not affect participants' stated intention to reduce meat consumption.

What, then, causes the null effect? Having ruled out several possible explanations for the observed null effect, we are led to conclude that people's decision to eat meat appears not to be subject to concerns about the associated CO₂ emissions. That is, even though we see that people are willing to invest in emission reduction when this willingness is elicited directly, their desire to curb emissions in meat consumption appears to be drowned out by the many other considerations that go into their consumption decision. If this is the reason behind the null effect, then we should be no more optimistic about finding an effect of information in still "wilder" settings. After all, we made sure that our information actually moved beliefs and we can be confident the climate impact of various consumption activities was a salient feature of the decision making environment.

5 Conclusion

We have used incentivized survey techniques to elicit both beliefs about the carbon impact of consumer products and the valuation of this impact. We find that most consumers underestimate the impact, but heterogeneity is large. While they are willing to pay to offset carbon emissions, this willingness is highly concave and varies by subgroups. We use these inputs in a simple structural model to predict the impact of information. In two experimental tests, we find no support for our predictions: despite a correction in the beliefs about beef meat, subjects are unresponsive in their valuations of beef products or their intentions to reduce consumption.

Our results show that correcting consumer beliefs does not necessarily lead to lower demand for carbon-intense consumer products, even in settings where misperceptions are large, and consumers indicate that they are interested in offsetting emissions. The results suggest that the climate impact of behavior is not a strong motivating force for most consumers in our experiment in everyday consumption decisions.

Our results also speak to the implications that can and cannot be drawn from existing evidence. First, we see our findings as consistent with those of studies that show the effects of climate labels, which are often small and short-lived. Our results suggest that behavioral effects from such labels are not primarily driven by changes in individual beliefs, but by other channels, such as an increase in the salience of the climate change phenomenon (Schulze Tilling, 2023), or social norms of mitigation, both of which were kept constant in our experiment. In addition, our representative sample differs from that in most previous studies, which often use university canteens, supermarkets or restaurants, that may attract a particular segment of the population.

Second, evidence of widespread misperception of the climate impact of different consumption behaviors has sometimes been used to argue that information campaigns can lead to meaningful change. We show that this conclusion may be too optimistic. Similarly, other papers have investigated attitudes toward climate change by using donation decisions, willingness to mitigate, and survey responses. The results from these papers may be important in their own right, but our results temper confidence that these measures translate directly into everyday behavior like food consumption.

In fact, the picture that emerges from our and other studies is that the immediate return on information provision policies does not justify their current popularity among policy makers. It suggests that relying on the good intentions of informed individuals will not by itself deliver the important changes that we need in our carbon consumption, and that we will need to rely on more systemic approaches (Chater and Loewenstein, 2022; Kaufmann and Kőszegi, 2023). Of course, our results leave open the possibility that other types of information provision in a different context will be more effective in changing behavior. Having more informed citizens may also have other beneficial effects through long-run reflective processes, for instance by increasing political support for a carbon or meat tax. Future research should help elucidate such mechanisms.

References

- Allcott, Hunt, and Dmitry Taubinsky.** 2015. “Evaluating Behaviorally Motivated Policy: Experimental Evidence From the Lightbulb Market.” *American Economic Review*, 105(8): 2501–2538.
- Andreoni, James, and John Miller.** 2002. “Giving according to GARP: An Experimental Test of the Consistency of Preferences for Altruism.” *Econometrica*, 70(2): 737–753.
- Andre, Peter, Teodora Boneva, Felix Chopra, and Armin Falk.** 2021. “Fighting Climate Change: The Role of Norms, Preferences, and Moral Values.” CESifo Working Paper No. 9175.

- Attari, Shahzeen Z., Michael L. DeKay, Cliff I. Davidson, and Wändi Bruine de Bruin.** 2010. “Public Perceptions of Energy Consumption and Savings.” *Proceedings of the National Academy of Sciences*, 107(37): 16054–16059.
- Azrieli, Yaron, Christopher P. Chambers, and Paul J. Healy.** 2018. “Incentives in Experiments: A Theoretical Analysis.” *Journal of Political Economy*, 126(4): 1472–1503.
- Bernard, René, Panagiota Tzamourani, and Michael Weber.** 2022. “Climate Change and Individual Behavior.” Deutsche Bundesbank Discussion Paper No. 01/2022.
- Bianchi, Filippo, Claudia Dorsel, Emma Garnett, Paul Aveyard, and Susan A. Jebb.** 2018. “Interventions Targeting Conscious Determinants of Human Behaviour to Reduce the Demand for Meat: A Systematic Review with Qualitative Comparative Analysis.” *International Journal of Behavioral Nutrition and Physical Activity*, 15(1): 1–25.
- Bilén, David.** 2022. “Do Carbon Labels Cause Consumers to Reduce Their Emissions? Evidence from a Large-Scale Natural Experiment.” Mimeo, Gothenburg University.
- Brunner, Florentine, Verena Kurz, David Bryngelsson, and Fredrik Hedenus.** 2018. “Carbon Label at a University Restaurant—Label Implementation and Evaluation.” *Ecological Economics*, 146: 658–667.
- Camilleri, Adrian R., Richard P. Larrick, Shajuti Hossain, and Dalia Patino-Echeverri.** 2019. “Consumers Underestimate the Emissions Associated with Food but Are Aided by Labels.” *Nature Climate Change*, 9(1): 53–58.
- Carlsson, Fredrik, Mitesh Kataria, and Elina Lampi.** 2022. “How Much Does It Take? Willingness to Switch to Meat Substitutes.” *Ecological Economics*, 193: 107329.
- Chater, Nick, and George Loewenstein.** 2022. “The i-Frame and the s-Frame: How Focusing on Individual-Level Solutions Has Led Behavioral Public Policy Astray.” *Behavioral and Brain Sciences*.
- d’Adda, Giovanna, Yu Gao, and Massimo Tavoni.** 2022. “A Randomized Trial of Energy Cost Information Provision Alongside Energy-Efficiency Classes for Refrigerator Purchases.” *Nature Energy*, 7(4): 360–368.
- Danz, David, Lise Vesterlund, and Alistair J. Wilson.** 2022. “Belief Elicitation and Behavioral Incentive Compatibility.” *American Economic Review*, 112: 2851–2883.
- Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva.** 2022. “Fighting Climate Change: International Attitudes Toward Climate Policies.” NBER Working Paper No. 30265.
- Diederich, Johannes, and Timo Goeschl.** 2014. “Willingness to Pay for Voluntary Climate Action and Its Determinants: Field-Experimental Evidence.” *Environmental and Resource Economics*, 57(3): 405–429.

- Dong, Diansheng, Christopher G. Davis, and Hayden Stewart.** 2015. “The Quantity and Variety of Households’ Meat Purchases: A Censored Demand System Approach.” *Agricultural Economics*, 46(1): 99–112.
- Ellis, Andrew, and David J Freeman.** 2020. “Revealing Choice Bracketing.” arXiv:2006.14869.
- Elofsson, Katarina, Niklas Bengtsson, Elina Matsdotter, and Johan Arntyr.** 2016. “The Impact of Climate Information on Milk Demand: Evidence from a Field Experiment.” *Food Policy*, 58: 14–23.
- European Commission.** 2020. “New Consumer Agenda: Strengthening Consumer Resilience for Sustainable Recovery.” COM(2020) 696 final. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020DC0696&qid=1605887353618>.
- Fairbrother, Malcolm.** 2022. “Public Opinion About Climate Policies: A Review and Call for More Studies of What People Want.” *PLOS Climate*, 1(5): e0000030.
- Fallucchi, Francesco, and Marc Kaufmann.** 2021. “Narrow Bracketing in Work Choices.” arXiv:2101.04529.
- Fisman, Raymond, Shachar Kariv, and Daniel Markovits.** 2007. “Individual Preferences for Giving.” *American Economic Review*, 97(5): 1858–1876.
- Fosgaard, Toke Reinholt, Alice Pizzo, and Sally Sadoff.** 2021. “Do People Respond to the Climate Impact of Their Behavior? The Effect of Carbon Footprint Information on Grocery Purchases.” IFRO Working Paper No. 2021/05.
- Hackett, M. J., and Nicholas. F. Gray.** 2009. “Carbon Dioxide Emission Savings Potential of Household Water Use Reduction in the UK.” *Journal of Sustainable Development*, 2(1): 36–43.
- Hulshof, Daan, and Machiel Mulder.** 2020. “Willingness to Pay for CO₂ Emission Reductions in Passenger Car Transport.” *Environmental and Resource Economics*, 75(4): 899–929.
- Jalil, Andrew J., Joshua Tasoff, and Arturo Vargas Bustamante.** 2020. “Eating to Save the Planet: Evidence from a Randomized Controlled Trial Using Individual-Level Food Purchase Data.” *Food Policy*, 95: 101950.
- Kanay, Aysegül, Denis Hilton, Laetitia Charalambides, Jean-Baptiste Corrége, Eva Inaudi, Laurent Waroquier, and Stéphane Cezera.** 2021. “Making the Carbon Basket Count: Goal Setting Promotes Sustainable Consumption in a Simulated Online Supermarket.” *Journal of Economic Psychology*, 83: 102348.
- Kaufmann, Marc, and Botond Köszegi.** 2023. “Understanding Markets with Socially Responsible Consumers.” Mimeo, Central European University.
- Klenert, David, Linus Mattauch, Emmanuel Combet, Ottmar Edenhofer, Cameron Hepburn, Ryan Rafaty, and Nicholas Stern.** 2018. “Making Carbon Pricing Work for Citizens.” *Nature Climate Change*, 8(8): 669–677.

- Kortelainen, Mika, Jibonayan Raychaudhuri, and Beatrice Roussillon.** 2016. “Effects of Carbon Reduction Labels: Evidence from Scanner Data.” *Economic Inquiry*, 54(2): 1167–1187.
- Lohmann, Paul, Elisabeth Gsottbauer, Anya Doherty, and Andreas Kontoleon.** 2022. “Do Carbon Footprint Labels Promote Climatarian Diets? Evidence from a Large-Scale Field Experiment.” *Journal of Environmental Economics and Management*, 114: 102693.
- Löschel, Andreas, Bodo Sturm, and Carsten Vogt.** 2013. “The Demand for Climate Protection—Empirical Evidence from Germany.” *Economics Letters*, 118(3): 415–418.
- Nemet, Gregory F., and Evan Johnson.** 2010. “Willingness to Pay for Climate Policy: A Review of Estimates.” La Follette School Working Paper No. 2010-011.
- Newell, Richard G., and Juha Siikamäki.** 2014. “Nudging Energy Efficiency Behavior: The Role of Information Labels.” *Journal of the Association of Environmental and Resource Economists*, 1(4): 555–598.
- Pace, Davide D., and Joël van der Weele.** 2020. “Curbing Carbon: An Experiment on Uncertainty and Information about CO₂ Emissions.” Tinbergen Institute Discussion Paper No. 2020-059.
- Padgett, J. Paul, Anne C. Steinemann, James H. Clarke, and Michael P. Vandenbergh.** 2008. “A Comparison of Carbon Calculators.” *Environmental Impact Assessment Review*, 28(2-3): 106–115.
- Perino, Grischa, Luca A. Panzone, and Timothy Swanson.** 2014. “Motivation Crowding in Real Consumption Decisions: Who Is Messing with My Groceries?” *Economic Inquiry*, 52(2): 592–607.
- Poore, Joseph, and Thomas Nemecek.** 2018. “Reducing Food’s Environmental Impacts through Producers and Consumers.” *Science*, 360(6392): 987–992.
- Rabin, Matthew, and Georg Weizsäcker.** 2009. “Narrow Bracketing and Dominated Choices.” *American Economic Review*, 99(4): 1508–1543.
- Rodemeier, Matthias, and Andreas Löschel.** 2020. “The Welfare Effects of Persuasion and Taxation: Theory and Evidence from the Field.” Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3587339>.
- Schlag, Karl H., and Joël van der Weele.** 2013. “Eliciting Probabilities, Means, Medians, Variances and Covariances without Assuming Risk Neutrality.” *Theoretical Economics Letters*, 3: 38–42.
- Schlag, Karl H., James Tremewan, and Joël J. van der Weele.** 2015. “A Penny for Your Thoughts: A Survey of Methods for Eliciting Beliefs.” *Experimental Economics*, 18(3): 457–490.
- Schulze Tilling, Anna.** 2023. “Changing Consumption Behavior with Carbon Labels: Causal Evidence on Behavioral Channels and Effectiveness.” Mimeo, University of Bonn.

- Shi, Jing, Vivianne H. M. Visschers, Michael Siegrist, and Joseph Arvai.** 2016. “Knowledge as a Driver of Public Perceptions about Climate Change Reassessed.” *Nature Climate Change*, 6(8): 759–762.
- Smith, Anna Jo Bodurtha, Imogen Tennison, Ian Roberts, John Cairns, and Caroline Free.** 2013. “The Carbon Footprint of Behavioural Support Services for Smoking Cessation.” *Tobacco Control*, 22(5): 302–307.
- Soregaroli, Claudio, Elena Claire Ricci, Stefanella Stranieri, Rodolfo M. Nayga Jr, Ettore Capri, and Elena Castellari.** 2021. “Carbon Footprint Information, Prices, and Restaurant Wine Choices by Customers: A Natural Field Experiment.” *Ecological Economics*, 186: 107061.
- Spaargaren, Gert, C. S. A. Van Koppen, Anke M. Janssen, Astrid Hendriksen, and Corine J. Kolfshoten.** 2013. “Consumer Responses to the Carbon Labelling of Food: A Real Life Experiment in a Canteen Practice.” *Sociologia Ruralis*, 53(4): 432–453.
- Taufique, Khan M. R., Kristian S. Nielsen, Thomas Dietz, Rachael Shwom, Paul C. Stern, and Michael P. Vandenbergh.** 2022. “Revisiting the Promise of Carbon Labelling.” *Nature Climate Change*, 12(2): 132–140.
- Tobler, Christina, Vivianne H. M. Visschers, and Michael Siegrist.** 2012. “Consumers’ Knowledge about Climate Change.” *Climatic Change*, 114(2): 189–209.
- UK BEIS.** 2020. “Greenhouse Gas Reporting: Conversion Factors 2019.” UK Department for Business, Energy & Industrial Strategy. <https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2019>.
- Vetter, David.** 2020. “Got Beef? Here’s What Your Hamburger Is Doing to the Climate.” *Forbes*, October 5. <https://www.forbes.com/sites/davidrvetter/2020/10/05/got-beef-heres-what-your-hamburger-is-doing-to-the-climate>.
- Visschers, Vivianne H. M., and Michael Siegrist.** 2015. “Does Better for the Environment Mean Less Tasty? Offering More Climate-Friendly Meals Is Good for the Environment and Customer Satisfaction.” *Appetite*, 95: 475–483.
- Vlaeminck, Pieter, Ting Jiang, and Liesbet Vranken.** 2014. “Food Labeling and Eco-Friendly Consumption: Experimental Evidence from a Belgian Supermarket.” *Ecological Economics*, 108: 180–190.
- Waxman, Henry A., and Edward J. Markey.** 2009. “American Clean Energy and Security Act of 2009.” Washington: US House of Representatives. <https://www.congress.gov/bill/111th-congress/house-bill/2454>.