Correcting Consumer Misperceptions about CO_2 Emissions

Taisuke Imai, Davide D. Pace, Peter Schwardmann, Joël J. van der Weele

Abstract: Policymakers frequently champion information provision about carbon impact on the premise that consumers are willing to mitigate their emissions 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 the carbon impact of different consumption behaviors and clear preferences for mitigation. Yet, in two separate experiments, we show that correcting beliefs has no effect on consumption in large representative samples. Our null results are well powered and informative, as we target information for maximal impact. These results call into question the potential of correcting carbon footprint misperceptions as a tool to fight climate change.

JEL Codes: C81, C93, D84, Q54

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

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, policymakers frequently stress the role of information about CO_2 emissions to consumers and

Taisuke Imai is at the University of Osaka and CESIfo. Davide D. Pace is at Ludwig-Maximilians-Universität Munich and CESIfo (davide.pace@econ.lmu.de). Peter Schwardmann is at Carnegie Mellon University and CESIfo. Joël J. van der Weele is at the Amsterdam School of Economics, Tinbergen Institute. 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 Ludwig-Maximilians-Universität Munich, University of Pittsburgh, BAdW, the University of Osaka, Düsseldorf Institute for Competition Economics, French Experimental Talks, CESifo Area Conference on Behavioral Economics 2022, Berlin Behavioral Economics Seminar, and BGSE/briq Applied Microeconomics Workshop, for helpful comments. Financial support from Deutsche Forschungsgemeinschaft through CRC
Dataverse data: https://doi.org/10.7910/DVN/8RWS8G

Received March 31, 2025; Accepted February 19, 2026; Published online July 16, 2026.

Journal of the Association of Environmental and Resource Economists, volume 13, number 5, September 2026.
 Published by The University of Chicago Press for the Association of Environmental and Resource Economists.
https://doi.org/10.1086/741602

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 United States, 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). 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, which presumes that people care about reducing emissions but 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 study, we test this argument empirically and show that it is flawed. 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 by eliciting their willingness to pay (WTP) 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 policymakers' argument.[2]

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

TRR 190 (Imai, Pace, Schwardmann), the Japan Society for the Promotion of Science (Kakenhi no. JP22K21358 and JP25H00388 to Imai), 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) is gratefully acknowledged. Edited by Teevrat Garg.

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

2. Note that Pace et al. (2025) use belief data from this climate survey in a purely descriptive manner to illustrate that consumers are uncertain about the emissions associated with common products.

the latest scientific estimates. The model allows us to predict the products for which information is likely to generate the biggest behavioral response. The predicted impact takes account of consumers' prior beliefs, their degree of uncertainty, and the (diminishing) responsiveness of their willingness to mitigate to emission size.

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 CO_2 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 nonexistent.

Our "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 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 and, depending on their choices, actually shipping the frozen meat to them. In four between-subject treatments, we varied the type of meat (beef vs. poultry) participants were offered 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 the carbon impact of beef, we find no change in the demand for either beef or chicken. This null result is informative. Our well-powered experiment targets a population that cares about avoiding emissions and, in the case of beef, a product for which misperceptions are large. The null result obtains for all subgroups in our sample and survives when we focus only on those participants 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 one-off statistical fluke. We also rule out behavioral channels like an intention-action gap.

Our second experiment measures changes in self-reported consumption. It gives some participants in our initial survey precise information about the carbon impact of beef, poultry, and other products. We contacted the same subjects two weeks later and elicited their recollection of the information and its impact on their consumption patterns. While subjects remembered some of the information and shifted their beliefs, we replicate the null result from the butcher experiment. The participants who received information about beef are as likely to report a reduction in consumption as participants who received information about other products. However, we find an effect of information provision per se: the participants who receive some information report a general consumption reduction compared to participants who receive none, irrespective of the content or target of information and participants' priors.

Our results indicate that there is a disconnect between knowledge or beliefs and actual behavior in everyday consumption decisions. While emission size plays an important role in abstract elicitations of green preferences, like our incentivized survey measure of willingness to mitigate, it does not necessarily factor into people's demand for consumer products. This refutes the implicit argument made by proponents of information campaigns to combat climate change, as well as the standard intuitions of economic decision-making embodied in our structural model. Furthermore, our findings show that information provision can have effects via channels other than beliefs, which may help explain the (small) effects of information in studies focusing on carbon labels.

These results matter for information policies on climate change. If beliefs are key, then education and belief correction have an important role to play, as current policy measures suppose. If, instead, the effect of information is driven by salience or perceived social norms, then policymakers will have to design interventions that effectively change the attentional, emotional, or social context at the point of purchase. This will likely involve a multitude of strategies alongside information provision, possibly including more visceral aspects of information labels. More generally, our null results suggest that better-informed consumers are unlikely to be a potent substitute for more systemic policies that affect supply chains and relative prices.

Our study contributes to the academic literature in three ways. First, we know of no other study that combines preference and belief elicitations about CO_2 emissions into a structural model and that estimates the quantitative effect of correcting misperceptions for different groups and products. Since different people react differently to information interventions and nudges (Allcott et al. 2025), this exercise ostensibly helps us select the right target and increase impact. Our interventions are then able to focus on a "best case" target, where information is most likely to work. The fact that the interventions do not work, and that the null result obtains in two experiments, is therefore a powerful signal that standard experimental elicitation of economic primitives has limited predictive power over actual behavior in the domain of climate mitigation.

Second, we add to a growing literature on carbon labeling. 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). Most of these papers find a small and short-lived effect of labels on behavior and, ultimately, emissions.[3] The effects documented in these papers may stem from various channels. Our study isolates the effect of belief movements and shows that such movements do not necessarily translate

3. While most studies in this literature focus on hypothetical choices, several papers have looked at real consumption choices in the context of restaurants, university canteens, or supermarkets,

into changes in behavior. Thus, our study suggests that any effects of information are due to other channels, such as increasing the salience of climate impact or norms around climate change. This evidence complements two contemporaneous papers suggesting that beliefs matter less than salience when it comes to carbon labels (Ho and Page 2025; Schulze-Tilling 2025). While these other studies achieve a high level of ecological validity, our study features more detailed belief measurements, information provision that targets only beliefs, and representative samples of consumers, thus producing a forceful null result on the effect of belief movement.

Indeed, our third contribution is to the measurement of (mis)perceptions and valuations related to climate change and to the quantification of their behavioral impact. Some articles elicit broad knowledge of the phenomenon of climate change and link it with measures of concern and policy support (e.g., Shi et al. 2016; Dechezleprêtre et al. 2025). Closest to our study, Camilleri et al. (2019) find a systematic underestimation of greenhouse gas emissions associated with food and electric appliances. Previous papers have used unincentivized point estimations to measure misperceptions. By contrast, our study administers incentivized elicitation of the entire belief distribution about the emissions of common products and activities, providing a much more complete picture of subjects' beliefs. To do so successfully in representative samples, we develop an intuitive and accessible elicitation interface.

We similarly contribute to the measurement of willingness to pay to mitigate climate impact. Previous literature uses mostly unincentivized surveys and contingent valuation methods (see Nemet and Johnson [2010] for a review). Some studies use incentivized revealed preference techniques to elicit WTP for a single amount of emissions (e.g., Löschel et al. 2013; Diederich and Goeschl 2014; Andre et al., forthcoming). Only recently, people have studied (incentivized) valuations as a function of emission sizes, which is necessary to properly evaluate the impact of belief changes (Berger and Wyss 2021; Pace et al. 2025; Rodemeier 2026). While our elicitation technique is similar to the measure in Berger and Wyss (2021), we go further by applying these measures to representative samples (as opposed to students) and by using interpolation techniques to derive a smooth willingness-to-mitigate curve for use in structural model estimation.

1. CLIMATE SURVEY

Our initial survey measures consumers' existing beliefs about CO_2 emissions generated in the production of common consumer goods, as well as their willingness to pay (WTP) to avoid CO_2 emissions. These quantities subsequently serve as inputs for a structural model that allows us to make predictions about the effect of information, as we explain

sometimes studying labels in combination with another information intervention, like posters (e.g., Lohmann et al. 2022; Fosgaard et al. 2024; Ho and Page 2025; Schulze-Tilling 2025). The literature also includes some null results for specific products like detergents (Kortelainen et al. 2016) and in abstract online experiments (Pace et al. 2025).

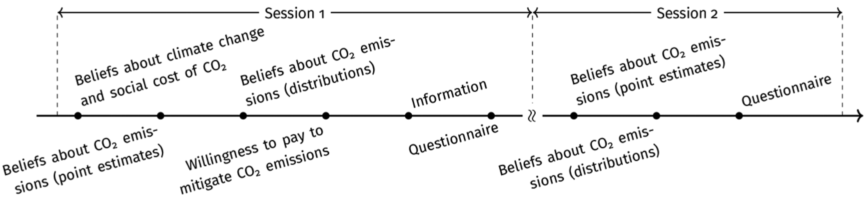

Figure 1. Timeline of the climate survey

in section 2. 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 3.3). 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 collected both point beliefs and belief distributions of CO_2 emissions from several common consumer products and activities. The last task elicited willingness to pay for mitigating CO_2 emissions.

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. Appendix A.1.2 (apps. A–D are available online) contains additional information about the steps we took to maximize the data quality.

1.1. Belief Elicitation

At the start of the survey, we elicited participants' beliefs about the CO_2 emissions generated by driving 1 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 (see table A.1; tables A.1–A.11, B.1–B.7 are available online). Table 1 presents the scientific estimates we used to incentivize the guesses together with their source. 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 CO_2 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 CO_2, we can convert all measures to the perceived grams equivalent (see table A.4 and fig. A.9; figs. A.1–A.17, B.1–B.6 are available online). Moreover, the model we describe in section 2 further mitigates any concern that systematic misperceptions about the CO_2 emissions associated with driving bias our predictions, because these predictions will be independent of the denomination of CO_2 emissions.

Table 1. List of Consumer Products and Actions

	Quantity	Emission Size: Estimate	Emission Size: Unit	Source
Beer	12 fl oz	1.46	mile	Poore and Nemecek (2018)
Phone call	1 hour	1.55	mile	Smith et al. (2013)
Microwave	1,000 W, 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	100 g	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)

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 CO_2 each of the 12 products in table 1 emitted relative to driving 1 mile by car. Participants answered all 12 questions on one page, and the order of the products was randomized across participants (fig. 2*A*). 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 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 et al. 2015).[4]

In order to understand the participants' confidence in their answers, we then elicited the subjective probability distribution about the size of CO_2 emissions. For each product, we presented five "bins" around the point estimate the participant reported in the first part and asked the participant to allocate 20 balls into these five bins. We informed participants that each bin represents an interval that might contain the scientific estimate. They were instructed to allocate the balls to represent their level of confidence that the estimate falls within each bin. Figure 2*B* provides an illustration. We incentivized the elicitation by randomly selecting one of the bins and scoring the answer according to a

4. We did not incentivize the questions about the CO_2 emissions and the social cost of driving 1 mile by a car, as we realized that answers to these questions can be straightforwardly obtained on Google.

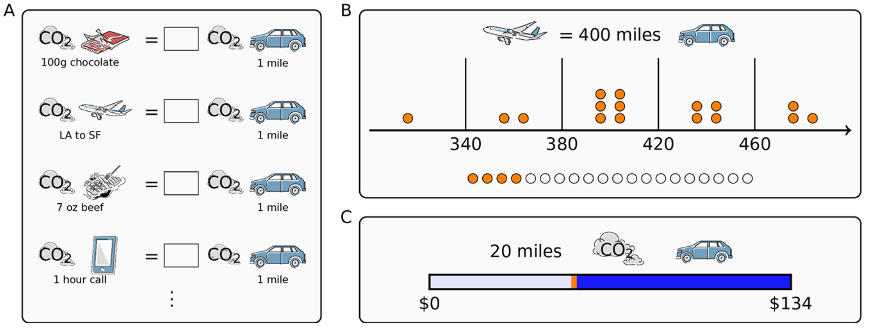

Figure 2. Illustration of the belief and willingness to mitigate (WTM) elicitation interface. *A*, Point-belief elicitation task; *B*, bins-and-balls belief elicitation task; *C*, WTM elicitation task. *B* shows an example in which a participant stated 400 in the previous point belief elicitation task and is now asked to allocate 20 balls into five bins, centered around this number. See appendix A.1.2 for screenshots of the interface.

randomized quadratic scoring rule. Theoretically, this mechanism ensures that participants truthfully reveal their belief about the probability 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.

1.2. Willingness to Mitigate

After the belief tasks, we elicited the participants' willingness to pay for mitigating CO_2 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, a charity that finances various projects to offset CO_2 emissions and offsets 1 ton of CO_2 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 CO_2 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 CO_2 using a slider between $0 and $134 (£100); see figure 2C.[5] The interface was designed to help participants make consistent choices. 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 app. A.1.2).

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

We incentivized the WTM with a Becker–DeGroot–Marschak mechanism, which means that reporting the true WTM is in the best interest of the participant. Specifically, the instructions explain to the participants that we will show them an amount of CO_2 emissions and that they will have to indicate their corresponding "requested bonus": the minimum payment they need to receive in order to accept causing those emissions. We use this framing to make the elicitation feel as close as possible to emitting CO_2 as a by-product of consuming a good. As standard in Becker–DeGroot–Marschak elicitations, the instructions further explain that we will compare the requested bonus with a randomly generated monetary amount between $0 and $134. If the randomly generated number is higher than the requested bonus, the participants will earn a bonus payment equal to the randomly generated number, and the CO_2 will be emitted. Otherwise, the participant will get no bonus, and no emissions will be generated. The emissions are generated by canceling the donations to Carbonfund.org. Incentive-wise, this mechanism is equivalent to endowing participants with an emission certificate they could either cancel or sell to someone who will use it. Under this representation, the Becker–DeGroot–Marschak elicits the participants' willingness to accept the emissions to take place. 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 the carbon offset receipts from the experiment.

1.3. Information Provision and Session 2

After the WTM elicitation, we conducted an information provision experiment and asked participants a series of questions about their demographic background, consumption habits (regarding the 12 products), and attitudes toward climate change. Additionally, we followed up with the participants two weeks later to assess whether the information had changed their behavior. The detailed description of these parts of the survey is postponed to section 3.3.

1.4. Implementation

We recruited 1,430 participants on Prolific between December 3 and 6, 2020, and 1,022 of those completed session 1 of the survey.[6] We restricted participation to US residents, and we aimed to collect a sample representative for age, gender,

6. We 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. To retain enough older people in the sample, we invited them to restart the experiment from the WTM instructions on December 7, 2020, and we gave them the solutions to two of the seven related comprehension questions. Of the 41 subjects that were allowed to restart the experiment, 22 completed it.

and ethnicity.[7] 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 shows the full 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. We also divided the instructions into five blocks. After each block, we asked participants to answer several comprehension questions. We did not allow subjects to continue the experiment until they answered all the questions in 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, we paid them a bonus. This incentive mechanism elicits truthful answers in experiments with multiple tasks (Azrieli et al. 2018). Participants received $10.05 for completing the study plus a variable bonus depending on their answers (mean = $2.67, SD = 4.31). The median survey completion time was 55 minutes.

1.5. Results

1.5.1. Beliefs

Participants estimated CO_2 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 3*A* plots them against scientific estimates of CO_2 emissions.[8] Median beliefs lie below the identity line for all but one (microwave) product, indicating that participants underestimated the size of CO_2 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 (light bulb vs. car). The fraction of participants who underestimated the size of emissions varies from 41% (microwave) to 92% (gas heating), with this fraction increasing with 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 by only 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.

7. We compared the demographic characteristics of study participants and information from US Census Bureau (2022) and found that our sample is not representative in the age dimension (table A.3).

8. We focus on median beliefs since there are several extreme outliers. Figure A.8 shows empirical CDFs of reported CO_2 emission sizes for each product. We find that reported beliefs vary across participants' demographic characteristics, although the associations are generally weak. Figure A.10 illustrates the distribution of reported beliefs for each product across participant categories (unconditionally). Quantile regression results indicate that median beliefs are correlated with participants' political views and with being 58 years of age or older, but not with other demographic characteristics (table A.5).

Table 2. Summary Statistics of Elicited (Point) Beliefs about CO_2 Emissions from 12 Consumer Products and Activities

			Belief			
Product	Emissions	Unit	Q1	Median	Q3	Underest.
Beer	1.46	miles	.50	1.20	6.00	.516
Phone call	1.55	miles	.40	1.00	5.00	.549
Microwave	1.76	miles	.80	2.15	10.00	.406
Milk	2.60	miles	.50	2.00	8.00	.570
Shower	3.90	miles	.50	1.50	5.00	.689
Egg	4.81	miles	.50	1.50	6.00	.697
Poultry	6.78	miles	.60	2.50	10.00	.676
Chocolate	16.03	miles	.40	1.20	8.00	.831
Coffee	44.41	miles	.50	2.00	10.00	.885
Beef	68.39	miles	1.00	5.00	20.00	.858
Flight	304.60	miles	10.00	150.00	600.00	.586
Gas heating	606.68	miles	3.00	20.00	100.00	.919
Car (drive 1 mile)	291.00	grams	5.03	85.05	402.99	.677

Note. The last column "Underest." shows the fraction of participants who underestimated the size of emissions.

Even though the participants misperceived the size of CO_2 emissions from each product, they had a good understanding of which products emit more CO_2. As figure 3*B* shows, the "true" ranking of emission sizes based on scientific estimates and the ranking "revealed" by each participant's estimate are positively correlated.

All the qualitative results of this section replicate if we express participants' beliefs in terms of grams of CO_2 using their beliefs about the CO_2 emissions linked with driving 1 mile by car. Figure A.9 shows that, since participants underestimate the grams of CO_2 emitted when driving, the underestimation is more severe if we express the beliefs in grams.

Taken together, the belief elicitation tasks in the climate survey suggest that consumers significantly underestimate the size of CO_2 emissions associated with common consumer products and activities, but they have more accurate perceptions about the ordinal ranking of CO_2 emissions.

1.5.2. Willingness to Mitigate

We now turn to participants' willingness to mitigate (WTM) CO_2 emissions. We elicited WTM for 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 CO_2 emissions, and they exhibit a concave pattern (fig. 3C, dark line). Moving from emissions equivalent to driving 5 miles to 20 miles, a fourfold 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

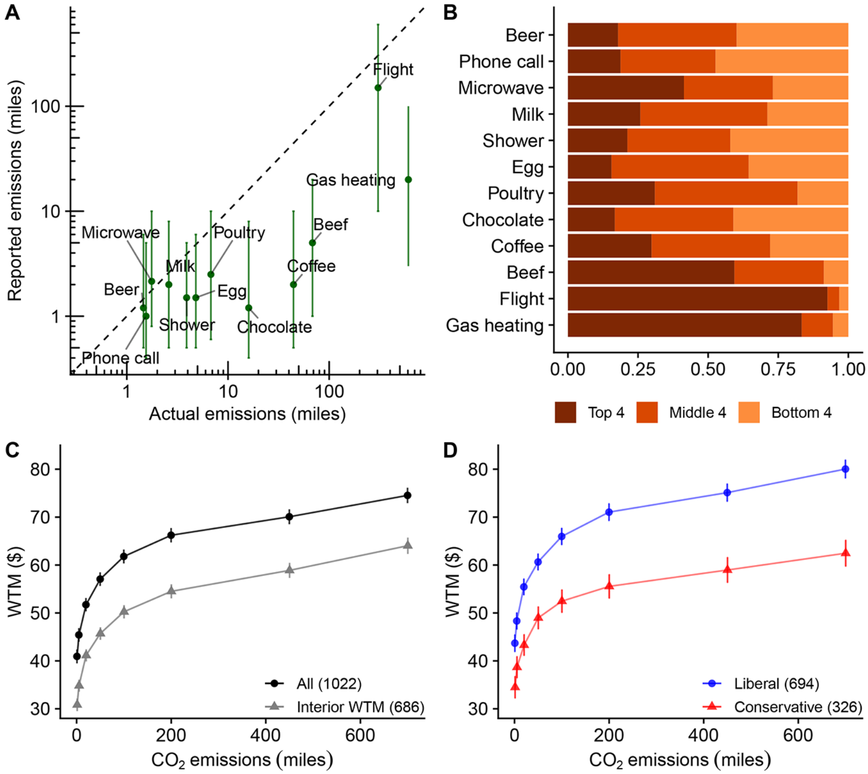

Figure 3. Beliefs and willingness to mitigate (WTM). *A*, Summary statistics of reported CO_2 emissions (median and interquartile range [IQR]). Axes are on a logarithmic scale. *B*, Ranking of reported emission sizes. Products are sorted by the true emission size from low to high. C, Concave WTM (mean and standard error of the mean [SEM]). *D*, WTM and political view (mean and SEM). In C and *D*, numbers in parentheses indicate the number of observations. In *D*, "somewhat liberal" and "somewhat conservative" are grouped into liberal and conservative, respectively.

the emission size increases. 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 (fig. 3C, light gray line). See tables A.6 and A.7 for summary statistics of WTMs and the number of "corner" observations for each level of emissions.

We observe correlations between WTM and some demographic characteristics. Participants who identified themselves as liberal on the political spectrum have uniformly higher WTM than conservative participants (fig. 3*D*). 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 (figs. A.11 and A.12).

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 CO_2 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 nonmonotonic. 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 appendix A.2.4 for details.

The concave WTM we document confirms findings in Pace et al. (2025) and is in line with the scope insensitivity documented in many domains, including offsets for emissions generated by delivery services (Rodemeier 2026), the contingent valuation of environmental public goods (Kahneman et al. 1999), donations for vaccines (Ziegler et al. 2024), and the number of people affected by a decision (Schumacher et al. 2017).

In the next section, we describe how to combine WTM and belief measures to predict the effect of information.

2. 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 consumer 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 et al. 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 CO_2 emissions $c \geq 0$. The consumer's utility from consuming the product is

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

where $w : \mathbb{R}_+ \to \mathbb{R}_+$ captures the psychological cost from CO_2 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 CO_2. 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.[9]

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 CO_2 emissions F. The climate survey we discussed above is designed to measure these two quantities as precisely as possible.

The WTMs stated by each participant provide information about w. Requesting a bonus of y_m to allow emitting CO_2 corresponding to emissions generated by driving m 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 app. 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 the subjective belief distribution F_k for each product k. See appendix A.3 for details.

2.1. Quantifying the Effect of Information

Given a WTM function w and a subjective belief distribution F about CO_2 emissions associated with a good or activity, we can calculate the expected WTM,

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

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 CO_2 emissions.

We model an information policy as a device that shifts consumer i's belief about CO_2 emissions associated with good k from F_{ik} to F_k^*, a degenerate distribution at

9. Narrow bracketing has also been documented in choices over monetary outcomes (Rabin and Weizsäcker 2009; Ellis and Freeman 2024) 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 US consumer, broad bracketing implies that consumers would be on a flat part of w.

the "true" size of CO_2 emissions.[10] 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 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 information campaign Δ_k has a simple interpretation: providing accurate information on the CO_2 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.

Second, because our model combines beliefs and willingness to mitigate CO_2 emissions that were both expressed as miles-driven-in-a-car equivalents, the unit of denomination of CO_2 emissions drops out of our prediction. This allows us to use an intuitive and common way of denominating CO_2 emissions while assuring that any systematic misperceptions about the climate impact of driving do not affect our predictions.

2.2. Predictions

We now calculate our measure of the effect of information 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 substantial variation in the effect of information across products. The model predicts a positive effect for five products (gas heating, beef, coffee, flights, and chocolate), no effect for three (showers, poultry, and eggs), and a negative effect for four (phone calls, milk, beer, and microwaves). The model's predictions find support in the empirical literature. For example, negative informational effects for electrical

10. Note 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 about the information or has difficulty in giving up her original belief.

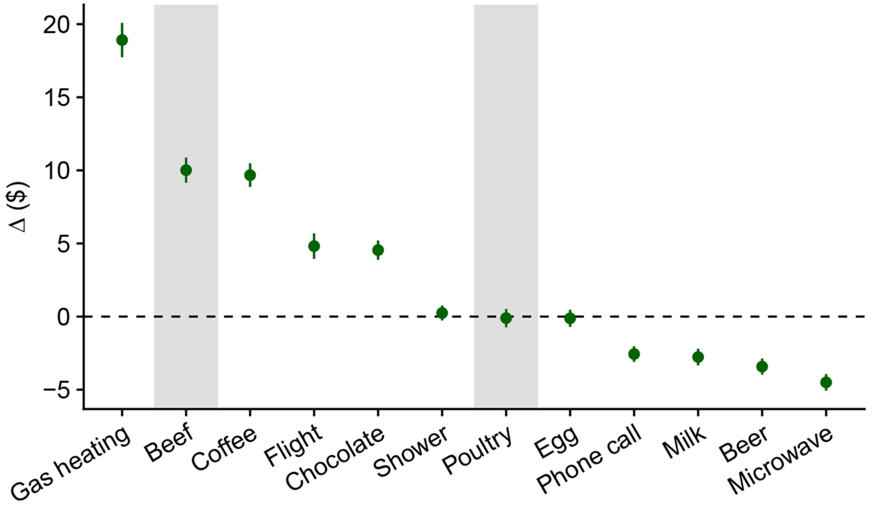

Figure 4. Predicted effect of information provision Δ_k for each product. The reference group G is the entire sample of survey participants ($N = 1{,}022$). Bars indicate standard error of the mean.

appliances have been documented in several studies (Rodemeier and Löschel 2020; d'Adda et al. 2022).

The model aggregates highly heterogeneous individual-level inputs, including WTM curves and belief distributions. Nevertheless, the predicted product-level informational effect is driven primarily by the product's actual emission level. We find more positive information effects for products with higher CO_2 emissions, and the ranking in figure 4 closely mirrors that in table 1 (see also fig. A.14).

Actual emissions matter for two reasons. First, emission size is correlated with the extent of underestimation, as shown in figure A.8. As a result, providing accurate information leads to more negative belief updates for low-emission products. Second, belief updates interact with the concavity of the WTM function. Products with relatively low emissions lie on a steeper portion of the WTM curve, so downward belief revisions have a larger effect on valuations. In contrast, correcting overestimations for high-emission products is less consequential, as these adjustments are smaller and translate into weaker valuation responses.

We also quantify the effect of information for various subgroups in the population, focusing on two meat products, beef and poultry, which will be the subject of the experiment in the next section. Figure A.15*A* quantifies Δ_k and illustrates the benefits of integrating preferences and beliefs over simpler approaches, such as merely targeting populations with a high willingness to pay. For example, the model predicts a larger effect for males than females (fig. A.15*c*) and for participants with conservative political views compared to those with liberal views (fig. A.15*d*). This is despite the fact that, in both cases, the former group has a lower WTM (see fig. A.11). 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 the information provided.

3. THE CAUSAL EFFECT OF CORRECTING CONSUMER BELIEFS

We now test the predictions we derive from our calibrated structural model in section 2. To this end, we compare the effect of correcting misperceptions 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 (Poore and Nemecek 2018). 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, as we show in figure 4. While the predicted effect of information on beef consumption is among the 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, is subject to correction through information. Thus, the main hypothesis that we test in our experiments is that information about carbon impact will have a bigger impact on consumer willingness to pay for beef products than for chicken products. We focus on the willingness to pay as it is the variable on which our model makes direct predictions.

We conduct two experiments. Our first and main experiment involves a real choice to buy meat from an online butcher. It is designed to identify the effect of correcting misperceptions while controlling for other possible effects of information, such as increased salience of emissions. Our second experiment, described in section 3.3, investigates the impact of information on self-reported behavior, elicited in the climate survey.

3.1. Butcher Experiment

Our experiment provides evidence on the role of information in an actual consumption decision. We recruited a new sample and 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 by a premium online butcher Porter Road; 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 United States 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 CO_2 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 meat. We pinned down participants' beliefs about the car CO_2 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. Information about car emissions was available for all treatments.

As an additional manipulation, we varied whether the participants were first offered 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 spillover effects, in which information about beef affects beliefs and WTP for poultry, and vice versa.

The timeline of the experiment is illustrated in figure 5. The experiment has two parts, one for 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 CO_2 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 CO_2 emitted by driving 1 mile by car.[11]

To help participants get a sense of the magnitudes of emissions, just before they could express their guesses, we informed them of how the CO_2 emissions from driving 1 mile by car compared with the emissions generated by the production and distribution of 12 fluid ounces 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 ±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.[12]

---

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

12. In 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.

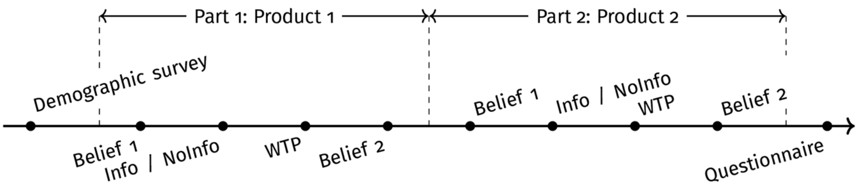

Figure 5. Timeline of the butcher experiment. WTP = willingness to pay

We then elicited participants' WTP using a two-stage multiple price list (MPL) with forced single switching.[13] 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 increments. 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.[14] The instructions encouraged participants to consider their own valuation of the meat bundle and to use it to make decisions.

After completing the MPL task, we asked the participants to guess the size of the emissions associated with the meat product they had the opportunity to purchase one more time. 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 information about the CO_2 emissions associated with the new meat bundle, along with all previously provided information. In the NoInfo treatment, 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

13. We 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.

14. We used a Becker–DeGroot–Marschak 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 received 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 switch, thereby enforcing consistency in choices.

asked for their contact information (home address and email) so we could deliver the meat product or, if any, the monetary bonus.

3.1.1. Implementation

We recruited participants on the platform, Lucid, between March 31 and April 15, 2022. When starting the experiment, the participants were unaware that it was about meat consumption or that they had the chance to receive a premium meat shipment. Hence, our sample comes from the population of US people completing surveys on Lucid, excluding only those participants who did not consume meat and those who lived outside contiguous US states due to shipment requirements by Porter Road;[15] 2,081 participants satisfied the preregistered inclusion criteria: 1,047 were assigned to the NoInfo treatment, and 1,034 were assigned to the Info treatment.[16] Participants are representative along gender and age (table B.2). Table B.1 shows that demographic characteristics are balanced across treatments.[17] Our sample is, on average, 46.8 years old (SD = 17.1), and 48.4% of the participants are 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 and 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 appendix B.2.

3.1.2. 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 for the beef bundle come from BeefFirst treatments ($N = 1{,}048$), and data for the poultry bundle come from PoultryFirst treatments ($N = 1{,}033$).

As in the climate survey discussed in section 1, participants exhibited a significant underestimation of the size of CO_2 emissions from beef and poultry. Figure 6 shows

15. To 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.

16. Number 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.

17. Table B.3 compares the Lucid and Prolific samples. We observe differences in the distribution of demographic characteristics between the two samples, particularly in age (the Lucid sample includes more older participants) and political views (the Lucid sample includes more conservative participants).

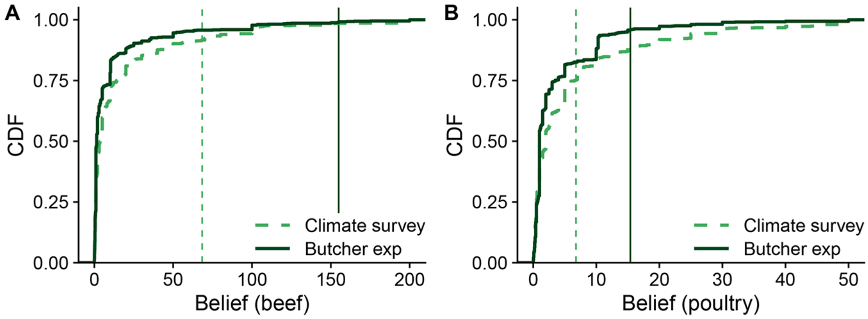

Figure 6. Empirical CDFs of beliefs about CO_2 emissions from two samples. *A*, Beef; *B*, poultry. The size of meat products for belief elicitation was 7 oz in the climate survey and 1 lb (=16 oz) in the butcher experiment. For the butcher experiment data, we focus on belief data from the first elicitation in the first part of the experiment. Vertical dashed lines correspond to the "true" size of CO_2 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).

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. As 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 CO_2 emissions across treatments. The distributions of prior beliefs (asked before WTP) show no differences between Info and NoInfo treatments for both meat products (fig. 7*A*, *B*). 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 7*C*, *D*. 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.[18]

Remember that our model in section 2 predicts that information reduces the valuation 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 by the data. Figure 8*A* shows the WTP for meat products by treatment. If anything, there is a small upward movement in the valuation of the beef package after information provision. Average WTPs are not significantly

18. If we allow a margin of ±10%, the number increases to 68.3% (351/514) for poultry.

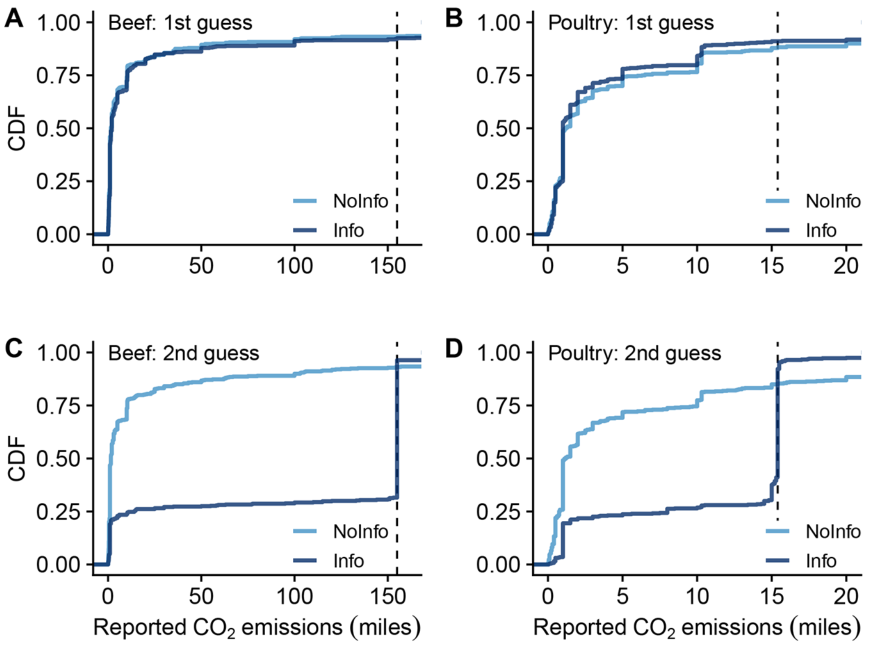

Figure 7. Beliefs about CO_2 emissions from two meat products. We focus on the data from the first part of the experiment (*A*, *C*, BeefFirst treatments; *B*, *D*, PoultryFirst treatments). Vertical lines correspond to the "true" size of CO_2 emissions (15.4 miles for poultry and 155 miles for beef).

different between treatments for both products (beef: $t(1{,}046) = -1.200$, $p = .230$; poultry: $t(1{,}031) = 0.938$, $p = .349$). Figure 8*B* and figure 8*C* 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.

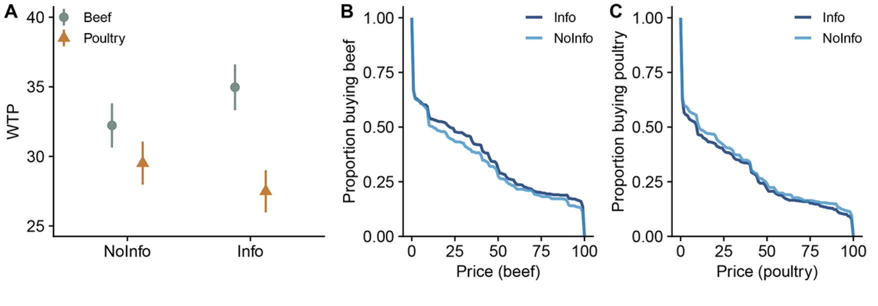

Figure 8. *A*, Average willingness to pay (WTP) for the first meat product. *B*, *C*, The proportion of participants buying the meat product at each price. We focus on the data from the first part of the experiment. In *A*, bars indicate SEM. Figure B.4 shows the CDFs of WTPs.

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.4). 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 among women and younger individuals. In 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.

In summary, our results do not support the hypothesis that correcting CO_2 misperceptions changes consumption decisions.

3.2. Interpretation of the Null Effect

We now turn to investigate possible reasons for the observed null effect of information about CO_2 emissions on the demand for meat. We focus on beef, where we predicted that information should affect willingness to pay negatively and decisively.

3.2.1. 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. To test for this channel, we look at the participants' stated intentions about future consumption in the postexperimental questionnaire. At the end of the experiment, we asked participants "Do you intend to reduce your beef/poultry consumption in light of its CO_2 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 (fig. 9; $\chi^2(4) = 0.964$, $p = .915$). Thus, the null effect of information does not stem from a failure to implement virtuous plans but from a failure to make such plans.

Indeed, we find evidence that participants simply do not plan to reduce beef consumption. We ask participants how hard it would be for them to reduce their beef consumption by half. We find that the null effect persists in the subsample that declares that such a reduction would not be so difficult (see fig. B.6).

3.2.2. Were Participants' Beliefs Insensitive to the Information Treatment?

As shown in figure 7, most participants become less optimistic about beef CO_2 emissions after seeing the information. 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 (col. 1), the null effect remains. Thus, insensitivity to the information cannot explain the null result.

Table 3. Interpretation of the Null Effect of Information on WTP for Beef

	BeefFirst: All (1)	NoInfo + Updaters (2)	Completed Sample (3)	Env-Conscious (4)	High Beef Eaters (6)	Beef Consumers (6)	Trusts Delivery (7)	PoultryFirst: All (8)
Info	2.743	.528	2.914	1.673	5.277		4.835	−.860
	(2.285)	(2.448)	(2.400)	(2.847)	(3.284)		(3.191)	(2.292)
Belief (poultry)			−.011					
			(.012)					
Difficulty (beef)						3.074***		
						(1.004)		
Constant	32.225***	32.225***	33.018***	32.912***	34.574***	25.580***	35.062***	33.080***
	(1.590)	(1.590)	(1.660)	(1.984)	(2.263)	(2.942)	(2.162)	(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	.001	.0001	.002	.001	.005	.010	.004	.0001

Note. The dependent variable is willingness to pay (WTP) for beef. Samples are as follows. Column 1, all participants in the BeefFirst treatments. Column 2, participants in the NoInfo treatment and those in the Info treatment who responded to information by updating their beliefs upward. Column 3, participants who completed both parts of the experiment. One subject who reported an extremely large belief ($\approx 1.46 \times 10^9$) about CO_2 emissions from poultry is excluded. Column 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?"). Column 5, participants in the BeefFirst treatments who self-reported consuming beef at least three times per week. Column 6, participants in the BeefFirst treatments who self-reported consuming beef. Column 7, participants in the BeefFirst treatments who expressed trust in us actually sending meat. Column 8, all participants in the PoultryFirst treatments. Robust standard errors are reported in parentheses.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

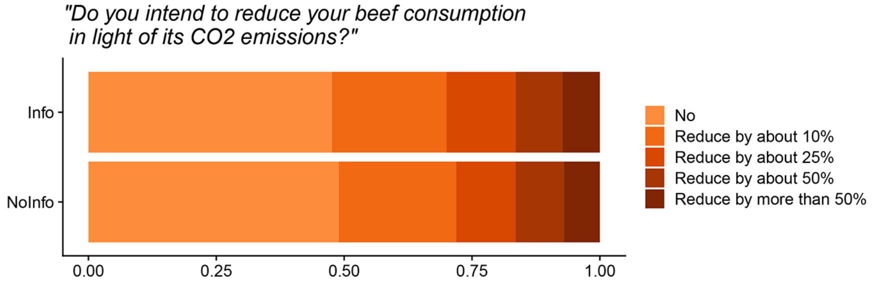

Figure 9. Distribution of responses to a survey question: "Do you intend to reduce your beef consumption in light of its CO_2 emissions?" We focus on 1,013 participants in the BeefFirst treatments. Participants responded on a 5-point Likert scale (1 = No. 2 = Yes, I am prepared to reduce my current consumption by about 10%. . . . 5 = Yes, I am prepared to reduce my current consumption by more than 50%).

3.2.3. 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 the WTP for beef insensitive to information.

We can address this point in several ways. The first is to directly control for this spillover in beliefs. In the BeefFirst treatment, we measured participants' beliefs about the CO_2 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.[19] 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, col. 3).

In addition, we can look at the case where beef is the second product that participants can buy. Here, by the time participants stated their willingness to pay for beef in the Info treatment, they had 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, col. 8).

Finally, we can analyze 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. Nevertheless, this information did not change their intended behavior.

19. About 63% of the participants in the BeefFirst, Info treatment (317/505) overestimated the size of CO_2 emissions from poultry (reported numbers above 15.4 miles; see fig. B.5).

3.2.4. Is Optimism Concentrated Among Unresponsive Participants?

One reason that information may have little impact on CO_2 emissions is that prior optimism about CO_2 emissions is concentrated among individuals whose behavior is unresponsive to belief changes, either because their willingness to mitigate is low or because they do not consume much meat in the first place (e.g., near vegetarians). Note that our structural model was explicitly designed to make predictions that take this mismatch into account, so our initial predictions, based on the climate survey, are not subject to this concern.

We perform two exercises to investigate this pattern in our data. First, we restrict our analysis to the 65% of participants who self-proclaim to be concerned about climate change. The degree of underestimation of beef emissions in this subsample is similar to that of the overall one (91.7% vs. 92.7% of participants underestimate beef emissions), indicating that concerned participants are not better informed. Moreover, the null effect persists in this restricted sample (table 3, col. 4). Second, we restrict our dataset to participants who consume meat at least three times per week (i.e., above-median frequency), thus ruling out near vegetarians. The null effect persists also in this restricted sample (table 3, col. 5).

3.2.5. Does the Information Cause Participants to Offset Emissions Outside the Experiment?

To offer a product that participants find appealing, we used high-quality meat. If participants find this a particularly attractive offer that they would not like to pass on, they may respond to the information by demanding less (low-quality) meat outside the experiment or offset the emissions in some other way. If this were the case, the information treatment might reduce emissions but not the willingness to pay for the meat we offer.

We already showed above that information did not change participants' stated willingness to buy beef in the future, casting doubt on this conjecture. This is all the more informative, since the price elasticity of demand for beef steaks is substantial—between −0.42 and −0.52 (Dong et al. 2015). Thus, if participants wanted to reduce emissions, one would expect them to do it via the elastic and immediate option of reducing beef demand, not by taking additional actions after the experiment.

3.2.6. What If People Chose Directly Between Beef and Poultry?

Many shopping environments offer a choice between variations of comparable products. Information about CO_2 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 comparing the valuations for beef and poultry in the Info treatment depending on whether they are offered as the first or the second product. In the latter case, information about a substitute product is available as well. We carry out this analysis in appendix B.3.6, but we do not find evidence of substitution across products.

3.2.7. Do People Care About Relative Emission Comparisons Rather than Absolute Emission Levels?

Our experiment is designed to correct individuals' absolute misperceptions about emission levels. However, it is possible that individuals primarily care about avoiding the most polluting products rather than the absolute amount of emissions. We refer to this theory as the "relative emission hypothesis." This hypothesis could account for our null result. Suppose that (1) individuals focus only on avoiding the most polluting products and (2) they already hold accurate beliefs about the relative pollution ranking of products. In that case, the information treatment would not be expected to change their WTP for a given product.

The first step in testing the relative emission hypothesis is to check how many participants knew, prior to the experiment, that poultry is less polluting than beef. In the NoInfo treatment, participants' reported beliefs suggest that 47.5% held an incorrect subjective ranking. This substantial share of misperceptions indicates that our information treatment had the potential to influence consumption behavior, even if participants primarily care about relative rather than absolute emissions.

However, we do not find evidence that providing information revealing the correct ranking affects behavior. First, we test whether ranking information reduces participants' WTP for beef. This test is possible because, when beef is presented as the second product available for purchase, participants in the Info treatment state their WTP after viewing a screen that displays the emission ranking of poultry and beef (see fig. B.1). We find no significant difference in WTP between the Info and NoInfo treatments (table 3, col. 8). Table B.7 replicates these results by focusing on the subset of subjects who initially wrongly believed that beer is more polluting than beef.

Second, we leverage responses to a question in the final questionnaire to check whether participants in the Info treatment are more likely to consider switching from beef to poultry to reduce their CO_2 footprint. Again, we find no evidence of such an effect: 34.8% and 35.7% of participants report an intention to consume more poultry in the NoInfo and Info treatment, respectively—a difference that is not statistically significant ($p = .793$, Fisher's exact test).

Overall, we find no evidence in favor of the relative emission hypothesis in the butcher experiment. However, some correlational results from the climate survey are consistent with the hypothesis (see app. A.5). We conclude that a promising direction for future research is to design experiments that test the hypothesis more directly.

3.2.8. Does Our Willingness to Pay Measure Suffer from Noise, Misinterpretation, or Lack of Trust?

A possible reason for the 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.

However, we find that WTP for beef is significantly correlated with participants' self-reported difficulty in reducing beef consumption if they had to (table 3, col. 6).

A related worry may be that despite our elaborate efforts to be credible, (some) participants did not believe that we would actually send them the meat they purchased and did not take the WTP elicitation seriously. 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 postexperimental survey.[20] The null effect persists in this restricted sample (table 3, col. 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 CO_2 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 the information reduces the intention to consume beef. However, as discussed above, we do not find evidence for this treatment effect.

3.2.9. *Is the Null Effect Explained by a Different Sample Composition in the Butcher Experiment Compared to the Climate Survey?*

Table B.3 compares the Lucid and Prolific samples. While there are some differences between the two samples (the Lucid sample skews older and more conservative), it is unlikely that these differences in sample composition explain the null result. As figure A.15 shows, our model predicts information to reduce WTP for beef in every single demographic group. As such, any reweighting of the climate survey sample would still predict a negative effect of information on WTP for beef.

A related concern is that the (unmeasured) WTM is lower in the Lucid sample. Maybe in the butcher experiment, participants are simply unwilling to make sacrifices for the climate. While we cannot fully exclude this possibility, we believe it is unlikely, given that 65% of the participants in the butcher experiment stated to be concerned about climate change (compared to 77% of concerned participants in the climate survey). Furthermore, the null result in the butcher experiment replicates if we restrict our sample to this 65% of concerned participants (table 3, col. 4).

20. Participants 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).

3.2.10. Was the Null Effect a Fluke?

Even relatively well-powered studies may sometimes result in erroneous null effects. Four results speak against this hypothesis. First, we can ask whether there is any correlational evidence that beliefs about CO_2 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 CO_2 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. Column 8 of table 3 shows that the null result persists in the second part. Of course, because these data 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.

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, a result we replicate with the survey experiment described in section 3.3.

3.2.11. What, Then, Causes the Null Effect?

Having ruled out several possible explanations for the observed null effect, we conclude that people's decision to eat meat is not influenced by beliefs about the size of its CO_2 emissions. That is, while the climate survey reveals that people are willing to pay more to directly mitigate larger quantities of emissions, the size of the emissions does not affect their consumption in a less stylized setting. This might be because many considerations, from one's children's food preferences to family recipes, weigh on food purchases, and so CO_2 emissions fail to rise to sufficient prominence in participants' minds when they make their decisions. It might also be that the concern for marginal emissions we find in our climate survey was an artifact of the elicitation method. The elicitation asks each participant to report their willingness to pay to avoid a series of increasing emission amounts. This type of elicitation could encourage participants to report increasing valuations, either to appear consistent or because the interface makes comparing emission amounts easier than in real life. Indeed, Rodemeier (2026) and Ho and Page (2025) find using between-subjects elicitation that individuals' (positive) willingness to pay to avoid emissions does not vary with emission amounts. Their results imply that people either do not care about marginal increases in emissions or cannot articulate their preferences in between-subject designs, which is consistent with the findings for our butcher experiment.

If individuals' valuations for emissions do not change with emission size, then we should be no more optimistic about finding an effect of information in even "wilder" settings. After all, the information effectively moved beliefs, and we can be confident that the climate impact of various consumption activities was a salient feature of the decision-making environment.

3.3. Survey Experiment

To lend robustness and refine the interpretation of the results of the butcher experiment, we now present additional evidence from a survey experiment that we conducted at the end of the climate survey of section 1. The survey experiment provided some participants with information about the CO_2 emissions associated with several consumption behaviors.

Our main outcome variable is participants' self-reported consumption two weeks after information was received, which differs from the willingness to pay we elicited in the butcher experiment in several key ways. First, the self-reported consumption captures past purchase decisions. This variable may be more intuitive for the participants compared to the willingness to pay, but it is further removed from the predictions of our structural model. Second, consumption was self-reported rather than incentivized. While self-reports have obvious drawbacks, they may be more responsive to information and allow a more sensitive test of the effects of information. Third, we measured self-reported consumption some time after the information provision, so we can speak about information retention.

The design of the survey experiment allows us to refine the interpretations of the null result in several ways. First, it allows us to check the results in the exact sample that we used for the model predictions in section 2, alleviating potential concerns about switching samples in the butcher experiment. Second, it allows us to test whether the null effect replicates for other products. Finally, the survey experiment allows us to separate the effect of information content from the broader effects of information provision. Recall that in the butcher experiment, all participants obtained information about some benchmark products, which kept the salience of climate impact and the presence of information constant. By contrast, the survey experiment features an additional control group that received no information at all. This condition may help reconcile our null result, which relates to information content, with the broader labeling effects of information provision in the literature.

3.3.1. Design of the Survey Experiment

The first part of the survey experiment (session 1) took place at the end of the climate survey. The second part (session 2) took place approximately two weeks later (between December 16 and 21, 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 subset of products randomly selected from the 12 products in our survey.[21] The information provided the latest scientific estimate of the carbon impact of each product and included a link to the source of the information.

To ensure participants engaged with the information, they were first required to repeat it back to us (fig. A.5). Subsequently, we presented them with a "surprise" short-term memory task where they had to recall the information again. Each correct answer earned an additional £0.20 bonus (fig. A.6). Finally, we asked several questions about participants' consumption of the products included in the survey. Among these questions, we inquired whether participants planned to reduce their consumption of each of the products due to their CO_2 emissions (Intentions). See appendix A.1.3 for the complete list of questions.

Session 2. We called back participants 13–18 days after the first session. Out of our original 1,022 participants, 946 (92.5%) completed the second session. At the 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. Finally, we asked the participants whether they changed their consumption of any of the products in light of the associated carbon emissions since the previous session (Consumption).

3.3.2. Results of the Survey Experiment

We first ask whether the information provided in session 1 affected beliefs two weeks after it was communicated. Consistent with the rest of the paper, our focus is on beef. Figure 10 shows CDFs of elicited beliefs prior to receiving information in session 1, and after having received information in session 2. Figure 10*A* shows that beliefs about beef do not significantly differ before any participants received information in session 1 (two-sample Kolmogorov-Smirnov test, $p = .151$). Figure 10*B* shows that information effectively shifted beliefs and reduced underestimation of CO_2 emissions (two-sample Kolmogorov-Smirnov test, $p < .001$). However, it is notable that only a minority of participants remembered the exact value, and a majority still underestimated the impact of beef. Figure A.16 shows results for all products, split by those who over- and

21. We implemented three different conditions: 322 participants received information about six of the 12 products we used in the survey, 344 received information about only three products, and the remaining 356 participants received no information. These different information treatments were designed to explore issues related to recall and information overload, which are not analyzed in this study.

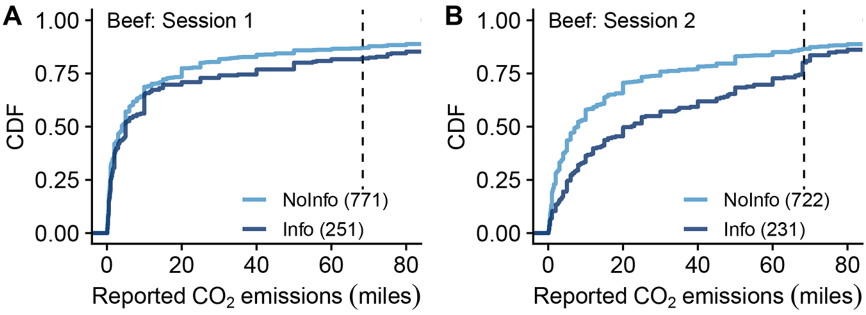

Figure 10. Impact of information on beliefs about the impact of beef and poultry. *A* shows beliefs from the original survey in session 1, before information was given. *B* shows beliefs in session 2. Vertical lines correspond to the "true" size of CO_2 emissions (68.39 miles). Numbers in parentheses indicate the number of observations.

underestimate, and shows that the beliefs generally moved in the direction of the correct answer after information was provided.

Next, we ask whether information affects self-reported behavior. We first focus on participants who received information about at least some of the products, making the survey experiment most comparable to the butcher experiment, where all participants were informed about the emissions associated with a beer and a domestic flight. For a specific product, such as beef, we thus compare participants who received information about beef and several other products with those who received information about products other than beef. As an outcome, we focus on self-reported decreases in consumption reported in session 2. However, all results remain robust when considering consumption intentions reported in session 1 instead (see the first two columns in table A.9).

Table 4 presents the results from linear regressions testing the effect of information on consumption. The first column focuses on beef consumption, replicating the finding of the butcher experiment by showing no significant effect of information on consumption. The second column examines all products. Given the general tendency of participants to underestimate the CO_2 impact of consumption behaviors, one might anticipate an overall reduction in consumption.[22] However, the small point estimate is not significantly different from zero, indicating that this expectation is not met.

We further investigate the role of information in correcting misperceptions by leveraging reported beliefs from session 1. We construct three belief-based proxies to predict the effect of information provided at the end of the session. First, we measure whether

22. We exclude observations featuring the product category "flight," as there was little opportunity to fly during the COVID-19 pandemic. Figure A.17 shows the effect product by product.

Table 4. Effect of Information on Consumption

	Consumption Reduction Beef (1)	Consumption Reduction All Products (2)	Consumption Reduction All Products (3)
Information	−.034	−.001	.050**
	(.042)	(.011)	(.021)
Observations	613	6,742	6,218
R^2	.198	.124	.121

Note. The dependent variable is self-reported consumption reduction of beef (col. 1) and of each of the 12 products in the survey (cols. 2 and 3). This variable is a dummy equal to 1 if the participant reported a reduction in the consumption of a given product. Information is a dummy variable equal to 1 if we provided information about emissions for a participant-product observation and 0 otherwise. The sample in cols. 1 and 2 includes all participants who received at least some information. The sample in col. 3 consists of participants who received no information and those participant-product observations for which we provided information. Robust standard errors in col. 1 and standard errors clustered at the individual level in cols. 2 and 3 are reported in parentheses.

* $p < .1$.

** $p < .05$.

*** $p < .01$.

participants underestimated the impact, which should lead to a larger consumption reduction. Second, we compute the degree of underestimation, which should lead to a larger consumption correction. Finally, we construct the predicted effect of information based on the model outlined in section 2, incorporating both the participant's belief and their willingness to mitigate. Table A.10 reports regressions of consumption reductions on these three measures, focusing exclusively on the product-participant pairs that received information about the impact. The results clearly indicate no significant effect of information using any of these belief-based measures. This further underscores that changes in beliefs do not translate into changes in consumption behavior.

Finally, we examine the condition in which participants were not informed about the impact of any product and compare it with participant-product pairs for which information was provided. The difference between these two groups captures the broader effects of information provision, which may include increasing the salience of climate impacts of CO_2 emissions, conveying normative expectations from the experimenter, and other changes in the perception of consumption decisions unrelated to emission size. Column 3 of table 4 shows a positive and statistically significant effect of information relative to this control group. This result is replicated when examining consumption intentions from session 1 (see col. 3 in table A.9). These findings suggest that information per se influences consumption behavior, even when changes in beliefs induced by the information do not.

3.3.3. Discussion

We replicate the main finding from the butcher experiment that correcting misperceptions about emission sizes has no impact on the consumption of beef (or any other product). This evidence is all the more striking, since it would have been easy for participants to self-report a socially desirable change in consumption after learning that a product is more polluting than expected. Moreover, we do observe a broader effect of information provision per se.

The results suggest that the effects of information provision in the current literature do not (just) arise from the content of the information but from the act of providing information itself. Information provision may increase the salience of climate change or may communicate a normative signal to participants that it is important they do something, or may make participants feel more capable of taking action.[23] More generally, these results suggest that climate labeling primarily affects consumption through channels other than correcting misperceptions of consumers. A better understanding of these alternative channels might be important in the design of effective information policies.

4. 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-intensive 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 in the general population.

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 primarily driven not by changes in individual beliefs but by other channels, such as an increase in the salience of the climate change phenomenon (Bettinger et al. 2022; Schulze-Tilling 2025) or social norms of mitigation. In addition, our representative sample differs from that in most previous studies, which

23. We find suggestive evidence that the information treatment makes participants feel more informed and that more informed participants reduce consumption more.

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 policies does not justify their current popularity among policymakers. 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 2023; Kaufmann et al. 2024; Hardaker et al. 2025). 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, Daniel Cohen, William Morrison, and Dmitry Taubinsky. 2025. When do nudges increase welfare? *American Economic Review* 115 (5): 1555–96.

Andre, Peter, Teodora Boneva, Felix Chopra, and Armin Falk. Forthcoming. Misperceived social norms and willingness to act against climate change. *Review of Economics and Statistics*.

Andreoni, James, and John Miller. 2002. Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica* 70 (2): 737–53.

Azrieli, Yaron, Christopher P. Chambers, and Paul J. Healy. 2018. Incentives in experiments: A theoretical analysis. *Journal of Political Economy* 126 (4): 1472–1503.

Berger, Sebastian, and Annika M. Wyss. 2021. Measuring pro-environmental behavior using the carbon emission task. *Journal of Environmental Psychology* 75:101613.

Bettinger, Eric, Nina Cunha, Guilherme Lichand, and Ricardo Madeira. 2022. When the effects of informational interventions are driven by salience: Evidence from school parents in Brazil. https://dx.doi.org/10.2139/ssrn.3644124.

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.

Chater, Nick, and George Loewenstein. 2023. The i-frame and the s-frame: How focusing on individual-level solutions has led behavioral public policy astray. *Behavioral and Brain Sciences* 46: E417.

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): 360–68.

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva. 2025. Fighting climate change: International attitudes toward climate policies. *American Economic Review* 115 (4): 1258–1300.

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): 405–29.

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. 2024. Revealing choice bracketing. *American Economic Review* 114 (9): 2668–2700.

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.

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): 1858–76.

Fosgaard, Toke R., Alice Pizzo, and Sally Sadoff. 2024. Do people respond to the climate impact of their behavior? The effect of carbon footprint information on grocery purchases. *Environmental and Resource Economics* 87 (7): 1847–86.

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.

Hardaker, Adam, Igor Asanov, František Bartoš, and Stephan B. Bruns. 2025. No evidence for effectiveness of behavioral interventions to mitigate climate change after adjusting for publication bias. I4R Discussion Paper Series no. 263.

Ho, Lisa, and Lucy Page. 2025. Got beef with beef? Evidence from a large-scale carbon labeling experiment. http://dx.doi.org/10.2139/ssrn.4803757.

Kahneman, Daniel, Ilana Ritov, and David Schkade. 1999. Economic preferences or attitude expressions? An analysis of dollar responses to public issues. *Journal of Risk and Uncertainty* 19 (1–3): 203–35.

Kaufmann, Marc, Peter Andre, and Botond Kőszegi. 2024. Understanding markets with socially responsible consumers. *Quarterly Journal of Economics* 139 (3): 1989–2035.

Kortelainen, Mika, Jibonayan Raychaudhuri, and Beatrice Roussillon. 2016. Effects of carbon reduction labels: Evidence from scanner data. *Economic Inquiry* 54 (2): 1167–87.

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): 415–18.

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.

Pace, Davide, Taisuke Imai, Peter Schwardmann, and Joël van der Weele. 2025. Uncertainty about carbon impact and the willingness to avoid CO_2 emissions. *Ecological Economics* 227:108401.

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): 106–15.

Poore, Joseph, and Thomas Nemecek. 2018. Reducing food's environmental impacts through producers and consumers. *Science* 360 (6392): 987–92.

Rabin, Matthew, and Georg Weizsäcker. 2009. Narrow bracketing and dominated choices. *American Economic Review* 99 (4): 1508–43.

Rodemeier, Matthias. 2026. Willingness to pay for carbon mitigation: Field evidence from the market for carbon offsets. *Review of Financial Studies* 39 (1): 1–29.

Rodemeier, Matthias, and Andreas Löschel. 2020. The welfare effects of persuasion and taxation: Theory and evidence from the field. http://dx.doi.org/10.2139/ssrn.3587339.

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): 457–90.

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.

Schulze-Tilling, Anna. 2025. The effectiveness of carbon labels. CRC TRR 224 Discussion paper no. 639.

Schumacher, Heiner, Iris Kesternich, Michael Kosfeld, and Joachim Winter. 2017. One, two, many: Insensitivity to group size in games with concentrated benefits and dispersed costs. *Review of Economic Studies* 84 (3): 1346–77.

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): 759–62.

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): 302–7.

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): 132–40.

UK BEIS (UK Business, Energy, and Industrial Strategy). 2020. Greenhouse gas reporting: Conversion factors 2019. UK Department for Business, Energy and Industrial Strategy. https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2019.

US Census Bureau. 2022. National population by characteristics: 2020–2022. US Census Bureau. https://www.census.gov/data/tables/time-series/demo/popest/2020s-national-detail.html.

Waxman, Henry A., and Edward J. Markey. 2009. American Clean Energy and Security Act of 2009. US House of Representatives, Washington, DC. https://www.congress.gov/bill/111th-congress/house-bill/2454.

Ziegler, Andreas G. B., Giorgia Romagnoli, and Theo Offerman. 2024. Morals in multi-unit markets. *Journal of the European Economic Association* 22 (5): 2225–60.