

Uncertainty about Carbon Impact and the Willingness to Avoid CO₂ Emissions *

Davide D. Pace^{1,5} Taisuke Imai^{2,5}
Peter Schwardmann^{3,5} Joël J. van der Weele⁴

¹Ludwig-Maximilians-Universität München

²Institute of Social and Economic Research, Osaka University

³Carnegie Mellon University

⁴Amsterdam School of Economics, Tinbergen Institute

⁵CESifo

December 1, 2023

Abstract

With a large representative survey ($N = 1,128$), we document that consumers are very uncertain about the emissions associated with various actions, which may affect their willingness to reduce their carbon footprint. We experimentally test two channels for the behavioural impact of such uncertainty, namely risk aversion about the impact of mitigating actions and the formation of motivated beliefs about this impact. In two large online experiments ($N = 2,219$), participants make incentivized trade-offs between personal gain and (uncertain) carbon impact. We find no evidence that uncertainty affects individual climate change mitigation efforts through risk aversion or motivated belief channels. The results suggest that reducing consumer uncertainty through information campaigns is not a policy panacea and that communicating scientific uncertainty around climate impact need not backfire.

*This paper supersedes a previous version circulated with the title “Curbing Carbon: An Experiment on Uncertainty and Information about CO₂ Emissions” by a subset of the authors. Financial support from Deutsche Forschungsgemeinschaft through CRC TRR 190 (Imai, Pace), the Bavarian Academy of Sciences and Humanities (Schwardmann), the NWO in the context of VIDI grant 452-17-004 (van der Weele), the Amsterdam Center for Behavioral Change of the University of Amsterdam, and the initiative, “A Sustainable Future,” at the University of Amsterdam (van der Weele), are gratefully acknowledged. The data collection was approved by the Ethics Committee Economics and Business (EBEC) of the University of Amsterdam, with request numbers EC 20191018121005, EC 20200810100845, and EC 20210713120725.

Climate action takes place against a backdrop of considerable uncertainty. Scientists face fundamental uncertainty about the exact contribution of different human activities to aggregate CO₂ emissions and about how emissions contribute to climate change and resulting extreme weather events^{1,2}. On top of that, non-experts may be uncertain about the best estimates of how a given behaviour maps onto CO₂ emissions. Understanding how uncertainty about CO₂ emissions affects climate action is therefore central to debates around the costs and benefits of communicating scientific uncertainty about climate change³⁻⁵ and around the effectiveness of carbon labels or other information campaigns^{6,7}.

This paper consists of three studies. Our first study investigates the degree of consumer uncertainty about emissions impact. We survey a large representative sample of the US population on their beliefs about the CO₂ emissions associated with several common consumer products. We elicit participants' entire belief distributions for each product using incentivized elicitation methods. The survey shows that people have a high degree of subjective uncertainty about the impact of all of these products, with only 6 per cent expressing high confidence in their guesses.

We then investigate two plausible mechanisms by which such uncertainty may causally affect emission mitigation behaviour. The first of these mechanisms is risk aversion. Uncertainty about the impact of emissions may reduce climate action if people do not place much value on avoiding high amounts of (potential) emissions. Theoretically, this requires that the marginal willingness to forgo personal benefits to avoid emissions declines with each additional unit of emissions. To test this mechanism, our second study uses large-scale online experiments in which participants' choices to buy an online convenience product affect real CO₂ emissions. We find that participants indeed have an increasing but concave willingness to mitigate (WTM) CO₂ emissions. However, contrary to the predictions of standard decision theory, we do not find an effect of uncertainty, as consumption of the polluting product is similar in treatments with and without uncertainty about emissions.

The second mechanism is the use of "moral wiggle room" that allows consumers to self-deceive into believing that emissions are small and consumption is harmless, as happens in other types of ethical decision-making⁸⁻¹¹. Such self-serving belief formation may also interact with carbon pricing: as CO₂-intensive products become more expensive, the temptation to form self-serving beliefs decreases⁹. In our third study, participants may again buy a virtual polluting product. They see a vague signal about the emissions associated with the product and have to update their beliefs about the emission size. The signal we use provides an opportunity for self-serving belief distortions, but only in a condition where the incentives to hold self-serving beliefs are known before the signal is seen. We find no evidence that uncertainty is exploited by

participants to develop self-serving beliefs that the emissions are low, nor do we find an interaction between prices and belief formation.

We contribute to the literature by documenting the extent of individuals' uncertainty about the carbon impact of their actions and by designing novel controlled experiments to study money-emission trade-offs. Our findings on subjective uncertainty complement previous evidence that people's best guesses about the size of CO₂ emissions are very heterogeneous and often too optimistic^{6,7}. Moreover, we are the first to study the explicit and isolated role of uncertainty in affecting climate action. In this way, we complement earlier studies on information provision, which focus on correcting the bias in consumer beliefs and have found conflicting evidence on carbon information provision^{6,7,12,13}. We show that the variance of beliefs has little impact on climate-friendly consumption. This demonstrates the limits of information campaigns to spur climate action, but it also suggests that communicating scientific uncertainty does not negatively affect consumption patterns.

In addition, our findings deliver new insights into the willingness to make sacrifices to reduce CO₂ emissions. Most studies in this literature have relied on hypothetical choices, which have been shown to lead to large overstatements in willingness to pay for offsets^{14,15}. A few studies, like ours, use incentivized willingness to pay to retire offsets or permits but have looked at a single emission quantity^{16–18}. We empirically show that the valuation of such offsets is concave in emission size, in line with concurrent evidence¹⁵. This suggests that future welfare measurements can take such nonlinearity into account when scientists estimate the welfare benefits of climate policies, study the demand for carbon offsets in particular industries like aviation¹⁹ or the optimal pricing of offsets more generally²⁰. Finally, we contribute to the literature on motivated cognition with our null results on self-serving environmental beliefs and on the interaction between beliefs and prices^{8,9,21,22}.

Results

Study 1: Subjective uncertainty about CO₂ emissions. We investigate subjective uncertainty with a representative sample of the US population ($N = 1,128$). Our survey elicited participants' beliefs regarding the carbon emissions associated with 12 common consumer products and activities. The elicitation process comprises two steps. First, we asked participants to indicate their best guess about the size of the emissions associated with each product. Then, we elicited the participants' entire subjective probability distribution regarding the emissions size of each product. To do so, we partitioned the interval of possible emissions into five bins centred around the participants' initial guesses. The three central bins have a width equal to 10% of the initial guess;




Figure 1: The percentage of participants who satisfy three subjective uncertainty criteria. A participant is *Not certain* if she puts at least one ball in a bin that does not contain her initial guess. She is *Less than 50% certain* if she puts less than 10 balls in the bin which contains her initial guess. She is *Extremely uncertain* if she puts four or fewer balls each in the bin that contains her initial guess and in the two adjacent bins. The three criteria are cumulative. The *Extremely uncertain* participants are also *Less than 50% certain*. The participants who are *Less than 50% certain* are also *Not certain*.

the two outer bins contain all the other possible emissions values. We asked the participants to allocate 20 balls across these bins to reflect the probability that the size of the emissions belongs to each of them. To incentivize truthful responses, we employed state-of-the-art techniques from experimental economics^{23,24} (see [Methods](#)).

We study participants' uncertainty with three (nonexclusive) nonparametric criteria based on the ball allocation. A participant is *Not certain* if she puts at least one ball in a bin that does not contain her initial guess. She is *Less than 50% certain* if she puts less than 10 balls in the bin which contains her initial guess. She is *Extremely uncertain* if she puts four or fewer balls in each of the three central bins, indicating less than 20% confidence in being (close to) the correct answer.

Figure 1 depicts the levels of uncertainty for each product. In each case, at least 94% of the participants are *Not certain*, at least 85% are *Less than 50% sure*, and at least 35% are *Extremely uncertain*. The figure also displays limited variability across products, indicating that people's uncertainty is not domain-specific. These results show that people are aware of having very limited knowledge about the carbon footprint of common products and activities.

Study 2: Does risk aversion towards carbon impact increase polluting activity? This experiment tests whether people are risk-averse towards reducing carbon emissions. In the experiment, participants had to work on an effortful and tedious en-

coding task. Before engaging in the task, subjects were offered the option to obtain a computer code that would do the encoding for them, reducing their burden. However, the code was associated with the release of CO₂ emissions into the atmosphere. Across conditions, we varied the uncertainty about the size of these emissions. In addition, we measured willingness to mitigate (WTM) CO₂ emissions via an incentive-compatible multiple price list. We define the WTM as the monetary amount a participant requires in order to accept the emission of a certain quantity of CO₂. For every subject, we elicit the WTM to avoid 0, 4, 8, 12, 16, or 20 kg of CO₂ emissions. The decisions were incentivized: participants could really accrue private benefits in exchange for increasing the CO₂ emissions in the atmosphere (see [Methods](#)).

According to Expected Utility Theory, the classical theory of decisions in economics, risk aversion is driven by a concave WTM, that is, a declining marginal willingness to make sacrifices for CO₂ reductions. Intuitively, an increase in the variance of the emissions will increase the probability of both very low and very high emissions. A consumer with a concave WTM will heavily discount the very high emissions, decreasing the subjective value of offsetting these emissions, even if the expected emissions remain constant. Consequently, consumers with a concave WTM should be more likely to engage in polluting actions as the (subjective) variance of the emissions grows.

To understand whether WTM is concave, Figure 2a shows the WTM curve aggregated over all subjects. The figure shows that subjects, on average, indeed display a diminishing WTM to reduce carbon emissions. They are willing to sacrifice about £2.7 to avoid 4 kilograms of CO₂ emissions, whereas to avoid 20 kilograms of emissions, they are only willing to forgo only £4. That is, the WTM increases by less than 50% when the amount of CO₂ increases by 500%. The graph shows that this effect is robust if we exclude subjects who have decreasing valuations (a possible sign of confusion) and who are top-censored (i.e. they select the maximum WTM of £7 at least once, which could produce concavity as an artefact). Supplementary Information A.2.2 contains further details about the variation in concavity across subjects. Supplementary Information A.2.4, instead, discusses how neither cognitive uncertainty²⁵ nor concave moral evaluations seem to explain the concavity of the WTM. It also discusses how, given the framing of the WTM elicitation, the concavity we found is more likely due to a marginally decreasing disutility from CO₂ emissions, which are seen as a loss, rather than decreasing marginal utility from implementing offsets, which are seen as a gain.

We now turn to our main outcome, the purchase of the polluting computer code, which simplifies the real effort task. Before purchase, participants were randomized into two treatments. In the *Information* treatment, the participants knew that the emissions from buying the product were equal to 4kg of CO₂. Instead, in the *Uncertainty* treatment, they knew that there was a 40% chance that the emissions were equal to

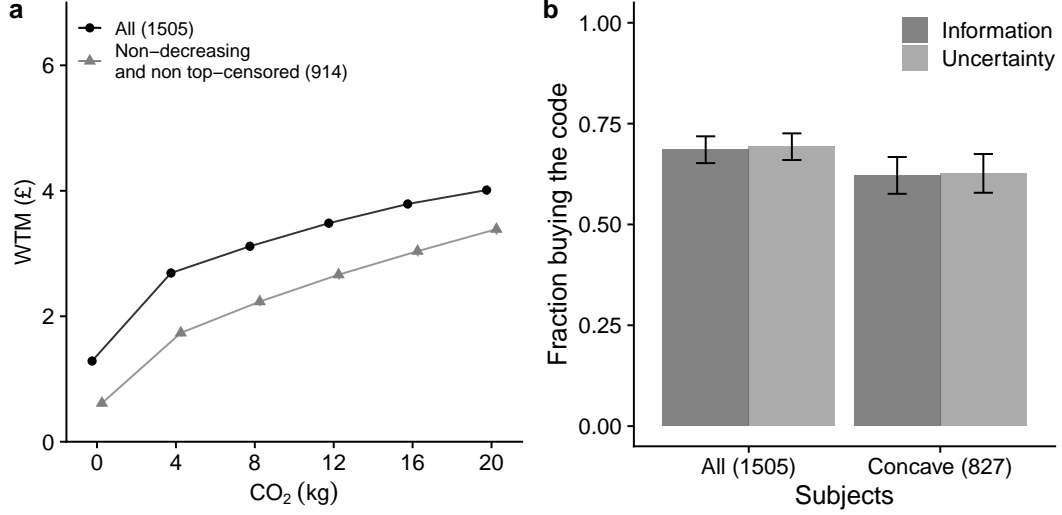


Figure 2: (a) Aggregate WTM. (b) Purchasing decision and uncertainty. *Notes:* In panel (b), the light-grey bars correspond to the group of subjects whose WTM exhibits concavity. Bars indicate 95% CI.

0kg, a 20% chance that the emissions were 4kg, and a 40% chance that they were 8kg. Thus, the expected emissions are the same in both cases, but the variance is larger in the *Uncertainty* treatment.

In line with the reasoning above and our finding that WTM is concave, we hypothesize that uncertainty increases the fraction of subjects who purchase the polluting product. We do not find support for this hypothesis. The left bars in Figure 2b show that purchasing decisions are similar in the two treatments (68.5% vs. 69.3%, $z = 0.3168$, two-sided $p = 0.7514$). It is possible, however, that there are offsetting effects for subjects with concave and convex WTM. The right bars in Figure 2b show the treatment difference only for participants with a concave WTM. Again, we find little evidence for the hypothesized effect (62.2% vs. 62.7%, $z = 0.1494$, $p = 0.8813$).

To provide further statistical backup, we run several regression models, in which we regress the purchasing decision on the treatment, several concavity scores, and the interaction of these two (Supplementary Table 4). We also include the average WTM and several demographic controls. We find little evidence for our hypothesis. Only a single specification of concavity score produces a significant interaction with the information treatment dummy, yet this effect is not robust to other concavity measures. However, we do find a statistically significant negative effect of the average WTM on purchasing decisions, showing that the WTM data is predictive of subjects' decisions to get the computer code. We also find some demographic effects, as young people are less likely to buy than old, women less than men, and left-wing less than right-wing (Supplementary Figure 2).

Study 3: Do participants form self-serving beliefs? We now turn to a second potential channel through which uncertainty may matter: the formation of motivated beliefs. As in study 2, participants made trade-offs between monetary gains and CO₂ emissions. These trade-offs were framed as consumer decisions: participants could buy a virtual product, represented as a nondescript wrapped package on the screen. The purchase increased their pay-off by the product value of £2, minus the price at which we, the experimenters, offered it. However, purchasing the product resulted in the emission of CO₂ into the atmosphere.

We conducted an online experiment ($N = 714$) that consisted of a *Motivated* and an *Unmotivated* treatment. Both treatments featured uncertainty about the size of the emissions associated with the product. Unlike in the risk aversion experiment discussed above, we did not state explicit probabilities, as there is evidence that more ambiguous settings are conducive to the formation of motivated beliefs^{26,27}. Participants could reduce uncertainty by engaging in an attentional task, somewhat akin to reading a product label, after which we elicited their beliefs about the emission size.

To study motivated beliefs, we vary whether participants complete the attentional task before or after knowing about the emission-money trade-off. Manipulating the timing of knowledge of the incentive scheme is a standard design feature in experiments studying motivated cognition^{28–30}. Thus, in the *Unmotivated* treatment, subjects were presented with the task before they knew any other details of the experimental design. In this way, participants have no self-serving motive to distort their attention or beliefs in the direction of their economic interest. By contrast, in the *Motivated* treatment, participants engaged in the task after reading the full experimental instructions. Therefore, they knew that higher numbers corresponded to higher CO₂ emissions and indicated a more “inconvenient” trade-off. We hypothesize that the latter treatment will lead to motivated beliefs, i.e. a lower estimate of the impact of the emissions.

Moreover, we hypothesized that motivated beliefs are more pronounced when the payoffs from the product are higher^{9,22}. To test this last hypothesis, we implemented three price treatments that varied the price of the product: a low price (£0.25), a medium price (£1), and a high price (£1.75). The information treatments are orthogonal to the price treatments, creating a 2×3 design (see [Methods](#) for details). While the trade-offs in our experiment are not about concrete, branded products, this simple design allows us to fully control the price of the products and the uncertainty about emissions while ensuring that beliefs remain, on average, correct. The average belief in the *Unmotivated treatment* is 62.7, not significantly different from 60—the real value of the emissions ($t(303) = 1.31$, 95% CI [58.1, 67.3], $p = 0.25$, two-sided).

We find no evidence of the formation of self-serving beliefs. The left panel of Figure 3 shows the distribution of beliefs in both treatments, where the spikes are driven by the

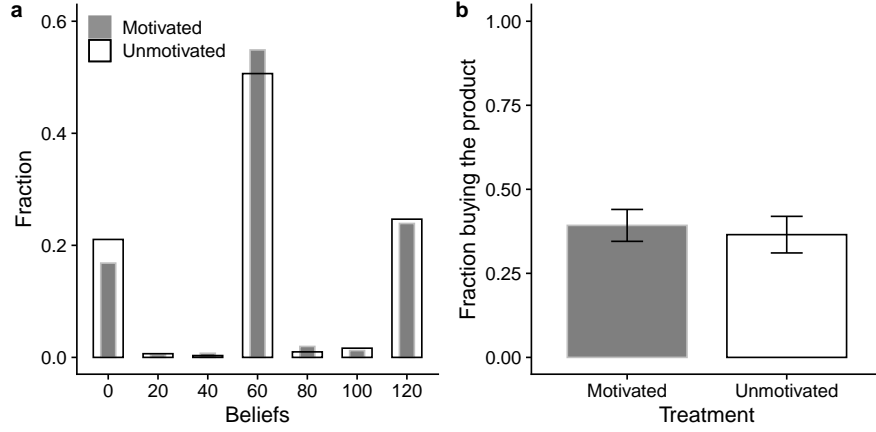


Figure 3: (a) Distributions of beliefs in the *Motivated* and *Unmotivated* treatment. (b) Fraction of participants buying the product in the *Motivated* and *Unmotivated* treatment. *Notes:* In panel (b), bars indicate 95% CI.

nature of the perceptual task (see [Methods](#)). We cannot reject the hypothesis that beliefs about emissions are the same in both treatments (Fisher’s Exact test, $p = 0.66$, two-sided). Moreover, Table 1 provides the results of regressing beliefs from both treatments on a dummy for the *Motivated* treatment controlling for individual characteristics. The coefficient in column (1) is positive, indicating that, if anything, participants in the *Motivated* treatment believe that emissions are larger. Hence, the direction of the effect is the opposite of what theories of self-serving beliefs and cognitive dissonance would predict. However, the coefficient is not significantly different from zero ($t(659) = 0.60$, 95% CI $[-4.36, 8.20]$, $p = 0.548$, two-sided). Furthermore, we find no evidence that the *Motivated* treatment causes people to spend less time looking at the information as we discuss in Supplementary Information A.3.2.

This null result obtains despite substantial ambiguity: only 51% of participants answered the belief question correctly in the *Unmotivated* treatment, even though they spent, on average, 50 seconds on the task screen. Moreover, the standard deviation of beliefs is 0.65 times the average, indicating that the degree of uncertainty in the experiment is sizable. Thus, there was room for motivated subjects to perceive emissions to be lower than they actually were.

We also test for differences in purchasing behaviour between the two treatments. The right panel of Figure 3 shows that behaviour is similar in the two treatments. Both a Fisher’s exact test and a t -test based on column (2) of Table 1 fail to reject that subjects are equally likely to buy the product in the *Motivated* and in the *Unmotivated* treatment (Fisher’s Exact test: $p = 0.483$; t -test: $t(695) = 1.07$, 95% CI $[-3.33, 11.34]$, $p = 0.285$; both tests are two-sided). By contrast, a higher product price has a strong independent impact on purchasing behaviour: Column (4) of Table 1 shows that an increase in the price of one pound leads to a 13 percentage point decline in purchases.

Table 1: Comparison between the *Motivated* and the *Unmotivated* treatments.

	Beliefs		Units	
	(1)	(2)	(3)	(4)
Motivated	1.921 (3.199)	−4.406 (6.226)	0.040 (0.037)	0.051 (0.074)
Price		−4.242 (4.127)		−0.128** (0.045)
Price × Motivated		6.172 (5.215)		−0.022 (0.060)
Controls	Yes	Yes	Yes	Yes
Observations	695	695	695	695
R^2	0.050	0.052	0.092	0.122

Notes: Models include observations from the *Motivated* and the *Unmotivated* treatments, where the latter is the baseline. The regressions exclude 19 participants for which we failed to record the demographic characteristics. Dependent variables are beliefs in columns (1) and (2) and purchasing decisions (1 if the participant purchased the product) in columns (3) and (4). Control variables: age, gender, student status, education (6 categories), frequency of car usage (5 categories), and nationality (27 categories). Robust standard errors are reported in parentheses. *: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$.

Discussion

We document that people are uncertain about the CO₂ emissions associated with common consumer products. Yet we find no evidence that uncertainty affects climate action through either risk aversion or the formation of motivated beliefs about the magnitude of emissions.

Our findings add to a growing literature on the effect of information provision on climate-related behaviours. While information campaigns are popular among politicians, previous studies have found mixed evidence on the effect of correcting beliefs about emissions: some labelling studies find small reductions in emissions from informing consumers about the carbon impact of e.g. meat products^{6,13,31,32}, whereas other studies do not find such an effect⁷. Our findings suggest that while consumer information may reduce uncertainty, this alone may be insufficient to spur voluntary reductions in emissions. At the same time, our results also suggest that scientists can be upfront about the uncertainty of their estimates without fear of providing an excuse for polluting behaviour.

Another set of implications stems from our finding that consumers’ willingness to mitigate the first units of CO₂ is much higher than for subsequent units. The first is that consumers’ behaviour will be sensitive to reference points and to the framing of emission impact. For instance, framing multiple emission events separately may lead to a higher willingness to avoid or offset them than framing them as a single event. Future research

could explicitly test this prediction. A second implication concerns the use of people’s WTM to offset emissions to calculate the benefits of climate policies. The literature has so far relied on linear extrapolations of the WTM for a single emission amount in order to compute the benefits of reducing one ton of CO₂ in the atmosphere^{16–18}. Our results strongly suggest that future studies should measure individual WTM to avoid several emission amounts and should use nonlinear models to produce accurate estimates of these benefits. It is reasonable to expect that the two implications above hold, albeit we did not find that the concavity of the WTM curve generates risk aversion toward emissions. The participants’ WTM is significantly associated with them buying the polluting product. Hence, the WTM is a meaningful predictor of consumption behaviour and environmental preferences.

One limitation of our research is that, while we do not find evidence for the formation of motivated beliefs about emission size, it is possible that consumers find other dimensions for motivated reasoning about climate change²¹. Further research could investigate these additional dimensions.

Methods

General remarks

The implementation of the CO₂ emissions in the experiments. In both of our experiments, participants were asked to make decisions that could result in the emission of CO₂. To ensure that these emissions are consequential, we prepared a monetary transfer to Carbonfund.org, an organization that offsets CO₂ emissions. Every time a participant made a decision resulting in CO₂ emissions, we decreased the amount of our transfer. We explicitly communicated this procedure to the participants. We always framed these decisions as choices between private benefits and emitting CO₂ to enhance the external validity and maximize the salience of the emissions. We took several measures to assure participants of the tangible nature of these CO₂ emissions associated with the purchase of the product. We emphasized the role of the no-deception policy in obtaining ethical approval for the experiment. Additionally, we promised participants to send them the invoice for the donation to Carbonfund.org (see Supplementary Information B.5) and actually did so. To ensure that the participants believed that their actions had environmental consequences, we asked them whether they believed that we would buy the CO₂ offsets as described in the instructions. We find that between 80% and 84% of the participants expressed trust in our intentions.

Data quality. We took several measures to maximize data quality as described in Supplementary Information B.4.

Study 1: survey about subjective uncertainty

Our survey measures consumers’ prevailing beliefs about the CO₂ emissions generated in the production of common consumer goods. The survey then continues with additional questions that we describe and analyze elsewhere⁷.

Our elicitation methods employed incentive-compatible payment schemes developed in the experimental economics literature while maintaining simplicity and participant-friendliness in both instructions and the interface to allow for a representative sample to participate. Below, we elaborate on each of the elicitation procedures in more detail. Supplementary Information B.4 presents additional information about the steps we took to maximize the data quality.

Belief elicitation

In the survey, we elicited participants’ beliefs about the CO₂ emissions generated by driving one mile by car. We then elicited beliefs about 12 common consumer products and activities listed in Table 2, including food items, household appliance use, and transportation. To enhance understanding, we provided participants with information about product specifications and the type of emissions considered by the estimate.^a We retrieved these estimates from top-tier academic journals or from the estimates the UK government uses for its environmental regulations. We disclosed these scientific sources only at the end of the experiment. To make the values more meaningful to subjects, we refrained from eliciting emissions in grams and instead asked about the number of miles one must drive by car to emit the same amount of CO₂ as the product in question, an approach in line with previous studies⁶.

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

^aParticipants could learn a detailed breakdown of the factors considered in the scientific calculations of CO₂ emissions size. See Supplementary Table 12.

Table 2: List of consumer products and activities.

	Quantity	Emission size		Source
		Estimate	Unit	
Beer	12 fl oz	1.46	mile	Poore and Nemecek ¹
Phone call	1 hour	1.55	mile	Smith et al. ³³
Microwave	1000W, 2 hour	1.76	mile	UK BEIS ³⁴
Milk	1 cup	2.60	mile	Poore and Nemecek ¹
Egg	6 eggs	4.81	mile	Poore and Nemecek ¹
Poultry meat	7 oz	6.78	mile	Poore and Nemecek ¹
Shower	Average usage	3.90	mile	Hackett and Gray ³⁵
Dark chocolate	100g	16.03	mile	Poore and Nemecek ¹
Coffee beans	1 lb	44.41	mile	Poore and Nemecek ¹
Beef	7 oz	68.39	mile	Poore and Nemecek ¹
Flight	SFO to LAX	304.60	mile	UK BEIS ³⁴
Gas heating	One month	606.68	mile	Padgett et al. ³⁶

the scientific estimate^{23, b}

To understand the participants’ confidence in their answers, we then elicited the subjective probability distribution of the size of CO₂ 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 told participants that each bin represents an interval that might contain the scientific estimate and that they should allocate the balls to represent their level of confidence that the estimate is, in fact, in that bin. A screenshot of the interface is in Supplementary Information B.1. We incentivized the elicitation by randomly selecting one of the bins and scoring the answer according to a randomized quadratic scoring rule. This mechanism encourages participants to truthfully reveal their belief that the scientific estimate falls in a particular bin³⁷. To keep things simple and avoid information overload, we did not provide participants with the exact details of the scoring rule but made them available with a mouse click. We told subjects that they would maximize their expected earnings by answering truthfully, an approach suggested by Danz et al.²⁴

Implementation We recruited 1,430 participants on Prolific (<https://www.prolific.co/>) between the 3rd and 6th December 2020, and 1,128 completed all the belief elicitation questions described in this paper. We restricted participation to US residents, and we aimed to collect a sample representative for age, gender, and ethnicity.^c Our sample

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

^cWe compared the demographic characteristics of study participants and information from US Census Bureau³⁸, and confirmed that our sample is representative for gender and ethnicity, but not

is, on average, 42.6 years old ($SD = 15.4$), and 48.5% of the participants self-identified as male. Supplementary Table 1 shows the demographic characteristics of the sample. In total, participants were asked to answer 21 comprehension questions.

At the end of the survey, we randomly selected one question from the entire study per participant. There was a 30% probability that one of the belief questions analyzed above was selected for payment. The average duration of the survey was 51 minutes, but all the questions described here were in the first half of the study.

Study 2: Risk Aversion experiment

The experiment consists of two parts. In the first part, we measure the participants' valuations of CO₂ emissions. In the second part, we ask participants whether they want to get a valuable but polluting product. The instructions for this experiment are in Supplementary Information C.2. We preregistered the experimental design and the hypotheses on AsPredicted.org (see Supplementary Information B.6).

Valuation of CO₂ emissions. We measured the participants' valuation of CO₂ emissions by presenting them with trade-offs between money and emissions. Specifically, participants were offered a choice between Option A and Option B. Opting for Option A meant forgoing any monetary gain but preventing the generation of CO₂ emissions. In contrast, choosing Option B allowed them to earn money but resulted in CO₂ emissions. For a given CO₂ emission level, the participants had to make 15 choices where the amount of money they could earn increased from £0 to £7 in 50 pence increments. These decisions were embedded in a Multiple Price List (see Supplementary Information C.2), which enforced a single switching point between Option A and Option B. This switching point gives us the participants' valuation for a given amount of CO₂ emissions. To gauge the participants' certainty regarding their valuations, we employed the “cognitive uncertainty” elicitation method developed by Enke and Graeber²⁵. This elicitation was skipped if the participants never switched from Option A to Option B. In total, the participants saw six Multiple Price Lists, each corresponding to one of the following emissions levels: 0 kg, 4 kg, 8 kg, 12 kg, 16 kg, and 20 kg. Half of the participants saw these lists in ascending order of emission size, while the other half saw them in descending order.

Consumption decision. In the second part of the experiment, the participants had the opportunity to receive a valuable yet polluting product, framed as a “computer code”, which significantly speeded up the completion of a time-consuming and laborious real-effort task. The task involved typing 15 strings, each consisting of 15 characters, in

for age (Supplementary Table 2).

reverse order (Supplementary Figure 6). The participants were required to transcribe these strings flawlessly to complete the task: any mistakes incurred resulted in an error message indicating the specific strings that required correction before they could proceed with the experiment. Those who obtained the computer code were presented with a task in which the code generated all the correct answers, reducing their task to simply clicking one button to submit them. To ensure task completion without interruption and to make the task sufficiently “annoying,” we implemented an attention check. The participants saw a warning sign every 30 seconds, and upon its appearance, they had a 5-second window to click on a button on the screen to confirm their active engagement with the task. The participants were excluded from the experiment if they failed to click on the button within the specified time window in more than four instances.

At first, all the participants were required to complete the real-effort task once without the help of the computer code. In the second round, they had the chance to obtain the code to speed up the completion time. On average, the participants spent 8 minutes and 40 seconds to complete the first task.

Obtaining the code came at the cost of emitting CO₂. There are two treatments that vary the information participants had about the size of the CO₂ emissions. In the *Information* treatment, participants were informed that purchasing the product would result in emissions equivalent to 4kg of CO₂. In contrast, in the *Uncertainty* treatment, participants were informed of a probability distribution: a 40% probability of emissions being 0kg, a 20% probability of emissions being 4kg, and a 40% probability of emissions being 8kg.

Survey. Before the participants complete (or are allowed to skip) the second real effort task, they complete a questionnaire. A battery of questions elicits people’s moral evaluations of emissions-money trade-offs. Participants indicate the morality acceptability of emitting 4, 12, and 20kg of CO₂ for either £1 or £5 using a Likert scale, where 1 indicates a decision that is “morally very inappropriate” and 7 a decision that is “morally very appropriate”. Further questions in the survey ask about attitudes toward climate change and demographic characteristics.

Sample and data collection. We recruited a total of 1,935 participants through the online platform Prolific.co on January 5th, 2023. Following the preregistration, 1,505 participants who successfully completed the final survey were included in the analysis. Among them, 753 participants were assigned to the *Information* treatment, while the remaining 752 were assigned to the *Uncertainty* treatment. Fifty per cent of the participants identified as females, and the average age is 39 years old (min = 18, max = 79, SD = 12.51). We restricted participation to individuals based in the UK. Subjects earned a fixed reward of £3, with the potential for a bonus payment

based on their decisions. On average, they earned £3.66, and they took about 34 minutes to complete the tasks. Following the participants’ decisions, we donated \$160 to Carbonfund.org to offset CO₂ emissions, resulting in a reduction of almost 13 metric tons of CO₂ in the atmosphere. The experiment maintained participant anonymity, as we identified participants only via alphanumerical IDs provided by the platform. Participants knew that they were taking part in an experiment.

Study 3: Motivated Belief experiment

In this experiment, we offered participants the opportunity to buy a single unit of a virtual product.^d If participants decided to purchase the product, they increased their pay-off by the product value of £2, minus the price at which we, the experimenters, offered it. Importantly, purchasing the product entailed the emission of CO₂ into the atmosphere. We framed the experiment as a market interaction, employing terminology such as “virtual product” and “price,” to make it closer to a real-life purchasing situation.

We informed participants that purchasing the product would result in CO₂ emissions equivalent to burning 60 litres of gasoline. This emission size had offset costs of £1.07, commensurate to the other payments in the experiments.^e

As we describe next, we orthogonally implemented three price treatments and two information treatments, resulting in a total of six treatments. Each participant was assigned to a single treatment. The instructions for this experiment are in Supplementary Information C.3. We preregistered the experimental design and the hypotheses on AsPredicted.org (see Supplementary Information B.6).

Uncertainty treatments. Our primary focus is on the formation and impact of beliefs about the CO₂ emissions associated with the product, as well as the role of emission information. To study this, we employed two treatments, called the *Motivated* and *Unmotivated* treatments, that varied the nature of uncertainty about the size of the emissions.

In both treatments, there was uncertainty about the size of the emissions. The participants had the opportunity to reduce this uncertainty by engaging in an attentional task designed to mimic real-life consumption scenarios, such as reading product labels or conducting an online search for information. The task involved examining a matrix

^dThis product is virtual in the sense that it exists only inside the experiment; it is not a physical product nor a service. Nevertheless, the product is valuable to the participants since their pay-off from the experiment increases if they “buy” it.

^eUsing a report from the US Environmental Protection Agency³⁹, we calculate that burning 60 litres of gasoline produces 140kg of CO₂ emissions. At the time of the experiment, Carbonfund.org offset 1 metric ton of CO₂ per every \$10 (or £7.9) it receives in donations, so offsetting the products’ emissions cost £1.07.

of numbers, with the most frequently appearing number representing the emission size, measured in terms of the CO₂ emissions generated by burning a litre of gasoline. We adapted this task from Sandro Ambuehl’s work⁴⁰, which shows that the information-gathering strategy in this task is influenced by incentives for subsequent decisions.^f Participants had up to one minute to engage with the task, after which we elicited their beliefs by asking them which number they believed was the most frequently found in the table. Providing the correct answer was rewarded with a bonus of £0.10, which incentivized participants to report the mode of their belief distribution^{23,g}.

These two treatments differ in the order in which we presented the attentional task and the information about the emission size. In the *Unmotivated* treatment, participants were presented with the task prior to receiving any instruction about the possibility of emitting CO₂. In this way, we eliminated any self-serving motives that might lead participants to distort their attention or beliefs in the direction of their economic interests. By contrast, in the *Motivated* treatment, participants engaged in the task after they had read the full experimental instructions. Consequently, they were aware that the correct answer to the task indicated the magnitude of the CO₂ emissions as well as of the surplus they could obtain from the product. This treatment enables us to test if the surplus from buying the product has a causal effect on participants’ belief formation in the attentional task and, in turn, on the product’s purchase.

At the end of the experiment, we collected demographic information using a survey.

Price treatments. We also investigate the relationship between purchasing decisions and prices. The effect of prices provides a natural quantitative benchmark for assessing the effects of beliefs and information, as price incentives are a primary tool used by policymakers and are often discussed in the context of reducing carbon emissions.

We implemented three price treatments that varied the price of the product: a low price of £0.25, a medium price of £1, and a high price of £1.75. Participants were informed that the price was randomly assigned and held no informational content regarding emission size. We made sure of the participants’ understanding of this aspect by asking them a comprehension question on the topic.

^fThe task can be found in Supplementary Information B.3. The matrix contained a total of 143 numbers drawn from the set {0, 20, 40, 60, 80, 100, 120}. The number 60, the most frequently occurring, appeared 35 times, with 0 and 120 being the next most frequent, each appearing 26 times. All other numbers appeared 14 times each.

^gNote that the experiment had a third treatment which gave people full information about the carbon impact of the product. This treatment is analyzed and described elsewhere⁴¹. We do not analyze it here, as the comparison between the info and the uncertainty treatments is complicated by the fact that we do not have full information about the belief distributions. Hence, we cannot be sure that the expected value of emissions is constant in the different treatments. Experiment 1 in this study, therefore, provides higher quality evidence on the effect of precise information on behaviour.

Sample and data collection. We recruited 714 participants using Prolific.co, an online platform, between 9th and 11th May 2019.^h Of those, 304 participants were assigned to the *Unmotivated* treatment (87 faced a £0.25 price, 107 a price of £1.00, and 110 a price of £1.75), and 410 participants were assigned to the *Motivated* treatment (146 faced a £0.25 price, 125 a price of £1.00, and 139 a price of £1.75). Demographic information for 19 participants was not successfully recorded (11 and 8 from the *Unmotivated* and *Motivated* treatment, respectively). These 30 subjects are included in the analysis when we run non-parametric tests, but they are excluded in the regression analysis, which includes the demographic controls.

Fifty per cent of the participants identified as females, 42% are students, and the average age is 29 years old. We accepted only EU nationals as participants. The most represented countries in our sample are the UK (34.7%), Poland (14.54%), and Portugal (12.1%). Participants earned a fixed reward of £1.60, with a potential bonus payment contingent on their decisions. On average, they earned £2.04, and they took less than 13 minutes to complete the tasks. We obtained participants’ demographic information, including gender, age, student status, and nationality, directly from Prolific.co. Following the participants’ decisions, we donated \$911.40 to Carbonfund.org to offset CO₂ emissions, resulting in a reduction of over 90 metric tons of CO₂ in the atmosphere.

The experiment maintained participant anonymity, as we identified participants only via alphanumerical IDs provided by the platform. Participants knew that they were taking part in an experiment.

References

- [1] Joseph Poore and Thomas Nemecek. Reducing food’s environmental impacts through producers and consumers. *Science*, 360(6392):987–992, 2018.
- [2] Robert S Pindyck. What we know and don’t know about climate change, and implications for policy. *Environmental and Energy Policy and the Economy*, 2(1): 4–43, 2021.
- [3] David V. Budescu, Han-Hui Por, Stephen B. Broomell, and Michael Smithson. The interpretation of IPCC probabilistic statements around the world. *Nature Climate Change*, 4(6):508–512, 2014.

^hResearchers have recently voiced concern that bots, instead of humans, might complete online experiments. To the best of our knowledge, no such concerns have been raised regarding Prolific.co. Nevertheless, we took two precautionary measures to further minimize the risk of undetected bots completing our experiment. First, we incorporated two “honey-pots,” which are questions that a human participant would not be able to see, but that a bot that reads the source code of the experimental program should identify as questions to be answered. Second, we kept track of the number of attempts participants needed to answer the comprehension questions correctly. We did not find any evidence that bots completed our study. Additional details are in Supplementary Information B.4.

- [4] Baruch Fischhoff and Alex L. Davis. Communicating scientific uncertainty. *Proceedings of the National Academy of Sciences*, 111(supplement_4):13664–13671, 2014.
- [5] Stephen B. Broomell and Patrick Bodilly Kane. Public perception and communication of scientific uncertainty. *Journal of Experimental Psychology: General*, 146(2):286, 2017.
- [6] Adrian R. Camilleri, Richard P. Larrick, Shajuti Hossain, and Dalia Patino-Echeverri. Consumers underestimate the emissions associated with food but are aided by labels. *Nature Climate Change*, 9(1):53–58, 2019.
- [7] Taisuke Imai, Davide D. Pace, Peter Schwardmann, and Joël J. van der Weele. Correcting consumer misperceptions about CO₂ emissions. CESifo Working Paper No. 10138, 2022.
- [8] Roland Bénabou and Jean Tirole. Mindful economics: The production, consumption, and value of beliefs. *Journal of Economic Perspectives*, 30(3):141–164, 2016.
- [9] Nina Hestermann, Yves Le Yaouanq, and Nicolas Treich. An economic model of the meat paradox. *European Economic Review*, 129:103569, 2020.
- [10] Rafael Di Tella, Ricardo Perez-Truglia, Andres Babino, and Mariano Sigman. Conveniently upset: Avoiding altruism by distorting beliefs about others’ altruism. *American Economic Review*, 105(11):3416–3442, 2015.
- [11] Christine L. Exley. Excusing selfishness in charitable giving: The role of risk. *Review of Economic Studies*, 83(2):587–628, 2015.
- [12] Matthias Rodemeier and Andreas Löschel. The welfare effects of persuasion and taxation: Theory and evidence from the field. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.3587339>, 2020.
- [13] Paul Lohmann, Elisabeth Gsottbauer, Anya Doherty, and Andreas Kontoleon. Do carbon footprint labels promote climatarian diets? Evidence from a large-scale field experiment. *Journal of Environmental Economics and Management*, 114:102693, 2022.
- [14] John A. List and Craig A. Gallet. What experimental protocol influence disparities between actual and hypothetical stated values? *Environmental and Resource Economics*, 20:241–254, 2001.
- [15] Matthias Rodemeier. Willingness to pay for carbon mitigation: Field evidence from the market for carbon offsets. IZA Discussion Paper 15939, 2023.

- [16] Gregory F. Nemet and Evan Johnson. Willingness to pay for climate policy: A review of estimates. Available at SSRN: <https://dx.doi.org/10.2139/ssrn.1626931>, 2010.
- [17] Andreas Löschel, Bodo Sturm, and Carsten Vogt. The demand for climate protection—Empirical evidence from Germany. *Economics Letters*, 118(3):415–418, 2013.
- [18] Johannes Diederich and Timo Goeschl. Willingness to pay for voluntary climate action and its determinants: Field-experimental evidence. *Environmental and Resource Economics*, 57(3):405–429, 2014.
- [19] Hannes Cordes, Stefan Baumeister, and Minna Käyrä. Factors influencing the willingness to pay for aviation voluntary carbon offsets: A literature review. *European Journal of Tourism Research*, 36:3602, 2024.
- [20] Marc N. Conte and Matthew J. Kotchen. Explaining the price of voluntary carbon offsets. *Climate Change Economics*, 1(02):93–111, 2010.
- [21] Shahzeen Z. Attari, David H. Krantz, and Elke U. Weber. Energy conservation goals: What people adopt, what they recommend, and why. *Judgment and Decision Making*, 11(4):342–351, 2016.
- [22] Peter Schwardmann. Motivated health risk denial and preventative health care investments. *Journal of Health Economics*, 65:78–92, 2019.
- [23] Karl H. Schlag, James Tremewan, and Joël J. van der Weele. A penny for your thoughts: A survey of methods for eliciting beliefs. *Experimental Economics*, 18(3):457–490, 2015.
- [24] David Danz, Lise Vesterlund, and Alistair J. Wilson. Belief elicitation and behavioral incentive compatibility. *American Economic Review*, 112:2851–2883, 2022.
- [25] Benjamin Enke and Thomas Graeber. Cognitive uncertainty. *Quarterly Journal of Economics*, 138(4):2021–2067, 2023.
- [26] Emily C. Haisley and Roberto A. Weber. Self-serving interpretations of ambiguity in other-regarding behavior. *Games and Economic Behavior*, 68(2):614–625, 2010.
- [27] Jan Engelmann, Maël Lebreton, Peter Schwardmann, Joël van der Weele, and Li-Ang Chang. Anticipatory anxiety and wishful thinking. Tinbergen Institute, 2019.
- [28] Linda Babcock, George Loewenstein, Samuel Issacharoff, and Colin Camerer. Biased judgments of fairness in bargaining. *American Economic Review*, 85(5):1337–1343, 1995.

- [29] Peter Schwardmann and Joël van der Weele. Deception and self-deception. *Nature Human Behaviour*, 3(10):1055–1061, 2019.
- [30] Uri Gneezy, Silvia Saccardo, Marta Serra-Garcia, and Roel van Veldhuizen. Bribing the self. *Games and Economic Behavior*, 120:311–324, 2020.
- [31] Anna Schulze Tilling. Changing consumption behavior with carbon labels: Causal evidence on behavioral channels and effectiveness. Mimeo, University of Bonn, 2023.
- [32] David Bilén. Do carbon labels cause consumers to reduce their emissions? Evidence from a large-scale natural experiment. Mimeo, Gothenburg University, 2022.
- [33] Anna Jo Bodurtha Smith, Imogen Tennison, Ian Roberts, John Cairns, and Caroline Free. The carbon footprint of behavioural support services for smoking cessation. *Tobacco Control*, 22(5):302–307, 2013.
- [34] UK BEIS. Greenhouse gas reporting: Conversion factors 2019, 2020. UK Department for Business, Energy & Industrial Strategy. <https://www.gov.uk/government/publications/greenhouse-gas-reporting-conversion-factors-2019>.
- [35] M. J. Hackett and Nicholas. F. Gray. Carbon dioxide emission savings potential of household water use reduction in the UK. *Journal of Sustainable Development*, 2(1):36–43, 2009.
- [36] J. Paul Padgett, Anne C. Steinemann, James H. Clarke, and Michael P. Vandenberg. A comparison of carbon calculators. *Environmental Impact Assessment Review*, 28(2-3):106–115, 2008.
- [37] Karl H. Schlag and Joël van der Weele. Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. *Theoretical Economics Letters*, 3:38–42, 2013.
- [38] US Census Bureau. National population by characteristics: 2020-2022, 2022. United States Census Bureau. <https://www.census.gov/data/tables/time-series/demo/popest/2020s-national-detail.html>.
- [39] United States Environmental Protection Agency. Emission facts: Average carbon dioxide emissions resulting from gasoline and diesel fuel. EPA420-F-05-001, February 2005. <https://nepis.epa.gov/Exe/ZyPURL.cgi?Dockkey=P1001YTF.txt>, 2005.

- [40] Sandro Ambuehl. An offer you can't refuse? Incentives change what we believe. Available at SSRN: <https://dx.doi.org/10.2139/ssrn.2830171>, 2016.
- [41] Davide D. Pace and Joël van der Weele. Curbing carbon: An experiment on uncertainty and information about CO₂ emissions. Tinbergen Institute Discussion Paper No. 2020-059, 2020.