Ecological Economics 227 (2025) 108401

Contents lists available at ScienceDirect

Ecological Economics

journal homepage: www.elsevier.com/locate/ecolecon

ANALYSIS

Uncertainty about carbon impact and the willingness to avoid CO_2 emissions[☆]

Davide D. Pace [a,e], Taisuke Imai [b,e], Peter Schwardmann [c,e], Joël J. van der Weele [d,*]

[a] Ludwig-Maximilians-Universität München, Germany
[b] Institute of Social and Economic Research, Osaka University, Japan
[c] Carnegie Mellon University, United States of America
[d] Amsterdam School of Economics, University of Amsterdam, Tinbergen Institute, Netherlands
[e] CESifo, Germany

ARTICLE INFO

Dataset link: https://osf.io/pu8q2/

JEL classification:
C91
D81
Q54

Keywords:
CO_2 emissions
Sustainable consumption
Economic experiments

ABSTRACT

Using data from a large representative survey, 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 then experimentally test two channels for the behavioral impact of such uncertainty, namely risk aversion about the impact of mitigating actions and the formation of motivated beliefs about this impact. In two novel 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.

1. Introduction

Climate action takes place against a backdrop of considerable uncertainty. Scientists face fundamental uncertainty about the speed of climate change and its effect on economic output (Berger et al., 2017; Pindyck, 2021; Barnett, 2023). On top of that, non-experts may be uncertain about how a given behavior maps onto CO_2 emissions. Understanding how uncertainty about CO_2 emissions affects climate action is therefore central to debates around the costs and benefits of communicating scientific uncertainty about climate change (Budescu et al., 2014; Fischhoff and Davis, 2014; Broomell and Kane, 2017) and around the effectiveness of carbon labels or other information campaigns (Camilleri et al., 2019; Imai et al., 2022).

However, it is unclear how such uncertainty about impact affects individual mitigation behavior. While there is a large literature on risky choices concerning monetary payoffs, mitigation behavior is different: given the small and delayed effect of individual choices on the climate, its effects are not felt physically. Instead, motivations are psychological in nature and depend on subjective beliefs, identity concerns, feelings of guilt, signaling, etc. Little is known about how such belief-based motivations for mitigation react to uncertainty and information provision about impact.

In this paper, we investigate the role of uncertainty in climate mitigation behavior. We first present evidence to motivate the importance of the research question. Using data from an incentivized survey in Imai et al. (2022), we show that people have little confidence in their estimates about the impact associated with common consumer products, and are indeed very uncertain about the CO_2 emissions.

We then present two novel experiments to investigate two potential ways in which uncertainty about carbon impact interacts with mitigation behavior. The first experiment studies whether people exploit uncertainty to form motivated beliefs that emissions are small and consumption is harmless, as happens in other types of ethical decision-making (Kunda, 1990; Bénabou and Tirole, 2016). Such motivated or

☆ This paper supersedes a previous version circulated with the title "Curbing Carbon: An Experiment on Uncertainty and Information about CO_2 Emissions" by a subset of the authors. Financial support from Deutsche Forschungsgemeinschaft, Germany through CRC TRR 190 (Imai, Pace), the Bavarian Academy of Sciences and Humanities, Germany (Schwardmann), the NWO, Netherlands in the context of VIDI grant 452-17-004 (van der Weele), the Amsterdam Center for Behavioral Change of the University of Amsterdam, Netherlands, and the initiative, "A Sustainable Future," at the University of Amsterdam, Netherlands (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.

* Corresponding author.
E-mail address: j.j.vanderweele@uva.nl (J.J. van der Weele).

https://doi.org/10.1016/j.ecolecon.2024.108401
Received 29 April 2024; Received in revised form 14 August 2024; Accepted 19 September 2024

self-serving beliefs can reduce guilt from emissions and hence obviate the need for behavior change. They can explain why people consistently underestimate the emissions associated with the products they consume (Camilleri et al., 2019; Imai et al., 2022) and the environmental and health consequences of CO_2 emissions (Semken, 2024). The formation of self-serving beliefs may also interact with carbon pricing: as CO_2-intensive products become more expensive, the temptation to form self-serving beliefs decreases (Hestermann et al., 2020; Serra-Garcia and Szech, 2022).

In our experiment, participants may buy a virtual product. The product is valuable, but buying it entails emitting CO_2. The participants receive 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 the participants to develop self-serving beliefs that the emissions are low, nor do we find an interaction between prices and belief formation.

The second experiment asks whether uncertainty affects behavior because people are risk averse over carbon impact. 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. A lot of work has focused on such risk aversion in monetary gambles, but we know little about the willingness to mitigate as a function of emissions amounts.

To test this mechanism, in our second experiment, a new set of participants is offered the choice to buy another valuable but polluting product. We find that participants have an increasing but concave willingness to mitigate (WTM) CO_2 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.

Our findings contribute to several strands of literature. First, we provide new evidence on the effect of information provision on climate-related behaviors. While information campaigns are popular among politicians, previous studies have found mixed evidence on the effect of correcting beliefs about emissions: some labeling studies find small reductions in emissions from informing consumers about the carbon impact of, e.g., meat products (Camilleri et al., 2019; Lohmann et al., 2022; Tilling, 2023), whereas other studies do not find such an effect (Imai et al., 2022). However, these studies do not isolate the effect of uncertainty, as they study interventions that also shift the expected size of the emissions and the beliefs about social norms. Other papers run public good games in the lab to study the effects of uncertainty about tipping points on individual contributions and climate negotiations (Barrett and Dannenberg, 2012, 2014). In these studies, uncertainty manipulation also shifts beliefs about others' contributions and the equilibrium of the game. Our experiments instead identify the isolated effects of uncertainty on behavior.

Second, we contribute to the literature on motivated cognition. Lind et al. (2019) and Stoetzer and Zimmermann (2024) show that people do not avoid free information about the environmental consequences of their actions, while Epperson and Gerster (2021) shows avoidance of information concerning animal welfare in the meat industry. We complement the existing evidence with a null results about self-serving beliefs distortion about the size of environmental externalities. In addition, we test, but we do not find support for the theoretical prediction that the size of motivated beliefs depends on the incentives to self-deceive, as theoretically shown in Schwardmann (2019) and Hestermann et al. (2020).

A final contribution is methodological. Our experimental protocol, in which participants can buy valuable but polluting products, can be adapted to investigate related questions about the behavioral economics of climate change. Indeed, Schöller and Ulmer (2023), He et al. (2024), and Woerner et al. (2023) run experiments whose design builds upon Experiment 1 in this paper.

The outline of the paper is as follows. Section 2 formulates an illustrative theoretical framework on how motivated beliefs and risk aversion affect consumers' actions. Section 3 reanalyzes the data from Imai et al. (2022) and provides the motivating evidence for the two novel experiments presented in Section 4. Section 5 discusses possible reasons for the null results, and additional results about visual attention and the psychological mechanisms that could generate a concave WTM. Section 6 discusses the implications of our findings.

2. Theoretical framework

The aim of this section is to show theoretically why uncertainty about emissions size might matter for consumer behavior. The proofs are in Online Appendix A.

Since mitigation activities do not have a direct, perceptible impact on the consumer, we model the utility function as dependent on beliefs about impact. Consider an agent whose utility is given by:

$$U(b \mid v, m) = bv - m\left[bE_{\hat{F}}[w(c)] - \sum_{c\in C}(\hat{f}_c - f_c)^2\right] - (1-m)bE_F[w(c)],$$

where $b \in \{0,1\}$ is an indicator equal to 1 if the agent buys the product, $v \in \mathbb{R}_+$ is her valuation of a polluting product, c is the amount of emissions associated with the product drawn from a finite set $C \subset \mathbb{R}_+$, $m \in \{0,1\}$ is an indicator variable equal to 1 if the environment allows for the formation of motivated beliefs, $w : \mathbb{R}_+ \to \mathbb{R}_+$ is a function mapping emission amounts to disutility from polluting, and $F, \hat{F} \in \Delta(C)$ are two discrete distributions on the space of possible emissions. Distribution F represents the objective probabilities while $\hat{F}$ represents the (possibly motivated) subjective beliefs of the agent at the moment of purchase. Let $\hat{f}_c$ denote the subjective probability attached to the amount of emission $c \in C$ and f_c the corresponding objective beliefs. Finally, E_F is an expectation operator with respect to distribution F.

When $m = 1$, the agent can choose a belief distribution $\hat{F} \neq F$ and experience an expected disutility from polluting equal to $E_{\hat{F}}[w(c)]$. To distort the beliefs, she needs to pay a cost $\sum_{c\in C}(\hat{f}_c - f_c)^2$. If $m = 0$ instead, the agent cannot distort her beliefs ($\hat{F} = F$). The psychological cost of polluting, in this case, is $E_F[w(c)]$, and there is no cost for belief distortions.

We interpret v as the material utility of consumption net of the (unmodeled) price. We interpret the disutility from polluting $E_{\hat{F}}[w(c)]$ as the psychological cost of imposing a negative externality on others.

There is a continuum of agents of mass 1 who share the same w and the same F but with individual valuation v_i. Let G be the cumulative distribution function over possible valuation $v \in \mathbb{R}_+$, where we assume positive density $g(v) > 0$ for all $v \in \mathbb{R}_+$. G and F are independent distributions.

Each agent makes two choices— she first chooses $\hat{F}$ and then decides whether or not to buy the product. Note, again, that the choice of $\hat{F} \neq F$ is possible only in the environment $m = 1$.

We use this framework to study how uncertainty about the size of the emissions changes people's behavior. First, we will show that people manipulate their beliefs to believe that the emissions are lower than they actually are and that these motivated beliefs lead more people to buy the product. Afterward, we will prove that the average population belief is more biased if people derive higher utility from buying the product. Finally, we will show under which conditions more people buy the product if there is more uncertainty about the size of the emissions. The experiments we present in the rest of the paper will test these theoretical predictions.

Proposition 1. *When $m = 1$, the following statements are equivalent:*

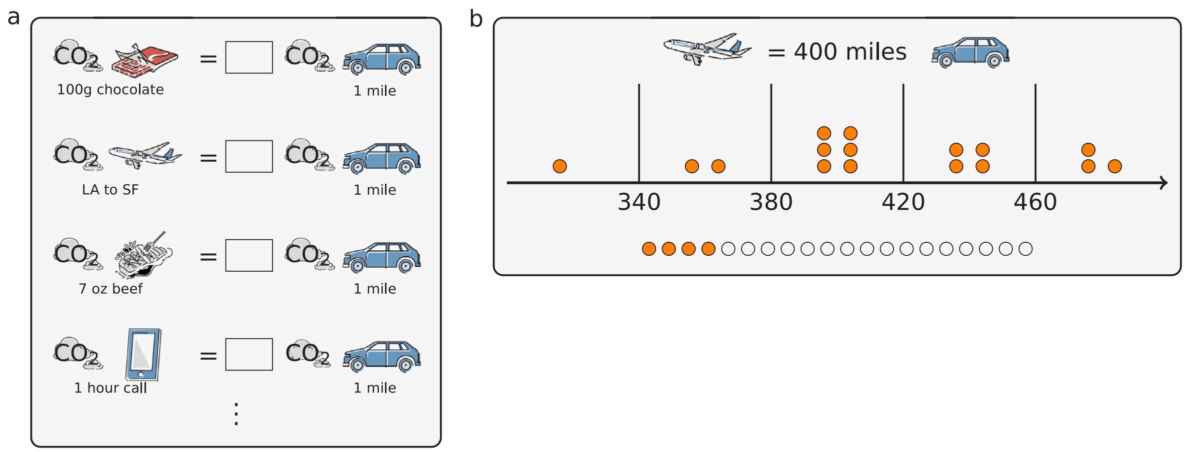

Fig. 1. Illustration of the belief and WTM elicitation interface. (a) Point-belief elicitation task. (b) Bins-and-balls belief elicitation task. *Notes*: Panel B shows an example in which a participant stated 400 in the previous point belief elicitation task and is now asked to allocate 20 balls into five bins, centered around this number. See Online Appendix C.1 for screenshots of the interface.
Source: Adapted from Imai et al. (2022).

1. *The agent buys the product, i.e.,* $b = 1$.
2. *The agent forms motivated beliefs, i.e.,* $E_{\hat{F}}[c] < E_F[c]$.

Intuitively, when an agent is able to form motivated beliefs and decides to purchase the product, it is beneficial for her to form motivated beliefs in order to alleviate the psychological cost associated with polluting. Conversely, if she chooses not to buy the product, there is no reason for her to incur the psychological cost of distorting her beliefs. As a result, her subjective belief distribution aligns with the objective one.

We will now present two corollaries of this result. For clarity, let ψ_1 and ψ_0 represent the fraction of agents who purchase the product when motivated beliefs are possible and not possible, respectively. The first corollary is:

Corollary 1.1. *When* $m = 1$, $\psi_1 > 0$ *if and only if* $E_G\left[E_{\hat{F}}[c]\right] < E_G\left[E_F[c]\right]$.

This corollary indicates that to observe motivated beliefs at the population level, there must be at least some agents purchasing the product in an environment where motivated beliefs are possible. This result is crucial because Experiment 1 will test for motivated beliefs, and it provides a condition to assess whether the model predicts the formation of motivated beliefs in our experimental setting.

The second corollary is:

Corollary 1.2. $\psi_1 \geq \psi_0$ *with strict inequality if and only if* $\psi_0 < 1$.

This corollary suggests that motivated beliefs are consequential as they lead to an increase in the purchase of polluting products— a prediction that our Experiment 1 will test. Intuitively, this occurs because motivated beliefs reduce the psychological cost of polluting, making it optimal even for participants with a relatively lower valuation of the product to purchase it. The presence of motivated beliefs does not affect the proportion of buyers only if everyone would have bought the product even in the absence of such beliefs.[1]

Next, we study how motivated beliefs within the population respond when individuals derive greater material utility from buying the product. This could occur due to a decrease in the product's price or an increase in its appeal. Our goal is to replicate, within our simple framework, the results of Hestermann et al. (2020), which suggest that motivated beliefs are larger when the surplus from consuming the product increases.

Proposition 2 (*Motivated Beliefs and Incentives*)**.** *Suppose two distributions of* v, G_1 *and* G_2, *are such that* $G_2(v) < G_1(v)$ *for all* $v \in \mathbb{R}_+$. *Motivated belief distortions in the population are higher under* G_2, *that is:*

$$E_{G_2}\left[E_{\hat{F}}[c]\right] - E_{G_2}\left[E_F[c]\right] < E_{G_1}\left[E_{\hat{F}}[c]\right] - E_{G_1}\left[E_F[c]\right].$$

Intuitively, when the value of the product is higher, a larger mass of agents buy the product, making it optimal for a larger number of them to develop motivated beliefs.

Next, we consider environments where motivated beliefs cannot emerge ($m = 0$) and examine a different behavioral channel, namely risk aversion over carbon impact. In this context, we show that uncertainty increases the mass of agents who buy the product if and only if the disutility from emissions is concave with respect to emission size.

Proposition 3 (*Risk Aversion*)**.** *Consider the case in which* $m = 0$. *Consider furthermore two distributions of emissions* F_1 *and* F_2 *such that* $E_{F_1}[c] = E_{F_2}[c]$ *and* $\mathrm{Var}_{F_1}(c) < \mathrm{Var}_{F_2}(c)$. *The fraction of agents buying the product is higher under* F_2 *than under* F_1 *if and only if the agents' disutility from emissions is represented by a concave* w.

Intuitively, an increase in the variance of emissions raises the probability of both low and high emissions. A consumer with a concave disutility for emissions will heavily discount the very high emissions, decreasing the subjective expected cost of polluting, even if the expected emissions remain constant. As a consequence, her utility from buying the product increases, making her more likely to buy it.

In Experiment 2, we will estimate the shape of the disutility from emissions and will empirically test Proposition 3.

3. Motivating evidence: Subjective uncertainty about CO_2 emissions

In this section, we provide motivating evidence behind our research questions. To do so, we reanalyze the Climate Survey in Imai et al. (2022). Whereas Imai et al. (2022) describes participants' point estimates, here we focus on the subjective uncertainty, which was also measured in the survey. The results show that subjective uncertainty about emissions is large and our research questions are relevant to consumer choice. Below, we briefly describe the methodology of the relevant parts of Imai et al. (2022), and Online Appendix D.1 presents the relevant instructions.[2]

[1] This case would realize if $w(c) = 0$ for all c and indifferent agents buy the product.

[2] We refer the reader to Imai et al. (2022) for a complete description and explanation of the survey design.

Belief elicitation. In the Climate Survey of Imai et al. (2022), we elicited participants' beliefs about CO_2 emissions generated 13 common consumer products and activities, including food items, household appliance use, and transportation (see Table C.1 in the Online Appendix). We first elicited a point estimate for the modal value of the emissions. To understand the participants' confidence in their answers, we then elicited the subjective probability distribution of the size of CO_2 emissions. For each product, we presented five "bins" centered around the point estimate the participant reported and asked the participant to allocate 20 balls into these five bins to represent the probability that the value of the emissions is in each of them. Fig. 1 provides an illustration of the elicitation screens (the screenshots of the interfaces are presented in Online Appendix C.1).

We incentivized the elicitation by randomly selecting one of the bins and scoring the answer against the most recent scientific estimates with a randomized quadratic scoring rule to determine. This mechanism encourages participants to truthfully reveal their belief that the scientific estimate falls in a particular bin (Schlag and van der Weele, 2013).

Implementation. We recruited 1,430 participants on Prolific (https://www.prolific.com) between the 3rd and 6th December 2020, and 1,128 completed all the belief elicitation questions described in this paper. After the belief elicitation, the survey continues with additional questions that we describe and analyze in Imai et al. (2022). We restricted participation to US residents, and we aimed to collect a sample representative for age, gender, and ethnicity.[3] Table B.1 in the Online Appendix shows the demographic characteristics of the sample.

Results. 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 point estimate. She is *Less than 50% certain* if she puts less than 10 balls in the bin which contains her point estimate: a ball allocation that indicates a probability of more than 50% that the true value is at least 5% away from the point estimate. She is *Extremely uncertain* if she puts four or fewer balls in each of the three central bins: a distribution which indicates a probability of at least 40% that the real value of the emissions is at least 15% away from her best guess.

We chose these non-parametric measures as we aim to verify whether participants are uncertain about the accuracy of their guesses. These measures allow immediate and assumption-free ways to check if people assign some probability mass to emission amounts that are far from their point estimates. The exact quantification of this uncertainty is beyond the scope of this paper.

Fig. 2 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*. These results show that people are aware of having very limited knowledge about the carbon footprint of common products and activities.

4. Experimental designs and empirical analysis

Having established that people can be very uncertain about the carbon impact of common products and activities, we now present the two experiments investigating the effects of this uncertainty on behavior. This section introduces the experiments sequentially. We preregistered Experiment 1 on the 8th May 2019 (https://aspredicted.org/2yz6-95w8.pdf) and Experiment 2 on the 11th October 2022 (https://aspredicted.org/6wjk-rdv7.pdf). Online Appendices D.2 and D.3 present the instructions. Online Appendix C.6 describes the steps we took to maximize the data quality.

[3] We compared the demographic characteristics of study participants and information from US Census Bureau (2022), and confirmed that our sample is representative for gender and ethnicity, but not for age (Table B.2 in the Online Appendix).

4.1. Experiment 1: Do participants form self-serving beliefs?

This experiment tests whether people exploit uncertainty to form self-serving beliefs that emissions are lower than they actually are. Fig. 3 displays the experimental timeline separately for each of the two treatments.

Consumption decision. In this experiment, we offered participants the opportunity to buy a single unit of a virtual product.[4] If the participants decided to purchase the product, they increased their payoff by the product value of £2, minus the price at which we, the experimenters, offered it. Importantly, purchasing the product entailed the emission of CO_2 into the atmosphere, equivalent to burning 60 liters of gasoline (see below for how we implemented these emissions). 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. At the end of the experiment, we collected demographic information using a survey.

As we describe next, we orthogonally implemented three price treatments and two information treatments, resulting in six treatments. Each participant was assigned to a single treatment.

The implementation of the CO_2 emissions. In the experiment, the participants were asked to make decisions that could result in the emission of CO_2. To ensure that these emissions are consequential, we prepared a monetary transfer to Carbonfund.org, an organization that offsets CO_2 emissions. Every time a participant made a decision resulting in CO_2 emissions, we decreased the amount of our transfer by £1.07, the amount needed to offset the CO_2 generated by burning 60 liters of gasoline, which are the emissions generated by buying the experimental product.[5] We explicitly communicated this procedure to the participants. To enhance the external validity and maximize the salience of the emissions, we always framed these decisions as choices between private benefits and emitting CO_2. We took several measures to assure participants of the tangible nature of these CO_2 emissions. 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 Online Appendix C.7) and actually did so. These measures were successful. 80% of the participants expressed trust in us following through with our promises of buying offsets as described in the instructions. Online Appendix B.4.2 shows that our results are robust to excluding the participants who do not trust us.

Uncertainty treatments. Our primary focus is on the formation and impact of beliefs about the CO_2 emissions associated with the product and 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 knew that buying the product generated CO_2 emissions equivalent to burning between 0 and 120 liters of gasoline. To learn more about the exact size of these emissions, the participants had to engage in an attentional task designed to mimic cognitively costly information-gathering processes, such as an online search. The task involved examining a matrix of numbers between 0 and 120, with the most frequently appearing number (60) representing the emission size,

[4] This product is virtual because 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 payoff from the experiment increases if they "buy" it.

[5] Using a report from the US Environmental Protection Agency (United States Environmental Protection Agency, 2005), we calculate that burning 60 liters of gasoline produces 140 kg of CO_2 emissions. At the time of the experiment, Carbonfund.org offset one metric ton of CO_2 per every \$10 (or £7.9) it receives in donations, so offsetting the products' emissions cost £1.07.

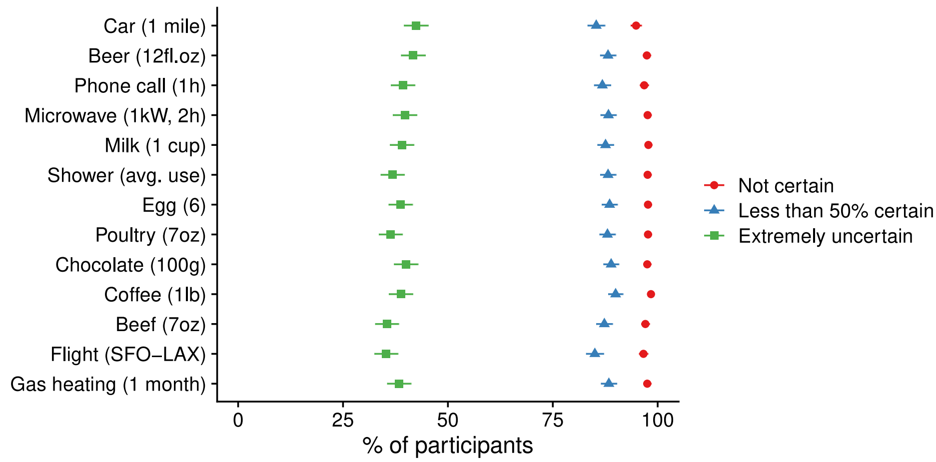

Fig. 2. The percentage of participants who satisfy three uncertainty criteria. *Notes*: See the text for the definitions of the criteria. *Source:* Data from the Climate Survey in Imai et al. (2022).

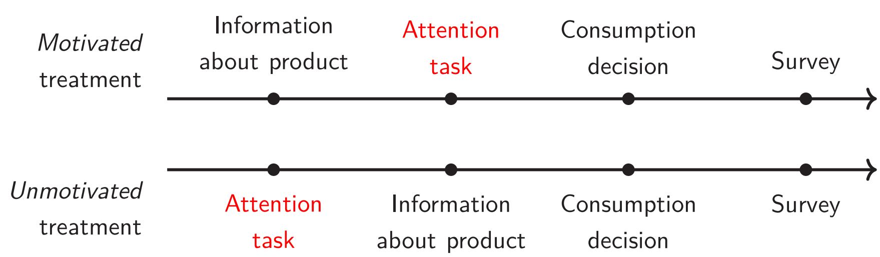

Fig. 3. Timeline of Experiment 1.

measured in terms of the CO_2 emissions generated by burning a liter of gasoline.[6] The 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 (Schlag et al., 2015).[7] We adapted this attention task from Ambuehl (2016), which shows that the information-gathering strategy in this task is influenced by incentives for subsequent decisions. We refrained from stating explicit probabilities, as there is evidence that more ambiguous settings are conducive to forming motivated beliefs (Chance and Norton, 2015; Engelmann et al., 2024).

The *Motivated* and *Unmotivated* 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 CO_2. They were instructed to find the most frequently occurring number without any knowledge about the meaning of the number within the experiment. In this way, we eliminated any self-serving motives that might lead participants to distort their attention or beliefs towards 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 of the surplus they could obtain from the product and that the correct answer to the task indicated the magnitude of the CO_2 emissions. Manipulating the timing of knowledge of the incentive scheme is a standard design feature in experiments studying motivated cognition (Babcock et al., 1995; Schwardmann and van der Weele, 2019). Consistently with Corollary 1.1, we hypothesize that the *Motivated* treatment would lead to motivated beliefs, i.e., to a lower estimate of the impact of the emissions. We also expect more people to buy the product according to Corollary 1.2.

Price treatments. We also investigate the relationship between purchasing decisions and prices. 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.

These treatments enable 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. Using a self-signaling model, Hestermann et al. (2020) predicts that when the surplus from consuming a polluting product is higher (low prices), people will distort their beliefs to a larger extent, leading to a further increase in consumption. Our model makes the same prediction (Proposition 2).

Sample and data collection. We recruited 714 participants using Prolific.co, an online platform, between the 9th and 11th of May 2019. 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

[6] The task can be found in Online Appendix C.2. The matrix contained 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.

[7] Note 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 in Pace and van der Weele (2020). We do not analyze it here, due to concerns about internal validity. In particular, the lack of full information about the belief distributions means we cannot be sure that the expected value of emissions is constant in the different treatments. Experiment 2 in this manuscript, therefore, provides higher quality evidence on the effect of precise information on behavior.

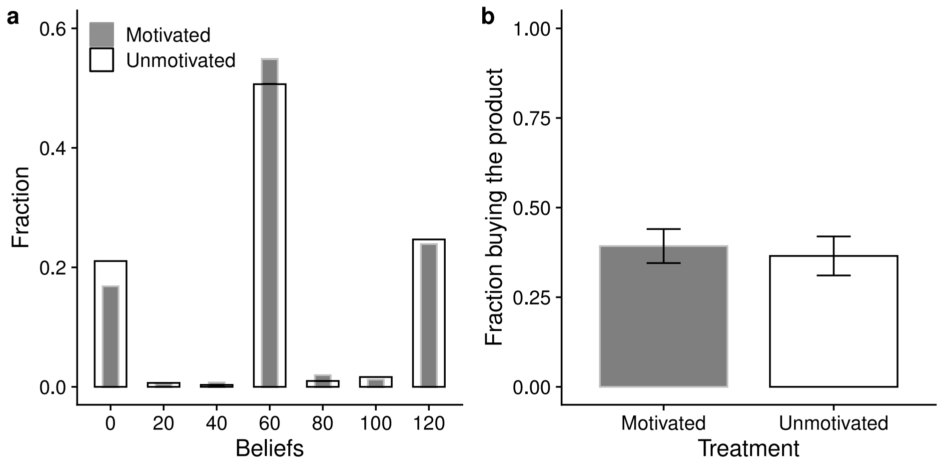

Fig. 4. (a) Distributions of beliefs about the correct answer in the attention task in the *Motivated* and *Unmotivated* treatment. (b) Fraction of participants buying the product in the *Motivated* and *Unmotivated* treatment. *Notes*: In panel (b), bars indicate 95% CI. Data from Experiment 1.

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 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 percent of the participants identified as females, 42% are students, and the average age is 29. We accepted only EU nationals as participants. 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 min to complete the tasks. We obtained some of the 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 CO_2 emissions.

Results. We find no evidence of the formation of self-serving beliefs. The left panel of Fig. 4 shows the distribution of beliefs in both treatments, where the spikes are driven by the nature of the perceptual task. 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 ($t(656) = 0.60$, $p = 0.548$, two-sided). Furthermore, contrary to Proposition 2, we do not find evidence that people are more likely to develop motivated beliefs when the product price is low. The coefficient for the interaction of price and *Motivated* treatment in Column (2) is not statistically significant ($t(654) = 1.18$, $p = 0.237$, two-sided). Finally, we find no evidence that the *Motivated* treatment causes people to spend less time looking at the information, as we discuss in Section 5.1.

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 s on the task screen. Thus, there was room for motivated subjects to perceive emissions to be lower than they actually were: harder tasks foster motivated beliefs as we discussed above.

We also test for differences in purchasing behavior between the two treatments. On average, 38% of the participants buy the product. The right panel of Fig. 4 shows that behavior is similar in the two treatments. Both a Fisher's exact test and a t-test based on Column (3) 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: $t(656) = 1.07$, $p = 0.285$; both tests are two-sided). By contrast, a higher product price has a strong independent impact on purchasing behavior: Column (4) of Table 1 shows that an increase in the price of one pound leads to a 13 percentage point decline in purchases.

The fact that some people buy the product in the *Motivated* treatment assures that our model predicts the emergence of motivated beliefs in the experiment.

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

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

Notes: The models include observations from the *Motivated* and the *Unmotivated* treatments, with the *Unmotivated* treatment as the baseline. Nineteen participants for whom demographic characteristics were not recorded are excluded from the regressions. The dependent variables are beliefs in columns (1) and (2), and purchasing decisions (an indicator variable equal to 1 if the participant purchased the product) in columns (3) and (4). Control variables include age, gender, student status, education (6 categories), frequency of car usage (5 categories), and nationality (27 categories). Robust standard errors are reported in parentheses. *: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$.

4.2. Experiment 2: Does risk aversion towards carbon impact increase polluting activity?

We now turn to a second potential channel through which uncertainty may matter: risk aversion. As we showed in Proposition 3, uncertainty about emissions might increase the propensity of buying polluting products if people's disutility from emitting CO_2 is concave in emission size.

We investigate the risk aversion channel with a new experiment since in Experiment 1 the participants face the same amount of uncertainty in every treatment. In the experiment, we first elicit the participants' valuation of CO_2 emissions. After that we ask participants whether they want to get a valuable but polluting product. The product

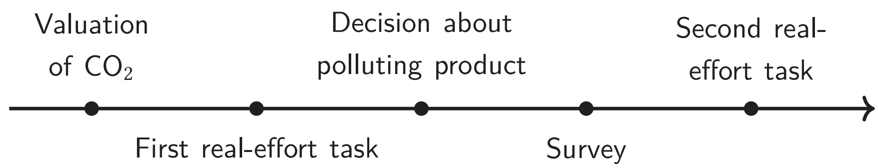

Fig. 5. Timeline of Experiment 2.

is valuable because it simplifies a real-effort task. Fig. 5 displays the experimental timeline.

Valuation of CO_2 emissions. We measured the participants' disutility from generating CO_2 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 CO_2 emissions. In contrast, choosing Option B allowed them to earn money but resulted in CO_2 emissions. For a given CO_2 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 Online Appendix D.3), which enforced a single switching point. This switching point gives the participants' valuation for a given amount of CO_2 emissions, which we call Willingness To Mitigate (WTM) in the text and $w(c)$ in the model. To gauge the participants' certainty regarding their valuations, we employed the "cognitive uncertainty" elicitation method developed by Enke and Graeber (2023). 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. To help participants understand the magnitude of the CO_2 emissions involved in the experiment, we wrote in the instructions that, on average, UK residents emit 14 kg of CO_2 per day (the sample of this experiment comes from the UK).

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 sped up the completion of a laborious real-effort task. This computer code is designed to mimic a convenience product like a tumble dryer which saves a few minutes of tedious work in exchange for emitting a few kilograms of CO_2.[8]

The task involved typing 15 strings, each consisting of 15 characters, in reverse order (Figure C.5 in the Online Appendix). The participants were required to transcribe these strings flawlessly to complete the task: any mistakes led to an error message specifying the strings needing correction before they could proceed with the experiment. 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 the button within the specified time window in more than four instances.

The participants were required to complete the real-effort task once without the help of the computer code— they spent on average 8 min and 40 s on this first iteration of the task. Then, before engaging with a second iteration, they had the chance to obtain the computer code. The code automatically generates all the correct answers for the task. If the participants get the code, they can complete the second iteration by simply clicking one button to submit the answers generated by the code.

Obtaining the code came at the cost of emitting CO_2— we implement these emissions using the same procedures as in Experiment 1 (see below for further details). There are two treatments that varied the information participants had about the size of the CO_2 emissions. In the *Information* treatment, participants were informed that purchasing the product would result in emissions equivalent to 4 kg of CO_2. In contrast, in the *Uncertainty* treatment, participants were informed of a probability distribution: a 40% probability of emissions being 0 kg, a 20% probability of emissions being 4 kg, and a 40% probability of emissions being 8 kg. Online Appendices C.4 and C.5 present the information and decision screens.

Differently from Experiment 1, we informed the participants about the objective beliefs distribution of the emissions as we wanted to keep constant the (subjective) expected value of emissions in both treatments. Furthermore, the explicit mention of the beliefs distribution limits the scope for motivated beliefs (Engelmann et al., 2024). This feature allows a cleaner analysis of the relationship between the participants' behavior and their WTM. We also told participants that they: "will get more information about the exact emission size later in the experiment" and indeed revealed the emission size at the end of the experiment. This feature ensures that the participants will get to know about the size of the emissions they are generating and hence can experience the related psychological disutility from polluting.

Survey. After the participants made the decisions about the computer code but before they completed the second real-effort task, they filled in a questionnaire. A battery of questions elicits people's moral evaluations of emissions-money trade-offs. Further questions in the survey ask about attitudes towards climate change and demographic characteristics.

The implementation of the CO_2 emissions. To implement the CO_2 emissions, we follow the same procedures as in Experiment 1. In this case, 84% of the participants said they believed that we would buy the CO_2 offsets as described in the instructions. Appendix B.4.2 shows that our results are robust to excluding the participants who do not trust us.

Sample and data collection. We recruited 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 percent of the participants identified as females, and the average age is 39 ($\min = 18$, $\max = 79$, $SD = 12.51$). 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 min to complete the tasks. Following the participants' decisions, we donated $160 to Carbonfund.org to offset CO_2 emissions.

Results. To understand whether WTM is concave, Fig. 6a shows the WTM curve aggregated over all subjects. The figure shows that subjects, on average, display a diminishing WTM to reduce carbon emissions. They are willing to sacrifice about £2.7 to avoid 4 kilograms of CO_2

[8] In this experiment we used a convenience product instead of the virtual product of Experiment 1 to make the purchase more intuitive to participants. We also changed the size of CO_2 emissions to make them more comparable to everyday consumption decisions.

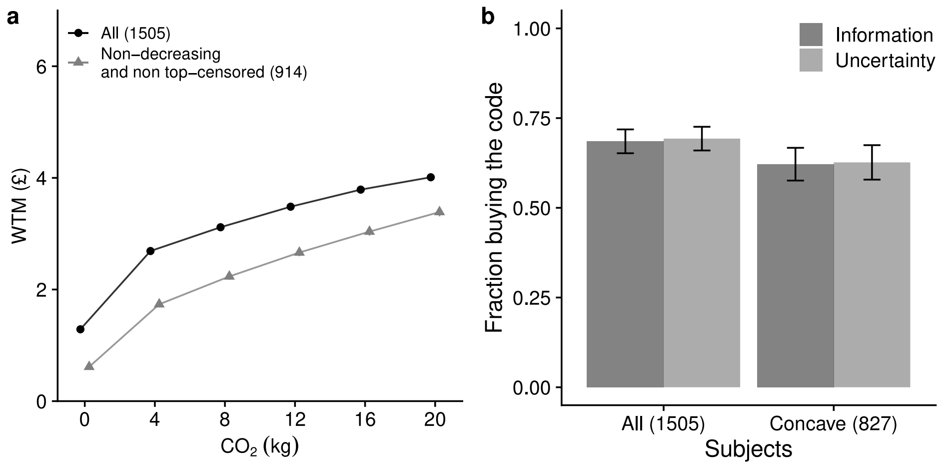

Fig. 6. (a) Aggregate WTM. (b) Purchasing decision and uncertainty.
Notes: In panel (b), the light-gray bars correspond to the group of subjects whose WTMs exhibit concavity. Bars indicate 95% CI. Data from Experiment 2.

emissions, whereas to avoid 20 kilograms of emissions, they are willing to forgo only £4. That is, the WTM increases by less than 50% when the amount of CO_2 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 artifact). Appendix B.3.2 contains further details about the variation in concavity across subjects. Section 5.2, instead, discusses how neither cognitive uncertainty nor concave moral evaluations seem to explain the concavity of the WTM.

We now turn to our main outcome, the purchase of the polluting computer code, which simplifies the real-effort task. On average, 69% of the participants buy the code.

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 Fig. 6b show that purchasing decisions are similar in the two treatments (68.5% vs. 69.3%, $z = 0.3168$, $p = 0.7514$, two-sided). It is possible, however, that there are offsetting effects for subjects with concave and convex WTM. The right bars in Fig. 6b 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%, $z = 0.1494$, $p = 0.8813$, two-sided).

To provide further statistical backup, we run several models, regressing the purchasing decision on the treatment, several concavity scores, and the interaction of these two (Table B.6 in the Appendix). 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 (Figure B.2 in the Appendix).

5. Discussion

In this section, we discuss whether our null results may be an artifact of the experimental design. Then, we document the effect of the *Motivated* treatment on information processing time in Experiment 1. Finally, we discuss possible cognitive mechanisms underlying the concavity results observed in Experiment 2.

5.1. Interpreting the null results

Are our null results informative about the public's reaction to climate information? One might worry that they are an artifact of the experimental setups. Here, we discuss a few potential issues.

Lack of individual impact. A possible driver for the null results in our experiments is that people perceive their actions to have a minimal impact on climate change. Indeed Semken (2024) shows that people underestimate by two orders of magnitude the environmental and health effects of CO_2 emissions and wrongly believe their individual actions to have limited effects if others do not behave sustainably.

Such an argument leads to the prediction that participants should disregard the CO_2 in the experiment and maximize their private utility. Yet, this prediction does not find support in the data. 30% of the participants declined to buy the computer code in Experiment 2, and 65% do not buy the product in Experiment 1. In Experiment 2, all participants but 19 (1.2% of the sample) have a positive WTM when emissions are equal to 4 kg (the lowest possible emission level in the experiment).

The participants' behavior indicates that they are willing to make sacrifices to reduce CO_2 emissions. Hence, it is incompatible with the hypothesis that the perceived limited impact of one's action is a driver of the null results.

Participants are insensitive to emission size. Another possible explanation for the null result is that participants dislike emitting CO_2, but they are unable to appreciate the difference between the various emission sizes. As a result, their behavior is not affected by the precise amount of CO_2 emitted. This insensitivity could account for the null result. If people do not care about the exact size of emissions, they have no reason to form motivated beliefs. Furthermore, if their valuation of emissions remains constant regardless of emission size they will not respond to the presence of uncertainty.

There are several pieces of evidence that point against this explanation of the null result. First, we can check if there is a relationship between the participants' beliefs about emissions and their probability of buying the product in Experiment 1. The insensitivity argument presented above predicts that the participants are equally likely to buy the product independently of their beliefs. To test this hypothesis, we focus on the *Unmotivated* treatment in which the participants solved the attention task *before* knowing what the answer to the task mean within the experiment. Table B.4 in the Appendix shows that the probability of buying the product decreases with the beliefs, a finding that contradicts the insensitivity hypothesis. While this evidence is correlational (the beliefs are endogenous), it is hard to construe a reason why the participants' answers to the task might be correlated with purchase behavior

Table 2
Time spent on the attention task.

	(1)	(2)	(3)
	Accuracy	Response time	
Motivated treatment	0.042	1.593	0.082
	(0.039)	(2.994)	(1.508)
Controls	Yes	Yes	Yes
Observations	694	694	632
R^2	0.052	0.067	0.079

Notes: The models are linear regressions. The dependent variables are: (1) a dummy variable indicating whether the participant correctly identified the most frequently occurring number (60) in the attention task, and (2–3) the time spent on the attention task (in seconds). The models include observations from both the *Motivated* and the *Unmotivated* treatments, with the *Unmotivated* treatment serving as the baseline. The first two columns include all participants for whom demographic data were recorded, except one for whom the program did not record the time spent on the task. The third column includes only participants who spent less than 70 s completing the attention task. Control variables include sex, age, student status, education (6 categories), frequency of car usage (5 categories), and nationality (27 categories). Robust standard errors are reported in parentheses. *: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$.

via omitted variables: in the *Unmotivated* treatment, the participants do not know anything about the purchase opportunity while completing the attention task.

Second, we can turn to Experiment 2 and check whether the participants' valuation of CO_2 emissions increases with emission size. Appendix B.3.2 shows that only 317 participants (21% of the sample) have a constant willingness to pay to avoid emissions, a finding inconsistent with the hypothesis that people are insensitive to emission size. The right bars in Fig. 6 show that these participants are not driving the results. The null result persists if we focus only on subjects with a strictly concave WTM, which are the ones that should increase consumption in the *Uncertainty* treatment according to our model

Effort in the attention task. In this analysis, we investigate a potential confound in Experiment 1. The *Motivated* treatment may give participants more reason to pay close attention to the task in the *Motivated* treatment as they know that the task is informative about the size of the emissions. If the participants indeed put more effort into the task, this could explain why we do not find motivated beliefs in the experiment.

To evaluate this, we look at participants' efforts in two ways. First, we check if the participants are more likely to give the correct answer in the task in *Motivated* rather than in the *Unmotivated* treatment. We find that 55% of participants gave the right answer in the *Motivated* treatment, which is not significantly larger than the 51% that gave the right answer in the *Unmotivated* treatment (Fisher's exact test, $p = 0.29$, two-sided). Column (1) of Table 2 confirms this null result, regressing the participant's answers to the task on a dummy for the *Motivated* treatment and on demographic controls.

Second, we look at the time subjects spend completing the attention task. Not only does this variable offer a proxy for effort, but it is also of interest because recent findings in economics and neuroscience suggest that dwell time on a piece of information causally increases the weight given to that information in subsequent decisions (Pärnamets et al., 2015; Engelmann et al., forthcoming; Li and Camerer, 2022; Amasino et al., 2024). In the context of our study, we observed no substantial differences between the *Motivated* and *Unmotivated* treatments. In Column (2) of Table 2, we regress the time the participants spent on the task (in seconds) on a dummy for the *Motivated* treatment and on demographic controls. We observe that participants in the *Motivated* treatment spent 1.6 s more on the task. However, this difference is not statistically significant ($t(655) = 0.53$, 95% CI $[-4.29, 7.47]$, $p = 0.595$, two-sided). Column (3) confirms this finding, focusing solely on the 91% of participants who completed the task in less than 70 s. This subset represents individuals for whom we can be most confident that they did not take any breaks between receiving the information and providing their responses (the information was displayed for up to 60 s).

Overall, our data does not support the hypothesis that differences in effort across treatments explain the null result concerning motivated beliefs.

Preferences for consistency. Another concern is that the null result in Experiment 2 is due to preferences for consistency. A participant who indicates a positive WTM in the first part of the experiment might feel like he has to abstain from getting the computer code in the second to act consistently. However, preferences for consistency are unlikely to explain the null result. As our theoretical model shows a participant with a concave WTM and acting consistently with it should be more likely to buy the computer code in the uncertainty treatment. Yet, we do not find that this is the case. In addition, we took several steps in the experiment to make sure that the two decision environments were dissimilar. The Willingness to Mitigate (WTM) elicitation involved a Multiple Price List, while the purchase decision involved a Yes/No question. Moreover, the value of the computer code depends on the subjective evaluation of the cost of effort, meaning that even environmentally conscious participants might find it optimal (and acceptable) to buy the computer code if they particularly dislike the task. These two features ensure that there is no clear reason why a participant with a positive WTM should avoid buying the product for consistency with previous choices.

Ceiling effects. Finally, one might worry that our null results are due to ceiling effects. This is unlikely to be the case. In Experiment 1, 63.5% of participants do not buy the product in the *Unmotivated* treatment, indicating that purchase levels are far from the ceiling.

In Experiment 2 instead, 30% of the participants do not buy the computer code in the *Info* treatment, leaving ample room for uncertainty to increase demand. What is more, the null result replicates in a subsample where buying behavior in the *Info* treatment is considerably lower. In particular, in the final questionnaire, we asked participants to indicate how annoying the real-effort task was on a scale from 1 (not annoying at all) to 5 (absolutely annoying). Of the 58% of participants that answered three or less to this question, 38% do not buy the product in the *Info* treatment. Yet the null effect of uncertainty on behavior persists in this subsample (Fisher's exact test, $p = 0.62$, two-sided).

Overall, we can exclude that ceiling effects artificially generate the null effects we document in the paper.

5.2. Psychological mechanisms behind concavity of WTM

We empirically explore two potential psychological mechanisms that may give rise to a concave WTM curve. The first mechanism relates to individuals' inability to appreciate increasingly large (and unfamiliar) amounts of emissions. The second mechanism considers the possibility that the concavity in WTM arises from concave moral judgments about the acceptability of causing different levels of emissions. Our data does not support either of these two mechanisms.

Increasing cognitive uncertainty. People may perceive the questions involving larger emission quantities as more challenging due to the inherent complexity of visualizing the precise scale of higher levels of emissions. This heightened level of complexity can lead participants to experience greater cognitive uncertainty when deciding their WTM, making them less sensitive to variations in increases in emission sizes. This relation between cognitive uncertainty and valuation can generate a concave WTM curve within the framework of an "anchoring and adjustment" model, in which the weight attributed to the anchor increases with cognitive uncertainty (Enke and Graeber, 2023). The anchor, in this case, is the default behavior of not compensating for CO_2 emissions, as is typical in everyday life.

The anchoring and adjustment model predicts the concavity of WTM under two conditions: (a) individuals generally do not engage in emissions offsetting, making an anchor value of £0 a plausible assumption, and (b) cognitive uncertainty increases with emission size.

Table 3
Concavity of WTM and cognitive uncertainty.

	(1)	(2)
Emissions	0.0762***	0.0782***
	(0.0042)	(0.0047)
Cognitive uncertainty	0.0013	0.0011
	(0.0092)	(0.0099)
Cognitive uncertainty ×Emissions	0.0005	0.0004
	(0.0007)	(0.0008)
Baseline WTM	0.7876***	0.8093***
	(0.0251)	(0.0275)
Constant	−0.3045	−0.2408
	(0.3505)	(0.3740)
Controls	Yes	Yes
Observations	4,924	4,148
Clusters	1,231	1,037
R^2	0.5261	0.5401

Notes: The dependent variable is WTM. The first column includes only participants with uncensored WTM values for all six emission amounts. The second column further excludes participants who stated they offset all their emissions or that they "often" offset their emissions. Control variables include age, gender (male, female, other), political affiliation (5 categories), education (6 categories), income (7 categories), and time needed to complete the first real-effort task. Standard errors clustered at the individual level are reported in parentheses. *: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$.

We find empirical support for both of these underlying assumptions. Specifically, over 82.7% of our participants reported to "Never" or "Rarely" compensate for their emissions. Furthermore, in a regression that controls for demographic characteristics, we find that cognitive uncertainty increases with emission size ($t(1231) = 8.132$, 95% CI $[0.067, 0.109]$, $p < 0.001$, two-sided).

However, our analyses do not reveal significant evidence of a relationship between cognitive uncertainty and the concavity of the WTM curve. Table 3 presents the results of a regression in which the WTM is regressed on (a) cognitive uncertainty, (b) emission levels, and (c) the interaction between cognitive uncertainty and the emission levels (d) individual controls, including the WTM for emission levels of 4 kg. The interaction between the emissions and cognitive uncertainty is included to explore whether subjects who are more cognitively uncertain are less sensitive to an increase in emission amounts. This is a prediction of the anchoring and adjustment model, and it implies a negative coefficient for the interaction term. The regression includes the WTM for emission levels of 4 kg to control for the fact that participants with a higher WTM are mechanically more sensitive to an increase in the size of the emissions. Since we include this control, we only use the WTM for emission levels above 4 kg as the dependent variable. The results reported below would be unchanged if we dropped the control for the WTM for emission levels of 4 kg and included all the observations.

Column (1) of Table 3 shows that the coefficient for the interaction term is small in magnitude and insignificant ($t(1230) = 0.755$, 95% CI $[-0.001, 0.002]$, $p = 0.451$, two-sided). Column (2) restricts the sample to subjects who "Never" or "Rarely" compensate for their emissions and confirms the null result ($t(1036) = 0.582$, 95% CI $[-0.001, 0.002]$, $p = 0.561$, two-sided).

Another approach to assess the relationship between cognitive uncertainty and a concave WTM leverages individual variations in the increase of cognitive uncertainty as emission rises. The subjects for which the cognitive uncertainty increases the most should be more likely to have a concave WTM. To test if this relationship is supported in the data, we define $CU_j(e)$ as the cognitive uncertainty of participant j at emission level e. The increase in cognitive uncertainty can then be quantified as:

$$\Delta_{CU} = CU_j(\bar{e}) - CU_j(0),$$

where $\bar{e}$ denotes the highest emission level for which the participant reported an uncensored WTM. We regress the concave-WTM dummy on Δ_{CU} and find that there is no statistically significant correlation between the two variables ($t(1082) = -0.431$, 95% CI $[-0.004, 0.003]$, $p = 0.666$, two-sided).[9]

Table 4
Concavity of WTM and morality.

	(1)	(2)	(3)	(4)
Concavity of moral judgment (ϕ)	0.022	0.020	0.016	0.045
	(0.020)	(0.020)	(0.021)	(0.026)
Constant	0.212	0.202	0.233	0.153
	(0.121)	(0.125)	(0.128)	(0.143)
Controls	Yes	Yes	Yes	Yes
Observations	1,504	1,430	1,373	1,100
R^2	0.040	0.040	0.042	0.040

Notes: The dependent variable is `concave`, a dummy variable equal to 1 if the WTM curve is classified as either "concave" or "concave†" based on the classification discussed in Appendix B.3.2. The samples become increasingly restrictive from left to right. Column (1) includes all participants. Column (2) excludes participants who failed the attention check embedded in the moral judgment elicitation. Column (3) excludes participants whose $m_j(e)$ are decreasing in e. Column (4) excludes participants whose WTM curve is either decreasing or non-monotonic. Control variables include age, gender (male, female, other), political affiliation (5 categories), education (6 categories), income (7 categories), and time needed to complete the first real effort task. Robust standard errors are reported in parentheses. *: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$.

Based on these two analyses, we conclude that there is insufficient support for the idea that cognitive uncertainty is a driver of concavity in the WTM curve.

Concave moral valuations. Another potential psychological channel that may explain a concave WTM relates to concave moral judgments. Individuals might perceive emitting 4 kg of CO_2 as considerably morally worse than emitting 0 kg, while the moral distinction between emitting 4 kg and emitting 20 kg might seem relatively minor. Such concave moral evaluations might, in turn, influence and shape the participants' WTM.

Let $\mu_j(e,k) \in \{1,2,\ldots,7\}$ denote the moral evaluation assigned by participant j to emitting $e$kg of CO_2 in exchange for £k, where the range spans from "morally very appropriate" (1) to "morally very inappropriate" (7). These evaluations are collected for each $e \in 4, 12, 20$ and $k \in 1, 5$. We aggregate these moral judgments by computing their average over the two values of k, yielding $m_j(e) = (\mu_j(e,1) + \mu_j(e,5))/2$. This composite measure is labeled as "Morality". Finally, we compute the variable ϕ_j as:

$$\phi_j = m_j(12) - \frac{m_j(4) + m_j(20)}{2}.$$

A positive value of ϕ_j indicates that the moral valuation of participant j is concave. The average ϕ is 0.107, which is positive and statistically significant ($t(1504) = 7.08$, 95% CI $[0.078, 0.137]$, $p < 0.001$, two-sided), suggesting that moral judgments are indeed concave.

To investigate whether the presence of concave moral valuations is linked to a concave WTM, we regress the "concavity" dummy, which is equal to 1 if a participant exhibits a concave WTM, on the variable ϕ_j. The results presented in Table 4 indicate that concavity in moral valuations has limited predictive power regarding the concavity of WTM. Column (1) includes all the observations and shows no significant relationship between the two variables ($t(1482) = 1.078$, 95% CI $[-0.018, 0.062]$, $p = 0.281$, two-sided). Columns (2) to (4) shows that the result is robust to excluding inattentive participants.

[9] We follow the classification of individual WTM curve discussed in Appendix B.3.2. Note that the concave-WTM is a dummy variable taking a value of 1 when the WTM curve is characterized as either "concave" or "concave†" in the classification. In this analysis, we excluded participants whose WTM curves were classified as decreasing or non-monotonic. Additionally, participants with only an uncensored WTM value at $e = 0$ were also excluded, as Δ_{CU} is undefined for this subgroup.

6. Conclusions and implications

We investigate the role of uncertainty about the impact on climate mitigation efforts. We document that people are uncertain about the CO_2 emissions associated with common consumer products, complementing previous evidence that people's best guesses about the size of CO_2 emissions are often too optimistic (Camilleri et al., 2019; Imai et al., 2022). 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 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 behavior.

Another set of implications stems from our finding that consumers' willingness to mitigate the first units of CO_2 is much higher than for subsequent units, a finding in line with concurrent evidence (Rodemeier, 2023). The first implication is that consumers' behavior should 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 to compute the benefits of reducing one ton of CO_2 in the atmosphere (Nemet and Johnson, 2010; Löschel et al., 2013; Diederich and Goeschl, 2014). Our results strongly suggest that future studies should measure individual WTM to avoid several emission amounts and use nonlinear models to estimate these benefits accurately. 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 towards emissions. The participants' WTM is significantly associated with their probability of buying the polluting product. Hence, the WTM is a meaningful predictor of consumption behavior and environmental preferences.

CRediT authorship contribution statement

Davide D. Pace: Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Taisuke Imai:** Writing – original draft, Visualization, Validation, Investigation, Formal analysis, Data curation. **Peter Schwardmann:** Writing – review & editing, Writing – original draft, Resources, Project administration, Methodology, Funding acquisition, Conceptualization. **Joël J. van der Weele:** Writing – review & editing, Writing – original draft, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

All data and code necessary to reproduce the analysis are available on the Open Science Framework repository (https://osf.io/pu8q2/).

Appendix A. Supplementary data

The mathematical proofs, additional empirical analyses, and the experimental and survey materials can be found in the online appendix at: https://doi.org/10.1016/j.ecolecon.2024.108401.

References

Amasino, Dianna R., Pace, Davide D., van der Weele, Joël J., 2024. Fair shares and selective attention. Am. Econ. J. Microecon. 16 (4), (forthcoming) https://www.aeaweb.org/articles?id=10.1257/mic.20220275.

Ambuehl, Sandro, 2016. An offer you can't refuse? Incentives change what we believe. Available at SSRN: https://doi.org/10.2139/ssrn.2830171.

Babcock, Linda, Loewenstein, George, Issacharoff, Samuel, Camerer, Colin, 1995. Biased judgments of fairness in bargaining. Amer. Econ. Rev. 85 (5), 1337–1343.

Barnett, Michael, 2023. Climate change and uncertainty: an asset pricing perspective. Manage. Sci. 69 (12), 7562–7584.

Barrett, Scott, Dannenberg, Astrid, 2012. Climate negotiations under scientific uncertainty. Proc. Natl. Acad. Sci. 109 (43), 17372–17376.

Barrett, Scott, Dannenberg, Astrid, 2014. Sensitivity of collective action to uncertainty about climate tipping points. Nature Clim. Change 4 (1), 36–39.

Bénabou, Roland, Tirole, Jean, 2016. Mindful economics: the production, consumption, and value of beliefs. J. Econ. Perspect. 30 (3), 141–164.

Berger, Loïc, Emmerling, Johannes, Tavoni, Massimo, 2017. Managing catastrophic climate risks under model uncertainty aversion. Manage. Sci. 63 (3), 749–765.

Broomell, Stephen B., Kane, Patrick Bodilly, 2017. Public perception and communication of scientific uncertainty. J. Exp. Psychol.: Gen. 146 (2), 286–304.

Budescu, David V., Por, Han-Hui, Broomell, Stephen B., Smithson, Michael, 2014. The interpretation of IPCC probabilistic statements around the world. Nature Clim. Change 4 (6), 508–512.

Camilleri, Adrian R., Larrick, Richard P., Hossain, Shajuti, Patino-Echeverri, Dalia, 2019. Consumers underestimate the emissions associated with food but are aided by labels. Nature Clim. Change 9 (1), 53–58.

Chance, Zoe, Norton, Michael I., 2015. The what and why of self-deception. Curr. Opin. Psychol. 6, 104–107.

Diederich, Johannes, Goeschl, Timo, 2014. Willingness to pay for voluntary climate action and its determinants: field-experimental evidence. Environ. Resour. Econ. 57 (3), 405–429.

Engelmann, Jan, Hirmas, Alejandro, van der Weele, Joël J., 2024. Individual and contextual effects of attention in risky choice. Exp. Econ. (forthcoming).

Engelmann, Jan B., Lebreton, Maël, Salem-Garcia, Nahuel A., Schwardmann, Peter, van der Weele, Joël J., 2024. Anticipatory anxiety and wishful thinking. Amer. Econ. Rev. 114 (4), 926–960.

Enke, Benjamin, Graeber, Thomas, 2023. Cognitive uncertainty. Q. J. Econ. 138 (4), 2021–2067.

Epperson, Raphael, Gerster, Andreas, 2021. Information avoidance and moral behavior: experimental evidence from food choices. Available at SSRN: https://doi.org/10.2139/ssrn.3938994.

Fischhoff, Baruch, Davis, Alex L., 2014. Communicating scientific uncertainty. Proc. Natl. Acad. Sci. 111 (supplement_4), 13664–13671.

He, Shutong, Blasch, Julia, Robinson, Peter John, van Beukering, Pieter, 2024. Social comparison feedback in decision-making context: environmental externality levels and psychological traits matter. Ecol. Econom. 216, 108047.

Hestermann, Nina, Le Yaouanq, Yves, Treich, Nicolas, 2020. An economic model of the meat paradox. Eur. Econ. Rev. 129, 103569.

Imai, Taisuke, Pace, Davide D., Schwardmann, Peter, van der Weele, Joël J., 2022. Correcting consumer misperceptions about CO_2 emissions. CESifo Working Paper No. 10138.

Kunda, Ziva, 1990. The case for motivated reasoning. Psychol. Bull. 108 (3), 480–498.

Li, Xiaomin, Camerer, Colin F., 2022. Predictable effects of visual salience in experimental decisions and games. Q. J. Econ. 137 (3), 1849–1900.

Lind, Jo Thori, Nyborg, Karine, Pauls, Anna, 2019. Save the planet or close your eyes? Testing strategic ignorance in a charity context. Ecol. Econom. 161, 9–19.

Lohmann, Paul, Gsottbauer, Elisabeth, Doherty, Anya, Kontoleon, Andreas, 2022. Do carbon footprint labels promote climatarian diets? Evidence from a large-scale field experiment. J. Environ. Econ. Manag. 114, 102693.

Löschel, Andreas, Sturm, Bodo, Vogt, Carsten, 2013. The demand for climate protection–empirical evidence from Germany. Econom. Lett. 118 (3), 415–418.

Nemet, Gregory F., Johnson, Evan, 2010. Willingness to pay for climate policy: a review of estimates. Available at SSRN: https://doi.org/10.2139/ssrn.1626931.

Pace, Davide D., van der Weele, Joël, 2020. Curbing carbon: an experiment on uncertainty and information about CO_2 emissions. Tinbergen Institute Discussion Paper No. 2020-059.

Pärnamets, Philip, Johansson, Petter, Hall, Lars, Balkenius, Christian, Spivey, Michael J., Richardson, Daniel C., 2015. Biasing moral decisions by exploiting the dynamics of eye gaze. Proc. Natl. Acad. Sci. 112 (13), 4170–4175.

Pindyck, Robert S., 2021. What we know and don't know about climate change, and implications for policy. Environ. Energy Policy Econ. 2 (1), 4–43.

Rodemeier, Matthias, 2023. Willingness to pay for carbon mitigation: field evidence from the market for carbon offsets. IZA Discussion Paper 15939.

Schlag, Karl H., Tremewan, James, van der Weele, Joël J., 2015. A penny for your thoughts: a survey of methods for eliciting beliefs. Exp. Econ. 18 (3), 457–490.

Schlag, Karl H., van der Weele, Joël, 2013. Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. Theor. Econ. Lett. 3, 38–42.

Schöller, Vanessa, Ulmer, Clara, 2023. Can monetized carbon information increase pro-environmental behavior? Experimental evidence. Ecol. Econom. 206, 107759.
Schwardmann, Peter, 2019. Motivated health risk denial and preventative health care investments. J. Health Econ. 65, 78–92.
Schwardmann, Peter, van der Weele, Joël, 2019. Deception and self-deception. Nat. Hum. Behav. 3 (10), 1055–1061.
Semken, Christoph, 2024. The marginal impact of emission reductions: estimates, beliefs and behavior. Mimeo, Universitat Pompeu Fabra.
Serra-Garcia, Marta, Szech, Nora, 2022. The (in)elasticity of moral ignorance. Manage. Sci. 68 (7), 4815–4834.
Stoetzer, Lasse S., Zimmermann, Florian, 2024. A representative survey experiment of motivated climate change denial. Nature Clim. Change 14 (2), 198–204.
Tilling, Anna Schulze, 2023. Changing consumption behavior with carbon labels: causal evidence on behavioral channels and effectiveness. Mimeo, University of Bonn.
United States Environmental Protection Agency, 2005. 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?Dockey=P1001YTF.txt.
US Census Bureau, 2022. National population by characteristics: 2020–2022. United States Census Bureau. https://www.census.gov/data/tables/time-series/demo/popest/2020s-national-detail.html.
Woerner, Andrej, Imai, Taisuke, Pace, Davide, Schmidt, Klaus, 2023. How to increase public support for carbon pricing. CRC TRR 190 Discussion Paper No. 489.