Understanding the Effectiveness of Carbon Labels

Anna Schulze-Tilling, University of Glasgow

First version: May 2023

This version: November 4, 2026 Click here for the most recent version

Behavioral climate interventions are often promoted as politically feasible alternatives to carbon taxes, yet we lack a way to value them in tax-equivalent terms. I show how the effectiveness of carbon labels can be benchmarked against an emissions-based price change resembling a local carbon tax. In German student canteens, a framed field experiment yields a local benchmark of approximately €155/t CO_2e. Natural field evidence suggests a similar benchmark. A third experiment investigates the roles of salience and misperception correction in driving label effectiveness, informing a structural comparison of consumer welfare under labels and carbon pricing.

JEL Classification: C93, D83, H23, Q54, Q58

Keywords: carbon labels, climate policy, behavioral interventions, salience, consumer welfare

Ethics approval and pre-registration. The German Association for Experimental Economic Research approved the experiments (#BnyKWhHQ; #RSzBqAcE). Preregistrations are in the AEA RCT Registry (#AEARCTR-0007858; #AEARCTR-0008435) and at AsPredicted (#95108).

Funding. This work was supported by the German Research Foundation (CRC TRR 224, Project B07; Excellence Strategy EXC 2126/1-390838866) and the European Research Council (Horizon Europe HEUROPE-2022-ADG grant 101097578, MemDec).

Acknowledgements. I am indebted to Lorenz Götte, Sebastian Kube, and Thomas Dohmen for continuous support and thank Hunt Allcott, Peter Andre, Sanchayan Banerjee, Simon Block, Stefano DellaVigna, Eugen Dimant, Katrin Gödker, Raphael Epperson, Ximeng Fang, Nicola Gennaioli, Holger Gerhardt, Luca Henkel, Michael Kosfeld, Chris Krekel, Matthew Levy, Paul Lohmann, Julia Mink, Davide Pace, Eugenio Proto, Simone Quercia, Alex Rees-Jones, Matthias Rodemeier, Paul Schäfer, Ganga Shreedhar, Egon Tripodi, Marie Claire Villeval, Jakob Wegmann, and audiences at the LSE, Bonn, Maastricht, the UIB, Geneva, PSE, Glasgow, briq, MCEE, AERE, SABE-IAREP, EAERE, EEA, VfS, and the ECBE Conference for comments. I especially thank the Studierendenwerk Bonn for non-financial support through its cooperation with the field experiment.

1 Introduction

Many economists agree that implementing carbon taxes is the most effective policy to combat climate change (European Association of Environmental and Resource Economists, 2019).[1] However, their implementation is not always feasible. For instance, even modest levies on food, a sector contributing an estimated 26% to 34% of global greenhouse gas emissions (Crippa et al., 2021; Poore and Nemecek, 2018), face considerable political resistance.[2] Behavioral climate interventions such as carbon labels have emerged as a promising policy tool where carbon pricing is difficult to implement. Carbon labels are included in the EU's Farm-to-Fork Strategy and are increasingly voluntarily adopted by producers. Yet we lack an evidence-based benchmark of the labels' effectiveness, comparing them with a carbon tax on the same products, faced by the same consumers. To gauge their true policy value, we need to understand: What per-tonne local tax would generate a comparable emissions reduction to the labels, and how do the two interventions compare in consumer welfare?

This paper develops an empirical approach for expressing the effect of behavioral climate interventions in tax-equivalent terms, and applies it to carbon labels in the student canteens in Bonn, Germany. Experiment 1 identifies, in a framed field experiment, the emissions-based price change that most closely reproduces the labels' emissions effect. This benchmark resembles a carbon tax implemented locally in the same consumption setting. Experiment 2 examines whether a similar estimate emerges from ordinary, recurring purchases. Experiment 3 provides experimental evidence on the role of misperception correction and increased salience in driving the labels' effect on consumer choices and elicits consumers' willingness to pay (WTP) for label presence. The tax

1. I use the term carbon tax to refer to price increases passed on to consumers that reflect the carbon emissions associated with a product, whether imposed through a direct tax (e.g., Germany's fuel tax) or through market mechanisms like the EU Emissions Trading System (EU ETS). Such carbon costs are typically embedded in the final price and are not explicitly communicated to consumers at the point of purchase as a distinct "carbon tax". In the empirical benchmark, the tax is local to the canteen setting: it is applied to the same meals and fully reflected in the prices paid by consumers. Supplementary Material S.3 discusses how this comparator relates to a tax imposed further upstream or across a broader market.

2. Dechezleprêtre et al. (2022) show opposition to a regulation of the food sector in global survey data. Additionally, Douenne and Fabre (2020) document considerable opposition to meat taxes in France, and Fesenfeld (2023) outlines the political challenges of implementing such taxes in Germany. Dietary shift towards lower-emission foods would substantially reduce food emissions (Grummon et al., 2023; Kim et al., 2020; Poore and Nemecek, 2018; Scarborough et al., 2023).

benchmark, the reduced-form mechanism evidence, and WTP for label presence are obtained directly from the experiments. I then use a structural model to quantify the relative contribution of salience and misperception correction and to decompose the resulting consumer-welfare effects.

In Experiment 1 ($N = 287$), I elicit participants' WTP for typical student canteen meals with and without seeing carbon labels. Participants receive the meal corresponding to their purchase decisions immediately after the experiment. The label-induced change in WTP identifies the local carbon tax that most closely reproduces the labels' aggregate emissions effect for the same participants, meals, and prices. Experiment 2 uses guest-level canteen purchase data to analyze a 7-week carbon labeling intervention in one of Bonn's canteens, using difference-in-differences to compare changes in purchasing behavior to those in Bonn's other canteens. I then compare the estimated label effect with demand responses associated with naturally occurring variation in relative meal prices. Within each experiment, the label effect and the relevant price response are estimated for the same choice environment and population.

The two approaches yield similar estimates. Experiment 1 identifies a local carbon-tax-equivalent benchmark of approximately €155 per tonne of CO_2e, with rates from €146 to €164 per tonne generating the same closest attainable emissions change (95% percentile interval for the €155 midpoint: €46.3–€242.6/t). These estimates imply an economically meaningful benchmark relative to EU ETS auction prices averaging approximately €75/t in the first two quarters of 2026 and the €55–65/t 2026 auction price corridor for Germany's carbon-pricing system.[3] Experiment 2 estimates convey a tax benchmark of €133/t, a reduction in meat choice of 5%, and a reduction in emissions of approximately 3%.[4] The similarity between Experiment 1 and Experiment 2 estimates suggests that the framed-field benchmark is informative for ordinary, recurring purchases in the Bonn canteens and, in turn, strengthens confidence in the benchmarking method.

3. These figures refer to 2026 values (see European Commission, 2026; European Energy Exchange, 2026). The experiments were conducted in 2021 and 2022. The average EU-ETS price in 2022 was €80, and €53 in 2021 (German Environment Agency (UBA), 2024; International Carbon Action Partnership (ICAP), 2025). The German carbon tax was €30 in 2022, and €25 in 2021 (Wettengel, 2024). These systems do not price the full agricultural greenhouse-gas footprint of food, although Germany's national system covers fuel use in agriculture.

4. Comparable field studies report emission reductions of 4.3% in British student canteens (Lohmann et al., 2022) and 3.6% in a Swedish student canteen (Brunner et al., 2018).

Experiment 3 ($N = 444$) provides experimental evidence on two channels through which carbon labels affect choices. In this experiment, before seeing the carbon labels, participants first estimate the carbon emissions caused by each of the meals, allowing me to relate their later response to the labels to prior misperceptions. A separate attention treatment makes carbon emissions salient without revealing the emissions of the respective meals. Label responses are more negative when participants previously underestimated a meal's emissions. However, WTP also falls when participants overestimated emissions. Moreover, increasing attention alone reduces WTP for high-emission meals. These comparisons show that the labels' effectiveness cannot be explained solely by their correction of misperceptions and that salience is an important driver.

I separately elicit participants' WTP for seeing carbon labels when making their meal choices. Fewer than 5% are willing to pay to avoid the labels, and mean WTP for label presence is approximately €0.20 per meal. This provides incentivized evidence that consumers value access to the labels.

I then use a simple discrete-choice model to quantify the behavioral channels and organize the consumer welfare comparison. The model represents meal choice as a function of taste, price, prior emissions beliefs, environmental preferences, and the salience of emissions. Under the baseline structural model, I estimate that increasing attention accounts for a substantial share of the label-induced emissions reduction. For consumer welfare, the model distinguishes between two components. The first component is the welfare change generated by altered meal choices, which I calculate using the model parameters estimated from participants' behavior. Under the baseline assumptions, both carbon labels and the benchmarked local carbon tax increase consumer welfare through altered meal choices, by similar amounts. The second component is the residual value consumers attach to label presence after accounting for expected improvements in meal choice. This value can be positive or negative and may reflect, for example, pure information value or curiosity, increased "warm glow" when consumers are reassured that the choice they make is low-carbon, or general annoyance with the labels.[5] I estimate

5. Allowing an intervention to affect welfare both through a choice-mediated component and through direct psychological or informational effects is standard in this literature and is less restrictive than setting the latter component to zero; in Butera et al. (2022), for example, public recognition generates image utility conditional on a given level of performance, so welfare can change even if behavior does not (see also Allcott and Kessler, 2019; Thunström, 2019). My analysis is limited to the treated consumers' own welfare and excludes spillovers through

this component by subtracting the model-predicted choice-mediated benefit of adding labels from participants' WTP for label presence. The estimate is positive and increases the welfare advantage of labels relative to the tax.

My quantitative estimates contribute to assessing the policy relevance of carbon labels in student canteens, which is a particularly promising setting due to the ease of implementation and the large student population.[6] Although the numerical estimate is local, the comparable magnitudes across the framed and natural field experiments, alignment with previous estimates of label and price effects, and the absence of consistent heterogeneity by demographic characteristics make it a useful reference point for comparable institutional food settings. The benchmarking method itself applies more broadly and can be re-estimated for other products and populations.[7] This paper yields three broader contributions.

First, this paper shows the value of benchmarking behavioral climate interventions against carbon taxation. The emission reductions I estimate are similar to those reported in earlier studies. Lohmann et al. (2022) and Brunner et al. (2018) find reductions of 4.3% and 3.6%, respectively, while I estimate reductions of 4% in the representative canteen simulation in Experiment 1 and 3% in Experiment 2. This suggests that carbon labels are not unusually effective in my setting. Previous studies have sometimes characterized emissions reductions of this magnitude as modest (see, e.g., Brunner et al., 2018), and the broader literature has questioned whether behavioral interventions can make a meaningful contribution to climate mitigation (Chater and Loewenstein, 2023; Nisa et al., 2019). Yet, assessing the policy relevance of these effects requires under-

other consumers' responses, such as utility from knowing that labels change others' choices. This matches the experimental setting, in which decisions were private and made in isolation.

6. In Germany alone, 2.9 million individuals were classified as students in 2021 (Federal Statistical Office (Germany), 2023), with approximately 54% dining in student canteens at least once per week (Federal Ministry of Education and Research (Germany), 2023). It is relatively cost-efficient to implement carbon labels in student canteens since a single facility often caters to many guests, which means that calculating carbon labels on a single menu impacts many guests. For example, student canteens in Bonn centralize their menu planning and cater to over 2,000 guests daily. Further, most student canteens in Germany operate as non-profit institutions aimed to serve students. This may make the introduction of carbon labels feel like a more natural step than in the case of commercial restaurants, especially if desired by a majority of the student population.

7. Supplementary Material S.3 distinguishes the portability of the benchmarking method from the external validity of the numerical €155/t estimate. It discusses comparable canteen and wider food-choice settings, reports exploratory heterogeneity analyses by demographic and attitudinal characteristics, and considers longer-run and larger-scale implementation.

standing how difficult it would be to achieve comparable reductions through alternative policy instruments. My tax-equivalent benchmark provides such a comparison: in this setting, achieving an emissions reduction similar to that of the labels would require a substantial local carbon tax. This shows that a seemingly modest percentage change may therefore represent a strong policy response when consumption choices are difficult to shift. More broadly, my results illustrate that, without a benchmark for consumers' responsiveness to alternative policy instruments, the percentage reduction alone is not sufficient evidence that an intervention is ineffective or not worth implementing. Dismissing such effects on this basis risks rejecting meaningful and feasible emission reductions, particularly in settings where carbon pricing faces political resistance. The benchmarking approach, particularly the framed field experiment, can readily be applied to other consumption settings and interventions.

Second, this study provides the first experimental evidence that salience is an important channel through which carbon labels affect choices and that correcting misperceptions alone cannot account for their effect. This helps reconcile previously mixed results. For instance, Imai et al. (2026) find no treatment response in a design that isolates the misperception-correction channel. My results show that this null finding does not imply that carbon labels are ineffective when they also make emissions salient at the point of choice.

Third, I provide incentivized evidence that consumers value the presence of carbon labels. Mean WTP for seeing the labels is approximately €0.20 per meal, while very few participants prefer to avoid them. I incorporate this directly elicited value into the structural welfare analysis, which separates welfare changes arising from altered meal choices from the value of label presence independent of choice. The resulting estimates suggest that the latter component is positive and economically important.

Correspondingly, this paper contributes to three key strands of literature. First, my findings inform our understanding of the effectiveness of carbon labels on food products. Previous field studies (most closely related are Bilén, 2022; Brunner et al., 2018; Ho and Page, 2025; Lohmann et al., 2022) report emission reductions ranging from less than 1% to around 5%. Differences in estimates across studies may reflect differences in setting (e.g., how substitutable high- and low-carbon options are) or population (e.g., the baseline carbon intensity of diets or concern for emissions). Findings also differ by

study design, with online studies (e.g. Banerjee et al., 2023; Imai et al., 2026; Lohmann, Gsottbauer, et al., 2024) sometimes finding smaller and sometimes larger effect sizes depending on incentivization and study design.[8] This paper adds evidence from a German canteen, and the first benchmark and welfare comparison to a carbon tax.

More broadly, my method for establishing a tax benchmark offers a framework for comparing behavioral interventions to price instruments in a consistent, internally valid way. Rather than relying on elasticity estimates elicited in a different context and population, my benchmarking approach compares a single population's responses to both carbon labels and price changes in a single setting. The resulting carbon-tax equivalent therefore applies to the same population and choice setting as the underlying label estimate. This avoids combining a label estimate from one population and choice setting with a price response estimated in another (also see Lusk and Tonsor, 2016). Further, the benchmarking procedure in the framed field experiment does not rely on price variation, avoiding concerns about the endogeneity of prices. This approach can be extended to benchmark other interventions aimed at reducing carbon emissions, including those in energy, water, and household consumption (relating to e.g., Allcott, 2011; Brent, Cook, and Olsen, 2015; Tiefenbeck et al., 2018), as well as food purchasing behavior (relating to e.g., Jalil, Tasoff, and Bustamante, 2023; Lohmann, Gsottbauer, et al., 2024).

Second, my findings on the relative role of salience and misperception correction, providing the first experimental evidence on the role of salience in driving carbon-label effectiveness, contribute to a growing literature on salience and attention in consumption decisions (e.g., Busse et al., 2013; Chetty, Looney, and Kroft, 2009; Taubinsky and Rees-Jones, 2018), particularly in environmental contexts (Allcott and Taubinsky, 2015; Rodemeier and Löschel, 2022; Tiefenbeck et al., 2018). I quantify the relative contributions of salience and misperception correction, offering new insight into the cognitive mechanisms behind behavioral policy tools.

Finally, I contribute to the literature on the welfare implications of nudges and information policies. One strand applies generalized frameworks to evaluate a range of interventions across settings (e.g., Allcott et al., 2025; Hahn et al., 2024; List et al., 2023), while another develops more detailed models for specific interventions (e.g., Allcott and

8. See Lohmann, Pizzo, et al. (2024) for a recent meta-analysis of demand-side interventions for food consumption. See Supplementary Material S.3 for exploratory heterogeneity analyses by demographic and attitudinal characteristics.

Kessler, 2019; Andor et al., 2023). My paper fits within the latter strand, offering a detailed welfare evaluation of carbon labels relative to taxes. However, my findings may also inform the former strand, for example by providing empirical guidance on modeling assumptions concerning behavioral channels and consumer welfare. In particular, the dominant role of salience and the presence of psychological effects independent of behavior change may generalize to other labeling interventions, such as those related to health or animal welfare.

The remainder of this paper proceeds as follows. Section 2 describes Experiment 1, which benchmarks the effectiveness of carbon labels using a framed field experiment. Section 3 presents Experiment 2 and examines whether a similar benchmark emerges from ordinary canteen purchases and field price variation. Section 4 examines the behavioral channels driving the observed effects, drawing on data from Experiment 3. Section 5 presents reduced-form evidence from all three experiments concerning consumers' preferences toward the presence of the label. Section 6 introduces a simple model of the labels' effects on behavior and welfare, structurally estimated using Experiment 3. Section 7 concludes.

2 Experiment 1: Benchmark identification in the framed field experiment

This section presents the framed field experiment method I use for benchmarking the effectiveness of carbon labels relative to a local carbon tax applied to the same meals. The framed field experiment provides precision and control. Subsection 2.1 describes the experimental design and data, and subsection 2.2 shows the empirical strategy and results.

2.1 Experimental design

Overview. I designed the framed field experiment to benchmark the effectiveness of carbon labels against a local carbon tax. The key design feature is that I observe participants' willingness to pay (WTP) for the same meals with and without carbon labels, allowing both sides of the benchmark to be estimated for the same participants and meal pairs. Participants make repeated, incentive-compatible meal purchase decisions, with one decision implemented and the corresponding lunch provided immediately after the experiment, at lunchtime that same day. WTP is elicited relative to a cheese sandwich,

which serves as a constant alternative lunch. Participants first value meals without labels and are then randomly assigned either to see carbon labels or to repeat the decisions without labels, allowing me to separate label effects from effects of repeated elicitation. The labels report greenhouse gas emissions using quantitative and traffic-light information. The experiment also separately elicits participants' WTP to see or avoid carbon labels in subsequent meal choices.

Experiment timeline. The experiment timeline is shown in Figure 1. First, participants receive an explanation of the WTP elicitation and answer comprehension questions.[9] Next, they indicate their baseline WTP for four canteen meals, incentivized by a Becker–DeGroot–Marschak (BDM) mechanism.[10] To create buffer time (approx. 5 minutes) before the second WTP elicitation, participants answer incentivized guessing questions on unrelated topics (e.g., the length of a popular running route in Bonn).[11] In the second WTP elicitation, the framing of the decisions depends on the randomly assigned treatment:

- Control: Decisions are identical to the baseline elicitation.
- Label: Participants see carbon labels as they make their decisions.

For additional insights, WTP is elicited a third time,[12] with altered treatment conditions:

- Participants previously in the Label condition are in the Offset condition: Participants are informed that meal emissions (meal or sandwich) will be offset. This serves as a robustness check for Experiment 3 and is detailed in Supplementary Material S.4.5, with reduced-form results in Table S.9.
- Half of the participants previously in the Control condition receive the Label condition, and half of the participants previously in the Control condition repeat the Control condition.

9. Participants must answer correctly to proceed. Those who made more than five unsuccessful attempts at the comprehension check are excluded, as preregistered.

10. See Supplementary Material S.4.4 for details.

11. Among the 287 participants in the Experiment 1 analysis sample, average completion time for the guessing block was 5.1 minutes (median: 4.7 minutes).

12. Analyses control for third-round data. Main results replicate using only the first two rounds; see Online Appendix Table OA.2, Column (2).

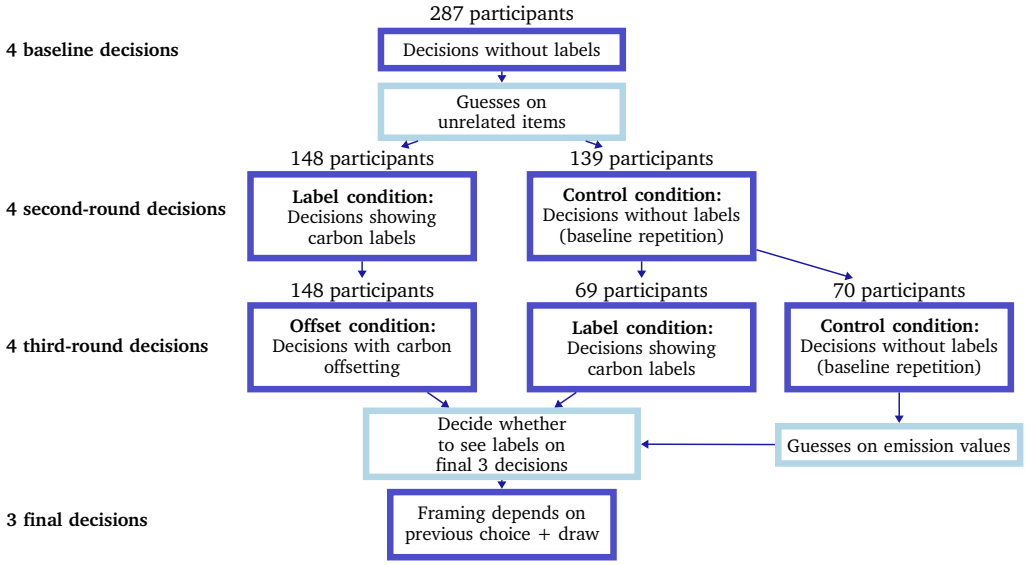


Figure 1. Experiment 1 schedule and participant numbers after data cleaning

Each round includes four meal purchasing decisions (12 total). After the third round, participants who have not previously seen labels estimate meal emissions.[13] All participants then guess the calories of the four meals and the cheese sandwich.[14] Three final decisions involve previously unseen meals. Before seeing these meals, participants indicate whether they want to see carbon labels for these decisions and state their WTP to enforce their choice, incentivized via a BDM mechanism.[15] Finally, participants answer questions on environmental attitudes and psychology.[16]

Details on the meal purchasing decisions. In each of the 15 decisions, participants first choose between a specific meal and a cheese sandwich. An example from the Label condition is shown in Figure 2. The left option changes across decisions to cycle through four meals, while the right option (cheese sandwich) remains constant.[17] In the repeated rounds, non-vegetarians evaluate two meat and two vegetarian meals, while

13. Used in Figure 6. As these questions occur after the third WTP elicitation, they do not affect results in this section.

14. I collect these to test whether the labels affect participants' perception of the nutritional content of meals. I do not find evidence of this being the case (Table S.7).

15. See Supplementary Material S.4.4 for details. Results are discussed in Section 5.

16. I collect these for the heterogeneity analyses in Supplementary Material S.3.3.

17. Participants always make their choice relative to the cheese sandwich as the alternative lunch to reflect the real-world fact that humans must eat and deciding not to eat one item implies needing to eat an alternative item. By directly eliciting WTP relative to an alternative lunch, I observe this alternative item, giving me more precision and control when interpreting results. Labels are always either shown or not shown for both options. The cheese sandwich option remains constant to ease comparison across WTP elicitations. If a participant rather likes or dislikes cheese sandwiches, this will affect all their decisions in the same manner, still enabling a clean comparison between WTP values. To control for left-right effects, positioning is reversed in half of the sessions. Meal order is randomized.

vegetarians evaluate the two common vegetarian meals plus two additional vegetarian meals.[18]

After selecting a preferred option, a second window appears where participants indicate how much of their experiment payment they would forego to secure their preference. The question wording in this second step differs depending on whether a participant selected the meal or sandwich as their preferred option. If they preferred the meal, they are asked to indicate their WTP to receive it instead of the sandwich. If they preferred the sandwich, they are asked to indicate their WTP to receive it instead of the meal. Responses are made on a slider in five-cent intervals up to €3.00, as this is the maximum price of any canteen meal and thus a natural upper bound. WTP values at the interval boundaries (±€3.00) occurred in less than 4% of observations. The WTP elicitations are incentive compatible and payoff consequential, mitigating potential experimenter demand effects.[19] This procedure effectively captures participants' WTP for the meal *relative to the cheese sandwich*, allowing for negative values: If a participant states they prefer the sandwich and have a positive WTP to secure their choice, I multiply this amount by -1 and interpret it as negative WTP for the meal.[17]

In the four baseline decisions, participants see only the meal name and main ingredient, without carbon labels. This mirrors how meals are typically displayed on the student canteen website. The four second-round and four third-round decisions resemble the baseline, except for framing differences across treatments. In the LABEL condition, emission labels are added to both meal options (see Figure 2 for an example). Carbon labels display greenhouse gas emissions (in kg), a traffic-light-style ordinal rating, and the equivalent driving distance (in km) to produce the same emissions (see Figure 2). I co-designed the labels with Bonn's student canteens to reflect a format they would realistically implement. A similar design was in fact adopted when the canteens later introduced permanent carbon labels (see Supplementary Material S.5.9). Combining quantitative and ordinal information has been shown to be effective in prior studies

18. Meals are detailed further below and in Supplementary Material S.4.2. Participants with stricter dietary requirements (vegan, gluten-intolerant, lactose-intolerant, or halal) are excluded.

19. Details on the BDM mechanism, the experiment procedure and experiment meals are shown in Supplementary Material Section S.4, and the distribution of baseline WTP values is shown in Section S.1.3.

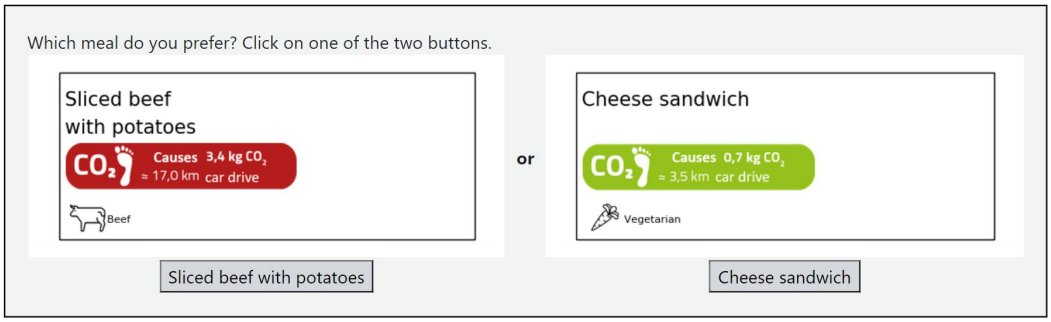


Figure 2. Meal purchase decision example: Decisions with labels

(Taufique et al., 2022; Potter et al., 2021).[20] In the CONTROL condition, framing remains unchanged. In the OFFSET condition, participants are informed that the emissions of the consumed meal option will be offset through a donation to the nonprofit carbon offsetting service Atmosfair. The OFFSET condition is not further discussed here, but details are in Supplementary Material S.4.5, with results in Table S.9.

Participants and data. 304 participants from the BonnEconLab, the University of Bonn's behavioral experimental lab, took part in one of eight experimental sessions between October 26 and November 5, 2021. I preregistered the experiment design and main outcomes shown in this section under AEARCTR-0008435 (Schulze Tilling, 2021a).[21]

Participants are informed that the experiment is conducted online but that they must visit campus directly afterward to collect their cash payment and lunch. They receive no further details on the experiment's purpose. The experiment is conducted using oTree software (Chen, Schonger, and Wickens, 2016).

Meals are provided by the student canteen. While all experiment meals are regularly offered by the canteen, they are not available on experiment days, meaning the canteen prepares them exclusively for participants. Participants receive a warm, ready-to-eat meal based on their online experiment choices and can consume it on-site. Fewer than 4% of participants did not collect their meal and cash payment after the experiment.

I exclude the 3% fastest participants and those who made more than five unsuccessful attempts at the comprehension check, as preregistered.[22] One incomplete response is also dropped, leaving 287 participants who were computer-randomized into treatments.

20. Meal emissions were calculated using the application Eaternity Institute (2020). The student canteen provided meal recipes for these calculations.
21. See Table S.9 in the Supplementary Material for all preregistered main results.
22. See Schulze Tilling (2021a). Dohmen and Jagelka (2024) find that fast respondents are more likely to give random answers.

Participants are on average 24 years old; 67% are female, 80% are students, and 25% are vegetarians. The sample is roughly representative of regular student canteen guests. 22% of baseline WTP values are 0 (indicating indifference between the meal and sandwich), 17% are negative (preference for the sandwich), and the rest are positive, with some bunching around €1. Less than 4% of WTP values are at the boundaries of the -€3 to €3 interval.[23]

2.2 Estimation strategy and results

I first estimate how carbon labels affect participants' WTP for high- and low-emission meals and then identify the local carbon tax that most closely reproduces the labels' emissions effect.

Effect of carbon labels on willingness to pay. For each participant-meal combination, I am interested in the within-subject *change* in participant *i*'s WTP between baseline and a subsequent elicitation round. Using the change directly as the outcome variable controls for individual- and meal-specific factors (e.g., taste, hunger, mood) that should remain constant across both elicitation rounds.[24] The Control condition captures changes due to repeated elicitation, while the Label condition additionally captures the effect of seeing carbon labels. Because WTP is measured relative to the cheese sandwich, I allow the label effect to differ depending on whether the meal has higher or lower emissions than the sandwich:

$$\Delta WTP_{ijm} = \beta_1 High_m + \beta_2 Low_m + \gamma_1 (Label_{ij} \times High_m) + \gamma_2 (Label_{ij} \times Low_m) + \rho ThirdRound_j + \varepsilon_{ijm} \quad (1)$$

where ΔWTP_{ijm} is the change in WTP for meal m by individual i in round j, relative to their baseline WTP for the same meal. $High_m$ and Low_m indicate whether meal m

23. See Supplementary Material S.1.1 for a randomization check, Supplementary Material S.1.2 for statistics on representativeness, Supplementary Material S.1.7 for descriptive split-sample estimates for students and non-vegetarians, and Supplementary Material S.1.3 for the full baseline WTP distribution.

24. As a robustness check, I estimate an alternative specification using WTP as the outcome and including participant-by-meal fixed effects. This yields similar results (Supplementary Material S.1.8).

has higher or lower emissions than the sandwich.[25] The coefficients β_1 and β_2 capture changes under repeated elicitation, while γ_1 and γ_2 identify the additional effect of seeing the labels. $ThirdRound_j$ controls for decisions made in the third elicitation round.[26] Standard errors are clustered at the participant level.

Relative to the Control condition, carbon labels increase WTP for lower-emission meals by €0.14 and reduce WTP for higher-emission meals by €0.31 (Table OA.2). The corresponding changes in the Control condition are comparatively small and move in the opposite direction.

Benchmarking labels against a local carbon tax. I next ask what emissions-based price adjustment would generate an emissions reduction similar to that caused by the labels. Figure 3 illustrates the intuition for higher-emission meals. Carbon labels shift the empirical demand curve downward relative to baseline WTP. A carbon tax can produce a similar shift in effective demand by increasing the price of a higher-emission meal relative to the sandwich. The benchmark asks how large this price adjustment must be to reproduce the labels' aggregate emissions effect.

The emissions effects of both interventions depend on the available meals, their prices, and which consumers are close to a choice threshold. I therefore calculate the benchmark under several choice environments, beginning with a simultaneous choice among all experimental meals and then adjusting the offer to more closely resemble the canteen setting.

The simulations use the four experimental meals for which I observe WTP both with and without labels. In the main specification, non-vegetarians choose among all four warm meals and the cheese sandwich, while vegetarians choose among the two vegetarian meals and the sandwich.[27] For each participant, I infer the option with the highest net valuation, defined as WTP relative to the sandwich minus the corresponding relative price. The sandwich's relative net valuation is normalized to zero, so it is selected when all available warm meals have negative net valuations. The main simulation uses

25. For non-vegetarians, three of four meals had higher emissions; for vegetarians, two of four. See Supplementary Material S.4 for details.

26. Dropping round 3 decisions yields similar results; see Online Appendix Table OA.2, Column (2). The categorization into $High_m$ and Low_m is exhaustive and I thus do not additionally include a separate $Label_{ij}$ indicator.

27. Vegetarians evaluated four vegetarian meals in the experiment; the simulations use only the two that were also offered to non-vegetarians.

the 217 participants who experienced labels, and uses canteen prices at the time of the experiment.

I simulate choices under three states. Baseline choices use participants' baseline WTP, while label choices use WTP elicited with labels. For the tax counterfactual, I start from baseline WTP and prices and change each warm meal's price relative to the sandwich according to the difference in their emissions. Specifically, I add to the warm meal's relative price the tax rate multiplied by its emissions difference from the sandwich. A meal that emits more than the sandwich therefore becomes relatively more expensive, while a meal that emits less becomes relatively cheaper. For example, at €155/t, a meal emitting 1.4 kg CO_2e has its price relative to the 0.7 kg sandwich increased by about 11 cents. I then calculate mean emissions under baseline, labels, and each tax rate and identify the tax rates whose simulated emissions reduction most closely matches that generated by the labels. Online Appendix OA.1.2 formalizes the simulation and reports additional implementation details.

Because choices are discrete, a range of tax rates can generate the same closest attainable emissions change. In the main specification, this matching set is €145.8–€164.3/t CO_2e. I use its midpoint, €155.0/t, as a representative scalar benchmark; its participant-bootstrap 95% percentile interval is €46.3–€242.6/t.

As a robustness check, I use the Control condition to estimate the WTP changes that arise when the same choices are repeated without labels, and subtract these changes from treated participants' WTP under labels. I also remove the common effect of making a decision in the third elicitation round. The resulting benchmark is slightly higher, at €167.4/t (Online Appendix Table OA.5).

Table 1 examines how the benchmark changes with the choice environment. Column (1) reports the simultaneous-menu specification described above. Column (2) more closely approximates the canteen's daily menu structure, where guests typically choose between one meat and one vegetarian main. I therefore pair the experimental meat and vegetarian meals to form four possible menus and weight these menus according to how frequently the corresponding types of meat and vegetarian meals were offered during the Experiment 2 labeling period. Column (3) uses the same weighted menu structure but removes the sandwich as an outside option. Online Appendix OA.1.2 provides the menu weights and other details.

Across the three scenarios, the labels reduce emissions by approximately 4–12%. Emissions savings differ across scenarios because of differences in the available options and substitution possibilities. For instance, the high-emission beef meal is available in every non-vegetarian choice set in scenario (1) but in only some of the choice sets in scenario (2), and vegetarians in scenario (3) have no margin for adjustment. Because these offer constraints limit responses to both the labels and a carbon tax, the emissions-matching range varies less across scenarios. This illustrates how benchmarking the labels against a tax provides a more targeted measure of their relative effectiveness within a given choice environment.

Table 1. Carbon-tax benchmark across alternative choice environments

	Meals on offer		
Outcome	(1) All options	(2) Canteen split	(3) Canteen split, no outside option
Baseline emissions (g)	943.4	790.1	883.6
Labels: Δ emissions (g)	-113.5	-33.3	-77.9
Labels: Δ emissions (%)	-12.0	-4.2	-8.8
€155/t: Δ emissions (g)	-101.4	-32.3	-79.0
Matching range (€/t)	[146, 164]	[146, 164]	[153, 170]

Notes: All simulations are based on baseline WTP and WTP with labels elicited from participants in the LABEL condition. In Column (1), non-vegetarians choose among the sandwich and all four warm meals, while vegetarians choose among the sandwich and the two vegetarian meals. Column (2) offers one meat and one vegetarian option only, together with the sandwich. Meat-vegetarian offer pairings are weighted in the aggregation to reflect the canteen offer. Column (3) removes the sandwich from this offer-pattern simulation. Column (1) contains 217 simulated choices from 217 participants; Columns (2) and (3) each contain 772 simulated choices from 217 participants. Matching ranges are rounded to the nearest euro and are sets of equally close point estimates. See Online Appendix Section OA.1.2 for details.

The right panel of Figure 3 illustrates the demand shift generated by the €155/t benchmark.[28] Comparing the two panels, the labels and the tax produce broadly similar shifts in demand, although the two curves are not identical. This is to be expected: the €155/t benchmark is chosen because it is constructed to come closest to reproducing the labels' overall emissions reduction at the observed study-period prices.[29] Nevertheless, this illustration shows how the comparison provides a useful sense of scale. The label-

28. For this illustration, I assume that the carbon cost is passed on to consumers. It does not matter whether it is added at checkout or imposed earlier in the supply chain, as long as it produces the same change in the final prices paid by consumers. Supplementary Material S.1.9 provides further intuition.

29. A complementary calculation summarizes the label-induced WTP response as an emissions gradient rather than matching simulated emissions at a particular set of prices. It yields a very

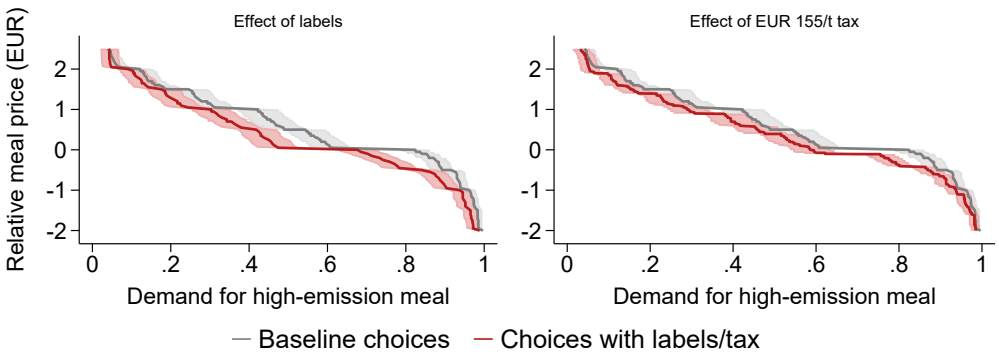


Figure 3. Demand curve shifts for high-emission meals

Notes: Pooled empirical demand curves for high-emission meals ($N = 603$ participant–meal observations) using the full sample of participants in the LABEL condition. At each five-cent price, demand is the share of participant–meal observations with WTP at least as high as that price. Gray lines show baseline WTP; red lines show WTP with carbon labels (left) and net WTP after applying a carbon tax of €155 per tonne to both the meal and the sandwich (right). Shaded areas are pointwise 95% percentile intervals from 999 bootstrap replications that resample participants.

induced shift in the left panel of Figure 3 is difficult to assess on its own, whereas the tax comparison shows how large it is relative to a familiar policy instrument. The key experimental feature enabling this benchmarking analysis is the incentive-compatible elicitation of WTP, a design feature which is portable to other interventions and domains.

3 Experiment 2: Field-based assessment of the benchmark

This section examines whether a similar benchmark emerges from ordinary canteen purchases. In subsection 3.1, I estimate how the labels affect purchases over seven weeks. In subsection 3.2, I compare this effect with demand responses associated with variation in the canteens' prices to establish a field-based benchmark. The resulting estimate is similar to that in Experiment 1, providing field evidence that a tax-equivalent effect of similar magnitude emerges for ordinary, recurring choices in the Bonn canteens and providing an independent check on the WTP-based method.

3.1 The effect of labels

3.1.1 Experimental design. To evaluate the effectiveness of carbon labels in a natural field setting, I partnered with Bonn's three student canteens. Key features of the exper-

similar benchmark of €156/t CO_2e (95% interval: €101–215/t); see Online Appendix OA.1.4 and Supplementary Material Section S.1.9.

imental design are summarized below; further details are provided in Supplementary Material S.5. The experiment and primary outcomes were preregistered.[30]

(1) The experiment follows a difference-in-differences design with three phases: (i) pre-intervention (4 weeks, no labels), (ii) intervention (7 weeks, labels added in the treatment canteen but not in the two control canteens), and (iii) post-intervention (3 weeks, no labels).

(2) The three canteens are located over 1.7 km apart, and only a minority of guests regularly visit more than one canteen, with no evidence of treatment-induced sorting (see Supplementary Material Section S.5.3 and Table S.27). My main analysis is intent-to-treat assigning treatment status based on guests' most-frequented canteen pre-intervention.

(3) Carbon labels resemble those in Experiment 1, combining quantitative and ordinal information. In the treatment canteen, labels are shown on the online menu, digital billboards, and paper signage at meal counters (examples in Figure 4).[31]

(4) Labels are applied only to the main meals, of which usually two are offered daily, one of these vegetarian and one meat-based. Labels are not applied to sides or desserts, for reasons of clarity and feasibility.[32] The main meals account for roughly 70% of total lunchtime emissions. During the sample period, the meat option always had higher or roughly identical emissions compared to the vegetarian option.

(5) The main analysis is based on canteen purchasing data. Because 69% of purchases are made with personalized payment cards, I can track individuals' purchasing behavior over time.

(6) Pre- and post-intervention surveys (pre: N>1,700, completed and identifiable post-survey respondents: $N = 1{,}004$) additionally capture demographic data and opinions. Survey responses are linkable to the purchasing data. See Supplementary Material S.5.8 for details.

3.1.2 Estimation and data. I estimate difference-in-differences models that compare purchases in the treatment and control canteens before, during, and after the labeling

30. See AsPredicted#95108. Preregistered outcomes include meat/vegetarian consumption, green-labeled meal uptake, greenhouse gas emissions, and canteen visits during and after the intervention. Full analyses are reported in Online Appendix OA.3.1.

31. Emissions are based on canteen recipes and the Eaternity Institute (2020) database, consistent with Experiment 1.

32. Details on the canteen setting are in Supplementary Material S.5.

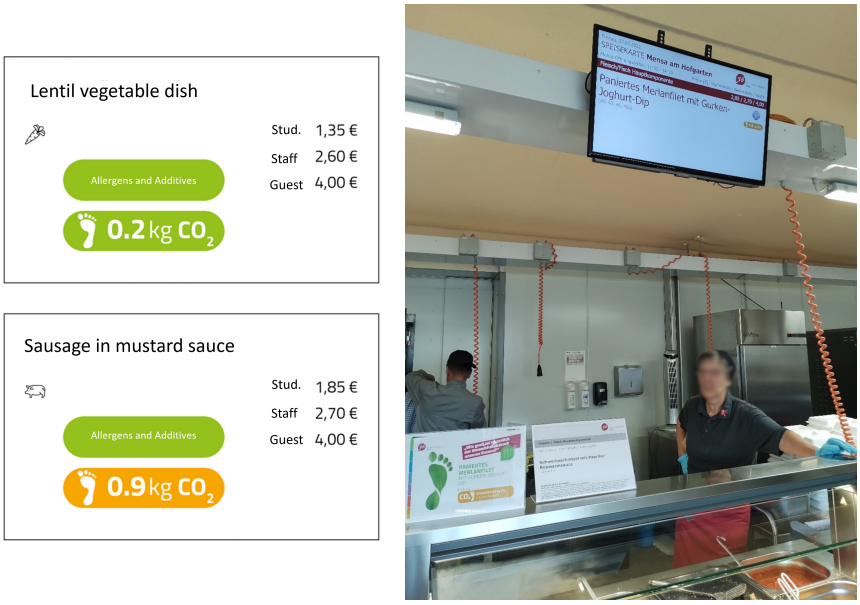


Figure 4. Labels in the canteen

Notes: Labels online (left, menu translated from German) and in the student canteen (right).

intervention.[33] The main outcome is whether a guest chooses the higher-emission meat option rather than the vegetarian option, a choice that is available every day. Alternative outcomes based on the emissions and displayed color of the chosen meal are reported in Online Appendix OA.3.1.

I use transaction data from April 4 to July 8, 2022, after which irregular canteen closures prevent a clean comparison. I exclude purchases by Ukrainian refugees who received free meals starting in week 9 of the study period, week 5 (which coincided with a "Healthy Campus" event in all canteens),[34] and the final week before the irregular closure. The final sample contains 124,784 transactions from over 2,000 guests. For each transaction, I observe the meal, price, location, time, and whether the guest is a student or employee.

33. Treatment was introduced in one canteen and compared with two control canteens. The treatment canteen was designated by the canteen operator based on operational aspects unrelated to potential outcomes. Causal identification hinges on the assumption of no contemporaneous treatment-canteen-specific shocks uncontrolled for. The small number of canteens reflects issues of practicality and is comparable to previous studies, with the design of Lohmann et al. (2022) involving five and Brunner et al. (2018) involving a single canteen in a before and after comparison. Experiment 2 provides complementary evidence to Experiment 1, in which treatment is assigned at the participant level. Supplementary Material S.5 further describes the setting, Online Appendix OA.3.1 reports robustness checks, and Supplementary Material S.3.3 discusses the scope of the field comparison.

34. The event featured an additional, rarely chosen ($\approx$ 10% of sales) premium-priced meal and light promotional activities. Since it affected all canteens, it should not confound the treatment effect unless its impact differed across locations; I omit it conservatively.

Approximately 69% of transactions are linked to personalized payment cards. For the intent-to-treat (ITT) analysis, I retain guests who visited any canteen at least five times before the intervention and made at least 80% of those visits to the same canteen. Treatment status is assigned solely from this pre-intervention behavior. See Supplementary Material Sections S.5.5–S.5.6 for alternative thresholds and exclusions. In this sample, I can identify days on which a guest's card was used to purchase more than one main meal. I exclude these days, because it is unclear which meal the guest personally consumed.

Figure 5 shows weekly difference-in-differences estimates in the ITT sample. The specification controls for week, day of the week, the exact meat and vegetarian main meals offered, and whether a second meat or vegetarian main is offered in the guest's ITT canteen. Because canteen visits are irregular, I do not include guest fixed effects for the week-by-week visualization. The estimates show similar pre-intervention trends and relatively stable negative effects on meat consumption during the intervention period. While the labels' effects seem to persist in the two weeks following label removal, post-intervention effects do not seem to persist into the following semester (see Figure S.6). Such a short-term persistence of effects might reflect a temporarily increased salience of carbon emissions in guests' consumption decisions (aligning with the findings of Byrne et al., 2024).

To estimate the average treatment effect used in the subsequent comparison with price responses, I additionally include guest fixed effects. This specification estimates that the labels reduce the likelihood of purchasing the higher-emission meat meal rather than the lower-emission vegetarian option by approximately 2 percentage points, a 5% decline relative to baseline meat consumption in the treatment group. The estimate is virtually identical when the specification is restricted to purchases made by students, excluding employees (Online Appendix Table OA.17, Panel B, Columns (3)–(4)). The estimate is also similar in the full sales sample, across progressively richer control sets, and when either control canteen is omitted (Online Appendix Tables OA.17 and OA.19). A separate analysis estimates an approximately 3% reduction in emissions from purchased main-meal components (Online Appendix Section OA.3.1).[35]

35. The main emissions specification excludes three days with unusually different meal offers across canteens. Including all sample days yields an estimated reduction of approximately 1.3%, which is less precisely estimated.

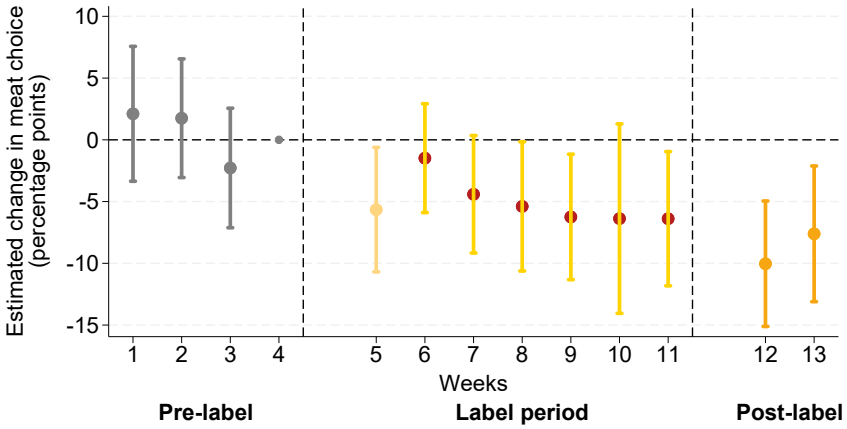


Figure 5. Event study: Difference-in-differences estimates

Notes: Difference-in-differences estimates of the likelihood of consuming the meat option (percentage points), using week 4 as the baseline. Weeks 1–4 are pre-intervention, weeks 5–11 intervention, and weeks 12–13 post-intervention. Week 5 is shown but excluded from the average treatment-effect estimation because it coincided with a university-wide event. The specification uses the ITT sample and controls for week, day of the week, the exact meat and vegetarian main meals offered in the guest's ITT canteen, and ITT canteen second meat and vegetarian offers. Standard errors are clustered at the guest level; bars show 95% confidence intervals. Full coefficients are reported in Supplementary Material Section S.2.2.

3.2 The effect of a local carbon tax

3.2.1 Setting. To benchmark the effect of carbon labels against a potential carbon tax in the natural field setting, I examine how canteen guests respond to price variations and compare these responses to the labeling effects estimated in the previous subsection.

If the student canteens were to implement an emissions-based relative-price adjustment resembling a local carbon tax, the most straightforward way to do so would be to increase the price difference between meat and vegetarian options without raising the overall price level.[36] Depending on the design of the tax, one could either apply a uniform price increase to all meat meals or vary the magnitude of the price adjustment by meat type.

To approximate demand reactions to such a local carbon tax, I exploit naturally occurring variation in the daily price difference between the main meat and vegetarian

36. The main meat meals offered in the canteens emit on average 1.59 kg CO_2e, compared with 0.45 kg for the vegetarian meals. Since the canteens are non-profit institutions mandated to provide affordable meals to students (see Section S.5 for details on the institutional set-up), raising the general price level would be inconsistent with their mission. The most practical way for them to introduce a carbon tax would therefore be to apply a linear tax based on emissions while simultaneously reducing the price of vegetarian meals, thereby generating a relative price increase for meat meals.

options. Price variation arises from daily meat–vegetarian meal pairings and nonuniform price increases across meals in October 2022. Because the canteens are non-profit institutions, the observed price differences primarily reflect ingredient and operating-cost variation rather than strategic pricing.

3.2.2 Estimation and results. For this analysis, I use student purchases from April 2022 to July 2023, excluding May 1–June 2, 2023.[37] I exclude meals outside the standard two main options, 0.5% of purchases for which the daily offer or price is unclear, and non-student guests since they face different prices. The resulting sample contains one meat-minus-vegetarian price difference for each canteen and day.

I estimate linear probability models relating meat choice to the meat-minus-vegetarian price difference. One specification pools the response for poultry, pork, and beef, while a second estimates separate responses by meat type. Fish is included as a separate interaction but is not used as a causal price-response estimate because it is offered almost exclusively on Fridays. Table 2 includes controls for the meat category and type of meat meal on offer. This specification controls for an array of meat meal characteristics, while still preserving 57% of the original variation in the price difference.[38] Meal pairings and prices were determined by the canteen operator rather than experimentally randomized. Interpreting the coefficients as causal price responses requires that, conditional on the included controls, the remaining variation in relative prices is uncorrelated with unobserved determinants of meal choice. Supplementary Material Section S.2.7 discusses this assumption and the institutional features supporting its plausibility.

Table 2 reports standard errors two-way clustered by calendar date and exact meat–vegetarian meal pair, allowing demand shocks to be correlated across purchases on the

37. This larger data set combines two data sets provided by the student canteen: one covering April to July 2022 (the main data set for this paper) and another covering August 2022 to July 2023. Data from May 1 to June 2, 2023 are excluded due to another intervention during this period. Unfortunately, individual guests cannot be linked across the two data sets, and emissions data are only available for meals served from April to July 2022.

38. Online Appendix Table OA.20 shows that including more granular controls preserves the general coefficient pattern, but naturally decreases the remaining variation in prices from which coefficients are identified. Price responses remain negative across alternative time periods reported in Supplementary Material Section S.2.7. Specifications that leave less identifying price variation sometimes yield smaller absolute coefficients; taken at face value, these would imply larger rather than smaller tax-equivalent benchmarks. Supplementary Material Table S.31 shows alternative standard errors.

Table 2. Comparison of effects: labels vs. an emissions-based relative-price adjustment

	Meat meal chosen	
	(1) Pooled poultry, pork, and beef	(2) By meat type
Price difference (in €)	-0.101***	
	(0.018)	
Price difference (in €) x Poultry		-0.058**
		(0.024)
Price difference (in €) x Pork		-0.132***
		(0.025)
Price difference (in €) x Beef		-0.202***
		(0.042)
Weekly time controls	Yes	Yes
Control for meat meal type	Yes	Yes
Predicted effect: €155/t CO_2e (pp.)	-1.7	-2.5
Observations	465,328	465,328

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable equals one when a transaction contains the meat meal and zero if the vegetarian meal is chosen. Coefficients are in probability units, so 0.01 corresponds to one percentage point. Linear probability regressions use student purchases from April 2022–July 2023, excluding all canteen observations from May 2023 and June 1–2, 2023 because of another canteen intervention during this period (a vegetarian month in the treatment canteen). Price difference is the price of the meat main minus the price of the vegetarian main. Meat-type indicators refer to the selected meat offer. Both columns include week and day-of-week fixed effects and canteen indicators. Treatment-canteen interactions control separately for the label period, the period after label removal through July 2022, and the period following another intervention in the treatment canteen, after June 2, 2023. Both columns also include broad meat-type controls and six mutually exclusive recipe-informed meat-offer categories: breaded or coated; minced, formed, or sausage; whole cut or roast; curry or international; goulash, ragout, or pan; and pasta or composed dishes. Both columns additionally include separate price-difference interactions for fish and lamb. Standard errors are two-way clustered by calendar date and exact meat–vegetarian meal pair. See Online Appendix Table OA.20 for alternative control sets, Supplementary Material Section S.2.7 for alternative time periods, and Supplementary Material Table S.31 for alternative standard errors. Table OA.21 shows the calculations underlying the predicted effect of the €155/t midpoint.

same date and when the same meal pair recurs. Column (1) implies that a €0.10 increase in the relative price of poultry, pork, or beef reduces the probability of choosing meat by 1 percentage point. Column (2) shows the strongest responses for pork and beef and a smaller response for poultry. The corresponding own-price elasticity is -0.5 for poultry, pork, and beef pooled; the type-specific elasticities are -0.3 for poultry, -0.7 for pork, and -1.1 for beef. These estimates are of a broadly similar order to those in Roosen, Staudigel, and Rahbauer (2022) and Liu, Ansink, and Bouma (2025).[39]

Applying the €155/t benchmark to the effective offer-share-weighted emissions gap of 1.11 kg CO_2e, the pooled coefficient predicts a 1.7-percentage-point decrease in meat consumption under a general meat tax; the corresponding predictions are 1.6 and 1.9 percentage points at the nearby €145/t and €165/t rates that bracket the Experiment 1 €146–164/t matching set. A tax differentiated by meat type, using the poultry, pork, and beef coefficients and weighting by offer shares, predicts a 2.5-percentage-point decrease at €155/t and decreases of 2.4 and 2.7 percentage points at the two bracketing rates.[40] These resulting effects are close to the approximately two-percentage-point label effect estimated both for the general canteen sample and the student-only canteen sample (Online Appendix Table OA.17), providing complementary field evidence consistent with the Experiment 1 benchmark. Direct inversion of the student-only label estimate gives €194/t for the general meat tax and €133/t for the differentiated tax.[41]

4 Experiment 3: Behavioral channels

Having established that carbon labels generate an economically meaningful demand response, I next examine why consumers respond. Experiment 3 focuses on two mechanisms that are particularly relevant for carbon labels and can be examined separately in the experimental design: correcting misperceptions about emissions and increasing

39. Supplementary Material Section S.2.6 provides a more extensive discussion of the price setting and elasticities.

40. Predictions are calculated as follows: For the general tax, I multiply the tax rate by the average meat-minus-vegetarian emissions gap and apply the pooled poultry, pork, and beef price coefficient. For the differentiated tax, I apply the poultry, pork, and beef coefficients to their type-specific emissions gaps and weight the resulting effects by canteen-day offer shares. In both estimations, fish retains its offer share but is assigned a zero response. Online Appendix Table OA.21 shows the calculation of the predicted change in meat choice under a €155/t tax differentiated by meat type.

41. I focus on the student-only sample for the main inversion because the price response is estimated using student purchases. Inversion of the estimate including all guests yields €195/t and €134/t, respectively. Supplementary Material Table S.33 reports 95% bootstrap intervals.

the salience of emissions at the moment of choice. These mechanisms are not intended to provide an exhaustive account of all channels through which labels may affect behavior. Rather, the experiment asks whether either can explain an important part of the observed response. To study them, I elicit participants' prior emissions beliefs and introduce an ATTENTION condition that makes emissions salient without revealing the emissions of the decision meals. This section presents reduced-form evidence on both mechanisms; Section 6 uses the same experimental variation to quantify their relative importance within the structural model.

Subsection 4.1 describes the experimental design, subsection 4.2 the experimental data, and subsection 4.3 the estimation strategy and results.

4.1 Experimental design

Overview. Experiment 3 follows a structure similar to Experiment 1, with two additions designed to study the two mechanisms. First, participants estimate the carbon emissions of the meals before seeing the labels, allowing me to relate their subsequent response to their prior beliefs. Second, the experiment includes an ATTENTION condition that increases the salience of emissions without revealing the emissions of the decision meals.

Experiment timeline. The experiment timeline is visualized in Figure S.11 and follows a structure similar to Experiment 1, with one key distinction. After the baseline WTP elicitation, participants in Experiment 3 guess the carbon footprints of various meals, including the meals on which they make purchasing decisions and six additional meals (see Figure 6 for a list). Participants receive emissions information for a single reference meal (Red Thai Curry with pork and rice, emitting 1.7 kg CO_2e, not used in the WTP elicitations) to guide their guesses. An example guessing screen is displayed in Figure S.15. These guessing tasks are incentivized and timed.[42] Participants do not receive feedback on whether their guesses were correct.

42. Participants answer each of the ten guessing questions on separate screens, presented in random order. The emissions of the reference meal are consistently displayed. See Supplementary Material S.4.7 for screenshots of the guessing instructions. The emission guessing questions take the place of the unrelated guessing questions in Experiment 1, i.e. Experiment 3 participants do not additionally respond to unrelated guessing questions as Experiment 1 participants do. Participants have 60 seconds per guess and receive €0.10 for each guess within 30% of the correct value.

The experiment then varies by treatment group, assigned via computer randomization. All participants again indicate their WTP for the four meals, but decision framing differs by treatment condition:

- Attention: The emission guessing questions themselves presumably direct participants' attention towards emissions, without revealing the emissions of the respective meals. To elicit the effect of this increase in attention on participants' choices, I elicit participants' WTP for the four meals again after the guessing task, with the WTP elicitation mirroring the baseline elicitation.
- Attention+Label: The WTP elicitation mirrors the WTP elicitation in the Label condition in Experiment 1 (see Figure 2), making participants both attentive and informed.
- Attention+Offset: Participants in the Attention+Offset condition are informed during the WTP elicitation that meal emissions will be offset, rendering choices carbon-neutral.[43]

To increase statistical power and gather additional insights, WTP for the same meals is elicited a third time,[44] with altered treatment conditions:

- Participants previously in the Attention+Label condition switch to the Attention+Offset condition, and vice versa.
- Participants in the Attention condition remain in the Attention condition.

Subsequent procedures, the design of meal purchase decisions and the incentivization of the WTP elicitations are as in Experiment 1.

Participants and set-up. 476 participants from the BonnEconLab pool at the University of Bonn take part in one of 12 experimental sessions between June 22 and July 8, 2021. The experimental design, sample restrictions, and comparison between the At-

43. In this condition, the meals receive labels that indicate carbon neutrality through CO_2e offsetting. The Offset condition is detailed in Supplementary Material S.4.5 and results are in Table S.9. This condition supports the structural estimation in Section 6 as discussed in Section 6.2.

44. Analyses control for third-round elicitations. Main results are consistent when using data from only the first two rounds.

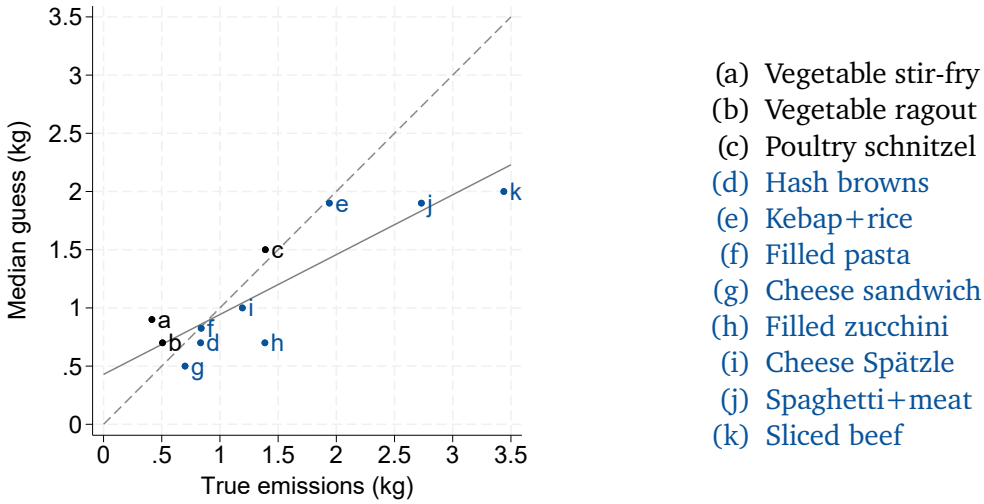


Figure 6. Median guess of the emissions caused by a given meal. Guesses are plotted against calculated emissions. Guesses closer to the dashed line are closer to the correct emission values. Black points indicate meals whose median guessed emissions exceed calculated emissions; blue points indicate meals whose median guessed emissions are below calculated emissions. Values are based on guesses made by the participants in Experiment 3, and 70 participants in the "Control, then Control" group in Experiment 1. The meal "Spaghetti with meat" was only guessed by the 70 participants of Experiment 1 guessing emissions. Graph based on 5,210 observations from 514 participants.

tention+Label and Attention conditions were preregistered.[45] Participant recruitment and set-up are as in Experiment 1.

4.2 Data

I exclude the 3% fastest participants and those who made more than five unsuccessful attempts at the comprehension check, as preregistered. The remaining 444 participants are computer-randomized into treatments. The sample is on average 26 years old, 55% female, 69% students, and 24% vegetarians.[46]

45. See Schulze Tilling (2021b). The preregistration refers to the Attention condition as the control group: these participants completed the emissions-guessing task but did not receive carbon labels. All preregistered main analyses are in Tables S.10 and S.11; Figure 6 was preregistered as an additional analysis. Figure 7a and Table OA.7 were not preregistered.

46. See Supplementary Material S.1.1 for a randomization check, Supplementary Material S.1.2 for statistics on representativeness, Supplementary Material S.1.7 for descriptive split-sample estimates for students and vegetarians, and Figure S.2 for the baseline WTP distribution.

4.3 Estimation Strategy and Results

Correcting Misperceptions as a Treatment Channel. This section provides reduced-form evidence on whether treatment effects align with a correction of misperceptions as the main mechanism driving effects. As descriptive evidence, Figure 6 shows that participants tend to underestimate emissions for high-emission meals (blue dots) and overestimate emissions for low-emission meals (black dots).[47] Misperception correction on its own predicts that WTP for a meal relative to the sandwich should fall following underestimation of its relative emissions and rise following overestimation. I use data from the Attention+Label condition to examine whether the observed responses follow this pattern.

$$\begin{aligned} \Delta WTP_{ijm} \;=\; & \beta_1 (AttLabel_{ij} \times Under_{im}) \\ & + \beta_2 (AttLabel_{ij} \times Over_{im}) + \rho\, ThirdRound_j + \varepsilon_{ijm}. \end{aligned} \tag{2}$$

where ΔWTP_{ijm} is the change in willingness to pay (WTP) for meal m by individual i in round j, relative to their baseline WTP for the same meal. As in Experiment 1, WTP for meal m is relative to the cheese sandwich.[48] $AttLabel_{ij}$ is an indicator for decisions in the Attention+Label condition. Data from other conditions are not included in this estimation. $Under_{im}$ is an indicator for whether the respective participant underestimated the difference in emissions between meal m and the cheese sandwich. This is determined by comparing the participant's estimated difference in emissions with the true difference.[49] Correspondingly, $Over_{im}$ is an indicator for whether the respective difference is overestimated.[50] $ThirdRound_j$ controls for whether round j is the third decision round.[51]

47. Further statistics on under- and overestimation appear in Supplementary Material S.1.10, including accuracy in meal ranking by carbon footprint.

48. See Section 2.2 and Supplementary Material S.1.8 for details on the dependent variable.

49. This refers to the signed difference, not the absolute difference. For example, if a meal generates 0.2 kg more emissions than the cheese sandwich, but a participant estimates it to produce 0.3 kg less, this constitutes an underestimation. In the main specification, I focus on the difference in emissions between meal m and the cheese sandwich because WTP is also measured relative to the sandwich. However, this implies that the relative error combines beliefs about both options. Online Appendix Table OA.7, Column (2), repeats the analysis conditioning only on participants' estimates of meal emissions and yields similar results.

50. In 55 instances (5% of observations) guesses are correct. Also in these instances, participants on average adjust WTP downward; see Online Appendix Table OA.7, Column (3).

51. This accounts for mixing data from second- and third-round decisions, as some participants experienced the Attention+Label condition in round two and others in round three. Excluding third-round observations yields similar results (Online Appendix Table OA.7, Column (4)).

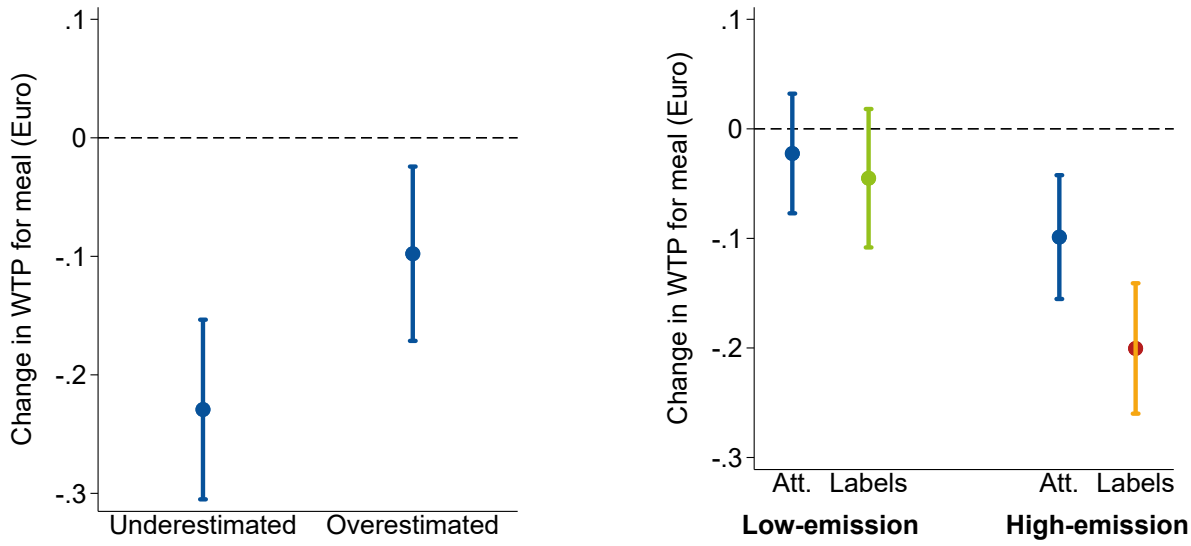
Figure 7a presents OLS estimates of equation 2. For participant-meal combinations in which participants underestimated meal m's emissions relative to the cheese sandwich, WTP decreases by approximately €0.23 when shown carbon labels. For participant-meal combinations in which participants overestimated meal m's emissions relative to the cheese sandwich, WTP decreases by approximately €0.10 when shown carbon labels. While the first coefficient aligns with a correction of misperceptions driving the labels' effect, the second does not; we would have expected an increase, not a decrease in WTP. The overall pattern is consistent with a correction of misperceptions playing *some* role, since the reduction is €0.131 larger following underestimation than following overestimation (underestimation minus overestimation: −€0.131, SE = €0.035, $p < .001$).[52] However, the negative response following overestimation shows that misperception correction cannot be the only channel driving treatment effects.

Increasing Salience as a Treatment Channel. This subsection examines whether treatment effects can be explained by an increased salience of carbon emissions. To estimate the magnitude of an attention effect, I analyze data from the Attention and Attention+Label conditions. Specifically, I estimate:

$$\begin{aligned}\Delta WTP_{ijm} = \beta_1(Attention_{ij} \times High_m) + \beta_2(Attention_{ij} \times Low_m) \\ + \delta_1(AttLabel_{ij} \times High_m) + \delta_2(AttLabel_{ij} \times Low_m) + \rho ThirdRound_j + \varepsilon_{ijm} \quad (3)\end{aligned}$$

where ΔWTP_{ijm} is defined as above. $Attention_{ij}$ and $AttLabel_{ij}$ are mutually exclusive indicators for decisions in the Attention and Attention+Label conditions, respectively. $High_m$ indicates that meal m has higher emissions than the cheese sandwich, while Low_m indicates that it has lower emissions. The four treatment-by-emissions-category coefficients report round-adjusted changes from baseline WTP in each treatment condition. The additional effect of displaying labels beyond increasing attention is $\delta_1 - \beta_1$ for high-emission meals and $\delta_2 - \beta_2$ for low-emission meals.

52. Note that prior misperceptions are not experimentally assigned, so this pattern is best interpreted as descriptive, not causal evidence on the role of misperceptions.



(a) Label effects by prior under- or over-estimation

(b) Attention versus Attention+Label

Figure 7. Reduced-form evidence on the behavioral channels in Experiment 3

Notes: Panel (a) plots OLS estimates of Equation 2 using the Attention+Label condition. It excludes 55 participant–meal combinations with exactly correct emissions-difference guesses. The underestimation-minus-overestimation comparison is −€0.131 (SE = €0.035, $p < .001$). Panel (b) plots estimates of Equation 3 using the Attention and Attention+Label conditions. The Attention+Label-minus-Attention comparison is −€0.102 (SE = €0.037, $p = .007$) for high-emission meals and −€0.023 (SE = €0.039, $p = .566$) for low-emission meals. Both specifications control for third-round decisions, cluster standard errors at the participant level, and display 95% confidence intervals. Complete regression results are reported in Online Appendix Tables OA.7 and OA.8.

Figure 7b illustrates the four coefficient estimates. Simply directing attention to carbon emissions reduces WTP for high-emission meals by approximately €0.10.[53] Providing labels in addition reduces WTP by a further €0.10. The corresponding low-emission estimates are small. These findings suggest that increased attention accounts for an important share of the treatment effect. Section 6 quantitatively assesses the relative contributions of attention and misperception correction to the label's overall effect.

5 Consumer Preferences for Carbon Labels

This section presents experimental evidence on consumer preferences for carbon labels in consumption decisions. Section 5.1 examines findings from Experiments 1 and 3, while Section 5.2 discusses results from Experiment 2.

53. Interpreting this change as an attention effect assumes that repeated elicitation would not otherwise affect WTP. Experiment 1's Control condition does not show a significant effect of repeated elicitation. Table S.9 combines Experiments 1 and 3 data for joint analysis.

5.1 Evidence from the Framed Field Experiments

In Experiments 1 and 3, participants make a payoff-consequential choice about whether carbon labels are displayed during their final consumption decisions. Using an incentive-compatible BDM mechanism, I elicit how much of their experiment payment they are willing to forgo to implement their preferred label condition.[54] Depending on the experimental draw, participants' bid can affect both the information displayed during their meal choices and their final payment. I record WTP to avoid labels as negative WTP for label presence. About 50% of participants submit a zero bid, indicating no positive monetary value for enforcing their preference. Fewer than 5% prefer labels to be absent and are willing to pay to avoid them. The remaining participants are willing to pay for labels, with 21% of all participants willing to pay €0.50 or more. WTP values are similar across treatment groups, with no clear differences by prior exposure to the labels. [55] Mean WTP for label presence is approximately €0.20. While Thunström (2019) finds that calorie labels impose greater psychological costs on individuals with low self-control, I find no evidence of such a correlation, as shown in Table S.22. However, I find a weak positive correlation between preferences for the labels and perceived social norms for reducing carbon emissions in food choices, as well as between preferences for the labels and self-reported willingness to use the information.

5.2 Evidence from the Natural Field Experiment

Following Experiment 2, student canteen guests participated in a follow-up survey assessing their preferences for permanent carbon label installation. As detailed in Supplementary Material S.5.8, the survey was framed as an opportunity for guests to provide feedback on various aspects of the canteen, with carbon labels being one of multiple topics. Carbon labels were not mentioned in the survey advertisement. Survey respondents were aware that their responses could impact student canteen policies, incentivizing them to report truthfully.

Among the 258 completed follow-up respondents who made at least one card-recorded main-meal purchase at the treatment canteen during the label period, reported

54. Both meal valuations and WTP for label presence were elicited using the same €0–€3 range, with BDM prices drawn uniformly in five-cent increments. I retained the same range for label presence to avoid introducing a different price scale during the experiment. However, elicited valuations may depend on the chosen range and price distribution, which I do not vary experimentally. See Supplementary Material S.4.4 for further details.

55. See Figure OA.1 for WTP distribution and Table OA.15 for differences across treatments.

noticing the labels, and answered both policy-support questions, 75.6% supported making the labels permanent, 16.7% were unsure, and 7.8% opposed the measure. By comparison, a revenue-neutral carbon tax of unspecified magnitude (described as making higher-emission meals more expensive and lower-emission meals cheaper) received 62.0% support, with 14.0% unsure and 24.0% opposed. While both measures have majority approval, opposition to the tax is more than three times as high as opposition to the labels. This difference suggests that carbon labels face markedly less resistance, making their implementation more politically feasible.

6 Structural Model

To formalize how the two behavioral mechanisms shown in Section 4 drive responses to carbon labels, I first introduce a simple discrete-choice model of meal choice and use Experiment 3 to identify its parameters. I then extend the model to consumer welfare and use participants' willingness to pay for label presence to identify an additional welfare component. Finally, I use the estimated model to compare the effects of carbon labels and a carbon tax on choices, emissions, and consumer welfare.

6.1 Basic model

There is a finite set of meals $\mathcal{M}$. Participant i chooses a meal $m \in \mathcal{M}$ that maximizes her *perceived utility*

$$u_i(m) = v_{im} - p_m - \theta\gamma e_{im}. \tag{4}$$

Here, v_{im} is the *consumption utility* of meal m that is independent of emissions,[56] p_m is the *price* of meal m, and e_{im} is the consumer's *estimate of emissions* caused by meal m at the moment of choice.[57]

56. One can think of v_{im} as a "catch-all" factor that captures the influence of all other meal attributes aside from carbon emissions that may influence utility. For the purposes of this paper, v_{im} can be treated as an exogenously given participant–meal-specific value. Alternatively, it can be derived from observable attributes x_m and an unobservable taste component ε_{im}, so that $v_{im} = \beta^T x_m + \varepsilon_{im}$.

57. Similar to Imai et al. (2026), I assume that perceived utility is additively separable in v_{im} and perceived environmental guilt.

The *salience* of carbon emissions, $\theta \in [0,1]$,[58] and *environmental guilt per perceived kg of emissions*, $\gamma \geq 0$,[59] jointly determine the weight placed on perceived emissions. At $\theta = 0$, emissions receive no weight in the choice; I normalize full attention to emissions as $\theta = 1$.[60]

The consumer's prior estimate of emissions caused by meal m is denoted by e_{im}^{prior}, which may differ from the true emissions, e_m^{true}. If the consumer receives information about true emissions, her updated estimate is

$$e_{im}^{\text{info}} = (1-\kappa)e_m^{\text{true}} + \kappa e_{im}^{\text{prior}}. \tag{5}$$

The parameter $\kappa \in [0,1]$ therefore captures the persistence of prior beliefs: $\kappa = 0$ corresponds to fully adopting the provided information, while $\kappa = 1$ corresponds to no updating.[61]

6.2 Identification and estimation using Experiment 3

The meal purchasing decisions used in Experiments 1 and 3 each have a binary choice set $\mathcal{M} = \{m, o\}$, where m is the warm meal and o is the outside option, the cheese sandwich. The structural estimation uses Experiment 3, which additionally elicits participant-specific prior emissions beliefs and contains the treatment states needed to identify θ, γ, and κ.

The WTP elicited in the meal purchasing decisions is the maximum surcharge the participant is willing to incur to receive meal m rather than option o. Setting perceived utility from the two options equal at this indifference surcharge gives model-predicted relative WTP

$$WTP_{im} = v_{im} - v_{io} - \theta\gamma(e_{im} - e_{io}).$$

58. I use a similar formulation as in the literature on attentiveness to taxes and resource consumption (Byrne et al., 2024; Chetty, 2009; DellaVigna, 2009). In the framework of Bordalo, Gennaioli, and Shleifer (2022), one reason why emissions may not be fully salient is a lack of prominence, as a meal's emissions are usually not featured at the moment of choice.

59. The model can equivalently be formulated in terms of warm glow from choosing a lower-emission option rather than guilt from choosing a higher-emission option. This changes the interpretation of γ but not the relative-utility comparisons used below.

60. For any alternative normalization $x > 0$ under full attention, one could redefine $\theta = \theta/x$ and $\gamma = \gamma x$.

61. The formulation leans on the evidence-informed framework proposed by Epstein, Noor, and Sandroni (2008) to model non-Bayesian updating. Bouchaud et al. (2019) use the same updating rule to study under-reaction in financial markets.

For compactness, define

$$\Delta v_{im} = v_{im} - v_{io}, \qquad \Delta e^{\text{prior}}_{im} = e^{\text{prior}}_{im} - e^{\text{prior}}_{io}, \qquad \Delta e^{\text{true}}_{m} = e^{\text{true}}_{m} - e^{\text{true}}_{o}.$$

The baseline structural specification allows Δv_{im} and $\Delta e^{\text{prior}}_{im}$ to vary freely across participants and meals, while treating θ, γ, and κ as common across participants.

At baseline, participants rely on their prior emissions beliefs and emissions receive the baseline salience weight θ:

$$WTP^{\text{B}}_{im} = \Delta v_{im} - \theta\gamma\Delta e^{\text{prior}}_{im}. \tag{6}$$

The Attention treatment does not reveal the emissions of the decision meals but makes emissions fully salient, setting $\theta = 1$:

$$WTP^{\text{A}}_{im} = \Delta v_{im} - \gamma\Delta e^{\text{prior}}_{im}. \tag{7}$$

The Attention+Label treatment both makes emissions fully salient and provides information about true emissions. Using the updating rule in Equation 5,

$$WTP^{\text{A+L}}_{im} = \Delta v_{im} - \gamma\Big[(1-\kappa)\Delta e^{\text{true}}_{m} + \kappa\Delta e^{\text{prior}}_{im}\Big]. \tag{8}$$

Finally, in the baseline specification I assume that the Attention+Offset treatment removes the perceived environmental consequences of either option. The emissions term therefore drops out of perceived utility:

$$WTP^{\text{A+O}}_{im} = \Delta v_{im}. \tag{9}$$

I estimate the model from three within-participant, within-meal WTP contrasts: $(\text{A}+\text{L})-\text{B}$, $\text{A}-\text{B}$, and $(\text{A}+\text{O})-(\text{A}+\text{L})$. Elicited WTP in treatment state s is

$$\widetilde{WTP}^{s}_{im} = WTP^{s}_{im} + \eta^{s}_{im}, \qquad s \in \{\text{B}, \text{A}, \text{A}+\text{L}, \text{A}+\text{O}\},$$

where η^{s}_{im} captures elicitation noise, transitory valuation changes, and other deviations from the common-parameter approximation. Differencing within participant and meal

removes the non-emissions valuation Δv_{im}, while participant-specific prior emissions beliefs remain in the estimating equations.

I estimate the resulting system by generalized method of moments (GMM). Identification uses the variation in prior and true emissions differences that enters the model: for each of the three WTP contrasts, I require the remaining model residual to be orthogonal to these two emissions differences, yielding six moment conditions in total. I do not impose a parametric distribution on the residuals. The second-step weighting matrix is constructed from participant-level sums of the moment contributions to account for dependence among observations from the same participant. The estimator imposes $\gamma \geq 0$ and $\theta, \kappa \in [0, 1]$.

Online Appendix OA.2.2 reports the estimating equations and further implementation details. To reduce the influence of extreme belief errors, I exclude, separately for each meal, the 10% of observations with the largest misperception of the emissions difference between meal and cheese sandwich. I report percentile confidence intervals from 999 bootstrap replications that resample participants and re-estimate the system.

6.3 Extension to consumer welfare

For the welfare analysis, I distinguish the perceived utility that governs choice from the *true utility* used to evaluate the resulting choice.[62] I define the consumer's *true utility* from meal m as

$$u_i^{\text{true}}(m) \;=\; v_{im} - p_m - \gamma e_m^{\text{true}}. \tag{10}$$

Thus, the welfare criterion evaluates the environmental consequence of a meal using true emissions and the full weight γ, rather than the consumer's potentially misperceived or inattentively weighted emissions beliefs.[63]

Let $P \in \{0, 1\}$ indicate whether labels are present, and let $m_i(P)$ denote the meal participant i chooses under label-presence state P. I decompose the utility effect of label presence as

$$U_i(P \;=\; 1) - U_i(P \;=\; 0) \;=\; u_i^{\text{true}}\big(m_i(1)\big) - u_i^{\text{true}}\big(m_i(0)\big) + F. \tag{11}$$

62. This analysis concerns consumer welfare. A comprehensive social-welfare assessment would additionally need to account for the environmental benefits of reduced emissions and the costs of implementing the interventions.

63. This welfare criterion treats misperceptions and inattention as affecting decisions but not the environmental component of true utility. One interpretation is that consumers may experience ex-post regret once the emissions consequences of their consumption become known.

The difference between the first two terms is the welfare effect mediated by a change in the meal selected in the presence of labels. F captures any additional utility effect of label presence that is not mediated by such a change.

6.4 Identification of the welfare component

I use participants' WTP for the presence of carbon labels to identify F. The label-presence elicitation itself introduces participants to carbon emissions and the possibility of seeing carbon labels. I therefore take the relevant no-label counterfactual to be the ATTENTION state: If labels are absent in the choices following the label-presence elicitation, the participant is attentive to emissions but relies on her prior beliefs and chooses m_i^A; if labels are present, she chooses m_i^L. I assume that elicited WTP for label presence reflects the participant's subjective expectation of the corresponding true-utility difference, plus the additional component F:[64]

$$\widetilde{WTP}_i^P = \mathbb{E}_i\left[u_i^{\text{true}}(m_i^L) - u_i^{\text{true}}(m_i^A)\right] + F + \eta_i^P. \tag{12}$$

Here, $\mathbb{E}_i$ denotes participant i's subjective expectation at the time of elicitation. Because participants have not yet seen the three final meals or the BDM draws that will determine the implemented choices, this expectation is over the future meal-choice problems they may face. I approximate this expectation using the model-implied choice benefit from labels for meals participants encountered earlier in the experiment. For each meal, I calculate the gain as a function of the possible BDM price, average it over the BDM price distribution, and then average the resulting meal-specific expected gains across meals within participants. I estimate F as the average difference between participants' WTP for label presence and this predicted benefit.[65] Online Appendix OA.2.2.4 provides details. If participants' subjective expectations are well represented by this approximation, F can be interpreted as a direct psychological or informational utility effect of label

64. Allowing label presence to affect welfare both through its expected choice-mediated benefit and through an additional psychological or informational component is less restrictive than setting the latter component to zero. Related welfare analyses similarly allow interventions to affect welfare through channels beyond the behavioral outcomes they change. Butera et al. (2022), for example, allow public recognition to generate image utility conditional on performance, while Allcott and Kessler (2019) use recipients' valuations of home-energy reports to account for costs and benefits not captured by energy savings. Here, F captures the corresponding residual value of label presence beyond its expected effect on meal choice.

65. I apply the same trimming rule concerning extreme belief errors as in the main model estimation.

presence that is not mediated by a change in the selected meal. More generally, the estimated F is the residual value of label presence after subtracting the model-implied choice-mediated component and can also absorb error in this empirical approximation.

6.5 Results

Estimated model parameters. I first estimate the basic model based on choice data alone, i.e. excluding data on WTP for label presence later used to estimate F. The main estimator places θ at its lower bound of zero (95% participant-bootstrap interval: [0.000, 0.273]), implying that emissions receive no weight in baseline choices at the point estimate. I estimate $\kappa = 0.162$ [0.000, 0.550], so the point estimate implies substantial updating toward the emissions information, although the interval is wide. Environmental guilt is estimated at $\gamma = €0.101$ per kg CO_2e [0.078, 0.131]. Online Appendix Table OA.9 reports the basic-model estimates and bootstrap intervals; Table OA.10 compares fit under alternative assumptions about baseline attention and belief updating; and Table OA.11 reports sensitivity to relaxing the parameter restrictions.

Comparison of interventions: emissions reduction. While section 4 provides reduced-form evidence on the channels driving the labels' effects, the stronger assumptions underlying the structural model allow me to quantify the relative importance of each channel. For this purpose, I simulate how Experiment 3 participants would choose under different interventions using their baseline meal preferences, prior emissions beliefs, and the estimated parameters θ, γ, and κ. I use the same simultaneous four-meal-plus-sandwich choice environment and canteen prices as in the main Experiment 1 benchmark.[66] The interventions I consider are:

(1) Knowledge intervention: Providing emissions information without increasing attention. Note that this is a fictional intervention for the purpose of decomposing the treatment effects of the labels. In practice, it is not possible to provide information without increasing attention.

(2) Attention intervention: Increasing attention without providing emissions information.

(3) Label intervention: Combining both knowledge and attention effects.

66. Details are shown in Online Appendix OA.2.2.5.

(4) Carbon tax: Applying the €155/t midpoint of the Experiment 1 benchmark and the tax that most closely matches the label-induced emissions reduction in the Experiment 3 simulation, with lump-sum redistribution.

(5) Meat ban: Eliminating meat options from the menu.

Table 3 presents the simulation results. Because the main estimate places baseline salience θ at zero, the pure knowledge intervention has no effect at the point estimate. In the baseline structural simulation, the model predicts that the attention intervention reduces emissions by 54.4 grams per meal, equivalent to approximately 73% of the 74.1-gram reduction predicted for the label intervention.[67] The remaining 27% reflects the additional effect of emissions information once attention has been increased. The participant-bootstrap interval for the attention share is wide, ranging from approximately 28.4% to 170.4%. Thus, attention appears quantitatively important, while the exact channel allocation is imprecisely estimated (Online Appendix OA.2.2).

The baseline estimation assumes that participants do not consider meal emissions in the Attention+Offset treatment. Supplementary Material S.4.5 provides reduced-form evidence indicating that the offsetting treatment indeed substantially reduces environmental guilt, although it may not eliminate it entirely. Supplementary Material Section S.3.8.2 thus repeats the structural estimation assuming that only 70% of guilt is removed. The point estimates of γ, θ, and κ are all higher under this specification. Point predictions for the Attention, Label, and meat-ban interventions remain essentially unchanged, while the knowledge-only prediction changes because estimated baseline salience moves away from zero.

Comparison of interventions: consumer welfare. The welfare columns of Table 3 report the *choice-mediated* component of the welfare decomposition in Equation 11; they do not use participants' WTP for label presence or estimate F. Carbon labels improve choice-mediated consumer welfare by an average of 0.32¢ per meal, compared with 0.22¢ under a €155/t tax and 0.32¢ under a tax at the Experiment 3 emissions-matching tax rate of €106/t. The choice-mediated welfare changes generated by labels and the two tax comparisons are therefore similar at the point estimates, while the

67. Online Appendix Table OA.12 compares model-predicted WTP changes with reduced-form estimates from the experimental data. The model reproduces the average Attention responses reasonably closely, but predicts a smaller reduction in WTP for higher-emission meals under Attention+Label than the reduced-form estimate (12 cents versus 20 cents). These reduced-form mean responses are not directly targeted by the GMM moment conditions.

Table 3. Simulated effects of canteen interventions

Intervention	(1) Δ GHGE (g)	(2) [95% CI]	(3) Δ welfare (cents)	(4) [95% CI]
Knowledge	0.0	[0.0, 38.0]	0.00	[−0.06, 24.79]
Attention	−54.4	[−94.1, −13.7]	−0.06	[−0.74, 0.33]
Labels	−74.1	[−120.1, −20.5]	0.32	[−0.18, 18.85]
Tax, €155/t	−108.5	[−159.7, −61.2]	0.22	[−0.16, 0.62]
Tax, €106/t (GHGE match)	−74.6	[−115.9, −34.1]	0.32	[0.01, 0.65]
Meat ban	−215.4	[−277.2, −156.3]	−3.98	[−5.91, −2.32]

Notes: Model-implied changes in greenhouse gas emissions and choice-mediated welfare resulting from different interventions. The simulation includes the 430 of 444 Experiment 3 participants who retain at least one estimation observation in the structural parameter estimation, but restores all pre-trimming baseline WTP records for those participants so that they face the complete menu. Non-vegetarians choose among all four warm meals and the cheese sandwich; vegetarians choose among the two vegetarian meals and sandwich. Every participant receives equal weight. Intervals use 999 participant-level bootstrap replications that re-estimate the structural system and reapply equal participant weights within each draw. The GHGE-matching rate minimizes the absolute point-sample difference between tax- and label-induced emissions changes using the same continuous threshold-search procedure as Experiment 1, but based on Experiment 3 data and participants' simulated reactions to carbon labels from the structural model. Rates of €102.3–109.5/t yield the same closest attainable change, so the table reports their midpoint, €106/t. For both tax rows, each bootstrap replication recalculates outcomes at the tax rate reported in the table; the tax rates themselves are not re-estimated within bootstrap replications. Welfare excludes F, the residual value of label presence.

meat ban decreases welfare. The welfare bootstrap distributions in Column (4) have wide upper tails. These are largely attributable to participants with unusually inaccurate prior emissions beliefs. Online Appendix Table OA.13 repeats the simulations using only participants for whom none of the relevant meals is excluded by the meal-specific 10% trimming rule used in the structural estimation; the resulting welfare intervals are substantially narrower.

Separately, among Experiment 3 participants who previously experienced labels, I estimate the residual value F at €0.214 (95% participant-bootstrap interval: [0.170, 0.264]; Online Appendix Table OA.16). This component is not included in Table 3; accounting for it therefore increases the estimated welfare advantage of labels relative to the tax.

As a robustness check, I relax the assumption that the Attention and Label treatments purely increase salience and correct misperceptions without imposing psychological costs. Instead, I allow the interventions to impose an additional psychological cost

based on perceived emissions, doubling the weight on emissions in the choice equations. Under this assumption, the estimated welfare change under labels excluding *F* is -4.51 cents per choice (Supplementary Material Section S.3.8.1). The separately estimated residual value is 21.02 cents, resulting in a total welfare increase of 16.51 cents. I further assess the model's sensitivity to its parameter restrictions and outlying emissions beliefs in Online Appendix Section OA.2.2, to observed parameter heterogeneity in Supplementary Material Section S.3.6, and to alternative treatment assumptions in Supplementary Material Section S.3.8.

7 Discussion

This paper demonstrates how benchmarking behavioral climate interventions against a carbon tax can make their effectiveness more interpretable and provides the first such benchmark for carbon labels. Experiment 1 identifies a local carbon-tax equivalent of about €155/t CO_2e for the canteen meals studied, while Experiment 2 obtains a similar estimate from ordinary, recurring canteen purchases using a different empirical approach. Experiment 3 shows that solely increasing the salience of emissions considerably changes choices and that correcting misperceptions alone cannot explain the full label effect, suggesting that a substantial part of the labels' effectiveness is driven by salience. The structural model quantifies the relative contributions of these channels, attributing approximately 73% of the effect to salience at the point estimate. Under the model's baseline assumptions, the changes in meal choices induced by labels and by the local tax generate similar changes in consumer welfare. Participants exhibit an incentivized WTP of approximately €0.20 per meal for label presence, most of which the model attributes to a value of label presence that is not mediated by improvements in meal choice.

Carbon taxation remains the most effective instrument for reducing emissions, and its implementation is essential for meeting long-run climate targets (European Association of Environmental and Resource Economists, 2019). Yet political resistance makes food taxation difficult to implement, even at modest levels.[68] While carbon labels may provide a politically feasible alternative, their effectiveness is often evaluated based solely on the reduction in emissions they produce, an approach which can be misleading. An emissions reduction of a few percentage points may indeed reflect an ineffective

68. See Dechezleprêtre et al. (2022), Douenne and Fabre (2020), and Fesenfeld (2023) for evidence on the unpopularity of agricultural taxation policies.

intervention, but it may also reflect that the intervention is tested in a context in which behavior is difficult to change. In my study setting, the substantial €155/t benchmark supports the latter interpretation: the labels generate a meaningful demand response in a domain in which consumption patterns are difficult to shift (also see Guthrie, Mancino, and Lin, 2015, on the difficulty of changing food choices). While labels will not suffice as a stand-alone response to the climate crisis,[69] they can still deliver a meaningful reduction within the range of politically viable interventions. Based on the evidence, one might also be tempted to conclude that policy and research should primarily focus on settings where behavior is easier to shift. However, the food sector contributes an estimated 26% to 34% of global greenhouse gas emissions (Crippa et al., 2021; Poore and Nemecek, 2018); not targeting food consumption would make the $2°C$ target difficult to impossible to reach.

The preferred point estimates are approximately €155/t in Experiment 1 and €133/t in Experiment 2, but the benchmark is estimated with uncertainty and is specific to the local Bonn canteen context. The exact midpoint should therefore be transferred to other contexts cautiously. However, emissions reductions of a similar order in other student canteens (Brunner et al., 2018; Lohmann et al., 2022), together with price-response estimates close to those in related work (Roosen, Staudigel, and Rahbauer, 2022), suggest that neither component is unusual in comparable institutional settings. The Bonn estimate may therefore provide a useful reference point for workplace and other institutional canteens.

To provide insights on less closely related settings, the Experiment 1 benchmarking method can be easily ported to other contexts. Supplementary Material S.3 provides guidance as well as exploratory heterogeneity analyses providing suggestive evidence on the role of demographic characteristics and environmental attitudes in shaping consumer responses to the labels. Choice sets and substitution possibilities also shape the absolute responses to both labels and prices. Importantly, however, estimating both responses within the same setting holds these common features of the choice environment fixed. Comparing locally estimated benchmarks across contexts can therefore help as-

69. Estimates for the social cost of carbon range from 50 USD (€49) per tonne of CO_2e and lower (some scenarios in Barrage and Nordhaus, 2024), through €160 per tonne (e.g. Rennert et al., 2022), to substantially higher estimates (Bilal and Känzig, 2024; Moore et al., 2024). My €155/t benchmark is estimated with considerable uncertainty (95% percentile interval for its midpoint: €46.3–€242.6/t).

sess whether the relative effectiveness of labels varies systematically with population characteristics.

Experiment 3 shows that intervention effectiveness likely hinges on the labels attracting consumers' attention at the point of choice. The relevance of the salience mechanism in driving labels' effectiveness is likely to extend beyond food whenever emissions are not already prominent at the point of choice. Although its quantitative importance will depend on the setting, there is no obvious reason to expect the underlying mechanism to be specific to food choices. Importantly, taxes and labels work through different margins: taxes change relative prices, while labels change information and attention. They may therefore generate similar aggregate emissions reductions while affecting different consumers and producing different distributional and welfare consequences.

Institutional context is also likely to matter for policy implementation. Canteens are particularly well suited to the introduction and study of carbon information. They serve large populations, centralize menu planning, and allow carbon information to be introduced incrementally by individual providers. In Bonn, the canteens permanently introduced carbon labels in early 2025 by linking emissions values from the Eaternity Institute (2020) database to their internal menu systems (Studierendenwerk Bonn, 2025, and author's conversations with the canteens). Evidence over longer horizons and after canteen or producer adjustments would further inform large-scale implementation, as discussed in Supplementary Material S.3.

More broadly, the results illustrate the value of evaluating interventions relative to realistic policy alternatives. Without a benchmark, the same percentage change may appear small to one reader and meaningful to another. A comparison within the same population and choice setting gives the effect an economic reference point and helps distinguish the intervention-specific response from the general responsiveness of demand. The same principle can be applied to other behavioral climate interventions and other policy domains.

References

Allcott, H. 2011. "Social norms and energy conservation." *Journal of Public Economics* 95 (9-10): 1082–95.

Allcott, H., D. Cohen, W. Morrison, and D. Taubinsky. 2025. "When do "nudges" increase welfare?" *American Economic Review* 115 (11): 3448–93.

Allcott, H., and J. B. Kessler. 2019. "The welfare effects of nudges: A case study of energy use social comparisons." *American Economic Journal: Applied Economics* 11 (1): 236–76.

Allcott, H., and D. Taubinsky. 2015. "Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market." *American Economic Review* 105 (8): 2501–38.
Andor, M. A., L. Goette, M. K. Price, A. Schulze-Tilling, and L. Tomberg. 2023. "Differences in how and why Social Comparison and Real-Time Feedback impact resource use: Evidence from a field experiment." Working Paper 31845. National Bureau of Economic Research.
Banerjee, S., M. M. Galizzi, P. John, and S. Mourato. 2023. "Sustainable dietary choices improved by reflection before a nudge in an online experiment." *Nature Sustainability* 6 (12): 1632–42.
Barrage, L., and W. Nordhaus. 2024. "Policies, projections, and the social cost of carbon: Results from the DICE-2023 model." *Proceedings of the National Academy of Sciences* 121 (13): e2312030121.
Bilal, A., and D. R. Känzig. 2024. "The Macroeconomic Impact of Climate Change: Global vs. Local Temperature." Working Paper 32450. National Bureau of Economic Research.
Bilén, D. 2022. "Do carbon labels cause consumers to reduce their emissions? Evidence from a large-scale natural experiment." Working Paper. University of Gothenburg.
Bordalo, P., N. Gennaioli, and A. Shleifer. 2022. "Salience." *Annual Review of Economics* 14 (1): 521–44.
Bouchaud, J.-P., P. Krueger, A. Landier, and D. Thesmar. 2019. "Sticky expectations and the profitability anomaly." *Journal of Finance* 74 (2): 639–74.
Brent, D. A., J. H. Cook, and S. Olsen. 2015. "Social comparisons, household water use, and participation in utility conservation programs: Evidence from three randomized trials." *Journal of the Association of Environmental and Resource Economists* 2 (4): 597–627.
Brunner, F., V. Kurz, D. Bryngelsson, and F. Hedenus. 2018. "Carbon label at a university restaurant–label implementation and evaluation." *Ecological Economics* 146: 658–67.
Busse, M. R., N. Lacetera, D. G. Pope, J. Silva-Risso, and J. R. Sydnor. 2013. "Estimating the effect of salience in wholesale and retail car markets." *American Economic Review* 103 (3): 575–79.
Butera, L., R. Metcalfe, W. Morrison, and D. Taubinsky. 2022. "Measuring the welfare effects of shame and pride." *American Economic Review* 112 (1): 122–68.
Byrne, D. P., L. Goette, L. A. Martin, A. Miles, A. Jones, S. Schob, T. Staake, and V. Tiefenbeck. 2024. "How nudges create habits: Theory and evidence from a field experiment." Working Paper 3974371. SSRN.
Chater, N., and G. Loewenstein. 2023. "The i-frame and the s-frame: How focusing on individual-level solutions has led behavioral public policy astray." *Behavioral and Brain Sciences* 46: e147.
Chen, D. L., M. Schonger, and C. Wickens. 2016. "oTree—An open-source platform for laboratory, online, and field experiments." *Journal of Behavioral and Experimental Finance* 9 (C): 88–97.
Chetty, R. 2009. "Sufficient statistics for welfare analysis: A bridge between structural and reduced-form methods." *Annual Review of Economics* 1 (1): 451–88.
Chetty, R., A. Looney, and K. Kroft. 2009. "Salience and taxation: Theory and evidence." *American Economic Review* 99 (4): 1145–77.
Crippa, M., E. Solazzo, D. Guizzardi, F. Monforti-Ferrario, F. Tubiello, and A. Leip. 2021. "Food systems are responsible for a third of global anthropogenic GHG emissions." *Nature Food* 2 (3): 198–209.
Dechezleprêtre, A., A. Fabre, T. Kruse, B. Planterose, A. Sanchez Chico, and S. Stantcheva. 2022. "Fighting climate change: International attitudes toward climate policies." Working Paper 30265. National Bureau of Economic Research.
DellaVigna, S. 2009. "Psychology and economics: Evidence from the field." *Journal of Economic Literature* 47 (2): 315–72.
Dohmen, T., and T. Jagelka. 2024. "Accounting for Individual-Specific Reliability of Self-Assessed Measures of Economic Preferences and Personality Traits." *Journal of Political Economy Microeconomics* 2 (3): 399–462.

Douenne, T., and A. Fabre. 2020. "French attitudes on climate change, carbon taxation and other climate policies." *Ecological Economics* 169: 106496.
Eaternity Institute. 2020. *Eaternity database.*
Epstein, L. G., J. Noor, and A. Sandroni. 2008. "Non-Bayesian updating: a theoretical framework." *Theoretical Economics* 3 (2): 193–229.
European Association of Environmental and Resource Economists. 2019. *Economists' Statement on Carbon Pricing.*
European Commission. 2026. *Price of CBAM certificates.* Directorate-General for Taxation and Customs Union.
European Energy Exchange. 2026. *EEX successfully conducts first nEHS auction.* Press release.
Federal Ministry of Education and Research (Germany). 2023. *The student survey in Germany: 22nd Social Survey - Executive Summary.*
Federal Statistical Office (Germany). 2023. *Hochschulen - Hochschulstatistik.*
Fesenfeld, L. 2023. "The political economy of taxing meat." *Nature Food* 4 (3): 209–10.
German Environment Agency (UBA). 2024. *New record revenue in emissions trading: More than 18 billion euros for climate protection.*
Grummon, A. H., C. J. Lee, T. N. Robinson, E. B. Rimm, and D. Rose. 2023. "Simple dietary substitutions can reduce carbon footprints and improve dietary quality across diverse segments of the US population." *Nature Food* 4 (11): 966–77.
Guthrie, J., L. Mancino, and C.-T. J. Lin. 2015. "Nudging consumers toward better food choices: Policy approaches to changing food consumption behaviors." *Psychology & Marketing* 32 (5): 501–11.
Hahn, R. W., N. Hendren, R. D. Metcalfe, and B. Sprung-Keyser. 2024. "A welfare analysis of policies impacting climate change." Working Paper. National Bureau of Economic Research.
Ho, L., and L. Page. 2025. "Got Beef with Beef? Evidence from a Large-Scale Carbon Labeling Experiment." Working paper.
Imai, T., D. D. Pace, P. Schwardmann, and J. J. van der Weele. 2026. "Correcting consumer misperceptions about CO2 emissions." *Journal of the Association of Environmental and Resource Economists* 13 (5): 1399–435.
International Carbon Action Partnership (ICAP). 2025. *EU Emissions Trading System (EU ETS).*
Jalil, A. J., J. Tasoff, and A. V. Bustamante. 2023. "Low-cost climate-change informational intervention reduces meat consumption among students for 3 years." *Nature Food* 4 (3): 218–22.
Kim, B. F., R. E. Santo, A. P. Scatterday, J. P. Fry, C. M. Synk, S. R. Cebron, M. M. Mekonnen, A. Y. Hoekstra, S. De Pee, M. W. Bloem, et al. 2020. "Country-specific dietary shifts to mitigate climate and water crises." *Global Environmental Change* 62: 101926.
List, J. A., M. Rodemeier, S. Roy, and G. K. Sun. 2023. "Judging nudging: Understanding the welfare effects of nudges versus taxes." Working Paper 31152. National Bureau of Economic Research.
Liu, Z., E. Ansink, and J. Bouma. 2025. "Plant-based meat complements." Discussion Paper 2024-046/VIII. Tinbergen Institute. Revised April 2025.
Lohmann, P. M., E. Gsottbauer, A. Doherty, and A. Kontoleon. 2022. "Do carbon footprint labels promote climatarian diets? Evidence from a large-scale field experiment." *Journal of Environmental Economics and Management* 114: 102693.
Lohmann, P. M., E. Gsottbauer, J. Farrington, S. Human, and L. A. Reisch. 2024. "Choice architecture promotes sustainable choices in online food-delivery apps." *PNAS Nexus* 3 (10): pgae422.
Lohmann, P. M., A. Pizzo, J. M. Bauer, T. M. Khanna, M. Callaghan, J. Minx, and L. A. Reisch. 2024. "Demand-side interventions for sustainable food systems: a meta-analysis of food-policy interventions targeting food consumption and waste behaviours." Working Paper 4811931. SSRN.
Lusk, J. L., and G. T. Tonsor. 2016. "How meat demand elasticities vary with price, income, and product category." *Applied Economic Perspectives and Policy* 38 (4): 673–711.

Moore, F. C., M. A. Drupp, J. Rising, S. Dietz, I. Rudik, and G. Wagner. 2024. "Synthesis of evidence yields high social cost of carbon due to structural model variation and uncertainties." Working Paper. CESifo.
Nisa, C. F., J. J. Bélanger, B. M. Schumpe, and D. G. Faller. 2019. "Meta-analysis of randomised controlled trials testing behavioural interventions to promote household action on climate change." *Nature Communications* 10 (1): 4545.
Poore, J., and T. Nemecek. 2018. "Reducing food's environmental impacts through producers and consumers." *Science* 360 (6392): 987–92.
Potter, C., A. Bastounis, J. Hartmann-Boyce, C. Stewart, K. Frie, K. Tudor, F. Bianchi, E. Cartwright, B. Cook, M. Rayner, et al. 2021. "The effects of environmental sustainability labels on selection, purchase, and consumption of food and drink products: a systematic review." *Environment and Behavior* 53 (8): 891–925.
Rennert, K., F. Errickson, B. C. Prest, L. Rennels, R. G. Newell, W. Pizer, C. Kingdon, J. Wingenroth, R. Cooke, B. Parthum, et al. 2022. "Comprehensive evidence implies a higher social cost of CO2." *Nature* 610 (7933): 687–92.
Rodemeier, M., and A. Löschel. 2022. "Information nudges, subsidies, and crowding out of attention: Field evidence from energy efficiency investments." Working Paper 4213071. SSRN.
Roosen, J., M. Staudigel, and S. Rahbauer. 2022. "Demand elasticities for fresh meat and welfare effects of meat taxes in Germany." *Food Policy* 106: 102194.
Scarborough, P., M. Clark, L. Cobiac, K. Papier, A. Knuppel, J. Lynch, R. Harrington, T. Key, and M. Springmann. 2023. "Vegans, vegetarians, fish-eaters and meat-eaters in the UK show discrepant environmental impacts." *Nature Food* 4 (7): 565–74.
Schulze Tilling, A. 2021a. *Additional data collection to: Quantifying the role of greenhouse gas emissions in consumption choice.*
Schulze Tilling, A. 2021b. *Quantifying the role of greenhouse gas emissions in consumption choice.*
Studierendenwerk Bonn. 2025. *CO_2-Ampel in den Mensen.*
Taubinsky, D., and A. Rees-Jones. 2018. "Attention variation and welfare: theory and evidence from a tax salience experiment." *Review of Economic Studies* 85 (4): 2462–96.
Taufique, K. M., K. S. Nielsen, T. Dietz, R. Shwom, P. C. Stern, and M. P. Vandenbergh. 2022. "Revisiting the promise of carbon labelling." *Nature Climate Change* 12 (2): 132–40.
Thunström, L. 2019. "Welfare effects of nudges: The emotional tax of calorie menu labeling." *Judgment and Decision Making* 14 (1): 11.
Tiefenbeck, V., L. Goette, K. Degen, V. Tasic, E. Fleisch, R. Lalive, and T. Staake. 2018. "Overcoming salience bias: How real-time feedback fosters resource conservation." *Management Science* 64 (3): 1458–76.
Wettengel, J. 2024. *Germany's carbon pricing system for transport and buildings.*

Online Appendix

Understanding the effectiveness of carbon labels

Contents

OA.1 Experiment 1: Benchmark robustness OA-3

OA.1.1 Experiment 1: Main treatment-effect estimates OA-3

OA.1.2 Experiment 1: Menu-simulation details OA-3

OA.1.3 Experiment 1: Additional menu-simulation results OA-6

OA.1.4 Experiment 1: Alternative continuous-emissions benchmark OA-7

OA.2 Experiment 3: Mechanism and structural robustness OA-10

OA.2.1 Experiment 3: Results using alternative definitions OA-10

OA.2.2 Details of structural estimation and simulations OA-11

OA.2.2.1 Basic-model estimating equations OA-11

OA.2.2.2 Basic-model results and robustness OA-12

OA.2.2.3 Observed-versus-predicted WTP fit OA-14

OA.2.2.4 Welfare-extension estimating equation OA-15

OA.2.2.5 Policy simulations OA-17

OA.2.2.6 Sensitivity to outlying prior emissions beliefs OA-19

OA.2.2.7 Observed parameter heterogeneity OA-19

OA.2.2.8 Preferences for label presence OA-20

OA.2.2.9 Welfare-extension results and robustness OA-22

OA.2.2.10 Additional structural robustness OA-23

OA.3 Experiment 2: Field robustness OA-23

OA.3.1 Additional analyses on labeling effects OA-23

OA.3.2 Additional price-effect robustness and field-implied tax-equivalent rates OA-29

Table OA.1. Overview of the experiments and observational canteen data

Study or data source	Description
Experiment 1	*Sample:* 287 participants (October–November 2021). *Design:* Carbon labels versus repeated choices without labels. *Outcomes and analyses:* Within-participant label-induced changes in relative WTP and simulated meal choices with labels and taxes. *Role:* Establishes the local carbon-tax-equivalent benchmark.
Experiment 2	*Sample:* 124,784 transactions; ITT subsample of 1,918 regular guests (April–July 2022). *Design:* Treated canteen versus two control canteens before, during, and after labeling. *Outcomes and analyses:* Estimated label effects on meat-meal choice and GHGE. *Role:* Tests label effectiveness in ordinary, repeated purchases.
Observational canteen data	*Sample:* 465,000 student purchases (April 2022–July 2023). *Design:* Naturally occurring daily meat–vegetarian price differences. *Outcomes and analyses:* Meat-meal choice responses to relative-price differences and predicted effects of local carbon taxes. *Role:* Benchmarks the field label effect against emissions-based relative-price adjustments.
Experiment 3	*Sample:* 444 participants (June–July 2021). *Design:* Attention, Attention+Label, and Attention+Offset. *Outcomes and analyses:* Within-participant changes in relative WTP under attention and labels, prior emissions beliefs, and WTP for label presence. *Role:* Identifies behavioral channels and informs the structural welfare comparison.

Notes: The experiments are numbered by their role in the paper's argument rather than by chronological order: Experiment 1 establishes the benchmark, Experiment 2 examines it in natural field data, and Experiment 3 studies mechanisms and welfare. WTP denotes willingness to pay; ITT denotes intention to treat; GHGE denotes greenhouse gas emissions.

OA.1 Experiment 1: Benchmark robustness

OA.1.1 Experiment 1: Main treatment-effect estimates

Table OA.2. Within-subject change in WTP for meals

	Change in WTP compared to baseline	
	(1) All rounds	(2) Excluding round 3
High emission meal × Shown label	−0.31***	−0.34***
	(0.05)	(0.06)
Low emission meal × Shown label	0.14***	0.15**
	(0.04)	(0.06)
High emission meal	0.01	0.02
	(0.02)	(0.02)
Low emission meal	−0.06*	−0.05
	(0.03)	(0.03)
Control for third round	0.02	
	(0.03)	
Participants control	139	139
Participants treated	217	148
Observations	1,704	1,148

Notes: Column (1) is the main specification and uses all second- and third-round observations. Column (2) repeats the specification after excluding third-round observations. The third-round indicator is therefore included only in Column (1). Standard errors are clustered at the individual level.

OA.1.2 Experiment 1: Menu-simulation details

This subsection formalizes the menu simulations underlying Table 1. Let $d_i \in \{N, V\}$ denote participant i's dietary group (non-vegetarian or vegetarian). In the main specification in which meals are offered simultaneously, non-vegetarians choose from $\mathscr{K}_N = \{0, 1, 2, 3, 4\}$ and vegetarians from $\mathscr{K}_V = \{0, 1, 3\}$, where option 0 is the cheese sandwich, meals 1 and 3 are vegetarian, and meals 2 and 4 contain meat.

For each warm meal k, let w_{ik}^0 and w_{ik}^L denote participant i's baseline and label WTP relative to the sandwich, and let $p_k - p_0$ and $e_k - e_0$ denote the corresponding price and emissions differences. Normalizing the sandwich's relative net valuation to zero, the relative net valuation of warm meal k is

$$v_{ik}^{0} = w_{ik}^{0} - (p_k - p_0), \tag{OA.1}$$

$$v_{ik}^{L} = w_{ik}^{L} - (p_k - p_0), \tag{OA.2}$$

$$v_{ik}^{T}(\tau) = w_{ik}^{0} - (p_k - p_0) - \tau(e_k - e_0), \tag{OA.3}$$

where τ is the carbon-tax rate in euros per kg CO_2e. For state $s \in \{0, L, T(\tau)\}$,

$$c_i^s = \arg\max_{k \in \mathcal{K}_{d_i}} \{v_{ik}^s\}, \qquad v_{i0}^s = 0. \tag{OA.4}$$

where v_{i0}^s denotes the relative net valuation of the outside option, the sandwich. Exact ties are resolved using the same fixed deterministic ordering in every state: a warm meal is selected over the sandwich when both have zero relative net valuation, while the warm meal with the lower numerical identifier is selected when multiple warm meals share the highest net valuation.[70]

Giving each participant equal weight, mean emissions are

$$E^s = \frac{1}{N} \sum_{i=1}^{N} e_{c_i^s}. \tag{OA.5}$$

The label- and tax-induced emissions changes are therefore $\Delta E^L = E^L - E^0$ and $\Delta E^T(\tau) = E^{T(\tau)} - E^0$.

Finally, I identify all tax rates whose simulated emissions change is closest to the label-induced change:

$$\mathcal{T}^* = \arg\min_{\tau \geq 0} \left| \Delta E^T(\tau) - \Delta E^L \right|. \tag{OA.6}$$

Because choices change only when a participant crosses a participant-specific choice threshold, I calculate these thresholds exactly and evaluate emissions on the intervals between them. Consequently, $\mathcal{T}^*$ can itself be an interval. In the main simultaneous-menu specification it is €145.8–€164.3/t, with midpoint €155/t.

For inference on the simultaneous-menu benchmark in Table 1 Column (1), I repeat the complete simulation in each participant bootstrap replication, resampling separately within the randomized treatment-sequence strata represented in the treated sample.

70. In the 217-participant simultaneous-menu sample, the sandwich–warm-meal rule is invoked for five label-state choices and no baseline choices. Ties for the highest net valuation among two or more warm meals are resolved by meal number for five baseline choices and four label-state choices. These counts refer to the observed baseline and label states; ties under a carbon tax depend on the evaluated tax rate.

Table OA.3. Point estimate and bootstrap inference in the simultaneous choice set

Outcome	Unit	Estimate	Bootstrap standard error	95% interval
Baseline emissions	g	943.4	48.0	[855.6, 1046.3]
Labels: Δ emissions	g	-113.5	37.9	[-188.8, -40.5]
Labels: Δ emissions	%	-12.0	3.7	[-18.9, -4.6]
€155/t: Δ emissions	g	-101.4	31.0	[-167.6, -46.0]
Matching-set midpoint	€/t	155.0	50.1	[46.3, 242.6]
Emissions-matching range	€/t	[146, 164]	[52.1, 50.0]	[24.4, 258.6]

Notes: This table reports Column (1) of Table 1. Each of 999 participant-bootstrap replications resamples the 169 non-vegetarian and 48 vegetarian participants within randomized treatment-sequence strata and repeats the simulation; participants receive equal weight. Tax-effect inference evaluates €155/t. The midpoint row reports its bootstrap standard error and percentile 95% interval. For the matching range, the two standard errors refer to the lower and upper endpoints; the 95% interval is the outer envelope formed from the lower endpoint's 2.5th percentile and the upper endpoint's 97.5th percentile. Results are unadjusted; Table OA.5 reports the control-adjusted benchmark.

Columns (2) and (3) of Table 1 replace the simultaneous menu with offers that more closely resemble the canteen's daily menu structure. Each simulated menu contains one vegetarian and one meat main. I use the four experimental meals as proxies for four broad menu categories: meal 1, Italian vegetable ragout with pasta (approximately 0.5 kg CO_2e), represents a low-emission vegetarian main; meal 2, poultry schnitzel with fruity sauce and rice (approximately 1.4 kg CO_2e), represents a poultry main; meal 3, filled zucchini with potato croquettes (approximately 1.4 kg CO_2e), represents a higher-emission vegetarian main; and meal 4, sliced beef with potatoes (approximately 3.4 kg CO_2e), represents a beef main. These four meals generate four possible vegetarian–meat menu types: $(1,2)$, $(3,2)$, $(1,4)$, and $(3,4)$. Based on actual canteen prices at the time of the experiment, meal 1 is priced €1.25 above the cheese sandwich, while the other three meals are priced €1.55 above the sandwich.

I weight these menu types according to how often the corresponding offers occurred during the seven-week labeling period in Experiment 2. A vegetarian main is classified as low-emission if its emissions are below those of the 0.7 kg CO_2e cheese sandwich and as higher-emission otherwise. Of the 96 canteen-days with complete offer information, 44 feature either a poultry or beef main and can therefore be represented by the Experiment 1 meat meals. Among these 44 days, 33 pair poultry with a low-emission vegetarian main, none pair poultry with a higher-emission vegetarian main, eight pair beef with a low-emission vegetarian main, and three pair beef with a higher-emission

vegetarian main. The four menu types therefore receive weights of 75.0%, 0.0%, 18.2%, and 6.8%, respectively. Canteen-days featuring pork, fish, or another meat main are not used to construct these weights because Experiment 1 contains no corresponding meat meal.

For each dietary group, I first average simulated emissions across these four menu types using the above offer weights, giving each participant equal weight. I then combine the non-vegetarian and vegetarian averages using their shares in the Experiment 2 field sample, 78.0% and 22.0%, respectively. Column (2) retains the cheese sandwich as an outside option: non-vegetarians choose among the meat main, vegetarian main, and sandwich, while vegetarians choose between the vegetarian main and sandwich. Column (3) removes the sandwich, so non-vegetarians choose between the meat and vegetarian mains and vegetarians consume the vegetarian main.

OA.1.3 Experiment 1: Additional menu-simulation results

Table OA.4 repeats the simultaneous-menu simulation after removing the cheese-sandwich outside option.

Table OA.4. Point estimate and bootstrap inference in the simultaneous choice set without the outside option

Outcome	Unit	Estimate	Bootstrap standard error	95% interval
Baseline emissions	g	1130.1	58.6	[1013.8, 1242.3]
Labels: Δ emissions	g	-170.4	55.5	[-279.8, -62.5]
Labels: Δ emissions	%	-15.1	4.6	[-24.1, -5.9]
€113.5/t: Δ emissions	g	-169.0	39.4	[-247.4, -93.3]
Emissions-matching range	€/t	[113, 114]	[51.2, 51.5]	[56.9, 227.4]

Notes: This table repeats Table OA.3 after removing the cheese-sandwich outside option. Non-vegetarian participants choose among all four warm meals, while vegetarian participants choose between the two vegetarian meals. Bootstrap inference otherwise follows Table OA.3.

Table OA.5 reports the simultaneous menu simulation adjusted for possible effects of repeated WTP elicitation. I first estimate the average no-label change in WTP from the Control condition and subtract this predicted change from treated participants' observed WTP under labels. I additionally estimate a common third-round effect using both the Control and Label conditions and subtract this adjustment where applicable. The resulting label-induced emissions reduction is 128.8 grams (bootstrap SE 41.1).

The point-sample matching set is €164.3–170.6/t, with midpoint €167.4/t. Across 999 participant-bootstrap replications, the midpoint's percentile 95% interval is €46.3–€243.0/t. A conservative 95% outer bootstrap envelope for the complete matching set, formed from the 2.5th percentile of its lower endpoint and the 97.5th percentile of its upper endpoint, is €24.4–€258.6/t.

Table OA.5. Control-adjusted benchmark

Outcome	Unit	Estimate	Bootstrap standard error	95% interval
Baseline emissions	g	943.4	47.5	[849.5, 1044.9]
Labels: Δ emissions	g	-128.8	41.1	[-203.1, -39.2]
Labels: Δ emissions	%	-13.7	4.0	[-20.7, -4.3]
€155/t: Δ emissions	g	-101.4	31.3	[-169.9, -43.8]
Matching-set midpoint	€/t	167.4	51.2	[46.3, 243.0]
Emissions-matching range	€/t	[164, 171]	[52.8, 51.6]	[24.4, 258.6]

Notes: This table repeats Table OA.3 adjusting for possible effects of repeated WTP elicitation. Before simulating label choices, it deducts the estimated no-label coefficient for low-emission meals, −€0.055 (SE €0.036, $p = 0.129$); the corresponding coefficient for high-emission meals, €0.001 (SE €0.021, $p = 0.946$); and the third-round coefficient, €0.007 (SE €0.035, $p = 0.849$), where applicable. None of these coefficients is statistically significant at conventional levels. The adjustment regression and complete simulation are re-estimated in every participant-bootstrap replication. The midpoint and matching-range inference rows are defined as in Table OA.3.

OA.1.4 Experiment 1: Alternative continuous-emissions benchmark

An alternative way to summarize the WTP response is to regress participants' change in WTP for meal m directly on the difference in emissions between meal m and the sandwich. A carbon tax changes relative prices by the tax rate multiplied by this same emissions difference. The fitted label–emissions gradient can therefore be expressed in euros per tonne without selecting a particular menu and price configuration.

I estimate

$$\Delta WTP_{ijm} = \alpha + \beta_1 Emi_m + \delta_1 (Label_{ij} \times Emi_m) + \rho ThirdRound_j + \varepsilon_{ijm}, \qquad \text{(OA.7)}$$

where Emi_m is the emissions difference, in kg CO_2e, between meal m and the sandwich. The interaction coefficient δ_1 captures how the label response varies with this difference; multiplying $-\delta_1$ by 1,000 expresses the fitted gradient in euros per tonne of CO_2e. The above Equation OA.7 does not include a separate $Label_{ij}$ term. This constrains the

label-related component to be zero when a meal and the sandwich have identical emissions. This mirrors a pure emissions-based tax: if a meal and the sandwich have identical emissions, such a tax does not change their relative price. Column (2) of Table OA.6 relaxes this assumption and adds a separate $Label_{ij}$ term, allowing labels to shift WTP by a common amount as well as changing the emissions gradient.

Table OA.6. Within-subject WTP changes and the continuous-emissions benchmark

	Change in WTP compared to baseline	
	(1) No label-specific intercept	(2) With label-specific intercept
Emissions gap (kg CO_2e) × Shown label	−0.16***	−0.12***
	(0.03)	(0.03)
Shown label		−0.08**
		(0.03)
Emissions gap (kg CO_2e)	0.04***	0.02
	(0.01)	(0.01)
Control for third round	0.02	0.02
	(0.03)	(0.03)
Constant	−0.06***	−0.02
	(0.02)	(0.02)
Participants control	139	139
Participants treated	217	217
Observations	1,704	1,704

Notes: Both columns report continuous-emissions specifications and include the conventional regression constant shown in the table. Column (1) reports Equation OA.7, which has no label-specific intercept. Column (2) adds a Shown label indicator, allowing a common label-induced WTP shift in addition to the emissions gradient. The third-round coefficient is shown in both columns. Standard errors are clustered at the individual level. The Column (1) tax-equivalent benchmark is €156/t; its percentile 95% interval is €101–€215/t, based on 999 participant-bootstrap replications.

Table OA.6, Column (1), estimates $\delta_1 = -0.156$, implying €156/t CO_2e. A participant-level bootstrap gives a 95% interval of €101–€215/t. The point estimate lies within the €146–164/t simultaneous-choice-set benchmark in the main text, although the calculations answer slightly different questions. Column (2) shows how the fitted gradient changes when the restriction of a zero label effect at a zero emissions difference is relaxed. Adding a separate $Label_{ij}$ term allows the labels to generate a common shift in WTP for all warm meals relative to the sandwich, in addition to an emissions-related gradient. This is conceptually related to removing the sandwich as an outside option in the simulation (compare Table OA.4). When participants must choose among

warm meals, a common shift in the value of all warm meals relative to the sandwich does not affect which warm meal is selected; only differences across warm meals matter. Similarly, the label-specific intercept in Column (2) absorbs such a common shift, leaving the slope to capture how label responses vary with emissions differences. The estimated slope becomes less negative from −0.156 to −0.12, but remains broadly similar and lies within the 95% confidence interval of the baseline estimate.

Importantly, δ_1 should not be interpreted as the causal effect of a one-kilogram increase in relative emissions on the label response. Such an interpretation would require emissions differences across the experimental meals to be exogenous to other meal characteristics, whereas the meals differ along several dimensions. Instead, δ_1 provides a descriptive summary of how label-induced WTP changes vary with the emissions difference between the meal and the sandwich across the experimental meals. Expressing this gradient in euros per tonne gives a tax-equivalent measure: it corresponds to the emissions-based relative-price gradient generated by a carbon tax. Unlike the main benchmark, this calculation does not require specifying a particular menu or set of prices. For evidence on the behavioral channels underlying participants' responses to the labels, see Experiment 3.

OA.2 Experiment 3: Mechanism and structural robustness

OA.2.1 Experiment 3: Results using alternative definitions

Table OA.7. Under- and overestimation under alternative classifications

	Change in WTP compared to baseline			
	(1) Main: gap	(2) Meal emissions	(3) Gap, including exact	(4) Main: gap, no round 3
Underestimated × Shown label	−0.23***	−0.22***	−0.23***	−0.22***
	(0.04)	(0.04)	(0.04)	(0.04)
Overestimated × Shown label	−0.10***	−0.10***	−0.10***	−0.10**
	(0.04)	(0.04)	(0.04)	(0.04)
Correct estimate × Shown label			−0.19**	
			(0.08)	
Control for third round	0.05	0.05	0.05	
	(0.05)	(0.05)	(0.04)	
Participants	293	293	293	144
Obs. underestimate	555	646	555	269
Obs. overestimate	562	450	562	281
Obs. exact			55	
Observations	1,117	1,096	1,172	550

Notes: The analysis uses data from the Attention+Label condition. Column (1) is the main specification: it classifies under- and overestimation based on whether the emissions gap between the meal and sandwich was under- or overestimated. It thus combines participants' estimates of the meal and sandwich emissions. It drops 55 participant–meal combinations for which the estimated and actual emissions gaps are equal after rounding to the nearest 0.1 kg CO_2e. Column (2) instead classifies under- and overestimation based only on whether the respective meal's own emissions were under- or overestimated. It drops 76 combinations for which estimated and actual meal emissions are equal after the same rounding. Column (3) returns to the emissions gap but retains correct guesses as a separate category. Column (4) repeats Column (1) after excluding third-round observations. The underestimation-minus-overestimation contrasts are –€0.131 (SE = €0.035, $p < .001$), –€0.13 (SE = €0.04, $p = 0.001$), and –€0.13 (SE = €0.04, $p < .001$) in Columns (1)–(3), respectively. Columns (1)–(3) control for third-round decisions; Column (4) excludes them. Standard errors are clustered at the individual level.

Table OA.8. Within-subject WTP changes under attention and labels

	Change in WTP compared to baseline	
	(1) All rounds	(2) Excluding round 3
Low emissions × Attention	−0.02	−0.01
	(0.03)	(0.03)
Low emissions × Attention+Label	−0.05	−0.06
	(0.03)	(0.04)
High emissions × Attention	−0.10***	−0.09***
	(0.03)	(0.03)
High emissions × Attention+Label	−0.20***	−0.21***
	(0.03)	(0.04)
Participants attention	151	151
Participants label	293	144
Observations	2,380	1,180

Notes: Coefficients report changes from baseline WTP within the Attention and Attention+Label conditions. The additional effect of displaying labels is the difference between the Attention+Label and Attention coefficients for the relevant emissions category. Column (1) uses all observations and controls for third-round decisions; Column (2) excludes third-round observations. Standard errors are clustered at the participant level.

OA.2.2 Details of structural estimation and simulations

OA.2.2.1 Basic-model estimating equations. Section 6.2 summarizes the identification strategy. Here, I write out the three WTP-difference equations and the moment conditions used in estimation. Define

$$\Delta e_{im}^{\text{prior}} = e_{im}^{\text{prior}} - e_{io}^{\text{prior}}, \qquad \Delta e_{m}^{\text{true}} = e_{m}^{\text{true}} - e_{o}^{\text{true}}.$$

Differencing within participant and meal gives

$$\widetilde{WTP}_{im}^{\text{A+L}} - \widetilde{WTP}_{im}^{\text{B}} = \gamma \Delta e_{im}^{\text{prior}}(\theta - \kappa) - \gamma \Delta e_{m}^{\text{true}}(1 - \kappa) + \xi_{im}^{L}, \tag{OA.8}$$

$$\widetilde{WTP}_{im}^{\text{A}} - \widetilde{WTP}_{im}^{\text{B}} = \gamma \Delta e_{im}^{\text{prior}}(\theta - 1) + \xi_{im}^{A}, \tag{OA.9}$$

$$\widetilde{WTP}_{im}^{\text{A+O}} - \widetilde{WTP}_{im}^{\text{A+L}} = \gamma \Delta e_{m}^{\text{true}}(1 - \kappa) + \gamma \Delta e_{im}^{\text{prior}} \kappa + \xi_{im}^{O}. \tag{OA.10}$$

Here, $\xi^{L}_{im} = \eta^{\mathrm{A+L}}_{im} - \eta^{\mathrm{B}}_{im}$, $\xi^{A}_{im} = \eta^{\mathrm{A}}_{im} - \eta^{\mathrm{B}}_{im}$, and $\xi^{O}_{im} = \eta^{\mathrm{A+O}}_{im} - \eta^{\mathrm{A+L}}_{im}$. I do not specify a parametric distribution for these residuals. Let $\psi = (\gamma, \theta, \kappa)$. The basic estimator uses the six moment conditions

$$\mathbb{E}\big[\,\xi^{r}_{im}(\psi_0)\,\Delta e^{\mathrm{prior}}_{im}\,\big] = 0, \qquad \mathbb{E}\big[\,\xi^{r}_{im}(\psi_0)\,\Delta e^{\mathrm{true}}_{m}\,\big] = 0, \qquad r \in \{L, A, O\}. \tag{OA.11}$$

The estimator therefore uses two moments for each of the three WTP contrasts: orthogonality of the residual to the prior emissions difference and to the true emissions difference.

I estimate the equations jointly by constrained two-step GMM, using Experiment 3 data. To reduce the effect of outliers, I rank, separately by meal, the absolute discrepancy between each participant's prior and the true emissions difference relative to the cheese sandwich and drop observations in the highest decile. This leaves 1,056 observations from 284 participants for Equation OA.8, 1,104 observations from 146 participants for Equation OA.9,[71] and 1,056 observations from 284 participants for Equation OA.10. I treat participants as independent sampling units and construct the second-step weighting matrix from participant-level sums of the first-step moment contributions. I draw 999 participant-bootstrap resamples from this sample, keeping the trimming decisions fixed, and re-estimate the parameters and weighting matrix in each replication. The restrictions $\gamma \geq 0$ and $\theta, \kappa \in [0, 1]$ are imposed in the point estimate and in every bootstrap replication.

OA.2.2.2 Basic-model results and robustness. Table OA.9 reports the basic-model estimates. Column (1) is the main specification; Columns (2)–(4) impose, respectively, $\kappa = 0$, $\theta = 0$, and $\theta = \kappa = 0$. Because the main estimate of θ lies at zero, imposing $\theta = 0$ in Column (3) leaves the other estimates essentially unchanged.

71. The higher number of observations per participant stems from the Attention equation using the second- and third-round comparisons with baseline separately.

Table OA.9. Basic-model estimates

	(1)	(2)	(3)	(4)
θ	0.00 [0.00, 0.27]	0.00 [0.00, 0.24]	0 (fixed)	0 (fixed)
γ	0.10 [0.08, 0.13]	0.10 [0.07, 0.12]	0.10 [0.07, 0.13]	0.10 [0.07, 0.12]
κ	0.16 [0.00, 0.55]	0 (fixed)	0.16 [0.00, 0.53]	0 (fixed)
Basic-equation observations	3,216	3,216	3,216	3,216

Notes: Experiment 3; sample and estimation procedure are described in Section OA.2.2.1. Column (1) is the main specification; the remaining columns impose the indicated equality restrictions. Brackets report percentile 95% intervals from 999 participant-bootstrap replications. γ is measured in euros per kg CO_2e.

Table OA.10 separately assesses these restrictions and the fully attentive baseline restriction $\theta = 1$. For this comparison, I hold the main model's participant-clustered weighting matrix fixed across specifications, so differences in the GMM criterion reflect the imposed restrictions rather than changes in the weighting matrix. Imposing $\kappa = 0$ increases the objective by only 0.54 (chi-squared reference $p = 0.463$), whereas imposing $\theta = 1$ increases it by 48.92 (reference $p < 0.001$). The fit therefore deteriorates much more under a fully attentive baseline than under $\kappa = 0$. Because the restrictions involve parameter-space boundaries, these ordinary chi-squared probabilities serve only as reference values rather than exact tests.

Table OA.10. Model fit under basic-model parameter restrictions

Specification	J	ΔJ	df	Reference p
Main (no equality restriction)	5.14	0.00	0	—
$\kappa = 0$	5.68	0.54	1	0.463
$\theta = 0$	5.14	0.00	1	1.000
$\theta = \kappa = 0$	5.68	0.54	2	0.764
$\theta = 1$	54.07	48.92	1	0.000

Notes: Same sample and moment conditions as Table OA.9. All specifications use the main model's weighting matrix. ΔJ is relative to the main specification; df counts additional equality restrictions. Reference p-values use the chi-squared distribution and are approximate at parameter-space boundaries.

Table OA.11 isolates the role of the theoretical parameter restrictions while retaining the same participant-clustered two-step weighting procedure. The weighting matrix is re-estimated for each specification. Without the restrictions, the estimate of θ is negative; imposing $\theta \in [0, 1]$ places it at the lower boundary of zero.

Table OA.11. Sensitivity to removing parameter restrictions

Parameter	Constrained	Unconstrained	At bound
γ	0.101	0.095	0.0%
θ	0.000	-0.069	58.6%
κ	0.162	0.115	20.3%

Notes: Sample and two-step GMM procedure as in Table OA.9; the unconstrained column removes the theoretical parameter restrictions. At bound reports the percentage of constrained bootstrap estimates within 10^{-10} of zero for γ, or of either zero or one for θ and κ. γ is measured in euros per kg CO_2e.

OA.2.2.3 Observed-versus-predicted WTP fit. Table OA.12 compares observed Experiment 3 WTP changes with the corresponding structural predictions, separately for meals above and below the sandwich's emissions. These cell means are not directly targeted by the six GMM moment conditions in Equation OA.11, so the comparison provides an additional model-fit check. The model reproduces the average Attention responses reasonably closely. For higher-emission meals in the Attention+Label condition, however, the observed average WTP reduction is 20.0 cents, compared with a structural prediction of 12.0 cents. The parsimonious model therefore preserves the qualitative treatment pattern but understates the label response for higher-emission meals. Some discrepancy is unsurprising because the baseline model imposes common parameters across participants and meals, while the GMM moments use covariation with prior and true emissions differences rather than targeting these pooled means. Moreover, a pooled WTP discrepancy does not map one-for-one into simulated emissions: discrete choices depend on relative WTP within each menu and on which responses cross a choice threshold. For reference, the table also reports the corresponding observed Label and Offset-minus-Label changes from Experiment 1; structural predictions are unavailable there because those participants did not provide the prior-emissions beliefs required to construct them.

Table OA.12. Observed and structurally predicted WTP changes

Comparison	Meal emissions	Obs.	Observed	Structural	Residual [95% CI]
Experiment 3					
Attention	Below sandwich	352	-0.002	-0.016	0.015 [-0.035, 0.071]
Attention	Above sandwich	752	-0.065	-0.082	0.017 [-0.018, 0.049]
Attention+Label	Below sandwich	323	-0.014	0.014	-0.028 [-0.089, 0.032]
Attention+Label	Above sandwich	733	-0.200	-0.120	-0.080 [-0.132, -0.038]
Offset minus Label	Below sandwich	323	0.024	-0.014	0.039 [-0.005, 0.084]
Offset minus Label	Above sandwich	733	0.172	0.120	0.052 [0.012, 0.091]
Experiment 1					
Label	Below sandwich	265	0.086	–	–
Label	Above sandwich	603	-0.291	–	–
Offset minus Label	Below sandwich	177	-0.040	–	–
Offset minus Label	Above sandwich	415	0.256	–	–

Notes: Observed changes are in euros. In Experiment 3, the Attention and Attention+Label comparisons are relative to baseline WTP, and Offset minus Label is the change from Attention+Label to Attention+Offset. In Experiment 1, Label is relative to baseline WTP, and Offset minus Label compares each participant's offset WTP with the same participant's preceding label WTP. Higher- and lower-emission meals are defined relative to the cheese sandwich. Residual equals observed minus structural prediction; its percentile interval re-estimates the structural parameters in each participant-bootstrap replication. Experiment 1 participants in the Label and Offset conditions did not provide the prior-emissions guesses required for participant-specific structural predictions, so the corresponding structural and residual cells are not reported. The Experiment 3 rows report unadjusted means in the trimmed structural-estimation sample, rather than the round-adjusted regression coefficients reported in Table OA.8.

OA.2.2.4 Welfare-extension estimating equation. Equation 12 relates WTP for label presence to the expected benefit of improved meal choices. At elicitation, participants do not yet know the final meals or implementation prices. I approximate their expected benefit using the meals they encountered earlier and the experimental BDM price distribution.

Let $\psi = (\gamma, \theta, \kappa)$. For a given parameter vector, I approximate the participant–meal-specific non-emissions valuation from baseline WTP as

$$\widehat{\Delta v}_{im}(\psi) \;=\; \widetilde{WTP}^{\mathrm{B}}_{im} + \theta\gamma\Delta e^{\mathrm{prior}}_{im}.$$

At the true parameter values, this approximation equals $\Delta v_{im} + \eta^{\mathrm{B}}_{im}$. I assume that baseline elicitation errors have mean zero and use the resulting valuations without a separate correction for individual elicitation noise.

Using this approximated valuation, the model-predicted WTP under BASELINE, ATTENTION, and ATTENTION+LABEL is

$$\begin{aligned}
\widehat{\widetilde{WTP}}^{\mathrm{B}}_{im}(\psi) &= \widehat{\Delta v}_{im}(\psi) - \theta\gamma\Delta e^{\mathrm{prior}}_{im} = \widetilde{WTP}^{\mathrm{B}}_{im}, \\
\widehat{\widetilde{WTP}}^{\mathrm{A}}_{im}(\psi) &= \widehat{\Delta v}_{im}(\psi) - \gamma\Delta e^{\mathrm{prior}}_{im}, \\
\widehat{\widetilde{WTP}}^{\mathrm{A+L}}_{im}(\psi) &= \widehat{\Delta v}_{im}(\psi) - \gamma\left[(1-\kappa)\Delta e^{\mathrm{true}}_{m} + \kappa\Delta e^{\mathrm{prior}}_{im}\right].
\end{aligned}$$

Here, $\widetilde{WTP}$ denotes elicited WTP and $\widehat{WTP}$ denotes its model-predicted counterpart. The recovered valuation makes predicted baseline WTP equal elicited baseline WTP by construction. Under labels, emissions are displayed and beliefs update according to Equation 5; under attention alone, the participant relies on her prior emissions beliefs.

The welfare calculation uses the same true-utility criterion as the main text. Using these approximated valuations, I calculate the warm meal's true utility relative to the sandwich at equal prices as

$$a_{im}(\psi) = \widehat{\Delta v}_{im}(\psi) - \gamma\Delta e^{\mathrm{true}}_{m}.$$

At a relative price q, this utility difference becomes $a_{im}(\psi) - q$.

Let $q \in [-3, 3]$ denote the signed BDM draw. A positive q means that the sandwich is initially assigned and q is the price of switching to the warm meal; a negative q means that the warm meal is initially assigned and $-q$ is the price of switching to the sandwich. A participant with signed WTP w for the warm meal relative to the sandwich therefore receives the warm meal whenever $q \le w$. For the calculation, I approximate the experimental five-cent price grid by a continuous uniform signed draw on $[-3, 3]$. For each draw, I normalize the utility of the sandwich to zero. Expected true utility relative to this alternative is therefore

$$G_{im}(w;\psi) = \mathbb{E}_q\left[\mathbb{1}\{q \le w\}\left(a_{im}(\psi) - q\right)\right] = \begin{cases} 0, & w \le -3, \\ \dfrac{w+3}{6}\left(a_{im}(\psi) - \dfrac{w-3}{2}\right), & -3 < w < 3, \\ a_{im}(\psi), & w \ge 3. \end{cases} \tag{OA.12}$$

The first case corresponds to always receiving the sandwich, whose relative utility is normalized to zero. In the second case, the warm meal is received with probability $(w+3)/6$, and the average signed price conditional on receiving it is $(w-3)/2$. In the third case, the warm meal is always received; the signed price draw then averages to zero, so expected relative utility equals $a_{im}(\psi)$.

I evaluate this expected utility using $\widehat{WTP}_{im}^{\mathrm{A+L}}(\psi)$ and $\widehat{WTP}_{im}^{\mathrm{A}}(\psi)$. Their difference gives the expected choice benefit from adding labels for that participant and meal. I then average these benefits across the participant's earlier meals. Let $\mathscr{M}_i$ contain the distinct earlier meals retained after the trimming described in Section OA.2.2.1 for participant i. Giving each meal equal weight, the predicted average benefit is

$$\widehat{\Delta U}_i^{\mathrm{choice}}(\psi) \;=\; \frac{1}{|\mathscr{M}_i|}\sum_{m\in\mathscr{M}_i}\Big[G_{im}\Big(\widehat{WTP}_{im}^{\mathrm{A+L}}(\psi);\psi\Big)-G_{im}\Big(\widehat{WTP}_{im}^{\mathrm{A}}(\psi);\psi\Big)\Big].$$

The estimating equation compares this predicted benefit with the participant's elicited WTP for label presence:

$$\widetilde{WTP}_i^{P} \;=\; \widehat{\Delta U}_i^{\mathrm{choice}}(\psi)+F+\xi_i^{P}. \tag{OA.13}$$

The error ξ_i^P includes elicitation noise and any difference between the participant's expected benefit and the model's approximation. Interpreting the average residual as F requires this error to have mean zero. Further, I treat the additional utility represented by F as accruing for the implemented consumption choice. The elicitation does not separately identify benefits from viewing labels in decisions that are not selected for implementation.

I estimate $\widehat{\psi}$ using the six moment conditions in Equation OA.11, without label-presence WTP, and use these estimates to calculate each participant's predicted choice benefit $\widehat{\Delta U}_i^{\mathrm{choice}}(\widehat{\psi})$. I subtract this benefit from her WTP for label presence and estimate F by averaging the resulting residuals, giving each participant equal weight.

OA.2.2.5 Policy simulations. I simulate the interventions described in Section 6.5 using participants' elicited baseline WTP $\widetilde{WTP}_{im}^{\mathrm{B}}$, prior emissions beliefs, and the estimated parameters θ,γ,κ. The choice environment is the same simultaneous four-meal-plus-sandwich menu used in the main Experiment 1 benchmark. Non-vegetarians choose

among the four warm meals and the sandwich; vegetarians choose among the two vegetarian meals and the sandwich. The simulation uses the complete pre-trimming baseline records of the 430 Experiment 3 participants that contribute to the structural estimation, i.e. at least one of their participant-meal observations is retained by the 10% trimming rule.[72]

Baseline, Attention, and Label choices use the model-predicted thresholds $\widehat{WTP}^{B}_{im}(\psi)$, $\widehat{WTP}^{A}_{im}(\psi)$, and $\widehat{WTP}^{A+L}_{im}(\psi)$ defined above. These are the plug-in counterparts of Equations 6, 7, and 8. The counterfactual Knowledge intervention provides the same emissions information as the label while retaining baseline salience:

$$\widehat{WTP}^{K}_{im}(\psi) \;=\; \widehat{\Delta v}_{im}(\psi) - \theta\gamma\left[(1-\kappa)\Delta e^{\text{true}}_{m} + \kappa\Delta e^{\text{prior}}_{im}\right]. \tag{OA.14}$$

For each participant, I subtract each meal's canteen price relative to the sandwich from its model-predicted WTP and select the available option with the highest resulting value, normalizing the sandwich's value to zero.

The tax simulations use elicited baseline WTP $\widetilde{WTP}^{B}_{im}$ and increase the relative price of meal m by $\tau\Delta e^{\text{true}}_{m}$. I consider the €155/t midpoint of the Experiment 1 benchmark and the Experiment 3 tax rate whose emissions effect most closely matches the structurally predicted label effect. I identify the latter using the same exact choice-threshold procedure described in Section OA.1.2. Tax revenue is returned as a choice-independent lump-sum rebate, so it affects welfare but not choices. The meat-ban simulation removes the two meat meals from the choice set.

I aggregate choices, emissions, and choice-mediated welfare using equal participant weights. Welfare is the change relative to baseline in true utility of the selected meal, evaluated using Equation 10. I approximate the non-emissions valuations relative to the sandwich from baseline WTP as described in Section OA.2.2.4. Under a tax, the utility calculation uses tax-inclusive prices, and the lump-sum rebate returns the resulting tax revenue.

Each participant-bootstrap replication re-estimates the structural parameters and repeats the simulation using bootstrap multiplicities as participant weights.

72. The structural-estimation sample retains 430 of the 444 Experiment 3 participants; the meal-specific trimming of observations in the highest decile of the absolute discrepancy between prior and true emissions removes all relevant observations for the remaining 14 participants.

OA.2.2.6 Sensitivity to outlying prior emissions beliefs. The structural estimation excludes, separately for each meal, observations in the highest decile of the absolute discrepancy between prior and true emissions differences. The main simulation includes the 430 participants with at least one observation retained in the structural estimation and restores all their pre-trimming baseline WTP records, so that each participant faces the complete diet-specific menu. This is useful for representing their full choice sets, but a small number of particularly inaccurate beliefs can contribute to wide welfare tails when a bootstrap draw assigns positive weight to baseline salience. Table OA.13 therefore repeats the basic-model simulation for the 361 participants whose complete simulated menu remains in the estimation sample: all four warm-meal observations for non-vegetarians and both vegetarian-meal observations for vegetarians survive trimming.

In this sample, attention alone reduces emissions by 38.5 grams per meal, compared with 64.5 grams under labels. The implied attention share is therefore approximately 60%, compared with 73% in the main simulation. Because the structural parameter estimates are unchanged, this difference reflects the restriction of the simulation sample.

Table OA.13. Structural simulations excluding participants with outlying prior-belief errors

Intervention	Δ GHGE (g)		Δ welfare (cents)	
	Estimate	95% CI	Estimate	95% CI
Panel A: Main simulation				
Knowledge	0.0	[0.0, 38.0]	0.00	[-0.06, 24.79]
Attention	-54.4	[-94.1, -13.7]	-0.06	[-0.74, 0.33]
Labels	-74.1	[-120.1, -20.5]	0.32	[-0.18, 18.85]
Panel B: No outlying prior-belief errors				
Knowledge	0.0	[0.0, 16.9]	0.00	[-0.14, 0.00]
Attention	-38.5	[-73.2, -6.6]	0.19	[-0.01, 0.49]
Labels	-64.5	[-119.6, -15.9]	0.33	[0.07, 0.70]

Notes: Simulation and bootstrap procedures follow Table 3. Panel A uses the 430 participants retained in the main simulation; Panel B uses the 361 whose complete diet-specific meal set survives trimming. Both panels use the same parameter estimates in each replication. Emissions changes are in grams CO_2e per meal; welfare changes are in cents per meal and exclude *F*.

OA.2.2.7 Observed parameter heterogeneity. The baseline model assumes common γ, θ, and κ. I assess sensitivity to this assumption by re-estimating the structural model separately for vegetarians and non-vegetarians, students and non-students, participants below versus at or above the median of food-specific willingness to sacrifice for the environment, and participants below versus at or above the median of general willingness to

sacrifice for the environment. I also repeat the food-specific split among non-vegetarians only. Full subgroup parameter estimates and definitions are reported in Supplementary Material Section S.3.6.

The summary below reports how assigning participants the parameter estimates for their subgroup affects the structural simulations and their fit to reduced-form WTP changes. Across the four heterogeneity dimensions, the implied attention share ranges from approximately 63% to 80%, while the label-induced emissions reduction ranges from 53.7 to 85.8 grams per meal. The structural model also continues to understate the average Attention+Label response for higher-emission meals by about 9–10 cents. Thus, allowing for these observed forms of parameter heterogeneity changes the simulated magnitudes but does not remove the WTP-fit gap.

Table OA.14. Summary of observed-heterogeneity checks

Outcome across diet, student-status, food-specific-WTS, and general-WTS specifications	Range of estimates	Bootstrap comparison
Attention share of label-induced GHGE reduction	63%–80%	–
Label-induced GHGE reduction (grams per meal)	53.7–85.8	–
Reduced form minus structural WTP change: Attention, higher-emission meals (euro)	−0.002–0.004	Every 95% interval includes zero
Reduced form minus structural WTP change: Attention+Label, higher-emission meals (euro)	−0.098–−0.092	Every 95% interval excludes zero

Notes: For each heterogeneity dimension, the structural model is estimated separately for the two subgroups and participants are assigned the parameter estimates for their subgroup before choices are simulated. WTS denotes willingness to sacrifice for the environment. The fit comparisons use the same subgroup definitions for the reduced-form estimates and structural predictions. Full estimates and 999-replication participant-bootstrap intervals are reported in Supplementary Material Tables S.35 and S.36.

OA.2.2.8 Preferences for label presence. Figure OA.1 and Table OA.15 report the distribution of WTP for label presence and its variation across prior treatment sequences.

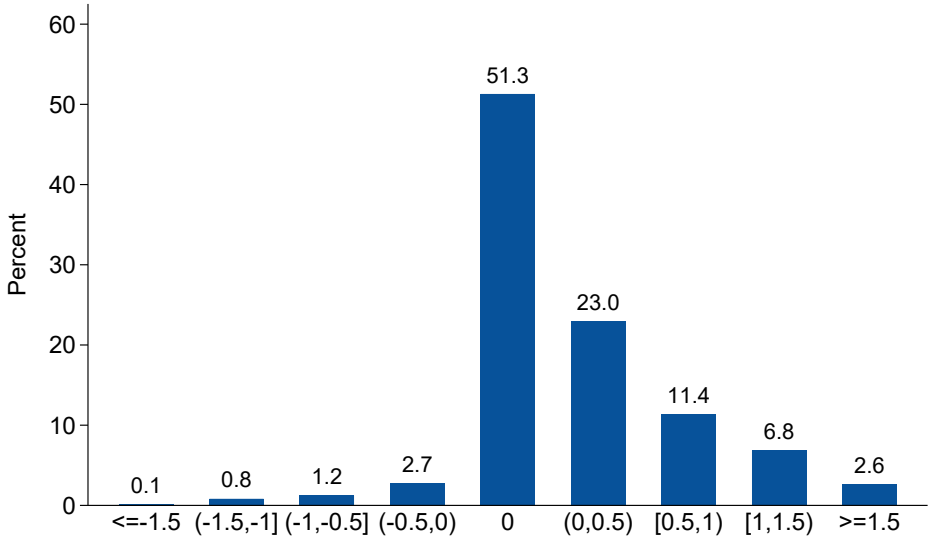


Figure OA.1. Distribution of WTP indicated to see carbon labels on the final three consumption decisions

Notes: In euros. Based on Experiments 1 and 3 and including all 731 participants.

Table OA.15. WTP for seeing carbon labels by treatment group

	WTP for labels
	(1)
Control, then Labels	−0.13
	(0.08)
Labels, then Offset	−0.11
	(0.08)
Attent, then Attent	−0.08
	(0.07)
Attention+Label, then Offset	−0.07
	(0.07)
Attention+Offset, then Labels	−0.04
	(0.07)
Constant	0.28***
	(0.06)
Participants control, then Control	70
Participants Control, then Labels	69
Participants Labels, then Offset	148
Participants Attention, then Attention	151
Participants Attention+Offset, then Labels	149
Participants Attention+Label, then Offset	144
Observations	731

Notes: Average deviation from WTP to see emissions labels in the Control–Control sequence. Heteroskedasticity-robust standard errors are reported.

OA.2.2.9 Welfare-extension results and robustness. Table OA.16 reports the participant-level estimates of F. In each of the same participant-bootstrap resamples used for the basic model, I use that replication's re-estimated parameters to recalculate expected choice benefits and average the resampled participants' residual valuations. The basic-model parameter restrictions continue to apply; F is unrestricted.

Column (1) uses participants who previously experienced Attention+Label, while Column (2) also includes participants from the Attention condition. I prefer Column (1) because participants who have seen the labels may form better-informed expectations about their value. Both columns use the basic-model estimates in Table OA.9, Column (1), with the corresponding replication-specific estimates used in the bootstrap.[73]

Table OA.16. Participant-level estimates of the residual value of label presence

	(1) Label-experienced	(2) All retained participants
Mean WTP for label presence (€)	0.214	0.210
Mean expected choice benefit (€)	0.000455	0.000464
Residual value F (€)	0.214	0.209
Bootstrap standard error of F	0.024	0.020
95% percentile interval for F	[0.170, 0.264]	[0.173, 0.249]
Participants	284	430
Unique participant–meal pairs	1,056	1,608

Notes: Column (1) includes the 284 participants who experienced Attention+Label before reporting their WTP for label presence, using the 1,056 participant–meal observations from this group that survive the structural estimation's meal-specific trimming rule. Column (2) additionally includes the 146 retained participants who experienced Attention but not Attention+Label, yielding 430 participants and 1,608 participant–meal observations. For these additional participants, I use one Attention observation per meal. For each participant, I calculate the model-implied Attention-to-Attention+Label benefit for every retained meal and average these benefits across meals. I subtract this participant-level expected choice benefit from her WTP for label presence and estimate F as the mean residual across participants. Both columns use the basic-model parameters from Table OA.9, Column (1); each bootstrap replication uses the corresponding re-estimated parameters and resampled participants. Differences are calculated before rounding, so the displayed components need not subtract exactly to the displayed F. F is measured in euros per choice.

73. The expected Attention-to-Attention+Label gain used to estimate F is approximately 0.046 cents per choice, while the difference between the Label and Attention welfare effects in Table 3 is 0.373 cents. The quantities need not coincide because they represent different choice problems. Table 3 restores all pre-trimming baseline meal records for the 430 participants with at least one observation retained in the structural estimation and simulates simultaneous choice from the complete menu at fixed canteen prices. Table OA.16 uses only participant–meal records retained after meal-specific trimming and averages binary meal-versus-sandwich gains over the experimental BDM price distribution.

OA.2.2.10 Additional structural robustness. I consider two alternative structural specifications in Supplementary Material Section S.3.8. First, I allow ATTENTION and ATTENTION+LABEL to impose an additional emissions-related psychological cost with the same functional form as the emissions term already entering perceived utility. The predicted emissions effects are unchanged at the point estimates, but welfare excluding F falls once the additional psychological costs are included. Under labels, welfare excluding F is -4.51 cents per choice, while the separately estimated residual value is 21.02 cents, giving a total welfare increase of 16.51 cents (95% participant-bootstrap percentile interval: $[11.37, 32.51]$).

Second, I relax the assumption that offsetting completely removes perceived environmental guilt and instead assume that 70% is removed. Under this specification, the estimates of γ, θ, and κ are higher. The predicted ATTENTION and meat-ban emissions effects are unchanged, the label-induced reduction changes from 74.1 to 69.4 grams per meal, and the KNOWLEDGE intervention increases emissions by 27.5 grams per meal. Supplementary Material Sections S.3.8.1 and S.3.8.2 report the full estimating equations, parameter estimates, policy simulations, and bootstrap intervals.

OA.3 Experiment 2: Field robustness

OA.3.1 Additional analyses on labeling effects

The pre-registration for Experiment 2 can be found on AsPredicted: #95108.

I preregistered to examine:

(1) **The effect of the labels on meat consumption.** This analysis is shown graphically in the main text, with regression estimates and robustness checks reported in Tables OA.17 and OA.18. Table OA.17, Panel A, applies the DiD model to the sales data and gradually adds controls. My full-sample specification in Col. 3 includes week and day-of-week fixed effects to flexibly control for time trends (e.g., seasonal or semester-specific effects), and indicators for the exact main meat and vegetarian meals on offer on a given day in a given canteen. Both main offers have the same exact dish identities across the three canteens for approximately 69% of transactions in the study sample. For the remaining transactions, at least one main offer differs across canteens. For instance, the treatment and the larger control canteen share the same main vegetarian option on 87% of the 54 days covered by the sample, and the

Table OA.17. Field estimates of the effect of carbon labels

Panel A: Full-sample meat-choice specifications

	Meat meal chosen		
	Base	+ Offer	+ Week FE
Treatment × Label period	-0.022***	-0.022***	-0.024***
	(0.007)	(0.007)	(0.007)
Treatment × Post period	-0.047***	-0.048***	-0.050***
	(0.009)	(0.009)	(0.009)
Week fixed effects	No	No	Yes
Offer controls	No	Yes	Yes
Observations	124,784	124,784	124,784

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Panel B: Progressive ITT meat-choice specifications

	Meat meal chosen			
	+ Offer	+ Week FE	+ Guest FE	Students
Treatment × Label period	-0.044***	-0.047***	-0.022**	-0.022**
	(0.014)	(0.014)	(0.010)	(0.011)
Treatment × Post period	-0.094***	-0.096***	-0.041***	-0.036**
	(0.020)	(0.020)	(0.014)	(0.014)
Offer controls	Yes	Yes	Yes	Yes
Week fixed effects	No	Yes	Yes	Yes
Guest fixed effects	No	No	Yes	Yes
Guests control	1,422	1,422	1,422	1,216
Guests treated	496	496	496	451
Observations	37,284	37,284	37,284	32,116

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Panel A uses all transactions. All three columns are linear probability models: Column (1) is the baseline difference-in-differences specification, Column (2) adds the exact meat and vegetarian main meals and indicators for second meat and vegetarian mains, and Column (3) adds week and day-of-week fixed effects. Specifications without week fixed effects include common label- and post-period indicators; week fixed effects absorb these indicators in the remaining specifications. Panel B assigns treatment from guests' pre-intervention visiting patterns. Column (1) includes offer controls, Column (2) adds week and day-of-week fixed effects, and Column (3) adds guest fixed effects. Column (4) repeats the guest-fixed-effects specification for student-card transactions. The dependent variable equals one for a meat meal and zero otherwise; coefficients are in probability units, so 0.01 corresponds to one percentage point. Standard errors are heteroskedasticity robust in Panel A and clustered at the guest level in Panel B.

same main meat option on 89% of these days. The meal indicators account for these differences. On occasion, a canteen offers a second vegetarian and/or a second meat main meal.[74] The specification therefore also controls for whether an additional vegetarian or meat option is offered in the relevant canteen-day. The estimates remain close to a 2 percentage point reduction, or 5% of baseline treatment-group meat consumption, as these controls are added. Col. (2) in Panel B uses the same set of controls to analyze the ITT sample, assigning treatment status based on which canteen guests frequent most in the pre-intervention period. My main ITT specification in Panel B Col. (3) includes guest fixed effects. Panel B Col. (4) shows that the estimate is unchanged when this specification is restricted to only students, i.e. excluding university employees. Table OA.18, Col. (5) estimates a binomial generalized linear model with a log link to estimate multiplicative effects of the labels on the probability of choosing a meat meal, yielding estimates in line with the main text. Supplementary Material Sections S.5.5 and S.5.6 probe robustness of the main estimates to alternative data-preparation methods.

In the main regressions, I use heteroskedasticity-robust standard errors for analyses using the full data set, and cluster at the individual level in the ITT specification. The identifying assumption is that there is no contemporaneous treatment-canteen-specific shock that I do not control for. The main ITT regression in Table OA.17, Panel B, Column (3), includes individual fixed effects, which absorb all time-invariant guest heterogeneity such as dietary preferences, demographics, or psychological traits. The combination of pre-intervention ITT assignment and guest fixed effects minimizes the influence of attendance-related changes in guest composition: treatment assignment remains fixed when guests change their visiting behavior, and guest fixed effects absorb stable differences in meat-choice propensity. Differential attendance can nevertheless affect the transaction-weighted estimate by changing guests' relative contribution or actual exposure. Table S.47 thus directly examines whether the labeling intervention impacted guests' likelihood of frequenting the canteens, and finds no evidence of a statistically significant change in visiting during the labeling period.

74. Specifically, 18% of observations pertain to canteen-day instances where a second vegetarian main was offered, and 7% to instances where a second meat main was offered.

(2) **The effect of the labeling intervention on guests' likelihood of choosing a green-, yellow-, or red-labeled meal.** Supplementary Material Table S.26 analyzes choices between the two main offers, conditional on the availability of the relevant color and meal type and controlling for the color of the other main offer. Green- and yellow-labeled vegetarian choices increase, while yellow-labeled meat and red-labeled choices decrease. These estimates describe substitution among the observed menu combinations; the columns use different samples, so their effects need not sum to zero.

(3) **The effect of the labeling intervention on greenhouse gas emissions.** Table OA.18, Column (1), analyzes GHGE directly and estimates an approximately 3% decrease; the remaining columns report alternative GHGE specifications. I report these results as secondary results and focus on meat consumption as the main outcome variable. During the sample period, the meat option always had higher or roughly identical emissions compared to the vegetarian option. I focus on meat consumption as the main outcome for two reasons. First, the initial difference in meat consumption levels across treatment and control canteens, together with a change in average meal emissions, can create misleading results. Average emissions of the meat meal offered decreased between the baseline and labeling period (from 1.8 kg to 1.4 kg CO_2e), while average emissions of the vegetarian meal remained around 0.4 kg. Together with different baseline meat-consumption levels (51% in the control group and 41% in the treatment group), such variation can generate misleading results when a naïve difference-in-differences strategy is applied. For instance, if the labels had no effect on behavior and meat consumption remained constant, a simple DiD would estimate a decrease of $0.4 \text{ kg} \times 0.51 = 0.20$ kg for the control group and $0.4 \text{ kg} \times 0.41 = 0.16$ kg for the treatment group, suggesting a spurious 40-gram *increase* in emissions attributed to the labels despite no change in behavior. The estimation in Table OA.18 therefore (i) controls for the specific main vegetarian and meat meal option offered on each day, as well as whether a second vegetarian or meat option is available, allowing these offer controls to differ between treatment and control canteens, and (ii) relies on a multiplicative model estimated via Poisson Pseudo-Maximum-Likelihood (PPML) as the main specification. Second, small differences in the daily offers of treatment and control canteens can lead to

large differences in estimated emissions reductions when the offered options differ substantially in emissions. The main-meal offers have identical emissions across the three canteens on 40 of 54 sample-period days (74%). Among the remaining non-extreme days, the largest cross-canteen difference is 0.7 kg CO_2e. On the 3 most extreme days—approximately 5% of sample-period days—the difference reaches up to 8.0 kg. I omit these three days from the main emissions analysis. On each of them, one or two of the three canteens offered a beef option while another canteen did not. Table OA.18 also reports the PPML specification using all sample days and a specification restricted to days on which both the vegetarian and meat main-meal offers have identical emissions across all three canteens. I estimate that emissions decreased by approximately 3% as a result of the labels, consistent with a back-of-the-envelope calculation relying on observed changes in meat consumption: In the pre-intervention period, the average emissions of meat meals in the treatment canteen were 1.9 kg, compared with 0.4 kg for vegetarian meals. Multiplying this 1.5 kg emissions gap by the estimated 2-percentage-point reduction in meat choice (1.5×0.02) implies a reduction of about 30 grams per meal, or 3% of baseline emissions of 1 kg, absent changes in menu composition.

Table OA.18. GHGE as outcome variable

	GHGE (% change)			GHGE (kg change)	Meat choice (% change)
	PPML	PPML full	PPML res.	OLS	Bin. GLM
Treatment × Label period	-3.061**	-1.329	-2.676	-0.039***	-4.759***
	(1.409)	(2.207)	(2.076)	(0.013)	(1.515)
Treatment × Post period	-4.580**	-3.230	-2.459	-0.046***	-11.892***
	(1.809)	(2.198)	(2.613)	(0.017)	(1.860)
Observations	118,125	124,784	93,989	118,125	124,784

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: Specifications analyze the effect of carbon labels on greenhouse gas emissions (GHGE). Columns (1)–(3) use Poisson Pseudo-Maximum-Likelihood (PPML); their coefficients are expressed as percentage changes, calculated as $100(\exp(\beta) - 1)$. Column (1), the main specification, omits the three sample-period days on which the maximum cross-canteen difference in main-meal emissions exceeds 1 kg CO_2e. Column (2) uses all sample days. Column (3) retains only days on which both the vegetarian and meat main-meal offers have identical emissions across all three canteens. Column (4) uses an additive OLS model with GHGE in kilograms as the outcome and applies the same three-day exclusion as Column (1). Columns (1)–(4) control for the specific vegetarian and meat main meal on offer and whether a second vegetarian or meat option is available, allowing these offer controls to vary by treatment canteen. Column (5) uses all sample days and a baseline difference-in-differences binomial generalized linear model with a log link, without offer controls; its coefficients, also expressed as percentage changes, estimate the effect of the labels on the probability of choosing a meat meal. All columns include treatment, label-period and post-period indicators and treatment–period interactions, but no week, calendar-date, day-of-week, or guest fixed effects. Heteroskedasticity-robust standard errors are reported in parentheses.

Experiment 2 compares one treated canteen with two control canteens. Table OA.19 examines whether the full-sample estimate is driven by changes in either control location. The estimate changes little when the treated canteen is compared separately with the larger or the smaller control canteen.

Table OA.19. Label effects using alternative control canteens

	Meat meal chosen		
	Both controls	Larger control	Smaller control
Treatment × Label period	-0.024***	-0.022***	-0.028***
	(0.007)	(0.008)	(0.010)
Treatment × Post period	-0.050***	-0.051***	-0.041***
	(0.009)	(0.010)	(0.013)
Week fixed effects	Yes	Yes	Yes
Offer controls	Yes	Yes	Yes
Canteen-days	162	108	108
Observations	124,784	99,058	64,817

Notes: The dependent variable equals one when a transaction contains the meat meal and zero otherwise. Coefficients are in probability units, so 0.01 corresponds to one percentage point. Column (1) uses both control canteens, Column (2) retains the larger control canteen, and Column (3) retains the smaller control canteen. Each column uses the specification and heteroskedasticity-robust standard errors of Table OA.17, Panel A, Column (3); Column (1) reproduces its coefficient and standard error. Observations and canteen-days are reported in the table. All inference is conditional on the observed canteens. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

OA.3.2 Additional price-effect robustness and field-implied tax-equivalent rates

The price coefficients in Table 2 are identified from variation in the meat–vegetarian price difference within broad meal categories and from non-uniform price changes over time. Table OA.20 documents these sources of variation and reports robustness to alternative control sets. The price response remains negative across these specifications, although more restrictive controls absorb a substantial share of the price variation from which the coefficients are identified. Supplementary Material Table S.29 further tests robustness by restricting the price-estimation sample to different time periods (before the October 2022 price change, within and outside of the label-study window).

Table OA.20. Comparison of effects: labels vs. "carbon tax" under alternative control sets

	(1) Broad type	(2) + Time	(3) + Meat cat.	(4) + Veg. cat.	(5) + Exact meat
PD x Poultry	-0.006	-0.033	-0.058**	-0.029	-0.058*
	(0.038)	(0.029)	(0.024)	(0.026)	(0.030)
PD x Pork	-0.047	-0.146***	-0.132***	-0.097***	-0.123***
	(0.030)	(0.029)	(0.025)	(0.024)	(0.044)
PD x Beef	-0.091	-0.181***	-0.202***	-0.150***	-0.100*
	(0.065)	(0.052)	(0.042)	(0.049)	(0.059)
PD x Fish	0.011	-0.002	-0.019	0.001	0.011
	(0.049)	(0.041)	(0.036)	(0.030)	(0.060)
Weekly time controls	No	Yes	Yes	Yes	Yes
Meat meal categories	No	No	Yes	Yes	No
Vegetarian meal categories	No	No	No	Yes	Yes
Exact meat meals	No	No	No	No	Yes
Remaining PD variation (%)	93.9	58.8	56.7	31.0	11.9
Observations	465,328	465,328	465,328	465,328	465,328

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable equals one when a transaction contains the meat meal and zero otherwise. Coefficients are in probability units, so 0.01 corresponds to one percentage point. Linear probability regressions use student canteen data from April 2022–July 2023, excluding all canteen observations from May 2023 and June 1–2, 2023. "PD" is the meat-minus-vegetarian price difference. All columns estimate separate price responses for poultry, pork, and beef, with fish and lamb price interactions included as controls. The lamb interaction is not displayed. Column (2) adds time, canteen, and intervention controls; Column (3) adds meat-offer categories; Column (4) also adds vegetarian-offer categories; and Column (5) replaces the meat-offer categories with exact meat-meal indicators while retaining the vegetarian-offer categories. The "Remaining PD variation" row reports residual PD variance as a share of raw variance. In corresponding pooled regressions, the PD coefficient for poultry, pork, and beef is −0.03 (SE 0.02), −0.09 (SE 0.02), −0.10 (SE 0.02), −0.08 (SE 0.02), and −0.09 (SE 0.03) under the five respective control sets. The first coefficient is not significant at the 10% level; the remaining four are significant at the 1% level. Standard errors are two-way clustered by calendar date and exact meat–vegetarian meal pair.

Table OA.21 shows the meat-type calculation underlying the predicted effect of the €155/t midpoint reported in Table 2. I apply the poultry, pork, and beef price coefficients to their type-specific emissions gaps, assign fish a zero response while retaining its offer share, and aggregate using canteen-day offer shares.

Table OA.21. Back-of-the-envelope effect at the €155/t midpoint, differentiating by type

Meat type	Emissions gap (kg CO_2e)	Price change (€)	Price coefficient (share/€)	Demand change (percentage points)	Offer share (%)	Contribution (percentage points)
Poultry	0.79	0.12	-0.058	-0.71	30.2	-0.21
Pork	1.10	0.17	-0.132	-2.25	34.6	-0.78
Beef	3.36	0.52	-0.202	-10.50	14.8	-1.56
Fish	0.10	0.02	0.000	0.00	20.4	0.00
				Total (weighted across types)		**-2.55**

Notes: The table illustrates the €155/t midpoint of the €146–164/t matching set. The GHGE gap is the paired difference between the main meat and vegetarian offers. Price changes equal this gap times €0.155/kg CO_2e. Price coefficients for poultry, pork, and beef are from Table 2, Column (2), and demand changes equal the price change times the coefficient. Fish is assigned a zero response but retains its offer share. Contributions weight each demand change by the type's share of canteen-days. Emissions gaps and offer shares use one observation per canteen-day in the label-study analysis sample (162 canteen-days; see Table S.46).

As a complementary summary, Supplementary Material Table S.33 reports the local tax-equivalent rates obtained by directly inverting the estimated field label and price effects. Using the student-only label estimate, which matches the population used to estimate the price response, the pooled estimate is €194/t CO_2e and the meat-type-differentiated estimate is €133/t. Using the label estimate for all guests instead gives €195/t and €134/t, respectively. Supplementary Material Section S.2.7 reports further specification and inference checks.

Supplementary Material

Understanding the effectiveness of carbon labels

Contents

S.1 Experiments 1 and 3: Further tables and figures S-3

S.1.1 Randomization checks S-3

S.1.2 Representativeness of the sample S-4

S.1.3 Descriptive statistics on baseline willingness to pay for meals S-7

S.1.4 Test for effect of labels on meal perception S-8

S.1.5 Distribution of individual label responses in Experiment 1 ..... S-8

S.1.6 Preregistered main effects in Experiments 1 and 3 ..... S-10

S.1.7 Results split by (non-) vegetarians and (non-) students S-15

S.1.8 Experiment 1: Alternative econometric specifications S-18

S.1.9 Experiment 1: Intuition behind expressing effect sizes in terms of a carbon tax. S-21

S.1.10 Experiment 3: Descriptives on under- and over-estimation S-24

S.1.11 Correlates of preferences for label presence S-26

S.2 Experiment 2: Further tables and figures S-27

S.2.1 Label-effect inference under guest and date clustering S-28

S.2.2 Time trends S-29

S.2.3 Label-color outcomes S-31

S.2.4 Guest numbers and canteen switching S-31

S.2.5 Field survey results S-33

S.2.6 Price elasticities S-34

S.2.7 Additional analyses on estimation of price effects S-35

S.2.7.1 Price-effect inference S-40

S.2.7.2 Direct field-implied tax-equivalent rates S-42

S.3 Applicability to other contexts S-42

S.3.1 Suitability of the study setting S-43

S.3.2 Contribution 1: Portability of the benchmarking framework S-44

S.3.3 Contribution 2: Scope and applicability of the local €155/t benchmark S-44
S.3.4 Contribution 3: Applicability of the mechanism findings S-48
S.3.5 Contribution 4: Applicability of the consumer-preference and welfare findings S-49
S.3.6 Structural-model fit and observed parameter heterogeneity S-49
S.3.7 Additional heterogeneity evidence S-53
S.3.8 Additional structural robustness S-58
S.3.8.1 Interventions impose additional psychological costs S-58
S.3.8.2 Offsetting does not completely remove perceived guilt S-59
S.3.9 Relation between environmental attitudes and baseline meat consumption S-61
S.4 Experiments 1 and 3: Details on the experimental set-up S-63
S.4.1 Pre-registration S-63
S.4.2 Meals used for elicitation S-63
S.4.3 Example of willingness-to-pay elicitation (baseline decision) S-65
S.4.4 Incentivization of elicitations S-65
S.4.5 Decisions under carbon offsetting S-67
S.4.6 Experiment 3 guessing screen example S-69
S.4.7 Experiment screens (English translation from German) S-70
S.5 Experiment 2: Details on the experimental set-up S-88
S.5.1 Pre-registration S-88
S.5.2 Canteen set-up in Bonn S-88
S.5.3 Canteen visiting patterns S-89
S.5.4 Carbon label calculation S-91
S.5.5 Data set construction: Full sample S-92
S.5.6 Data set construction: Intention-to-treat sample S-94
S.5.7 Descriptive statistics on meat consumption and average emissions . S-98
S.5.8 Survey accompanying natural field experiment S-100
S.5.9 Post-experiment label installation S-101
References for the Supplementary Material S-102

S.1 Experiments 1 and 3: Further tables and figures

S.1.1 Randomization checks

Tables S.1 and S.2 report baseline characteristics across the randomized treatment sequences. In Experiment 1, the Control–Control arm has a lower share of non-vegetarians than the Label–Offset arm (65.7% versus 80.4%). A joint test of the treatment-sequence indicators across gender, student status, employment, diet, and hunger does not reject equality across treatment sequences ($p = .481$); a complete-case test additionally including age yields $p = .307$. For Experiment 3, the corresponding p-values are .989 and .947. These joint tests do not establish exact equality, but they do not reject balance across the measured baseline characteristics. Because diet and student status are substantively relevant to baseline meal preferences and the scope for treatment response, Section S.1.7 also reports descriptive split-sample estimates for these groups.

Table S.1. Randomization Experiment 1

	Average value					
	(1)	(2)	(3)	(4)	(5)	(6)
	Age	Male	Student	Working status	Non-vegetarian	Hunger (1–10)
Control, then Control	-0.59	-0.00	0.08	0.05	-0.15	0.02
	(1.09)	(0.07)	(0.06)	(0.07)	(0.06)	(0.38)
Control, then Label	-0.80	-0.01	0.01	0.10	-0.08	-0.05
	(1.08)	(0.07)	(0.06)	(0.07)	(0.06)	(0.38)
Constant	24.60	0.33	0.78	0.58	0.80	5.16
	(0.62)	(0.04)	(0.03)	(0.04)	(0.04)	(0.21)
Control, then Control	60	70	70	70	70	70
Control, then Label	62	69	69	69	69	69
Label, then Offset	126	148	148	148	148	148
Observations	248	287	287	287	287	287

Standard errors in parentheses

Notes: Coefficients report differences relative to the Label–Offset sequence. No individual significance tests are highlighted. A joint test of both treatment-sequence indicators across gender, student status, employment, non-vegetarian status, and hunger yields $p = .481$. The corresponding complete-case test additionally including age yields $p = .307$. Age is missing for participants whose reported values were excluded as implausible. Summary statistics are reported in Table S.3.

Table S.2. Randomization Experiment 3

	Average value					
	(1)	(2)	(3)	(4)	(5)	(6)
	Age	Male	Student	Working status	Non-vegetarian	Hunger (1–10)
Attention+Offset, then Attention+Labels	0.04	-0.01	-0.00	0.00	0.03	0.27
	(0.88)	(0.06)	(0.05)	(0.05)	(0.05)	(0.29)
Attention+Labels, then Attention+Offset	-0.53	0.02	0.01	-0.04	0.04	0.10
	(0.89)	(0.06)	(0.05)	(0.05)	(0.05)	(0.30)
Constant	25.93	0.45	0.69	0.75	0.74	4.73
	(0.63)	(0.04)	(0.04)	(0.04)	(0.03)	(0.21)
Attention, then Attention	124	151	151	151	151	151
Attention+Label, then Attention+Offset	126	144	144	144	144	144
Attention+Offset, then Attention+Label	131	149	149	149	149	149
Observations	381	444	444	444	444	444

Standard errors in parentheses

Notes: Coefficients report differences relative to the Attention-then-Attention sequence. No individual significance tests are highlighted. A joint test of both treatment-sequence indicators across gender, student status, employment, non-vegetarian status, and hunger yields $p = .989$. The corresponding complete-case test additionally including age yields $p = .947$. Age is missing for participants whose reported values were excluded as implausible. Summary statistics are reported in Table S.4.

S.1.2 Representativeness of the sample

Tables S.3 and S.4 report descriptive statistics for Experiments 1 and 3. Table S.5 reports descriptive statistics elicited in a survey among student canteen guests, as described in section S.5.8. I use these survey data to assess the similarity of participants in Experiments 1 and 3 to the relevant student canteen guest population. In terms of age, participants in Experiments 1 and 3 are slightly older than the student canteen guests in my survey (average age of 24 and 26 vs. an average age of 23 in the survey). The proportion of males is slightly lower in Experiment 1 (33%) and slightly higher in Experiment 3 (45%) than in the survey (43%). The proportion of students is higher in the survey (95%) than in Experiments 1 and 3 (80% and 69%). However, it is likely that my survey over-proportionally surveyed student canteen guests who are students. In the student canteen purchase data analyzed in Experiment 2, 12% of guests paying with an individualized payment card are employees, 86% are students and 2% are non-student and non-employee.[75] 78% of guests in Experiment 2's ITT sample purchase

75. This is the only demographic characteristic reported in the student canteen purchase data. I thus rely on the survey data for the other characteristics.

meat at least once, while 22% purchase only vegetarian meals (Table S.6). These shares are close to the 75% and 76% of participants who report being non-vegetarian in Experiments 1 and 3. The measures are not identical, however: the experimental measure is self-reported dietary status, whereas the field measure classifies guests from their observed purchases during the study period.

The clearest difference between the experimental and comparison samples therefore concerns student status. Dietary composition appears similar when assessed using purchasing behavior, although exact comparability is limited by the different measures. Section S.1.7 repeats the main analyses by vegetarian and student status. These split-sample estimates are descriptive robustness checks. Their point estimates provide information about whether the main patterns appear in the respective subsamples, but differences in within-group significance do not constitute formal tests of treatment-effect differences across groups.

Table S.3. Socio-economic summary statistics for Experiment 1

Variable	Explanation	Mean	Std. Dev.
Age	Age of participant	24.27	6.9
Male	Dummy: 1 if participant is a man	0.33	–
Student	Dummy: 1 if participant is a student	0.80	–
Working	Dummy: 1 if participant is working in some form	0.62	–
Non-vegetarian	Dummy: 1 if participant eats meat	0.75	–
Hungry	Hunger on scale of 1 to 10 beginning experiment	5.16	2.58
N	287		

Notes: Table shows average socio-economic summary statistics for participants of Experiment 1.

Table S.4. Socio-economic summary statistics for Experiment 3

Variable	Explanation	Mean	Std. Dev.
Age	Age of participant	25.77	7.02
Male	Dummy: 1 if participant is a man	0.45	–
Student	Dummy: 1 if participant is a student	0.69	–
Working	Dummy: 1 if participant is working in some form	0.74	–
Non-vegetarian	Dummy: 1 if participant eats meat	0.76	–
Hungry	Hunger on scale of 1 to 10 beginning experiment	4.85	2.54
N	444		

Notes: Table shows average socio-economic summary statistics for participants of Experiment 3.

Table S.5. Socio-economic summary statistics for student canteen guests - survey data

Variable	Explanation	Mean	Std. Dev.
Age	Age of participant	22.88	3.87
Male	Dummy: 1 if participant is a man	0.43	0.50
Student	Dummy: 1 if guest is a student	0.95	0.23
N	1,439		

Notes: Statistics are based on responses to the first survey completed before the labels were introduced. I include only survey respondents who visited a student canteen at least once in the 14-week study period and paid with their individual payment cards. See S.5.8 for details on the survey design. To preserve anonymity (since I also asked these survey participants about their study field), I elicited age in intervals and assigned a representative value to each interval to estimate mean age. For 15% of respondents, I have the information that they are below 20, which I represent as age 18. For 54% of respondents, I have the information that they are between 20 and 23 (represented as 21.5), 21% are between 24 and 27 (represented as 25.5), 6% are between 28 and 31 (represented as 30), and 4% are 32 or older (represented as 35). Age is available for 1,437 respondents and gender for 1,413 respondents; 26 respondents either did not want to provide their gender or identified as non-binary.

Table S.6. Socio-economic summary statistics for student canteen guests - consumption data

Variable	Explanation	Mean	Std. Dev.
Student	Dummy: 1 if guest is a student	0.87	0.34
Non-vegetarian	Dummy: 1 if guest eats meat	0.78	0.41
N	1,918		

Notes: Statistics use the 1,918 guests and 37,284 transactions in Experiment 2's ITT sample. The sample includes individual payment-card guests with at least five pre-intervention visits and a clear pre-intervention canteen assignment, and applies the same transaction and study-period restrictions as Table OA.17, Panel B, Column (1). A guest is classified as non-vegetarian if they purchase meat at least once in this sample; guests whose observed purchases are all vegetarian provide the vegetarian population share used in the canteen-offer-weighted robustness in Online Appendix OA.1.2. Ukrainian refugees receiving meals for free are excluded.

S.1.3 Descriptive statistics on baseline willingness to pay for meals

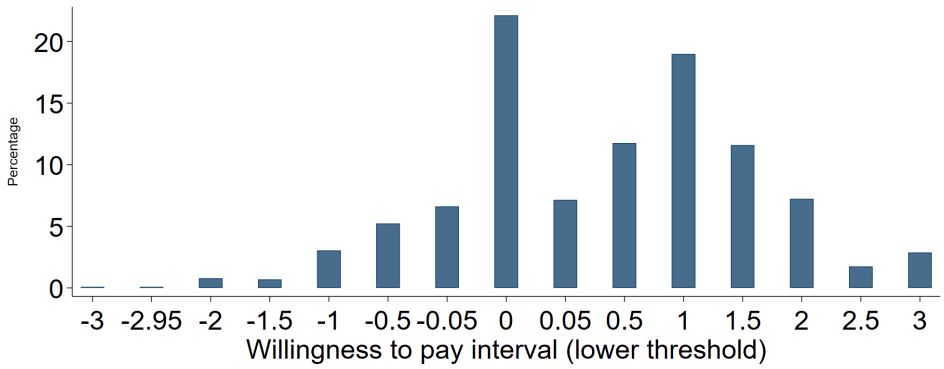


Figure S.1. WTP indicated for meals in the baseline purchase decisions in Experiment 1
Notes: $N = 1,148$ (287 participants making 4 baseline decisions each).

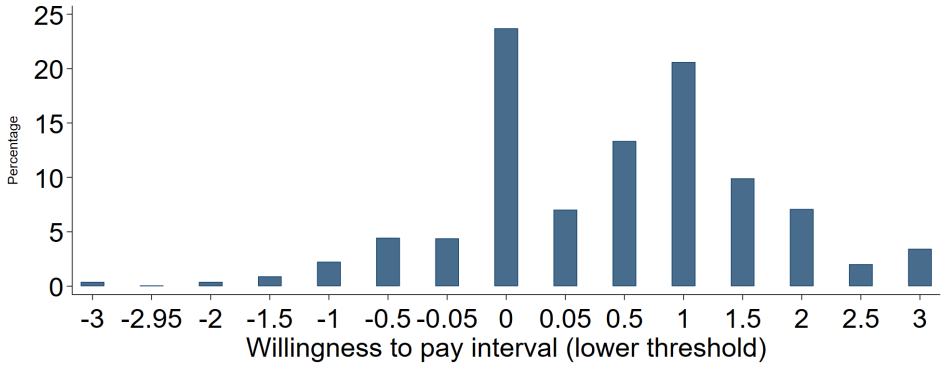


Figure S.2. WTP indicated for meals in the baseline purchase decisions in Experiment 3
Notes: $N = 1,776$ (444 participants making 4 baseline decisions each).

S.1.4 Test for effect of labels on meal perception

Table S.7. Test for effect of labels on perception of calories

	Guess of calories in				
	(1) Meat low	(2) Veg high	(3) Meat high	(4) Sandwich	(5) Veg low
Sees carbon labels	89.49	127.61	25.29	21.27	83.28
	(85.02)	(79.35)	(48.74)	(20.48)	(71.59)
Constant	608.00***	510.07***	708.04***	275.23***	521.23***
	(37.10)	(47.64)	(42.79)	(16.40)	(36.57)
Participants control	70	70	70	70	70
Participants treated	217	217	217	217	217
Observations	287	287	287	287	287

Notes: The analysis tests whether Experiment 1 participants who saw carbon labels perceived the nutritional quality of the experiment meals differently, using end-of-experiment calorie guesses as a proxy. Heteroskedasticity-robust standard errors are reported. See Section S.4.7 for the guessing screens.

S.1.5 Distribution of individual label responses in Experiment 1

For each of the 217 participants who experienced carbon labels in Experiment 1, I estimate a participant-specific no-intercept analogue of the continuous-emissions specification in Equation OA.7 and Table OA.6, using the participant's four label-round observations. I also construct control-adjusted coefficients by first subtracting, meal by meal, the no-label change in WTP predicted from the 139 control participants. The prediction includes a constant, the emissions-gap slope, and a third-round term. These estimates are best interpreted as participant-level label-response coefficients rather than individual causal effects because each treated participant's no-label counterfactual is predicted from average control-group behavior.

The estimates suggest that most participants react to the labels, but that there is heterogeneity in the strength of their response. As Table S.8 shows, 78.8% of control-adjusted coefficients are negative, 51.6% are −0.05 or stronger, and 29.0% are −0.30 or stronger. Without the control adjustment, 64.1% are negative; the shares at the stronger thresholds are very similar.

Figure S.3 shows the distribution of the unadjusted coefficients. It contains 26 participants with an exactly zero coefficient: none changed their WTP for any of the four meals between baseline and the label round. Of these 26 participants, 16 indicated a

Table S.8. Shares of participants with label-response coefficients at alternative thresholds

Estimated coefficient	Control-adjusted (%)	Unadjusted (%)
Negative	78.8	64.1
−0.05 or stronger	51.6	50.7
−0.10 or stronger	44.2	42.9
−0.15 or stronger	38.7	36.9
−0.20 or stronger	34.1	33.2
−0.25 or stronger	31.8	31.3
−0.30 or stronger	29.0	29.0

Notes: I estimate a participant-specific no-intercept analogue of the continuous-emissions specification in Equation OA.7 and Table OA.6. Each estimated participant-specific coefficient is the fitted change in WTP, in euros, associated with a one-kg increase in the meal's emissions relative to the sandwich. Coefficients are summarized in this table. Entries are percentages of the 217 participants shown carbon labels for whom I estimate a coefficient within a certain range. Rows are cumulative: for example, "−0.10 or stronger" includes every coefficient less than or equal to −0.10. More negative coefficients indicate a stronger emissions-related label response. The control-adjusted coefficients subtract predicted no-label behavior estimated using observations from the 139 control participants, including a constant, an emissions-gap slope, and a third-round term.

WTP of zero for all four meals in all three rounds, suggesting that at least some of these participants may have been disengaged from the experiment.

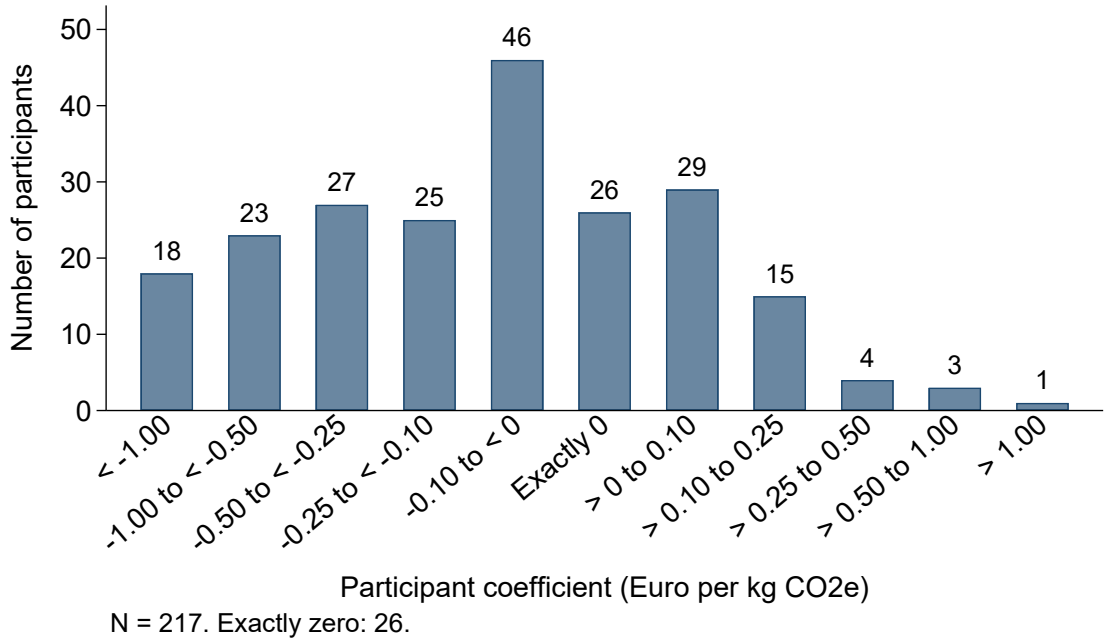


Figure S.3. Distribution of individual label-response coefficients in Experiment 1

Notes: Unadjusted participant-level no-intercept coefficients are estimated as in Table S.8. Categories are mutually exclusive, use the same cutoffs on either side of zero, and show exactly zero separately. The distribution contains all 217 participants who experienced carbon labels in Experiment 1.

S.1.6 Preregistered main effects in Experiments 1 and 3

Experiment 1 was preregistered under #AEARCTR-0008435, and Experiment 3 under #AEARCTR-0007858.

For Experiment 1, besides the analysis shown in the main text, I preregistered an analysis pooling all data from Experiments 1 and 3. Results are in line with those described in the main text and included here for completeness.

Table S.9 pools all data from Experiments 1 and 3. The baseline condition is the Control condition from Experiment 1, and rows 1 and 2 show that there is generally no effect of asking participants twice for their WTP, regardless of low- or high-emission meal. Rows 3 and 4 capture effects of participants seeing a carbon label, regardless of whether they are in a Label treatment, Attention+Label treatment, or an offset condition. I consider the offsetting notice participants in the offset conditions see to also be a carbon label and thus include them in this category. Rows 5 and 6 pick up differential effects of the offsetting treatments relative to the other labeling treatments. This is an increase in WTP for high-emission meals, in line with the labels removing some environmental guilt. Rows 7 and 8 pick up effects of being in an Attention condition (including Attention+Label and Attention+Offset), which decreases WTP for high-emission meals. Rows 9 and 10 pick up effects of being in an Attention condition *when any label is also present* (this includes both Attention+Label and Attention+Offset). These terms test whether "attention + label" is more or less than the sum of the separate attention and label effects. Rows 11 and 12 then capture any additional differential effect specific to the Attention+Offset treatment relative to Offset.

Table S.9. Experiments 1 and 3: Comparison of treatment effects

	Change in WTP compared to baseline
	(1)
Low	−0.05*
	(0.03)
High	0.02
	(0.02)
Low x Label	0.14***
	(0.04)
High x Label	−0.31***
	(0.05)
Low x Label x Offset	−0.03
	(0.04)
High x Label x Offset	0.23***
	(0.05)
Low x Attent	0.04
	(0.04)
High x Attent	−0.10***
	(0.03)
Low x Attent x Label	−0.16***
	(0.06)
High x Attent x Label	0.21***
	(0.06)
Low x Attent x Label x Offset	0.05
	(0.04)
High x Attent x Label x Offset	−0.07
	(0.05)
Control for third round	0.00
	(0.01)
Participants control	139
Participants label	217
Participants offset	148
Participants attent	151
Participants attent+offset	293
Participants attent+label	293
Observations	5,848

Notes: This pooled analysis combines Experiments 1 and 3 and was preregistered as part of Experiment 1. The dependent variable is the within-subject change in WTP relative to baseline. Standard errors are clustered at the individual level.

Tables S.10 and S.11 show the additional analyses I preregistered for Experiment 3. I preregistered to examine WTP for meals as the dependent variable, while including participant × meal fixed effects. As shown in section S.1.8, this approach is closely related to the change-in-WTP specification and produces similar estimates, although the two approaches weight baseline observations differently. The main text reports the change-

in-WTP specification because it is easier to interpret. Col. (1) of Table S.10 directly examines the effect of providing labels additionally to directing attention to carbon emissions, and the effect of offsetting relative to directing attention and providing labels. Similarly, Col. (2) performs a similar analysis interacting the emissions of each meal with treatments rather than using the Low and High indicators. Table S.11 further examines the effect of carbon offsetting relative to making participants attentive, excluding data from the ATTENTION+LABEL condition. Col. (2) examines the effect of directing attention and offsetting as a function of emissions guessed by participants.

Table S.10. Experiment 3: Analysis label and offsetting effects

	WTP	
	(1)	(2)
Low x Attent x Label	−0.02	
	(0.04)	
High x Attent x Label	−0.10***	
	(0.04)	
Low x Attent x Label x Offset	0.02	
	(0.02)	
High x Attent x Label x Offset	0.16***	
	(0.02)	
Low x Attent	−0.01	
	(0.03)	
High x Attent	−0.08***	
	(0.03)	
Emissions(kg) x Attent x Label		−0.01
		(0.03)
Emissions(kg) x Attent x Label x Offset		0.07***
		(0.02)
Emissions(kg) x Attent		−0.07***
		(0.02)
Attent x Label		−0.07**
		(0.03)
Attent x Label x Offset		0.06***
		(0.02)
Attent		−0.00
		(0.02)
Control for third round	−0.00	−0.00
	(0.01)	(0.01)
Constant	0.70***	0.70***
	(0.01)	(0.01)
Participant x Meal FE	Yes	Yes
Participants attent	151	151
Participants attent+offset	293	293
Participants attent+label	293	293
Observations	5,328	5,328

Notes: This Experiment 3 analysis was preregistered in Schulze Tilling (2021). The dependent variable is WTP for a meal and both columns include participant-by-meal fixed effects. Standard errors are clustered at the individual level.

Table S.11. Experiment 3: Analysis offsetting effects based on participants' emission guesses

	WTP	
	(1)	(2)
Low x Attent x Label x Offset	−0.07	
	(0.07)	
High x Attent x Label x Offset	0.07**	
	(0.03)	
Low x Attent	0.04	
	(0.05)	
High x Attent	−0.07***	
	(0.02)	
Guessed emissions(kg) x Attent x Label x Offset		0.06*
		(0.03)
Guessed emissions(kg) x Attent		−0.08***
		(0.02)
Attent x Label x Offset		−0.00
		(0.03)
Attent		0.01
		(0.02)
Control for third round	−0.01	−0.01
	(0.02)	(0.02)
Participant x Meal FE	Yes	Yes
Participants attent	151	146
Participants attent+offset	293	284
Observations	4,156	3,757

Notes: This Experiment 3 analysis was preregistered in Schulze Tilling (2021) and excludes the Attention+Label condition. Low- and high-emission categories and the continuous differences use participants' emission guesses. For each meal, Col. (2) excludes observations in the highest decile of the absolute error in the guessed meal–sandwich emissions difference, evaluated using the same true emissions difference in every round. Standard errors are clustered at the individual level.

S.1.7 Results split by (non-) vegetarians and (non-) students

The tables in this subsection report descriptive split-sample estimates by vegetarian and student status. The point estimates provide information about the robustness of the main patterns in these subsamples, but the tables are not designed as formal tests of treatment-effect differences across groups.

Experiment 1.

Table S.12. Experiment 1 analyses by vegetarian status

	Change in WTP compared to baseline	
	(1) Non-vegetarian	(2) Vegetarian
High emission meal x Shown label	-0.26***	-0.53***
	(0.05)	(0.11)
Low emission meal x Shown label	0.17***	0.11
	(0.06)	(0.07)
High emission meal	0.00	0.06
	(0.02)	(0.05)
Low emission meal	-0.10**	-0.02
	(0.05)	(0.04)
Control for third round	0.01	0.04
	(0.04)	(0.04)
Participants control	96	43
Participants treated	169	48
Observations	1,244	460

Notes: Each column repeats the specification in Table OA.2, Column (1), for the indicated subgroup, including the third-round indicator. Standard errors are clustered at the individual level.

Table S.13. Experiment 1 analyses by student status

	Change in WTP compared to baseline	
	(1) Student	(2) Non-student
High emission meal x Shown label	-0.29***	-0.41***
	(0.05)	(0.09)
Low emission meal x Shown label	0.15***	0.03
	(0.05)	(0.07)
High emission meal	-0.01	0.12**
	(0.02)	(0.06)
Low emission meal	-0.08**	0.08
	(0.03)	(0.07)
Control for third round	0.01	0.05
	(0.03)	(0.09)
Participants control	114	25
Participants treated	169	48
Observations	1,372	332

Notes: Each column repeats the specification in Table OA.2, Column (1), for the indicated subgroup, including the third-round indicator. Standard errors are clustered at the individual level.

Experiment 3.

Table S.14. Main under-/overestimation specification by vegetarian status

	Change in WTP compared to baseline	
	(1) Non-vegetarian	(2) Vegetarian
Underestimated × Shown label	−0.23***	−0.23***
	(0.04)	(0.08)
Overestimated × Shown label	−0.12***	−0.02
	(0.04)	(0.09)
Control for third round	0.05	0.05
	(0.05)	(0.10)
Participants	227	66
Obs. underestimate	451	104
Obs. overestimate	418	144
Observations	869	248

Notes: This table repeats the categorical specification from Table OA.7 by vegetarian status. Observations with exactly correct rounded emissions-gap guesses are excluded. Stars use ordinary two-sided *p*-values. Standard errors are clustered at the individual level.

Table S.15. Main under-/overestimation specification by student status

	Change in WTP compared to baseline	
	(1) Student	(2) Non-student
Underestimated × Shown label	−0.30***	−0.05
	(0.05)	(0.05)
Overestimated × Shown label	−0.12**	−0.05
	(0.05)	(0.06)
Control for third round	0.10*	−0.06
	(0.05)	(0.08)
Participants	203	90
Obs. underestimate	383	172
Obs. overestimate	391	171
Observations	774	343

Notes: This table repeats the categorical specification from Table OA.7 by student status. Observations with exactly correct rounded emissions-gap guesses are excluded. Stars use ordinary two-sided *p*-values. Standard errors are clustered at the individual level.

Table S.16. Main attention-versus-label specification by vegetarian status

	Change in WTP compared to baseline	
	(1) Non-vegetarian	(2) Vegetarian
Low emissions × Attention	−0.01	−0.04
	(0.04)	(0.04)
Low emissions × Attention+Label	−0.07	−0.01
	(0.04)	(0.05)
High emissions × Attention	−0.11***	−0.05
	(0.03)	(0.04)
High emissions × Attention+Label	−0.21***	−0.18**
	(0.03)	(0.08)
Participants attention	112	39
Participants label	227	66
Observations	1,804	576

Notes: This table repeats the categorical specification from Table OA.8 by vegetarian status, including the third-round indicator. Standard errors are clustered at the individual level.

Table S.17. Main attention-versus-label specification by student status

	Change in WTP compared to baseline	
	(1) Student	(2) Non-student
Low emissions × Attention	−0.03	−0.00
	(0.03)	(0.06)
Low emissions × Attention+Label	−0.05	−0.03
	(0.04)	(0.05)
High emissions × Attention	−0.08***	−0.14**
	(0.03)	(0.06)
High emissions × Attention+Label	−0.24***	−0.10**
	(0.04)	(0.05)
Participants attention	104	47
Participants label	203	90
Observations	1,644	736

Notes: This table repeats the categorical specification from Table OA.8 by student status, including the third-round indicator. Standard errors are clustered at the individual level.

S.1.8 Experiment 1: Alternative econometric specifications

Alternatively to the estimation approach described in Section 2.2, one could instead estimate the following specification:

$$WTP_{ijm} = \alpha_{im} + \beta_1(High_m \times Post_j) + \beta_2(Low_m \times Post_j) + \delta_1(High_m \times Post_j \times Label_{ij}) + \delta_2(Low_m \times Post_j \times Label_{ij}) + \rho ThirdRound_j + \varepsilon_{ijm} \quad \text{(S.1)}$$

This specification is more similar to a classic diff-in-diff approach. Instead of directly using the difference between indicated WTP for a meal and baseline WTP as the dependent variable (as in Equation 1), I use raw WTP of individual i in round j for meal m as the dependent variable. Accordingly, I also include observations from the baseline elicitation round in the regression.

α_{im} are individual and meal-specific fixed effects. These are 1,148 fixed effects in total: 287 participants × 4 meals. These fixed effects control for individual-specific baseline tastes. I use participant-by-meal rather than participant fixed effects because baseline valuations vary at the participant–meal level. Participant-by-meal fixed effects absorb each participant's time-invariant valuation of each meal, so the treatment coefficients are identified from within-participant, within-meal changes across rounds. A

participant fixed effect alone would absorb only the participant's average WTP across meals and would leave persistent participant–meal taste differences in the residual.

($High_m \times Post_j$) is an indicator variable for whether the meal causes higher emissions than the sandwich, interacted with the elicitation round $j > 1$, i.e. it being the second or third round of elicitations and not the baseline round. ($Low_m \times Post_j$) is the equivalent indicator for low-emission meals. Note that all meals are classified either as Low_m or $High_m$. The two variables thus together capture the $Post_j$ effect, and a separate $Post_j$ indicator would be dropped due to collinearity. I also do not include separate controls for Low_m and $High_m$ since meal characteristics are captured by the α_{im} fixed effects.

($High_m \times Post_j \times Label_{ij}$) interacts the high-emission and $Post_j$ indicator with an indicator for whether individual i saw carbon labels in round j. This describes the average causal effect of carbon labels on WTP for a meal that is high in carbon emissions. ($Low_m \times Post_j \times Label_{ij}$) describes the average causal effect of carbon labels on WTP for a meal that is low in carbon emissions. $ThirdRound_j$ is an indicator of whether it was the third round of decisions. Standard errors are clustered at the individual level.

Table S.18 shows regression results. They are very similar to those reported in the main text and Table OA.2.

Table S.18. Replication of Experiment 1 results with fixed effects

	WTP
	(1)
High x Post x Label	−0.30***
	(0.04)
Low x Post x Label	0.12***
	(0.04)
High x Post	0.01
	(0.02)
Low x Post	−0.05*
	(0.03)
Control for third round	0.01
	(0.03)
Constant	0.65***
	(0.01)
Participant x Meal FE	Yes
Participants control	139
Participants treated	217
Observations	2,852

Notes: The specification replicates the high- and low-emission analysis in Table OA.2, using WTP directly and participant-by-meal fixed effects. The small coefficient differences arise because this specification uses each participant-meal baseline observation once, whereas the change-score specification uses that baseline for each subsequent comparison. Standard errors are clustered at the individual level.

S.1.9 Experiment 1: Intuition behind expressing effect sizes in terms of a carbon tax

One of the main results in Section 2.2 is that carbon labels in Experiment 1 produce an emissions change similar to that generated by emissions-based relative-price adjustments in the €146–164 per tonne of CO_2e matching set, summarized by its €155/t midpoint. The comparison is fundamentally a consumer-side one. For a warm meal m relative to the cheese sandwich o, a carbon tax of τ changes the meal's relative net valuation by $-\tau(e_m - e_o)$. Carbon labels instead change the participant's WTP directly. Because both changes can be expressed on the same WTP scale, I can ask which tax-induced change in relative valuations most closely reproduces the emissions change generated by the labels.

Figure S.4 provides a stylized partial-equilibrium illustration of this comparison. A carbon tax can be drawn in two equivalent ways. In the common textbook representation, the tax is shown as an upward shift of the supply curve by the amount of the tax. Equivalently, one can hold the supply curve fixed and subtract the tax from consumers' willingness to pay, shifting the effective demand curve downward. The two representations impose the same wedge between the price paid by consumers and the price received by sellers and imply the same equilibrium quantity. I use the demand-side representation in the figure because it maps most directly onto the WTP measure observed in the experiment, making the comparison with the label-induced change in WTP easier to visualize. This is only a graphical representation; using the equivalent supply-side representation would leave the economic implications unchanged. The demand curve shift should not be interpreted as saying that a tax changes consumers' underlying tastes.

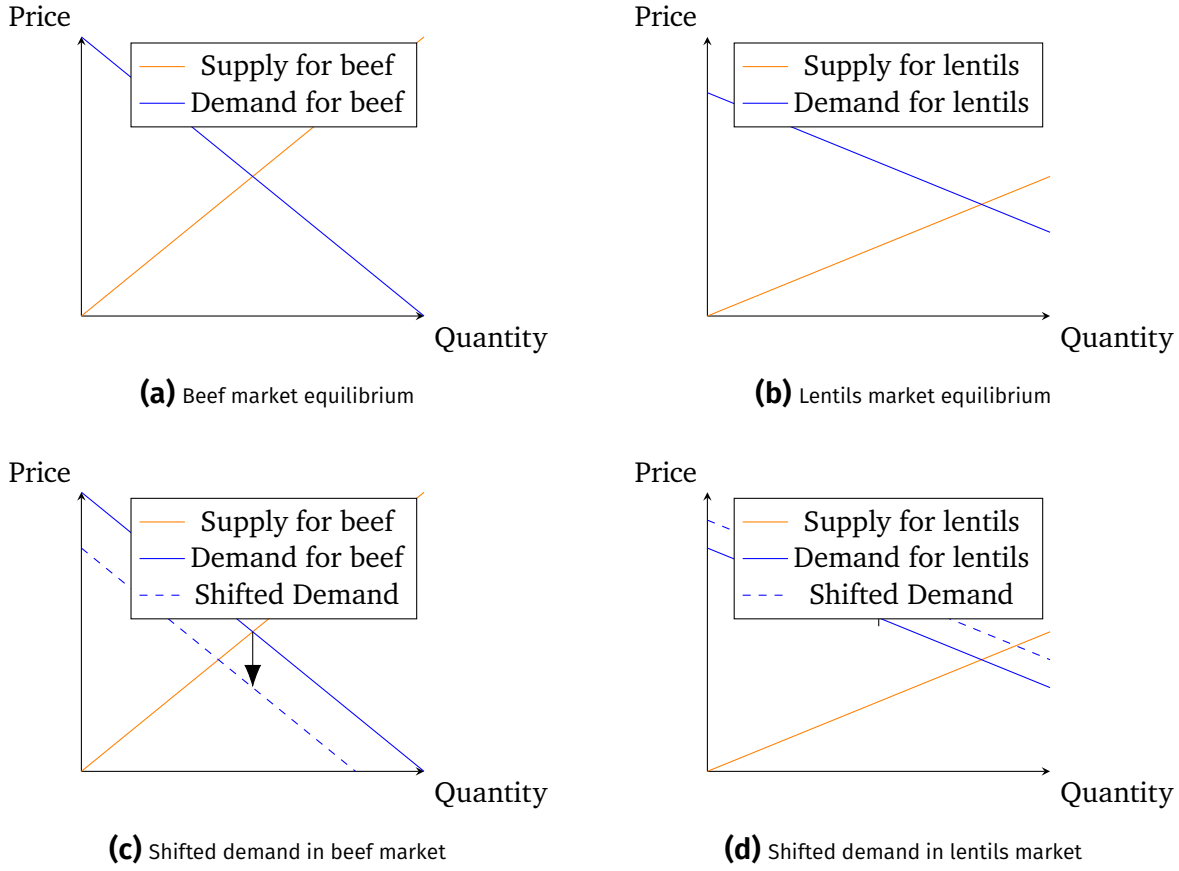
Panels (a) and (b) show stylized initial equilibria for a higher-emission good such as beef and a lower-emission good such as lentils. In Panel (c), carbon labels can reduce consumers' WTP for beef, shifting demand downward. A carbon tax can generate the same consumer-side shift in effective demand: for any given price received by sellers, consumers are willing to pay that price plus the tax only if their pre-tax WTP is sufficiently high. Both representations therefore reduce the quantity demanded.

For a lower-emission good, the relevant effect can go in the opposite direction relative to higher-emission substitutes. Carbon labels may increase its attractiveness, while an emissions-based tax raises the prices of higher-emission alternatives by more

and thereby shifts relative demand toward the lower-emission good. Thus, the lower-emission product need not be untaxed for its demand to increase; what matters for substitution is the change in relative prices across products.

Importantly, the demand shifts in Figure S.4 do not imply that the supply curve itself shifts. Holding the supply schedule fixed, a change in demand moves the market along the existing supply curve to a new equilibrium quantity. A shift of the supply curve would instead require a change in production costs, technology, capacity, or another determinant of supply. Such supply-side adjustments may matter over longer horizons, but they are outside the local short-run benchmark studied here and are discussed separately in Supplementary Material S.3.

The actual benchmark in Section 2.2 does not require estimating a supply curve or solving for a market equilibrium. I hold the canteen prices and available meals fixed and compare the choices implied by participants' WTP under labels with the choices implied by baseline WTP after applying different emissions-based relative-price adjustments. The tax benchmark is the rate for which these two different changes in consumer-side valuations generate the closest emissions effect.



(a) Beef market equilibrium

(b) Lentils market equilibrium

(c) Shifted demand in beef market

(d) Shifted demand in lentils market

Figure S.4. Comparison of supply and demand in beef and lentils markets

Notes: The figure uses a demand-side representation of a carbon tax: the tax is subtracted from consumers' willingness to pay while the supply curve is held fixed. An equivalent textbook representation shifts the supply curve upward by the tax wedge. The two representations imply the same equilibrium quantity and the same wedge between the price paid by consumers and the price received by sellers. A shift in demand causes movement along, not a shift of, the supply curve.

S.1.10 Experiment 3: Descriptives on under- and over-estimation

Table S.19. Under- and over-estimation of meal emissions

Meal	Relative emissions (kg CO_2e)	Under	Over	Correct	Total
Vegetable ragout	-0.2	31	249	13	293
Poultry schnitzel	0.7	47	163	17	227
Filled zucchini	0.7	249	33	11	293
Cheese Spätzle	0.5	31	24	11	66
Sliced beef	2.7	193	32	2	227
Vegetable stir-fry	-0.3	4	61	1	66
Total		555	562	55	1,172

Notes: Based on participants in the Attention+Label treatment. I show under- and overestimation for those meals that are used in the experiment decisions: In each decision round, non-vegetarians make decisions on the vegetable ragout, poultry schnitzel, filled zucchini, and sliced beef; vegetarians make decisions on the vegetable ragout, filled zucchini, cheese Spätzle, and vegetable stir-fry. Each participant therefore makes four meal decisions. Relative emissions are measured against the cheese sandwich (0.7 kg CO_2e). Under the high- versus low-emissions classification used in the main analyses, the vegetable ragout and vegetable stir-fry are low-emission meals because they emit less than the sandwich; the other four meals are high-emission meals. I classify a participant as underestimating relative emissions if their guessed meal-minus-sandwich emissions difference is lower than the actual difference, and as overestimating relative emissions if it is higher.

Table S.20. Number of under- and over-estimations per participant

Number overestimated	0	1	2	3	4	Total
Number underestimated						
0	0	0	0	2	10	12
1	0	1	21	54	0	76
2	1	24	128	0	0	153
3	4	31	0	0	0	35
4	17	0	0	0	0	17
Total	22	56	149	56	10	293

Notes: Based on participants in the Attention+Label treatment. I show under- and overestimation of the emissions caused by those meals that are also used in the experiment decisions. Relative emissions are emissions relative to the cheese sandwich (0.7 kg). I classify a participant as underestimating this amount if their guess for the meal's emissions minus their guess for the cheese sandwich is lower than the actual relative emissions. I classify a participant as overestimating this amount if their guess for the meal's emissions minus their guess for the cheese sandwich is higher than the actual relative emissions. Each cell shows the number of participants with the respective number of under- or over-estimations.

Table S.21. Number of participants who correctly guessed how the four decision meals rank relative to each other

Number of correctly ranked meals	Number of participants
0	66
1	35
2	135
3	48
4	9
Total	293

Notes: Table shows how many participants put meals on the correct rank, i.e. they guessed the lowest emission amount to the lowest-emitting meal, the highest emission amount to the highest-emitting meal, etc. For meat-eating participants, the meals on places 2 and 3 are very similar, so I count a meal as correctly ranked regardless of whether the ranking resulting from the participant's guesses is rank 2 or 3.

Table S.22. Correlation between WTP for seeing carbon labels and individual characteristics

	WTP for the presence of carbon labels				
	(1)	(2)	(3)	(4)	(5)
Perceived strength of social norms	0.01**				
	(0.01)				
In favor of labels in student restaurant		0.03***			
		(0.01)			
Self-reported willingness to use info			0.03***		
			(0.01)		
Self-reported confidence in own knowledge				−0.01	
				(0.01)	
Eating self-control					0.00
					(0.01)
Constant	0.15***	−0.03	0.03	0.18***	0.21***
	(0.03)	(0.06)	(0.04)	(0.02)	(0.02)
Observations	731	731	731	731	731

Notes: Exploratory correlational analysis. The dependent variable is WTP for seeing labels for the final three consumption decisions. Perceived social norms follow Krupka and Weber (2013); eating self-control follows Haws, Davis, and Dholakia (2016).

S.1.11 Correlates of preferences for label presence

Online Appendix Section OA.2.2.8 reports the distribution of WTP for label presence and its differences across prior treatment sequences. The exploratory tables below examine how this elicited preference correlates with participant characteristics and with participant-specific label responses.

Table S.23. Correlation between WTP for seeing carbon labels and treatment effect

	WTP for labels (EUR)
	(1)
Individual label-induced WTP slope (EUR/kg)	−0.40***
	(0.09)
Constant	0.13***
	(0.02)
Participants Control, then Labels	60
Participants Labels, then Offset	133
Participants Attention+Offset, then Labels	137
Participants Attention+Label, then Offset	130
Observations	460

Notes: Exploratory participant-level correlational analysis pooling Experiments 1 and 3. The dependent variable is each participant's WTP, in euros, to see carbon labels during the final three consumption decisions. The explanatory variable is a participant-specific measure of how meal WTP changed under labels as emissions increased relative to the cheese sandwich. I construct this measure from the participant's four meal decisions by estimating, without an intercept, the relationship between the label-induced change in meal WTP and the meal's emissions difference relative to the sandwich. The measure is expressed in euros per kg CO_2e. A more negative value indicates that the participant reduced meal WTP more strongly as relative emissions increased. The negative estimate reported in the table therefore indicates that participants with a stronger negative meal-choice response were willing to pay more to see the labels. To limit the influence of extreme participant-specific treatment estimates, I retain participants whose values lie between the sample's 5th and 95th percentiles, inclusive. Section S.1.5 provides further details and reports the distribution of participant-specific treatment estimates. Heteroskedasticity-robust standard errors are reported.

S.2 Experiment 2: Further tables and figures

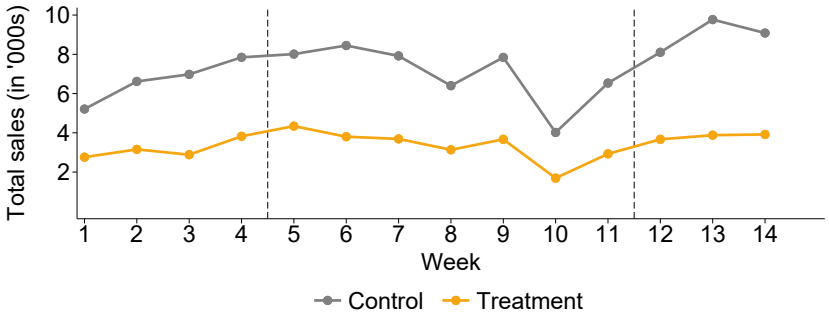


Figure S.5. Weekly student canteen sales of main meal components

Notes: Raw aggregate sales of main meal components, excluding sales to Ukrainian refugees; $N = 150{,}187$. Weeks 1–4 are the pre-intervention period (April 2022), weeks 5–11 are the intervention period (May to Mid-June 2022), and weeks 12–14 cover June 20–July 8, 2022. The drop in sales in week 10 is likely due to the one-week Pentecost holidays, during which no classes took place. Week 14 is excluded from the main analysis since it preceded an irregular two-week closure of the treatment canteen, and week 5 is excluded since it coincided with the Healthy Campus week. See main text for further details.

S.2.1 Label-effect inference under guest and date clustering

Table S.24 holds the ITT sample, outcome, controls, guest fixed effects, and point estimator fixed and changes only the covariance estimator. Two-way clustering by guest and calendar date allows residual dependence both across purchases by the same guest and across guests purchasing on the same date. The standard error rises from 0.010 under guest clustering to 0.012 under guest-and-date clustering; the corresponding 95% interval is $[-0.047, 0.003]$.

Table S.24. Label-effect inference under guest and calendar-date clustering

	Guest	Guest and date
Treatment × Label period	-0.022**	-0.022*
Standard error	0.010	0.012
95% interval	[-0.041, -0.003]	[-0.047, 0.003]
p-value	0.027	0.084
Observations	37,284	37,284
Guest clusters	1,918	1,918
Date clusters	–	54

Notes: Every column reports the treatment–label-period coefficient from the main ITT specification in Online Appendix Table OA.17, Panel B, Column (3). The estimation sample, outcome, controls, and guest fixed effects are identical across columns; only the covariance estimator changes. Column (1) clusters by guest, as in the main analysis. Column (2) clusters two ways by guest and calendar date. The latter permits arbitrary residual dependence across purchases by the same guest and across purchases made on the same date. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

S.2.2 Time trends

Table S.25. Regression coefficients for the event plot in Figure 5

	(1) Choice of meat
Treated × Week 1	2.105 (2.785)
Treated × Week 2	1.749 (2.450)
Treated × Week 3	-2.280 (2.467)
Treated × Week 5	-5.655** (2.569)
Treated × Week 6	-1.484 (2.246)
Treated × Week 7	-4.411* (2.425)
Treated × Week 8	-5.391** (2.668)
Treated × Week 9	-6.245** (2.589)
Treated × Week 10	-6.379 (3.911)
Treated × Week 11	-6.387** (2.770)
Treated × Week 12	-10.034*** (2.591)
Treated × Week 13	-7.610*** (2.801)
Guests control	1,422
Guests treated	496
Observations	40,952

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable is a 0/1 indicator for consumption of the meat option, multiplied by 100 so coefficients are in percentage points. The regression additionally includes week and day-of-week fixed effects; binary controls for the meat and vegetarian main meals assigned according to ITT classification; indicators for a second vegetarian or meat offer; and indicators for ITT assignment to the treated and smaller control canteens. Standard errors are clustered at the guest level.

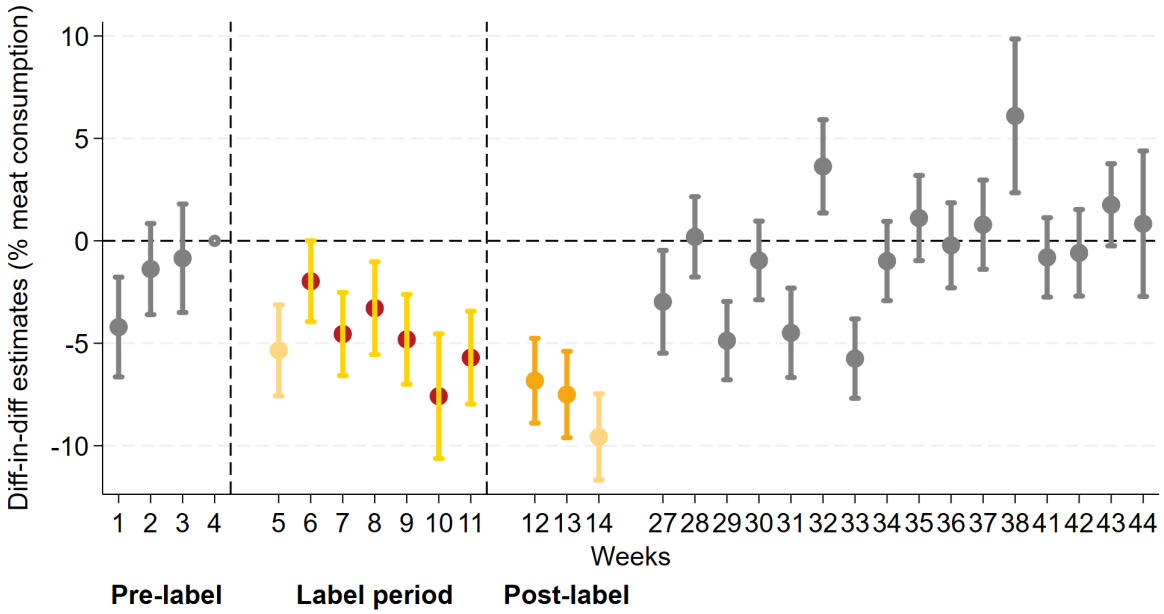


Figure S.6. Event study including data from the following semester

Notes: Difference in difference estimates of the likelihood of consuming the meat option (in percentage points), using week 4 of the pre-intervention phase as a baseline. Weeks 1–4 constitute the pre-intervention phase, while weeks 6–11 constitute the intervention phase, and weeks 12–14 the post-intervention phase. Week 14 is excluded from the main analysis since it preceded an irregular two-week closure of the treatment canteen, and week 5 is excluded since it coincided with the Healthy Campus week. See main text for further details. Weeks 27 onward pertain to the new semester. The regression specification closely follows the full-sample specification in Table OA.17, Panel A, Column (3). An ITT analysis as in the main text, Figure 5, is not possible using this data set, since individuals' anonymized ID numbers differ between the study period and the following semester. Weeks 15 to 26 are excluded due to the semester break. Weekly time controls, day-of-the-week controls, meat and vegetarian meal controls, and canteen controls are included. Bars indicate 95% confidence intervals.

S.2.3 Label-color outcomes

Table S.26. Choice among the two main offers, by label color

	Choice of meal with label color			
	Green	Yellow meat	Yellow veg.	Red
Treatment × Label period	0.021**	-0.037***	0.053***	-0.028***
	(0.009)	(0.012)	(0.012)	(0.009)
Treatment × Post period	0.025**	-0.066***	0.055***	-0.024**
	(0.012)	(0.015)	(0.015)	(0.012)
Observations	75,748	45,339	40,807	71,216

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable equals one when the purchased main meal has the color and meal type indicated by the column, and zero otherwise. Each column includes only purchases of the two main offers on canteen-days when the corresponding option is available; the main offers are the most frequently purchased vegetarian and meat meals. All specifications are linear probability models including indicators for the treated canteen, the label period, and the post period, treatment–period interactions, and indicators for the color of the other main offer. Yellow–yellow menus are retained in Columns (2) and (3). Coefficients are in probability units, so 0.01 corresponds to one percentage point. Standard errors are heteroskedasticity robust. Across canteen-days in the study sample, the main offers paired a green-labeled vegetarian and red-labeled meat meal (46%), a green-labeled vegetarian and yellow-labeled meat meal (15%), a yellow-labeled vegetarian and yellow-labeled meat meal (28%), or a yellow-labeled vegetarian and red-labeled meat meal (11%).

S.2.4 Guest numbers and canteen switching

(1) **The effect of the labels on guest numbers.** Figure S.5 shows that sales developed similarly in the treatment and control groups throughout the sample period. As an additional analysis, Table S.47 expands the ITT sample into a panel format by filling in zeros for days on which a given guest did not visit the canteen. I then repeat ITT analyses using a guest's decision to visit or not visit as the outcome variable. The coefficient during the labeling period is negative but insignificant, while the post-intervention coefficient is negative and significant. However, it is unlikely that the labeling intervention caused a decrease in canteen visits during the post period; this pattern more plausibly reflects end-of-semester differences in campus attendance, with treatment-canteen guests being less present during the final weeks.

Table S.27. ITT analysis restricted to guests' assigned home canteen

	Meat meal chosen			
	+ Offer	+ Week FE	+ Guest FE	Students
Treatment × Label period	-0.048***	-0.052***	-0.024**	-0.024**
	(0.014)	(0.015)	(0.010)	(0.011)
Treatment × Post period	-0.094***	-0.096***	-0.040***	-0.035**
	(0.020)	(0.020)	(0.014)	(0.015)
Offer controls	Yes	Yes	Yes	Yes
Week fixed effects	No	Yes	Yes	Yes
Guest fixed effects	No	No	Yes	Yes
Guests control	1,422	1,422	1,422	1,216
Guests treated	496	496	496	451
Observations	36,199	36,199	36,199	31,061

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: All specifications are estimated as linear probability models (LPMs). The dependent variable equals one when a guest chooses a meat meal and zero otherwise; coefficients are in probability units, so 0.01 corresponds to one percentage point. Columns (1)–(4) repeat Table OA.17, Panel B, Columns (1)–(4), retaining only purchases at the exact home canteen assigned from all of the guest's pre-intervention visits using the same assignment rule as the main ITT analysis. Column (1) includes offer controls and common label- and post-period indicators; Column (2) adds week and day-of-week fixed effects, which absorb these period indicators; Column (3) adds guest fixed effects; and Column (4) restricts the latter specification to student-card purchases. Standard errors are clustered at the guest level. Offer controls comprise indicators for the specific main vegetarian and meat meals and for whether a second vegetarian or meat option is available. None of these specifications includes date fixed effects.

S.2.5 Field survey results

Below I describe the results of surveys conducted in the control and treatment canteens pre- and post- intervention, as described in section S.5.8.

Did treatment canteen guests see the labels? Among the completed follow-up respondents, 371 made at least one card-recorded purchase of a vegetarian or meat main component at the treated student canteen during the label period. Of these, 260 (70.1%) report having seen the labels.

Did control canteen guests see the labels? Among the completed follow-up respondents, 422 made at least one card-recorded purchase of a vegetarian or meat main component at a control canteen, but not at the treatment canteen, during the label period. They might nevertheless have visited the treatment canteen without paying with their individual cards. Of these 422 respondents, 33 (7.8%) report having seen the labels.

Do canteen guests feel they reacted to the labels? Among the 260 completed follow-up respondents who made a card-recorded main-meal purchase at the treatment canteen during the label period and noticed the labels, 22.3% report not having incorporated the labels in their decisions, 34.6% report having incorporated them a little, 25.0% moderately, 14.6% relatively strongly, and 3.5% strongly.

Do canteen guests support the labels and a corresponding price policy? Among the 258 completed respondents who made a card-recorded main-meal purchase at the treatment canteen during the label period, noticed the labels, and answered both policy-support questions, 75.6% support permanent label installation, 16.7% are unsure, and 7.8% oppose it. A revenue-neutral price policy tied to the label colors receives 62.0% support, with 14.0% unsure and 24.0% opposed.

How do canteen guests make their consumption choices? Among completed respondents who answered this question, 34.4% report making their choice mainly using the information given on the canteen website, 31.3% mainly use the digital billboards, and 34.3% mainly decide by looking at the food counters. Figure 4 in the main text shows how the carbon labels were shown in each of these decision contexts.

S.2.6 Price elasticities

To ease comparison of the observed price effects with previous and concurrent literature, I calculate price elasticities (Table S.28). Related work uses household demand estimates to simulate emissions-based food taxes in the UK (Kehlbacher et al., 2016). Note that my analysis of price effects differs from conventional price elasticity estimation in two main ways:

- First, the variation in meat price I analyze is variation in the *price difference between the meat option and the vegetarian option*. Focusing on this relative price, rather than fluctuations in the absolute meat price, allows me to capture more directly the change in consumption behavior most relevant for a potential carbon tax. If the student canteen were to implement such a tax, the most straightforward way to do so would be to increase the price difference between meat and vegetarian dishes, rather than the overall price level.[76]
- Second, the resulting demand fluctuations I observe are changes in demand for the meat option *relative* to the vegetarian option. This outcome is directly analogous to my main outcome variable in the carbon-label analysis, facilitating a clean comparison between the effects of the labels and of a hypothetical tax.

These analysis choices make the comparison between carbon labels and a carbon tax more meaningful but may complicate direct comparison with elasticity estimates from the wider literature. Moreover, my elasticity estimates are based on the price of the meat meal component, but the ingredients for this component are rarely 100% meat. Filets, steaks, etc. are often filled with vegetables or cheese, sometimes breaded on the outside, and often served with sauce or garnish. My main pricing analysis takes this circumstance into account, since the carbon-tax simulations are based on the paired emissions difference between the canteens' main meat and vegetarian offers, which incorporates average meal composition and maps the tax into per-meal price changes on the relevant margin. However, this circumstance additionally complicates comparison with elasticities measured in a supermarket context and should be kept in mind in the following comparisons. I first compare my estimates to those reported by Roosen, Staudigel, and Rahbauer (2022), who use German household scanner data to estimate meat price elas-

76. The introduction of a linear carbon tax need not imply a general increase in the price level, since the canteen could simultaneously decrease the price of the vegetarian meal.

ticities. Their population is thus comparable to my canteen guest population. Roosen, Staudigel, and Rahbauer (2022) estimate uncompensated price elasticities of –0.861 (poultry), –0.966 (pork), and –0.960 (beef), and compensated elasticities of –0.771, –0.791, and –0.878, respectively. For low-income households, they report compensated elasticities of –0.591, –0.654, and –0.438. The pork elasticity in the student-canteen data (-0.7) lies between their compensated estimates for low-income households and the full sample. The beef elasticity (-1.1) is somewhat larger in absolute magnitude than their full-sample estimate, while the poultry elasticity (-0.3) is smaller in absolute magnitude. These comparisons indicate a broadly similar order of magnitude, although the canteen estimates concern substitution between a meat meal and its vegetarian alternative, whereas the household estimates concern quantities of meat purchased.

Liu, Ansink, and Bouma (2025), using store-level supermarket scanner data from the Netherlands, estimate uncompensated elasticities of –0.944 for poultry, –2.064 for pork, and –0.560 for beef.

I do not interpret a fish elasticity in my setting. Fish is almost exclusively served on Fridays and is more popular than other meats,[77] so its price interaction is difficult to interpret causally: it may capture residual Friday-specific demand, habitual Friday fish consumption,[78] or selection by pescetarians. I therefore retain the fish interaction as a regression control but assign fish a zero response in the carbon-tax calculations.

Lamb is observed as a single dish on 14 canteen-days across six dates. I retain a separate lamb indicator and price-difference interaction as controls, but do not report a lamb elasticity given the limited variation in lamb offers.[79] Lamb is absent from the label-study sample used to construct the emissions gaps and offer shares and therefore does not enter the carbon-tax calculation.

S.2.7 Additional analyses on estimation of price effects

This section shows additional checks to section 3.2 assessing the effect of a carbon tax in the natural field setting. For each canteen and date, I identify the main vegetarian

77. In the price-estimation sample, fish is chosen in 50.6% of transactions on days when it is the main meat/fish offer; the overall meat/fish choice share is 45.6%.

78. In many Christian traditions, fish is eaten on Fridays.

79. Excluding all lamb-offer canteen-days yields a pooled poultry, pork, and beef price coefficient of −0.102, compared with −0.101 when these observations are retained with separate lamb controls. The differentiated student tax-equivalent estimate rounds to €133/t in both specifications.

Table S.28. Estimated elasticities

	(1) Pooled	(2) By type
Poultry/pork/beef	-0.528***	
	(0.093)	
Poultry		-0.295**
		(0.123)
Pork		-0.710***
		(0.136)
Beef		-1.083***
		(0.225)
Observations	465,328	465,328

Notes: Elasticities estimated from the coefficients reported in Table 2, the average likelihood of choosing meat under the relevant offer, and the average price of the meat meal. The pooled elasticity covers poultry, pork, and beef. Fish and lamb price interactions are included as controls; their elasticities are not reported.

and meat meals by total transaction counts for each matched meal, pooling across payment amounts. I use the most frequent student payment for each selected meal among payments of at most €2.50. In the main specification, I include controls for broad meat type and six mutually exclusive meat-offer categories, time and canteen controls, and the intervention-period controls described in Table 2. I also include separate price interactions for fish and lamb, so the pooled coefficient describes poultry, pork, and beef. The identifying variation therefore comes from two sources:

- Meat offers within the same broad type and offer category can be paired with different vegetarian meals, generating variation in the relative price. Meat prices range from €1.85 to €2.50, and vegetarian prices from €1.35 to €2.40. Table OA.20, Column (4), additionally controls for six vegetarian-offer categories.
- Prices for vegetarian and meat meals increased by different amounts across dishes in October 2022. For dishes observed in the same canteen in both pricing periods, the average increase in the modal student price is €0.49 for meat/fish dishes (ranging from €0.35 to €0.65) and €0.33 for vegetarian dishes (ranging from €0.05 to

€0.85).[80] This changes relative prices even when the same meat–vegetarian pair recurs and supplies variation in addition to differences across daily meal pairings.

Figure S.7 displays the resulting distribution of relative prices before and after partialling out the controls. The controls leave 56.7% of the raw price-difference variance, and the 10th and 90th percentiles of the residual distribution are −€0.25 and €0.26. The figure therefore documents substantial conditional price support.

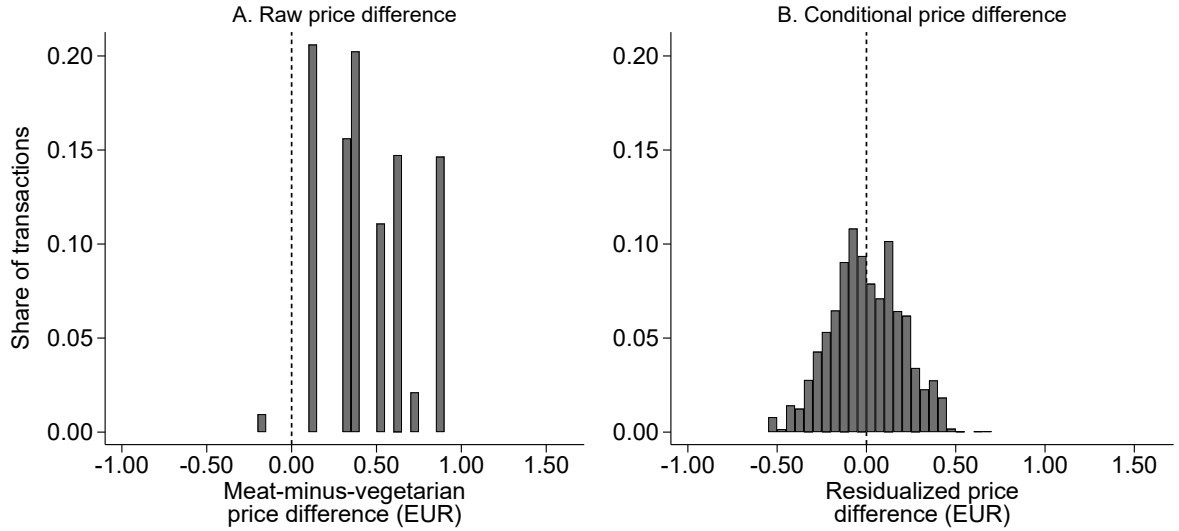


Figure S.7. Distribution of relative meal prices

Notes: Panel A shows the raw difference between the prices of the main meat and vegetarian meals. Panel B shows the residual after partialling out week, day-of-the-week, canteen, intervention-period, broad meat-type, and six recipe-informed meat-offer-category controls. Both panels use the 465,328 transactions in the price-regression sample and assign equal weight to each transaction; because the price difference is common within a canteen-day, this weights offers by transaction exposure. The sample contains 876 canteen-days. The bin width is €0.05, and the dashed vertical line marks zero.

As in other observational demand analyses, the central identification issue is to distinguish responses to prices from differences in other determinants of demand. The student canteen in Bonn is a non-profit organization aimed at serving the students. Prices are set with reference to ingredient costs and part of the canteens' fixed costs (Studierendenwerk Bonn, 2025). Posted prices use discrete price points and are adjusted periodically, rather than tracking changes in ingredient costs continuously.

80. Using the price-estimation sample, I compare each dish's most frequent daily student price in the same canteen before October 1, 2022 with its most frequent daily price from that date onward. The comparison includes 83 meat/fish and 55 vegetarian dish–canteen combinations observed in both periods and gives each combination equal weight. If prices tie in frequency, I use the lower price.

The canteens' cost-based pricing and social mission make demand-driven pricing, such as charging more simply because a dish is popular, less likely. However, a higher price for a meat meal could reflect more attractive ingredients, and these ingredients might also make the meal more popular. I therefore control for broad meat type and the six recipe-informed meat-offer categories in the main analysis. Similarly, the price increase occurred in a time of generally increasing price levels, and it is thus sensible to include time controls in the main analysis.

To provide more intuition on the variation in the data and test robustness, I provide the following additional analyses:

- Table OA.20 estimates the meat-type-specific responses from Table 2 with five control sets. Column (1) includes only broad meat-type controls. Column (2) adds week, day-of-the-week, canteen, and intervention-period controls. Column (3) adds the six meat-offer categories used in the main specification. Column (4) additionally controls for six vegetarian-offer categories. Column (5) instead controls for the exact meat meal and the six vegetarian-offer categories. The final row reports the residual variance in the price difference after partialling out each control set, relative to its raw variance.
- Table S.29 uses the main specification throughout. Column (1) restricts the sample to the period before the October price change, Column (2) to the exact sample period used for the label-effect estimation, and Column (4) to observations outside that period. Column (3) tests whether the pooled poultry/pork/beef price response differs in the treatment canteen. Column (5) estimates the price response using only the two control canteens. The pooled price response remains negative across all sample restrictions. The treatment-canteen interaction is small in magnitude and statistically insignificant (coefficient 0.014, SE 0.016), while the control-canteen-only coefficient is −0.117 (SE 0.024). The final row reports the residual variance in the price difference after partialling out the relevant controls, relative to its raw variance in each sample.
- Table S.30 tests whether the price response varies with the posted price of the vegetarian meal. It includes the centered vegetarian price and its interaction with the meat-minus-vegetarian price difference, and reports the implied price-difference co-

efficient at the first quartile, median, and third quartile of the vegetarian-price distribution.

Some of the more restrictive specifications yield smaller absolute price coefficients. In Table OA.20, these specifications also absorb a substantial share of the variation in relative prices from which the effect is identified. Importantly, using the smaller coefficients in the benchmark calculation would raise, rather than lower, the implied tax-equivalent rate.

Table S.29. Comparison of effects: labels vs. "carbon tax" robustness checks

	Meat meal chosen				
	(1) Pre-October	(2) Study period	(3) Treatment test	(4) Outside study	(5) Controls only
PD	-0.107***	-0.034	-0.106***	-0.089***	-0.117***
	(0.037)	(0.045)	(0.020)	(0.019)	(0.024)
PD x Treatment canteen			0.014		
			(0.016)		
Weekly time controls	Yes	Yes	Yes	Yes	Yes
Meat-offer controls	Yes	Yes	Yes	Yes	Yes
Remaining PD variation (%)	33.3	27.3	56.7	57.6	55.2
Observations	166,140	88,730	465,328	376,598	307,478

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable equals one when a transaction contains the meat meal and zero otherwise. Coefficients are in probability units, so 0.01 corresponds to one percentage point. Linear probability regressions use student canteen data from April 2022–July 2023, excluding all canteen observations from May 2023 and June 1–2, 2023. "PD" denotes the price difference between the main meat and vegetarian meals. Every column uses the main-text specification: week, day-of-the-week, canteen, and intervention-period controls; broad meat-type controls; six meat-offer categories; and separate fish and lamb price interactions. Column (1) uses observations before the October 2022 price change. Column (2) uses only the exact period and exclusions used for the label-effect estimation. Column (3) uses the full sample and interacts the pooled poultry/pork/beef PD measure with the treatment canteen. Column (4) uses observations excluded from the label-study sample. Column (5) uses only the two control canteens; treatment and intervention interactions are constant in that sample and are therefore omitted. The fish and lamb interactions and other controls are included but not displayed. "Remaining PD variation" is the residual variance after partialling out the controls, expressed as a percentage of raw PD variance within the corresponding sample. Standard errors are two-way clustered by calendar date and exact meat–vegetarian meal pair.

Table S.30. Price response and the vegetarian meal's base price

	Meat meal chosen			
	(1) Interaction model	(2) First quartile	(3) Median	(4) Third quartile
PD at mean veg. price	-0.109			
	(0.087)			
Vegetarian price (centered)	0.026			
	(0.094)			
PD x Vegetarian price	-0.079			
	(0.060)			
Implied PD coefficient		-0.088	-0.100	-0.128
		(0.087)	(0.087)	(0.089)
Vegetarian meal price (€)	1.85	1.60	1.75	2.10
Observations	375,863	375,863	375,863	375,863

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable equals one when a transaction contains the meat meal and zero otherwise. "PD" denotes the price difference between the main meat and vegetarian meals. The analysis excludes fish and lamb offers and therefore estimates the pooled response for poultry, pork, and beef. Column (1) includes the vegetarian meal price, centered at its transaction-weighted sample mean, and its interaction with PD. The PD coefficient in Column (1) is therefore evaluated at the mean vegetarian price. Columns (2)–(4) report linear combinations giving the implied PD coefficient at the first quartile, median, and third quartile of the vegetarian-price distribution; the corresponding prices are shown in the table. All specifications use the main controls for week, day of the week, canteen, intervention periods, broad meat type, and six meat-offer categories. Standard errors are two-way clustered by calendar date and exact meat–vegetarian meal pair.

S.2.7.1 Price-effect inference. The price difference is common to all transactions within a canteen-day, and repeated dishes may share demand shocks. Table S.31 holds the regression specification fixed and compares the headline two-way clustering by calendar date and exact meal pair with transaction-robust standard errors and alternative one-way clustering choices. The point estimate is unchanged by construction. Standard errors are larger when observations are allowed to be correlated within dates, dishes, or exact meal pairs than when transactions are treated as independent.

Table S.31. Price-effect inference under alternative clustering

Inference	Coefficient	Standard error	95% interval
Two-way calendar date and exact meal pair (main)	-0.101	0.018	[-0.136, -0.066]
transaction robust	-0.101	0.004	[-0.109, -0.093]
Canteen × calendar date	-0.101	0.013	[-0.127, -0.075]
calendar date	-0.101	0.014	[-0.129, -0.074]
exact meat dish	-0.101	0.018	[-0.137, -0.066]
exact vegetarian dish	-0.101	0.019	[-0.138, -0.064]
exact meal pair	-0.101	0.018	[-0.136, -0.067]

Notes: Every row reports the pooled poultry, pork, and beef price coefficient from the main specification; only the covariance estimator changes. The first row reports the two-way calendar-date and exact-meal-pair clustering used for the headline inference in Table 2; the second row reports transaction-robust standard errors. Canteen × calendar-date clustering allows errors to be correlated across transactions made in the same canteen on the same calendar date, reflecting that the price regressor is common within these cells. Calendar-date clustering allows common date-specific demand shocks across canteens, while dish and meal-pair clustering allow errors to be correlated when the same offer recurs.

Table S.32 compares guest-and-date clustering with date-and-meal-pair clustering using personalized-card purchases from April–July 2022, the source period covering the labeling experiment, within which individual guests' purchases can be linked. Both specifications use exactly the same 92,338 transactions and the same outcome, controls, and point estimator. The estimated price-response coefficient is −0.107; its standard error is 0.041 under guest-and-date clustering, compared with 0.049 under date-and-meal-pair clustering. Both estimates are statistically significant at the 5% level. This check supports the robustness of price-effect inference to guest-and-date clustering within this period.

Table S.32. Price-effect inference using linkable guest identifiers, April–July 2022

Inference	Coefficient	Standard error	95% interval	Observations
Date and meal pair	-0.107	0.049	[-0.205, -0.008]	92,338
Guest and date	-0.107	0.041	[-0.188, -0.026]	92,338

Notes: Every row reports the pooled poultry, pork, and beef price coefficient from the main specification, restricted to April–July 2022 purchases made with personalized payment cards. The observations, outcome, controls, and point estimator are identical across rows; only the covariance estimator changes. "Date and meal pair" denotes two-way clustering by calendar date and exact meat–vegetarian meal pair. "Guest and date" denotes two-way clustering by guest and calendar date, using guest identifiers that are linkable within this source period. Subject to the main price-analysis restrictions, the sample includes all purchases made with a personalized payment card, including all guests with at least one purchase.

S.2.7.2 Direct field-implied tax-equivalent rates. Table S.33 reports the tax rates at which the estimated price responses predict the same change in meat choice as the ITT label estimates.

Table S.33. Direct field-implied local tax-equivalent rates

	€/t CO_2e	
Definition	Students	All guests
Pooled	194 [9, 509]	195 [24, 502]
Differentiated by meat type	133 [6, 378]	134 [16, 355]

Notes: The rates invert the estimated field label and price effects. The student-only rates use the ITT label estimate and are the main inversion estimates because the price response is estimated from student purchases. The all-guests rates use the same ITT specification without restricting the sample to students. The pooled rate uses the common poultry, pork, and beef price coefficient. The differentiated rate uses the poultry, pork, and beef coefficients, assigns fish a zero response, and weights by offer shares. Brackets report 95% percentile bootstrap intervals for the corresponding sample. The intervals use 999 bootstrap replications. For both columns, guests are resampled within assigned treatment arm and the corresponding guest-fixed-effects label regression is re-estimated. Both combine the resulting label-effect draws with calendar-date resampling for the price coefficients, emissions gaps, and offer shares. The intervals propagate uncertainty from the respective observed sampling units, conditional on the three canteens and their realized treatment assignment. They do not capture treatment-effect heterogeneity across canteens.

S.3 Applicability to other contexts

All three experiments study meal choices in Bonn's student canteens, but the paper's findings differ in how readily they can be applied elsewhere. The benchmarking method is designed to be re-estimated in other settings, whereas the numerical €155/t benchmark is specific to the population and choice environment studied here. The mechanism and welfare results speak to broader questions about how carbon labels work, but their quantitative magnitudes may also depend on the consumers, products, and institutions involved.

This section discusses these different dimensions of external validity. I use List (2020) as a guide, focusing on how selection, the naturalness of the choices, and changes in scale or implementation may affect what can be learned from the Bonn setting. The three experiments provide complementary evidence for this purpose. Experiment 1 identifies the benchmark in a controlled but consequential choice environment, Experiment 2 examines similar behavior in ordinary repeated purchases in the

same local food market, and Experiment 3 studies the behavioral channels and welfare components underlying the response to the labels.

S.3.1 Suitability of the study setting

The EV Litmus Test proposed by List (2020) asks, in essence, whether the chosen setting is particularly well suited to providing an informative test of the question at hand. On this criterion, the Bonn setting has several advantages for the framed field experiments. Participants choose between familiar meals that recur in their ordinary lunch environment, make consequential decisions shortly before lunch, and receive one of the meals soon afterwards. WTP is elicited relative to a familiar alternative lunch, keeping the choice margin close to the substitution consumers ordinarily face. At the same time, the experimental setting permits controlled variation in the information shown to consumers and in the salience of emissions.

The cooperation with the canteens also makes it possible to compare the framed experiment with recurring purchases and naturally occurring price differences in the same institutional food market. Experiments 1 and 2 therefore share the institutional provider, menu structure, emissions calculations, and label design, while differing in experimental scrutiny, repetition, label exposure, and how choices are observed. The samples cannot be linked at the individual level, but they come from closely related local populations. Experiment 2 thus provides a useful check on whether the benchmark obtained from the framed choices is also visible in ordinary repeated purchases.

No single design would combine all of these features. A broader online sample, for example, would improve demographic coverage but make immediate implementation of a familiar meal difficult. A purely natural field design would preserve the ordinary purchase environment but would generally reveal only the resulting choices, rather than the underlying changes in valuations or the separate roles of attention and belief correction. The Bonn setting therefore provides a useful combination of consequential and familiar choices, close links to the natural consumption environment, and the experimental control needed to study the mechanisms behind the label response.

The framed-field participants were recruited from the Bonn experimental pool without participants being informed of the environmental focus of the study. Relative to the canteen survey sample, the experimental samples contain fewer students. The percentage of vegetarians is broadly similar to that inferred from canteen purchases (Sec-

tion S.1.2). Participants with strict dietary requirements that could not be accommodated by the experimental meal pairs were excluded, so the framed estimates speak most directly to consumers able to choose among common canteen menu alternatives. Attrition, comprehension exclusions, and meal non-collection are low. Experiment 2 uses ordinary purchases linked to personalized payment cards and focuses in its guest-level analyses on recurring users with a stable pre-intervention canteen affiliation. Treatment was implemented in one canteen, leaving variation across institutions unobserved.

S.3.2 Contribution 1: Portability of the benchmarking framework

Experiment 1 maps the label-induced change in relative WTP into choices at observed prices and then identifies the local carbon-tax rates that most closely reproduce the resulting emissions change. Because the label and price responses are evaluated for the same consumers and choice environment, the method does not require importing a price elasticity from another population or setting.

Applying the method elsewhere would require a consequential measure of relative WTP for products relevant to the target choice environment. The details of the elicitation and incentives can be adapted to the population and setting. The resulting estimate would then be a new local benchmark for that population, choice set, and price environment.

Experiment 2 provides a complementary check on the approach. A benchmark of a similar magnitude emerges from naturally occurring price differences and repeated purchases in the same consumption setting. Where sufficiently informative and plausibly exogenous price variation is available, such field variation can also provide a useful benchmark.

S.3.3 Contribution 2: Scope and applicability of the local €155/t benchmark

The closest test of how far the numerical benchmark travels is within the Bonn canteen setting itself. Experiment 1 participants are drawn from a population familiar with the canteens, while Experiment 2 observes ordinary purchases there. The two experiments differ substantially in how choices are made: Experiment 2 follows repeated purchases over seven weeks, outside experimental scrutiny, with incomplete exposure to the labels and choices potentially observed by others. Despite these differences, it yields a

benchmark of a similar magnitude. This does not show which of these features matters individually, but it shows that the Experiment 1 estimate remains visible when they change together.

Evidence from related work also suggests that neither side of the Bonn comparison is unusual for institutional food settings. The canteen studies by Brunner et al. (2018) and Lohmann et al. (2022) report label-induced emissions reductions of a similar order of magnitude, while the field price responses in this paper are broadly similar to meat-demand estimates in related work (Section S.2.6). This makes other student and institutional canteens the most natural settings in which to view the Bonn estimate as a useful reference point.

The benchmark may nevertheless vary across populations. Responses to labels or prices could differ with demographic characteristics, environmental preferences, or prior consumption patterns. Figures S.8– S.10 therefore provide exploratory evidence on heterogeneity in the label response, while Tables S.41– S.42 show corresponding differences in baseline food choices. In Experiment 1, point estimates suggest stronger label responses among participants with higher environmental preferences and greater openness to environmental food labels, but the environmental-preference gradients do not replicate consistently in Experiment 2. Estimates by student status and age likewise show no consistent pattern across the two experiments. Point estimates are larger among women in both experiments, although these subgroup comparisons are imprecise and should not be interpreted as causal effects of gender. Overall, I find no indication that the label response is specific to students, but the evidence also does not imply that its magnitude will be identical in substantially different populations.

The choice environment matters as well. Available substitutes and baseline consumption determine which consumers are close to a choice threshold and therefore affect the absolute response to both labels and prices. Because the benchmark compares those responses within the same choice environment, some of these constraints affect both sides of the comparison. This makes the tax-equivalent benchmark less mechanically tied to the choice environment than considering the absolute emissions reduction, but does not imply that it will necessarily remain unchanged when the choice environment changes.

Other student, workplace, and hospital canteens are therefore the closest applications: they share centralized menu design, recurring point-of-choice decisions, and rel-

atively small choice sets. Counter-service restaurants retain an immediate individual choice but may involve stronger branding and more competing information. Supermarkets and meal-kit platforms differ more substantially in the size of the choice set, the possibility of purchasing for others, and the delay between purchase and consumption.

Lohmann et al. (2024) study food-delivery app users in an incentivized online experiment, with participants selected to receive their orders contacted after the study to arrange delivery. Ho and Page (2025) study customers of a meal-kit provider who select meals at least one week in advance of delivery. The former find no statistically significant average effect of carbon labels on emissions, while the latter report a reduction of 0.6%, smaller than the reductions reported in the canteen studies discussed above. These comparisons alone do not establish why label effects differ across settings.

Estimating tax-equivalent benchmarks would help assess whether the smaller emissions reductions also imply that labels are less effective relative to carbon pricing. Even a small label-induced reduction may correspond to a substantial tax-equivalent rate if consumers are also relatively unresponsive to price changes. The findings of Lohmann et al. (2024) illustrate the relevance of this distinction: their prominently displayed emissions-based meat tax, which raised meat prices by 10% on average, also produced no statistically significant average emissions reduction. Although their menu-repositioning intervention did significantly reduce emissions, the small label effect does not by itself establish that carbon pricing would be a more effective alternative. Comparing responses to labels and price changes within each setting provides a more informative basis for assessing their relative effectiveness.

The €155/t estimate can provide a reference point for these settings, but the corresponding local tax-equivalent rate remains an empirical question.

Further scaling considerations. The numerical benchmark compares the response to carbon labels with the response to a local emissions-based relative-price adjustment on the same meals for the consumers and study period observed in Bonn. Even within a similar canteen setting, the benchmark may change as an intervention is maintained for longer or affects choices beyond the immediate purchase occasion. More generally, effects in large-scale nudge-unit trials can differ from those reported in academic studies (DellaVigna and Linos, 2022).

The first issue is duration. Experiment 1 measures an immediate response in a one-shot choice, while Experiment 2 follows repeated purchases over seven weeks and yields a similar benchmark. The data do not show how the label and price responses develop over longer periods. Attention to the labels may decline, while learning and changes in habits may persist. Consumers may also adjust gradually to price differences. Since the benchmark depends on both responses, its longer-run value may differ from the estimate obtained during the study period.

A second issue is whether effects carry over to other consumption decisions. The experiments observe lunch purchases rather than consumers' food choices over the full day. Consumers could compensate for a lower-emission lunch by consuming more meat later in the day, reducing the overall emissions benefit. Conversely, exposure to labels in the canteen could reduce emissions from later choices if consumers learn about relative emissions or continue to pay more attention to them. The post-intervention period in Experiment 2 provides some evidence on choices after the labels are removed, but studying broader spillovers would require data linking consumption across settings.

Scaling the intervention within institutional catering appears practically feasible. The federally funded "KlimaTeller in der Mensa" project supports 20 German Studierendenwerke (student services providers) in introducing automated CO_2e accounting and climate labeling (Bundesministerium für Umwelt, Klimaschutz, Naturschutz und nukleare Sicherheit, 2026), and several university catering systems have introduced carbon-footprint or climate information in their menus.[81] These initiatives rely on inputs similar to those used in Bonn—recipe-level emissions data, integration with menu planning, and display at the point of choice—and show that such systems can be introduced by individual institutional providers.

A different set of considerations would arise if either policy were implemented at a broader regulatory level rather than by the canteen itself. An upstream carbon tax, for example, need not translate one-for-one into the consumer-facing price differences studied here because of pass-through and other supply-side responses. Likewise, labels imposed further upstream or across a wider market could induce producers to change products, sourcing, or emissions intensities (see also Barahona, Otero, and Otero, 2023).

81. The precise labeling systems differ across providers. Some report emissions values and traffic-light scores for individual meals, while others identify particularly climate-friendly meals through the KlimaTeller label.

These responses may be important for broader regulatory policies, but they are outside the local comparison studied here, in which both the labels and the price adjustment are implemented at the canteen level. This is also a practically relevant scale for carbon labels, as several university catering providers have independently introduced carbon information into their menus.

S.3.4 Contribution 3: Applicability of the mechanism findings

The reduced-form evidence points to roles for both belief correction and attention in driving the labels' effectiveness. Participants who previously underestimated relative emissions react more strongly when shown the labels, which is consistent with belief correction contributing to the response. At the same time, participants who overestimated emissions also reduce their WTP when shown the labels, so correction of misperceptions alone cannot explain the effect. The Attention condition provides separate evidence on salience. After the emissions-guessing task, participants reduce their WTP for higher-emission meals even though the emissions of those meals are not revealed. Taken together, the reduced-form results indicate that making emissions more salient is an important part of the label response.

The structural model quantifies the importance of the two channels. At the point estimate, the attention intervention accounts for approximately 73% of the label-induced emissions reduction. This percentage should not be treated as a portable parameter: its participant-bootstrap interval is wide, and the estimated allocation depends on prior beliefs, environmental preferences, the meals being compared, and the structural assumptions. Allowing the structural parameters to differ by observed participant characteristics changes the magnitudes but continues to imply an important role for attention (Section S.3.6).

The qualitative implication may therefore carry beyond the Bonn canteens. The size of the attention channel, however, remains an empirical question in each setting. This interpretation is also consistent with Imai et al. (2026), who find no behavioral response in a design that isolates the correction of emissions misperceptions.

S.3.5 Contribution 4: Applicability of the consumer-preference and welfare findings

One concern with labels and other behavioral interventions is that their presence may itself impose psychological costs, even when they improve the choices consumers make. I find little evidence of strong aversion to the labels in this application. Across Experiments 1 and 3, fewer than 5% of participants prefer the labels to be absent and are willing to pay to avoid them, while mean WTP for label presence is approximately €0.20 per meal. This WTP is similar across experimental treatment histories (Table OA.15), including among participants who had and had not previously encountered the labels.

The structural analysis separates the welfare change generated by altered meal choices from the residual value of label presence, F. In the baseline simulation, the choice-mediated welfare changes generated by labels and the benchmark taxes are small and similar. Among Experiment 3 participants who had previously experienced the labels, the participant-level calculation gives a residual value of 21.4 cents per choice.

At the point estimate, this residual value is positive rather than indicative of a psychological cost that dominates the benefits of label presence. The same qualitative result holds in the robustness exercise that explicitly allows the Attention and Label interventions to impose additional emissions-related psychological costs: the estimated residual value exceeds those costs at the point estimate.

For consumers reasonably similar to the participants in Experiments 1 and 3 and facing a comparable label, these results provide evidence about the likely direction of the net welfare effect. The field survey points in the same direction descriptively: among respondents who purchased in the treatment canteen and noticed the labels, 75.6% support permanent installation and 7.8% oppose it (Section S.2.5). Policy support is not itself a welfare measure, but it provides complementary evidence that the labels were not broadly experienced as aversive. The numerical WTP, F, and welfare estimates should not be transferred mechanically to other populations or label designs.

S.3.6 Structural-model fit and observed parameter heterogeneity

Online Appendix Sections OA.2.2.3 and OA.2.2.7 summarize model fit and the implications of allowing parameters to vary with observed participant characteristics. The fuller tables below document the intervention-by-subgroup simulations and fit comparisons.

They support two complementary external-validity conclusions: attention remains an important component of the predicted label effect across the observed groups, but allowing for this heterogeneity does not close the model's WTP-fit gap for higher-emission meals under ATTENTION+LABEL.

The subgroup-specific parameter estimates underlying these exercises are reported below.

The baseline model assumes common γ, θ, and κ. I assess this assumption by re-estimating the structural model separately by diet, student status, food-specific willingness to sacrifice for the environment, and general willingness to sacrifice for the environment. Food-specific willingness to sacrifice is the reported importance of environmental aspects when choosing a dish in a restaurant or canteen. General willingness to sacrifice is the mean response to four statements completing the prompt "Keeping my greenhouse gas emissions as low as possible is something": "I am willing to do even if it involves higher costs and effort," "I am only willing to do if others do it as well" (reverse-coded), "that aligns with my principles," and "I feel personally responsible for." Both measures are expressed on a scale from 1 to 10, and observations at the median enter the upper group. Table S.34 reports the subgroup estimates. Estimates suggest a higher guilt parameter γ among vegetarians and among individuals with a higher food-specific willingness to sacrifice for the environment. The latter pattern remains pronounced when restricting the sample to non-vegetarians. This pattern does not replicate clearly for general willingness to sacrifice for the environment.

Table S.34. Structural parameters by observed participant characteristics

Group	N	γ	θ	κ
Panel A: Diet				
Vegetarian	102	0.200 [0.114, 0.305]	0.221 [0.000, 0.836]	0.094 [0.000, 0.605]
Non-vegetarian	328	0.100 [0.075, 0.133]	0.000 [0.000, 0.288]	0.258 [0.000, 0.659]
Panel B: Student status				
Student	298	0.101 [0.078, 0.139]	0.022 [0.000, 0.394]	0.000 [0.000, 0.419]
Non-student	132	0.109 [0.050, 0.176]	0.000 [0.000, 0.474]	0.607 [0.000, 1.000]
Panel C: Food-specific WTS				
Below median (< 6)	162	0.058 [0.034, 0.098]	0.077 [0.000, 0.748]	0.239 [0.000, 1.000]
At/above median (≥ 6)	268	0.134 [0.099, 0.177]	0.000 [0.000, 0.234]	0.084 [0.000, 0.476]
Panel D: General WTS				
Below median (< 7)	194	0.093 [0.060, 0.133]	0.170 [0.000, 0.523]	0.148 [0.000, 0.665]
At/above median (≥ 7)	236	0.127 [0.086, 0.165]	0.000 [0.000, 0.166]	0.247 [0.000, 0.691]
Panel E: Food-specific WTS, non-vegetarians only				
Below median (< 6)	151	0.057 [0.032, 0.097]	0.115 [0.000, 0.804]	0.214 [0.000, 1.000]
At/above median (≥ 6)	177	0.137 [0.099, 0.181]	0.000 [0.000, 0.227]	0.207 [0.000, 0.635]

Notes: Each row applies the estimation and bootstrap procedure of Table OA.9 to the indicated subgroup. The full-sample medians are 6 for food-specific and 7 for general willingness to sacrifice. These cutoffs remain fixed across subgroups, bootstrap replications, and simulations; the non-vegetarian food-specific split also uses 6. Observations at the median enter the upper group; missing attitude responses are excluded from the corresponding split. γ is measured in euros per kg CO_2e.

Table S.35 provides the full intervention-by-subgroup simulations underlying the summary in Online Appendix Table OA.14. These estimates illustrate how the model's policy implications vary across observed groups rather than establishing that the same magnitudes transport to other populations.

Table S.35. Simulated effects allowing observed parameter heterogeneity

Intervention	(1) Δ GHGE (g)	(2) [95% CI]	(3) Δ welfare (cents)	(4) [95% CI]
Panel A: Parameters heterogeneous by diet				
Knowledge	−0.4	[−2.5, 35.5]	0.01	[−0.05, 25.89]
Attention	−54.4	[−93.7, −13.0]	−0.08	[−0.81, 0.31]
Labels	−85.8	[−125.3, −21.2]	0.24	[−0.23, 14.56]
Tax, €155/t	−108.5	[−159.7, −61.2]	0.24	[−0.16, 0.65]
Meat ban	−215.4	[−277.2, −156.3]	−3.99	[−5.94, −2.29]
Panel B: Parameters heterogeneous by student status				
Knowledge	6.4	[0.0, 48.7]	0.95	[−0.10, 40.11]
Attention	−54.4	[−92.8, −11.4]	−0.07	[−1.04, 0.32]
Labels	−67.8	[−125.6, −14.1]	1.19	[−0.06, 40.09]
Tax, €155/t	−108.5	[−159.7, −61.2]	0.22	[−0.26, 0.64]
Meat ban	−215.4	[−277.2, −156.3]	−3.97	[−6.02, −2.25]
Panel C: Parameters heterogeneous by food-specific wts				
Knowledge	0.0	[−2.7, 39.6]	0.00	[−0.07, 29.17]
Attention	−34.4	[−77.8, −0.1]	−0.04	[−0.71, 0.30]
Labels	−53.7	[−95.5, −13.2]	0.28	[−0.19, 21.25]
Tax, €155/t	−108.5	[−159.7, −61.2]	0.04	[−0.54, 0.45]
Meat ban	−215.4	[−277.2, −156.3]	−4.23	[−6.30, −2.51]
Panel D: Parameters heterogeneous by general wts				
Knowledge	8.4	[−5.5, 27.1]	0.04	[−0.01, 15.03]
Attention	−47.6	[−88.5, −8.2]	0.02	[−0.58, 0.35]
Labels	−74.8	[−114.6, −19.6]	0.30	[−0.17, 5.26]
Tax, €155/t	−108.5	[−159.7, −61.2]	0.24	[−0.16, 0.68]
Meat ban	−215.4	[−277.2, −156.3]	−3.96	[−5.89, −2.25]

Notes: For each heterogeneity dimension, participants are assigned the structural parameter estimates for their subgroup before choices are simulated. Every participant receives equal weight, as in Table 3. Percentile intervals use 999 participant-bootstrap replications that re-estimate both subgroup models and repeat the simulation with the resulting subgroup-specific parameters. Welfare excludes F.

Table S.36 provides the corresponding full fit comparison. Its persistence across the observed groups cautions against interpreting subgroup heterogeneity as resolving the baseline model's understatement of the higher-emission Attention+Label response.

Table S.36. Reduced-form and structural WTP predictions with observed heterogeneity

Average change in WTP for higher-emission meals (euro)

Heterogeneity	Intervention	Reduced form	Structural	Difference
Diet	Attention	-0.081 [-0.129, -0.035]	-0.085 [-0.107, -0.059]	0.004 [-0.039, 0.041]
	Labels	-0.216 [-0.284, -0.156]	-0.123 [-0.152, -0.089]	-0.092 [-0.153, -0.047]
Student status	Attention	-0.081 [-0.128, -0.035]	-0.084 [-0.105, -0.055]	0.003 [-0.043, 0.039]
	Labels	-0.217 [-0.285, -0.158]	-0.121 [-0.151, -0.084]	-0.096 [-0.158, -0.054]
Food-specific WTS	Attention	-0.082 [-0.129, -0.037]	-0.080 [-0.100, -0.056]	-0.002 [-0.047, 0.036]
	Labels	-0.217 [-0.286, -0.157]	-0.120 [-0.149, -0.085]	-0.098 [-0.159, -0.053]
General WTS	Attention	-0.082 [-0.128, -0.035]	-0.081 [-0.101, -0.055]	-0.001 [-0.045, 0.038]
	Labels	-0.217 [-0.284, -0.158]	-0.121 [-0.150, -0.087]	-0.096 [-0.157, -0.051]

Notes: For each heterogeneity dimension, the structural parameters and reduced-form WTP coefficients are estimated separately for the two subgroups. Entries report average WTP changes for meals with emissions above those of the sandwich. Difference is reduced form minus structural; 95% percentile intervals are shown below each estimate. Intervals use 999 paired participant-bootstrap replications that re-estimate the subgroup-specific structural and reduced-form models.

S.3.7 Additional heterogeneity evidence

This subsection reports the exploratory heterogeneity estimates summarized in Section S.3.3. The demographic and attitudinal characteristics are not randomly assigned, several subgroup estimates are imprecise, and some Experiment 2 comparisons are restricted to guests who can be linked to the survey data. I therefore use these results to examine whether the label response differs systematically across observed groups, rather than to give a causal interpretation to the characteristics themselves.

To make the Experiment 1 estimates comparable across meals, I adapt the reduced-form WTP analysis used in Table OA.2. I express the treatment effect so that a more negative coefficient always indicates a shift away from the higher-emission option. When the warm meal has higher emissions than the sandwich, I retain the change in WTP for the warm meal relative to the sandwich; when the sandwich has higher emissions, I reverse its sign. I then estimate one label-response coefficient within each subgroup, pooling higher- and lower-emission warm meals. These coefficients are reduced-form WTP

responses and should not be interpreted as subgroup-specific carbon-tax-equivalent benchmarks.

Experiment 2 instead measures the label response as the change in the probability of choosing the meat meal. The outcomes in the two experiments are therefore in different units. The figures are most informative about whether the direction and relative pattern of subgroup differences are similar across the experiments, rather than about differences in coefficient magnitudes. Throughout, *N* denotes the number of food-choice observations entering the corresponding regression rather than the number of participants or guests.

Across these exploratory splits, I do not find a single demographic or attitudinal characteristic that consistently separates strong from weak label responses across Experiments 1 and 2. Some Experiment 1 point estimates suggest stronger responses among participants with greater environmental concern, and the estimates for women are larger in both experiments, but the subgroup comparisons are often imprecise and the environmental-attitude patterns do not replicate consistently across studies. These results therefore provide no clear evidence that the main label response is confined to a narrow observable subgroup.

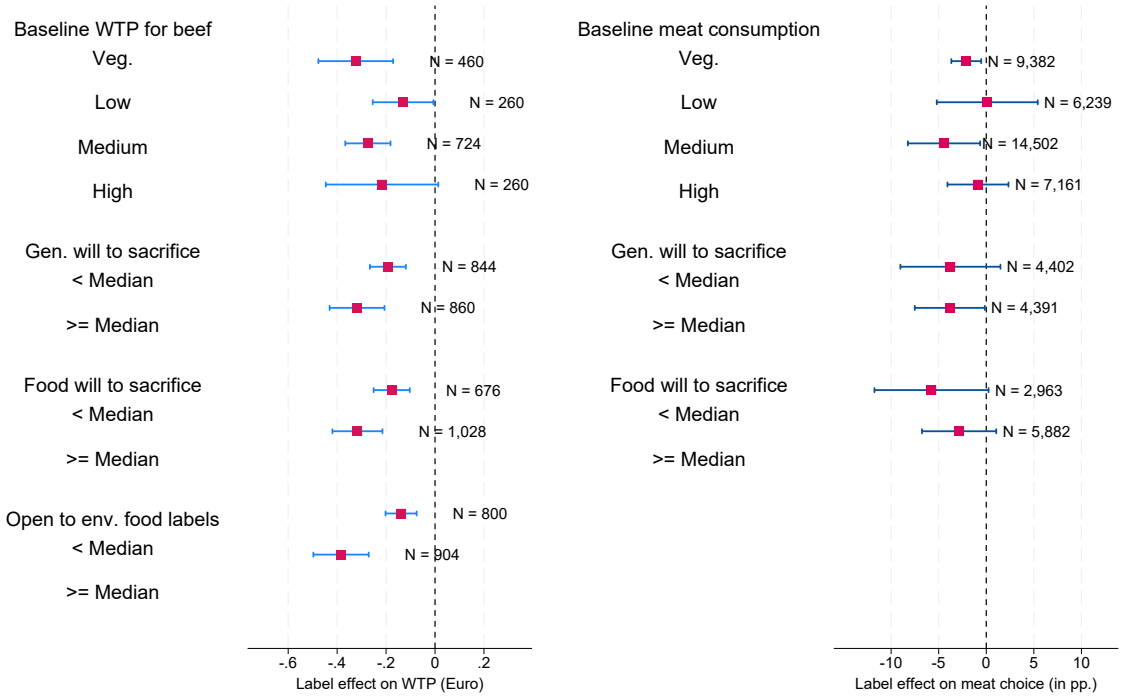
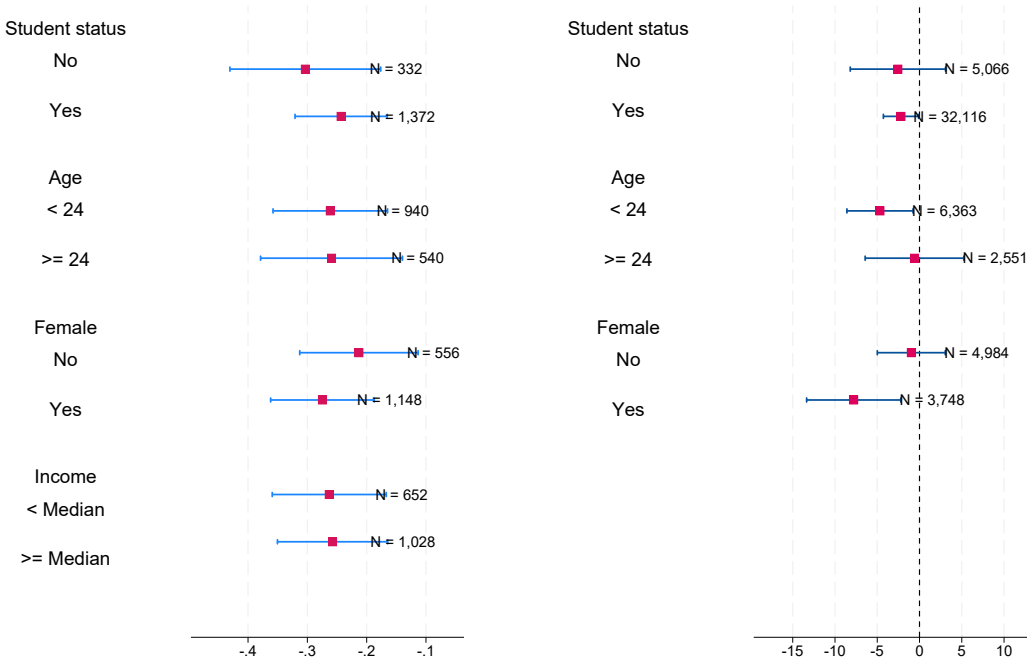


Figure S.8. Heterogeneity in label responses across studies

Notes: The heterogeneity dimensions are motivated by Ho and Page (2025), using the measures available in my experiments. The Experiment 1 panel reports the label effect on relative WTP for the higher-emission option, using the common sign convention described in the text. It estimates one common label coefficient, controlling for the emissions category and third-round decisions, with standard errors clustered at the participant level. The Experiment 2 panel reports percentage-point changes in the probability of choosing the meat meal and follows Table OA.17, Panel B, Column (3), with ITT assignment and guest fixed effects. For Experiment 1, baseline beef valuation is baseline WTP for the beef meal. General willingness to sacrifice for the environment is the mean of four 10-point items concerning willingness to bear costs and effort, conditional cooperation (reverse-coded), environmental principles, and personal responsibility. Food-specific willingness to sacrifice is the reported importance of environmental aspects when choosing a meal in a restaurant or canteen, measured on a 0–9 scale. For Experiment 2, baseline meat consumption is the share of pre-intervention purchases containing meat. General willingness to sacrifice is the mean of two seven-point pre-intervention items concerning more expensive air travel and environmental standards for new housing; food-specific willingness to sacrifice is the seven-point pre-intervention rating of the importance of environmental aspects in food choice. The Experiment 2 attitude comparisons use guests who can be linked to the survey data. More negative coefficients indicate a larger shift away from the higher-emission option. Bars show 95% confidence intervals using the regression's t distribution; N is the number of food-choice observations entering each regression.



Label effect on WTP (Euro)

Label effect on meat choice (in pp.)

Figure S.9. Heterogeneity in label responses by demographic characteristics

Notes: The Experiment 1 panel reports the label effect on relative WTP for the higher-emission option, using the common sign convention described in the text. It estimates one common label coefficient, controlling for the emissions category and third-round decisions, with standard errors clustered at the participant level. The Experiment 2 panel reports percentage-point changes in the probability of choosing the meat meal and follows Table OA.17, Panel B, Column (3), with ITT assignment and guest fixed effects. Age and gender in Experiment 2 are observed for guests who can be linked to the survey data; student status is observed directly in the purchasing data. Experiment 1 income and student status are self-reported. The age split is at 24 because Experiment 2 elicited age in relatively coarse categories. More negative coefficients indicate a larger shift away from the higher-emission option. Bars show 95% confidence intervals using the regression's *t* distribution; *N* is the number of food-choice observations entering each regression.

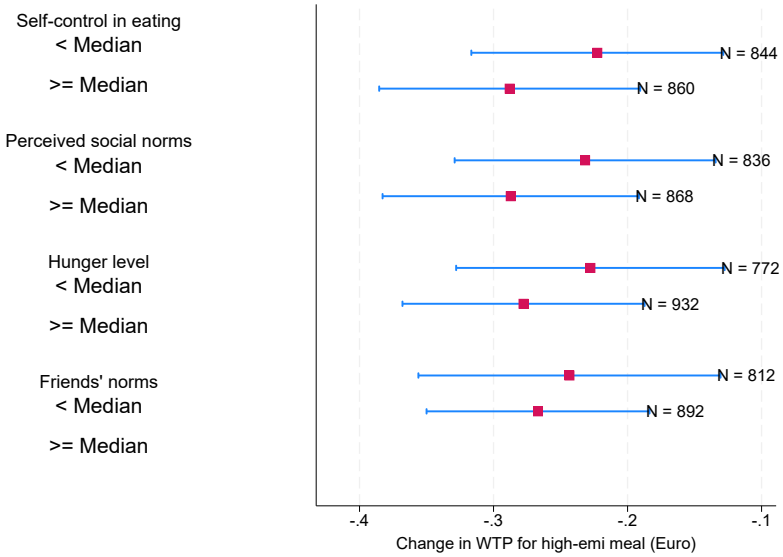


Figure S.10. Heterogeneity in label responses by psychological characteristics

Notes: The figure reports exploratory Experiment 1 heterogeneity estimates using the common sign convention described in the text, so that more negative coefficients indicate a larger shift away from the higher-emission option. It estimates one common label coefficient, controlling for the emissions category and third-round decisions, with standard errors clustered at the participant level. Eating self-control is measured using the nine-item scale of Haws, Davis, and Dholakia (2016). Perceived social norms are elicited using a procedure based on Krupka and Weber (2013); friends' norms summarize participants' perceptions of whether people important to them approve of and engage in emissions-reducing behavior. Hunger is self-reported on a 10-point scale at the beginning of the experiment. Comparable measures are not available in Experiment 2. Bars show 95% confidence intervals using the regression's *t* distribution; *N* is the number of food-choice observations entering the corresponding regression.

S.3.8 Additional structural robustness

This section reports two robustness exercises for the structural analysis. The first allows the Attention and Attention+Label interventions to impose additional emissions-related psychological costs. The second relaxes the baseline assumption that carbon offsetting fully removes perceived environmental guilt.

S.3.8.1 Interventions impose additional psychological costs. One might be concerned that the Attention and Attention+Label conditions impose additional psychological costs on consumers. To capture this possibility parsimoniously, I let the intervention impose an additional utility cost with the same emissions-based form as the term that already enters perceived utility in the corresponding treatment state. Thus, Attention adds a cost of $\gamma e_{ij}^{\text{prior}}$, while Attention+Label adds $\gamma[(1-\kappa)e_j^{\text{true}} + \kappa e_{ij}^{\text{prior}}]$. This introduces no additional parameter and effectively adds a second copy of the emissions-related disutility under these conditions.

The baseline true-utility criterion in Equation 10 is unchanged, but I subtract these intervention-generated costs separately when evaluating utility and welfare. In the choice equations, these additional costs enter through the difference between the warm meal's and the sandwich's costs. When evaluating welfare, I subtract the additional psychological cost of the selected option, including when the sandwich is selected.

I estimate the parameters from the meal-choice data alone. I then include the additional costs in expected utility under Attention and Attention+Label and estimate F using the participant-level averaging procedure in Section OA.2.2.4. This calculation retains any change in the sandwich's psychological cost between the two treatment states.

Table S.37 reports the parameter estimates and the separately calculated residual value F; Table S.38 reports the resulting policy simulations. Because the additional psychological cost doubles the emissions-related weight in the choice equations under Attention and Attention+Label, the observed WTP changes can be rationalized with a lower estimate of γ than in the baseline specification. The lower γ reduces the welfare gains from switching toward lower-emission meals, while the additional psychological costs directly reduce utility under Attention and Label.

The estimated welfare change under labels excluding F is -4.51 cents per choice. The separately estimated residual value is 21.02 cents, giving a total welfare increase of

16.51 cents (95% participant-bootstrap percentile interval: [11.37, 32.51]). The total is calculated within each bootstrap replication before forming the interval. Table S.38 also reports welfare changes under the two tax rates.

Table S.37. Structural estimates when interventions impose additional psychological costs

Parameter	Estimate	95% participant-bootstrap CI
γ	0.050	[0.039, 0.066]
θ	0.000	[0.000, 0.545]
κ	0.162	[0.000, 0.550]
F	0.210	[0.162, 0.262]

Notes: Sample, parameter bounds, and bootstrap procedure as in Table OA.9. The specification adds an intervention-generated psychological cost with the same emissions-based form as the term already entering perceived utility under Attention and Attention+Label, thereby doubling the emissions-related term in these two choice equations without introducing an additional parameter. F is calculated for label-experienced participants using the procedure of Table OA.16, with expected utilities adjusted for these additional psychological costs. γ is measured in euros per kg CO_2e and F in euros per choice.

Table S.38. Simulated effects when interventions impose additional psychological costs

Intervention	(1) Δ GHGE (g)	(2) [95% CI]	(3) Δ welfare (cents)	(4) [95% CI]
Knowledge	0.0	[0.0, 38.0]	0.00	[−0.03, 24.93]
Attention	−54.4	[−94.1, −13.7]	−5.22	[−7.49, −3.80]
Labels	−74.1	[−120.1, −20.5]	−4.51	[−6.20, 11.64]
Tax, €155/t	−108.5	[−159.7, −61.2]	−0.33	[−0.72, −0.06]
Tax, €106/t (GHGE match)	−74.6	[−115.9, −34.1]	−0.05	[−0.38, 0.08]
Meat ban	−215.4	[−277.2, −156.3]	−5.07	[−7.09, −3.36]

Notes: Sample, choice set, participant weights, and bootstrap procedure as in Table 3, using the specification in Table S.37. Emissions changes are in grams CO_2e per meal. Welfare changes are in cents per meal and include the added intervention-generated psychological costs, but exclude F. The €155/t row uses the midpoint of the Experiment 1 benchmark; the GHGE-matching row reports the tax rate that most closely reproduces the label-induced emissions change in this robustness specification.

S.3.8.2 Offsetting does not completely remove perceived guilt. The analyses above assume that the Attention+Offset treatment fully removes environmental guilt. Section S.4.5 provides descriptive evidence that offsetting removes at least part of participants' environmental guilt, but it may not eliminate it entirely. To examine sensitivity, I re-estimate the model and re-run the simulations assuming that offsetting removes 70% of perceived guilt (i.e., 30% remains). Formally,

Attention+Offset Treatment (30% of emissions considered):

$$WTP_{im}^{\text{A+O}} \;=\; \Delta v_{im} - 0.3\gamma \Delta e_{im}^{\text{prior}}. \tag{S.2}$$

with salience normalized to $\theta = 1$ under Attention and Attention+Offset.

Combining this expression with Equation 8 and allowing for elicitation error gives the estimating equation

$$\widetilde{WTP}_{im}^{\text{A+O}} - \widetilde{WTP}_{im}^{\text{A+L}} \;=\; \gamma(1-\kappa)\Delta e_{m}^{\text{true}} + \gamma(\kappa - 0.3)\Delta e_{im}^{\text{prior}} + \xi_{im}^{O}. \tag{S.3}$$

I replace Equation OA.10 with Equation S.3 and estimate the system as before. As shown in Table S.39, all three point estimates are higher: $\gamma = 0.136$, $\theta = 0.252$, and $\kappa = 0.380$. The bootstrap intervals for γ and κ exclude zero, while the interval for θ reaches the lower bound. Table S.40 shows the resulting predicted effects on choices, emissions, and consumer welfare. The Attention and meat-ban emissions effects are unchanged, while the label-induced reduction changes from 74.1 to 69.4 grams per meal. The Knowledge intervention now increases emissions by 27.5 grams per meal. The higher γ changes the welfare predictions: the Attention intervention reduces choice-mediated welfare by 0.49¢per meal, whereas the Label intervention and both tax rates remain slightly welfare-increasing at the point estimates.

When offsetting leaves some perceived guilt in place, the model fits the observed responses with a higher estimate of γ. Because the parameters are estimated jointly, the estimates of θ and κ also adjust.

Table S.39. Structural estimates under partial offsetting

Parameter	Estimate	95% participant-bootstrap CI
γ	0.136	[0.093, 0.184]
θ	0.252	[0.000, 0.491]
κ	0.380	[0.009, 0.671]

Notes: Sample, GMM procedure, parameter bounds, and bootstrap intervals as in Table OA.9. This specification assumes that offsetting removes 70% of perceived guilt. γ is measured in euros per kg CO_2e.

Table S.40. Simulated effects under partial offsetting

Intervention	(1) Δ GHGE (g)	(2) [95% CI]	(3) Δ welfare (cents)	(4) [95% CI]
Knowledge	27.5	[0.0, 56.2]	18.77	[−0.06, 84.97]
Attention	−54.4	[−94.1, −13.9]	−0.49	[−1.77, 0.27]
Labels	−69.4	[−126.0, −18.7]	0.09	[−0.92, 71.43]
Tax, €155/t	−108.5	[−159.7, −61.2]	0.24	[−0.29, 0.77]
Tax, €85/t (GHGE match)	−67.8	[−109.1, −32.8]	0.31	[−0.18, 0.72]
Meat ban	−215.4	[−277.2, −156.3]	−4.04	[−6.02, −2.22]

Notes: Sample, choice set, participant weights, units, and bootstrap procedure as in Table S.38, using the partial-offsetting parameters in Table S.39. Welfare reports choice-mediated effects and excludes *F*.

S.3.9 Relation between environmental attitudes and baseline meat consumption

Table S.41. Exp. 1 average baseline beef WTP by trait (only non-veg.)

Trait	Split	N	Mean	Std. Dev.
Gen. will to sacrifice	< Median	99	0.672	1.052
	≥ Median	116	0.556	1.061
Food will to sacrifice	< Median	73	0.879	0.934
	≥ Median	142	0.471	1.090
Open to env. food labels	< Median	91	0.638	1.113
	≥ Median	124	0.588	1.016

Notes: Table shows average baseline WTP for beef meals among Exp. 1 participants, split across different characteristics. Table excludes participants who report never consuming meat. In Exp. 1, baseline beef valuation is measured as baseline WTP for the beef meal. General will to sacrifice for the environment in Exp. 1 is measured as mean agreement with the statements (with the second weighted negatively): Keeping my greenhouse gas emissions as low as possible is something… I am willing to do even if it involves higher costs and effort.… I am only willing to do if others do it as well.… that aligns with my principles.… I feel personally responsible for.", all of which are measured on a 10-point scale. Food-specific will to sacrifice is measured after the experimental meal decisions by asking how much environmental aspects matter when choosing a dish in a restaurant, student canteen, or other canteen, on a 0–9 scale. Openness to environmental food labels is measured as approval of carbon labels being installed in the canteens, measured on a 10-point scale. *N* relates to the number of participants.

Table S.42. Exp. 2 average baseline meat consumption by trait (only non-veg.)

Trait	Split	N	Mean	Std. Dev.
Gen. will to sacrifice	< Median	131	62.646	28.515
	≥ Median	174	54.142	28.724
Food will to sacrifice	< Median	129	66.625	29.370
	≥ Median	178	51.254	26.719
Open to env. food labels	No/Don't know	22	65.708	28.845
	Yes	52	52.531	28.017

Notes: Table shows average pre-intervention meat consumption (percentage of visits in which meat is consumed) among Exp. 2 student canteen guests who responded to the baseline survey and visited the student canteens at least 5 times in the pre-intervention period, split across different characteristics. Table excludes participants who never consume meat in the canteens. In Exp. 2, I measure baseline meat consumption as the percentage of pre-intervention meals that contain meat. General will to sacrifice for the environment in Exp. 2 is measured as the mean of agreement with the statements "Flying should be made more expensive since it is bad for the environment" and "It should be forbidden to build new houses that do not adhere to current environmental standards" (survey respondents only, each measured on a 7-point scale). Food-specific will to sacrifice is measured as the self-reported importance of environmental aspects in food choice, pre-intervention (survey respondents only, measured on a 7-point scale). Openness to environmental food labels is measured as post-intervention approval of the labels being installed permanently (survey respondents in treatment canteen only, measured as Yes/No/Don't Know). *N* relates to the number of survey participants/canteen guests.

S.4 Experiments 1 and 3: Details on the experimental set-up

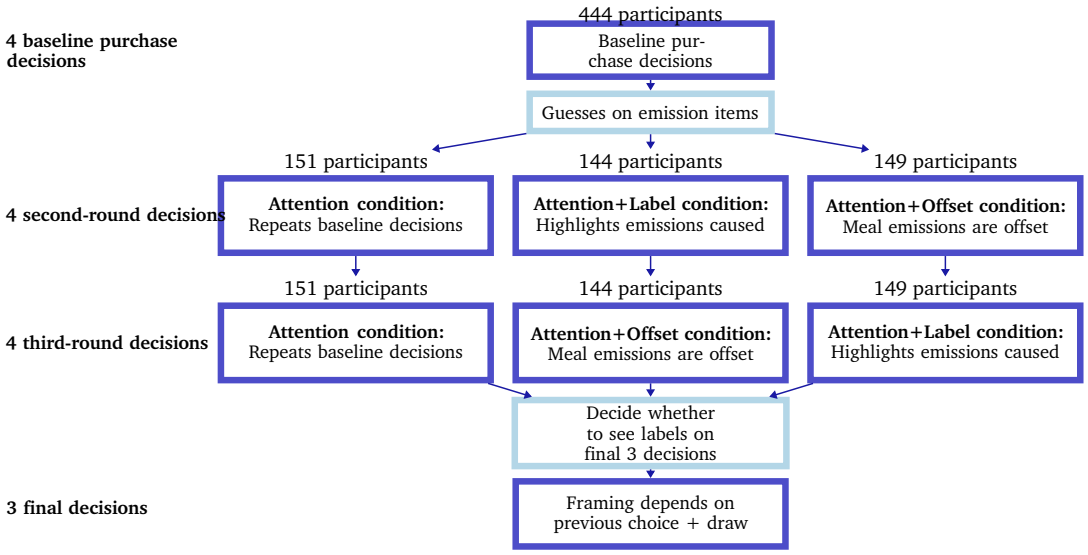


Figure S.11. Experiment 3 schedule and treatment groups

S.4.1 Pre-registration

I preregistered Experiment 3 on June 21st 2021 under #AEARCTR-0007858 and Experiment 1 on October 24th 2021 under #AEARCTR-0008435.

S.4.2 Meals used for elicitation

In the purchasing decisions in Experiments 1 and 3, participants make decisions on the same four student canteen meals. These are all meals which are regularly offered in the student canteen. Participants who indicate that they are not vegetarian decide on two vegetarian and two meat meals: Filled zucchini with potato croquettes (1.4 kg of emissions, colored yellow in the labels), Italian vegetable ragout with pasta (0.5 kg of emissions, colored green in the labels), poultry schnitzel with fruity sauce and rice (1.4 kg of emissions, colored yellow in the labels), and sliced beef with potatoes (3.4 kg of emissions, colored red in the labels). Participants who indicate they are vegetarian decide on four vegetarian meals: filled zucchini with potato croquettes (1.4 kg of emissions, colored yellow in the labels), Italian vegetable ragout with pasta (0.5 kg of emissions, colored green in the labels), cheese “Spätzle” with mushrooms (1.2 kg of emissions, colored yellow in the labels), and sweet-and-sour vegetable stir-fry with rice (0.4 kg of emissions, colored green in the labels). The cheese sandwich is the outside option to every choice, causes 0.7 kg of emissions, and is colored green on the labels.

I randomized the order in which meals appear (both in the decision and the emission estimating screens) to avoid order effects. The left-right positioning of the warm meal and cheese sandwich was reversed in half of the experimental sessions to avoid positioning effects.

S.4.3 Example of willingness-to-pay elicitation (baseline decision)

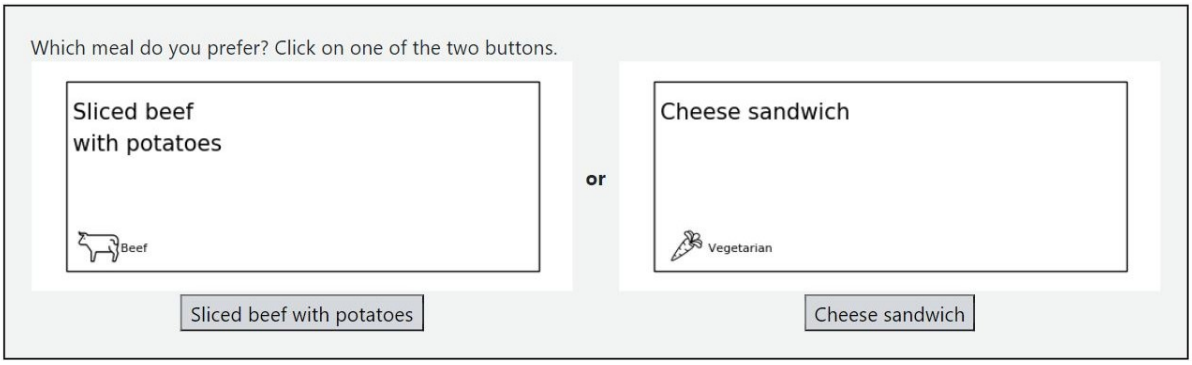


Figure S.12. Meal purchase decision example step 1

Notes: Step 1 of the purchasing decision. Depending on participants' choice, Step 2 (Figure S.13) asks for their WTP to receive or avoid the warm meal.

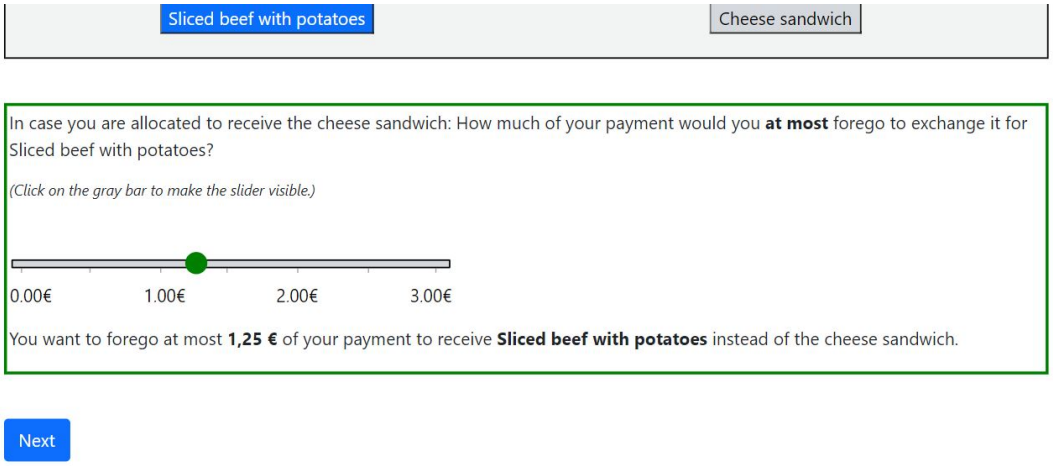


Figure S.13. Meal purchase decision example step 2

Notes: Step 2 of the purchasing decision, as shown to a participant preferring the meal in Step 1. The green window pops up underneath the slider depending on the selected meal option; Here the example shows the case of a participant preferring the beef. For a participant preferring the cheese sandwich, the wording is changed to "In case you are allocated to receive the Sliced beef with potatoes: How much of your payment would you at most forego to exchange it for the cheese sandwich? – You want to forego at most €1,25 of your payment to receive the cheese sandwich instead of the Sliced Beef with potatoes." Note that the €1.25 is an example of what a participant might indicate. The handle on the slider and the text below the slider only appear once participants click on it.

S.4.4 Incentivization of elicitations

The elicitation of participants' **WTP for consuming the meals** is incentivized with an adapted BDM mechanism: There is a 50% probability that the specific meal and a 50% probability that the cheese sandwich is randomly drawn as the default meal.



Figure S.14. Gazebo set-up on University campus

If the default meal and the preferred meal indicated in the first part of the decision (e.g. Figure S.12) coincide, the participant is given the preferred meal at zero price. If the two do not coincide, a price is randomly drawn at which the two options can be exchanged. Each value between €0.00 and €3.00 can be drawn with equal probability, in five-cent steps. If the WTP indicated by the participant in the second part of the decision (e.g. Figure S.13) is equal to or above the price drawn, the price is deducted from the participants' payment and participants are provided with the preferred option. If WTP is below the price drawn, participants are provided with the less preferred option, and no amount is deducted from participants' payments. The outcome lunch is provided to participants directly after the experiment, together with participants' payment in cash. The pay-out station is shown in Figure S.14. For this purpose, experiment participants are required to travel to the university campus immediately after completing the experiment. Fewer than 4% did not pick up their cash payment and meal. The incentivization structure was explained to participants, and they were required to pass an extensive comprehension check, which fewer than 4% of participants failed.

This **WTP for seeing labels elicitation** is incentivized with a similar BDM mechanism. There is a 50% probability that the default option is that choices are displayed with, and a 50% probability that the default option is that choices are displayed without labels. If the default display option and the preferred display option coincide, the preferred display option is implemented at zero price. If the two do not coincide, a price is randomly drawn at which the display option can be changed. Each value between €0.00 and €3.00 can be drawn with equal probability, in five-cent steps. If the WTP indicated

by the participant in the second part of the decision (similar to Figure S.13, with display options instead of meals) is equal to or higher than the price drawn, the preferred display option is implemented. The price drawn is only deducted from participants' payment if one of the final three meals is relevant for pay-out. If the WTP is lower than the price drawn, the less-preferred display option is implemented.

S.4.5 Decisions under carbon offsetting

In the ATTENTION+OFFSET condition in Experiment 3 and the OFFSET condition in Experiment 1, participants are informed that, if one of the decisions made in this treatment is implemented, the emissions of the meal provided to them (regardless of whether it is the warm meal or the cheese sandwich) are offset by the experimenter with a donation to Atmosfair. The example screens in Subsection S.4.7 show how this is communicated to experiment participants.

Towards the end of the experiment, after participants have completed all meal decisions, I elicit participants' attitudes towards the effectiveness of carbon offsetting and ask for participants' prior experiences with carbon offsetting. Tables S.43 and S.44 show descriptives pooled across Experiments 1 and 3. Across all participants, 73% had heard of carbon offsetting previously and 33% had used carbon offsetting themselves.

Table S.44 shows that participants broadly agree with carbon offsetting being effective (measured as agreement with the statement "Voluntary carbon offsetting is an effective climate protection measure"). They disagree with carbon offsetting replacing other climate protection measures (measured as agreement with the statement "If I offset emissions for a carbon-intensive activity such as a flight, it is okay to book another flight."). They agree that carbon offsetting should not replace other climate protection activities (measured as agreement with the statement "Carbon offsetting cannot replace personal efforts to protect the climate."). Interestingly, having experienced the ATTENTION+OFFSET or the OFFSET condition earlier in the experiment increases support for the second and decreases support for the third statement. Perhaps this is due to the information on carbon offsetting provided to them during the experiment (see section S.4.7). Carbon offsetting in the experiment occurs through the NGO Atmosfair, which does not engage in the often-criticized practice of rewarding credits for planting or protecting forests. Instead, Atmosfair supports the development and deployment of renew-

able energy in developing countries. See the explanations in the experiment screenshots (section S.4.7) and https://www.atmosfair.de/en/ for details.

These descriptive statistics convey that carbon offsetting likely removes a part of environmental guilt, but may not be removing it entirely. Accordingly, Supplementary Material Section S.3.8.2 re-estimates the model and repeats the simulation assuming that only 70% of environmental guilt is removed.

Table S.43. Familiarity with carbon offsetting

	Familiarity with offsetting	
	(1) Heard of	(2) Have used
In offset condition earlier in exp.	-0.04	-0.01
	(0.03)	(0.04)
Constant	0.75***	0.34***
	(0.03)	(0.03)
Observations	731	731

Notes: Participant responses to whether they had heard of and used carbon offsetting. The constant reports the mean among participants who had not experienced an offsetting treatment earlier in the experiment; the pooled percentages are reported in the surrounding text. Heteroskedasticity-robust standard errors are reported.

Table S.44. Beliefs on carbon offsetting effectiveness

	Beliefs about offsetting		
	(1) Effective	(2) Can replace	(3) Cannot replace
In offset condition earlier in exp.	0.15	0.44***	-0.50***
	(0.18)	(0.15)	(0.17)
Constant	5.55***	2.86***	8.14***
	(0.14)	(0.12)	(0.12)
Observations	731	731	731

Notes: Participant responses to three beliefs about carbon offsetting, recorded on a 10-point scale. Heteroskedasticity-robust standard errors are reported.

S.4.6 Experiment 3 guessing screen example

Figure S.11 reports the Experiment 3 timeline; Figure S.15 shows an example of the emissions-guessing screen.

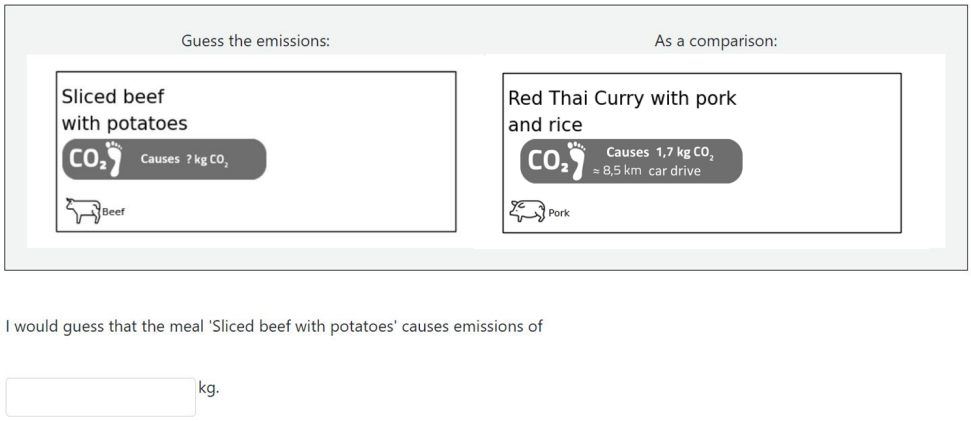


Figure S.15. Example guessing questions

S.4.7 Experiment screens (English translation from German)

The reproduced emission-guessing instruction and example screens below come from the 70-person emissions-guessing subgroup in Experiment 1, which answered eleven questions. Experiment 3 used the same procedure but omitted the spaghetti-with-meat item and therefore included ten questions.

Survey start screen

Welcome to the BonnEconLab online study. Please note that you may only take part in this study once. Furthermore, you may only take part if you have registered for the study in our participation database. Please complete this survey on your computer. Participation with mobile devices such as smartphones or tablets is not possible. The payout for this experiment will be done using your personal participant code: 12pI2q5vh Please write down your code! You will need approximately 45 minutes to process this survey. After fully completing the survey, you can collect your payout at our location on the central university lawn (see map below) until 2 p.m. today. You will not be able to receive your payout at any other time! In this experiment, your payout consists of several components:

- You receive exactly one dish (your lunch).
- You receive an experiment payment of €9.00 in cash.
- You may receive an additional payout of up to €1.60 in addition to the expense compensation. This depends on your answers in the marked part of the study.
- In addition, chance determines whether, depending on your answers in another (also clearly marked) part of the study, you will receive another additional payout of up to €1.10.

Payment will be made in the BonnEconLab pavilion on the central university lawn (Regina-Pacis-Weg). You will find us at the place marked with a blue arrow under a fabric gazebo.

Decision description screen - Screenshot 1/4

Description of upcoming decisions | Comprehension questions

The second part of the study is about to begin. Your decisions in this part of the study will affect your experiment payment and the meal you receive.

On this page you will find explanations and examples. On the following page we will check your understanding of these explanations. By clicking on the tab above you can switch between the two pages.

Once the comprehension questions have been answered correctly, you can proceed with further work on the survey.

How do your decisions affect your payout?

- In this experiment, your payout consists of three components:
  - You receive exactly one meal (your lunch).
  - You receive an experiment payment. At the moment, the experiment payment is €9.00. You will make a total of 15 decisions over the course of this study. For each of these decisions, you have the option of waiving part of the experiment payment (maximum €3.00). For that, you will receive a meal you prefer.
  - In two other parts of the study, you may receive an additional amount of up to €1.60 in addition to the experiment payment, depending on your answers. In addition, depending on your answers in a third part of the study, chance will determine whether you will receive an additional amount of up to €1.10. The relevant parts of the study are clearly marked.
- For each of the 15 decisions, indicate which of the two meals you prefer. Then specify the maximum amount of your experiment payment you would like to forgo in order to receive the preferred meal.
- Of the **15** decisions you make, only **one** is actually implemented.

The decision that is implemented shall be subject to the following:

- Chance decides whether you will receive your favorite meal for free:
  - Case 1 (50% probability): You will receive your favorite meal for free.
  - Case 2 (50% probability): You will be assigned the non-preferred meal first. In this case, specify the maximum amount of your experiment payment you would like to forgo in order to receive your favorite meal instead.

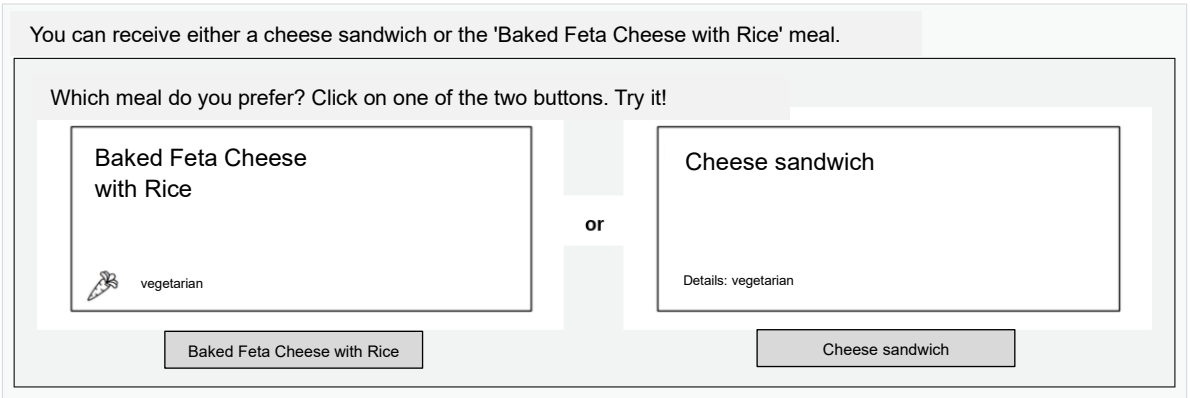
- If case 2 occurs, it is again a matter of chance:
  - A **surcharge** is determined at random. Any value between €0 and €3 (in 5 cent increments) is equally probable.
  - If the amount you have declared is more than the surcharge, you will receive your preferred meal. For this, the surcharge will be deducted from your experiment payment.
  - If the amount you specify is less than the surcharge, you will receive the non-preferred meal free of charge.

For the other 14 decisions which are not being implemented, the following rules apply:

- These decisions have no effect on the meal you receive.
- These decisions have no effect on your compensation.

You will not know which of the 15 decisions will be implemented until the end of the study. It is therefore in your best interest to make every decision carefully.

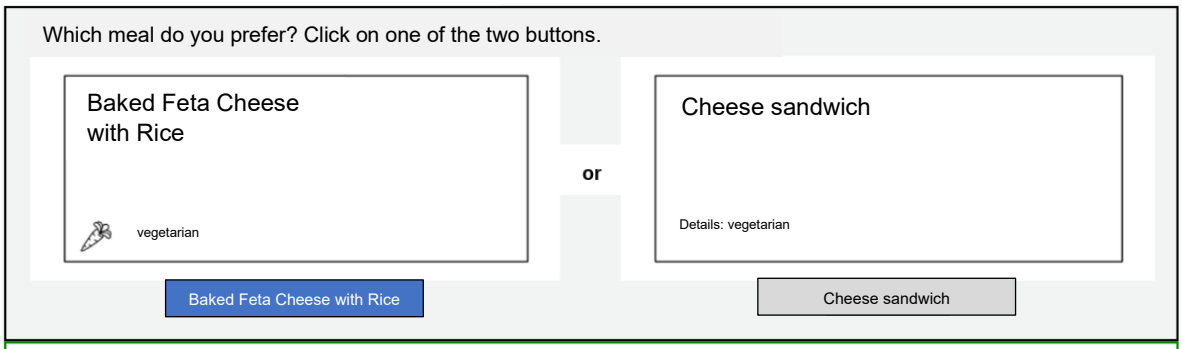
Example decision



Decision description screen - Screenshot 2/4

Example scenario 1

Assuming you made the following decision:



If you are given the cheese sandwich: What is the maximum amount of your experiment payment you would be willing to give up in exchange for Baked Feta Cheese with Rice?

(Click on the gray bar to make the slider visible).



You would like to give up a maximum of **€1.20** of your experiment payment to receive the meal **Baked Feta Cheese with Rice** instead of the cheese sandwich.

Here's what happens in this example (which you have no control over):

- You are first assigned your **less preferred meal**, the cheese sandwich.
- A **surcharge of €0.60** is randomly determined.

This means for you:

The surcharge with the amount of €0.60 is lower than the maximum amount of €1.20 you specified. You will receive the meal 'Baked feta cheese with rice'. For this, €0.60 will be deducted from your experiment payment.

Decision description screen - Screenshot 3/4

Example scenario 2

Assuming you made the following decision:



Here's what happens in this example (which you have no control over):

- You are first assigned your **less preferred meal**, the cheese sandwich.
- A **surcharge of €2.00** is randomly determined.

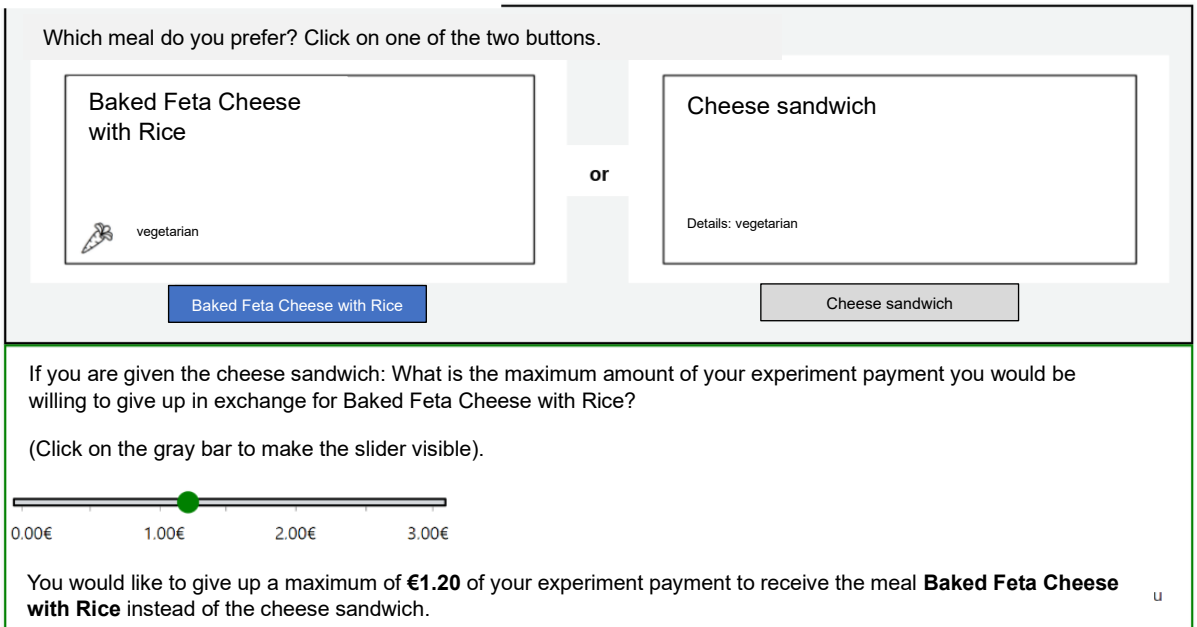
This means for you:

The surcharge with the amount of €2.00 is higher than the maximum amount of €1.20 you specified. You will receive the cheese sandwich. Therefore, nothing will be deducted from your experiment payment.

Decision description screen - Screenshot 4/4

Example scenario 3

Assuming you made the following decision:



If you are given the cheese sandwich: What is the maximum amount of your experiment payment you would be willing to give up in exchange for Baked Feta Cheese with Rice?

(Click on the gray bar to make the slider visible).

You would like to give up a maximum of **€1.20** of your experiment payment to receive the meal **Baked Feta Cheese with Rice** instead of the cheese sandwich.

Here's what happens in this example (which you have no control over):

- You are assigned your **preferred meal**, 'Baked feta cheese with rice', for free.

This means for you:

You receive the meal 'Baked feta cheese with rice'. Nothing will be deducted from your experiment payment.

Continue to the questions

You can always return to this page while answering the questions.

Comprehension questions - Screenshot 1/3

Description of upcoming decisions | Comprehension questions

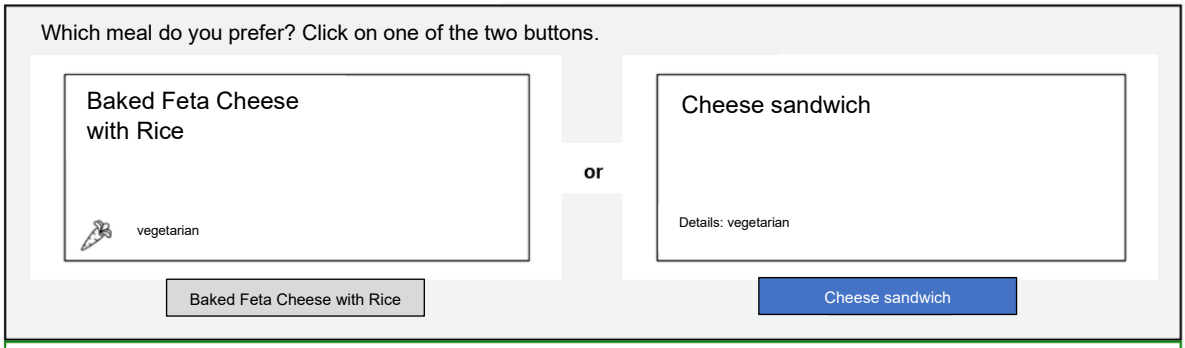
Comprehension questions

Please answer the following comprehension questions. If you want to look at the description of the survey again, you can switch back and forth between this page and the previous page by clicking on the tab at the top.

After correctly answering the comprehension questions, you can continue with the further processing of the survey.

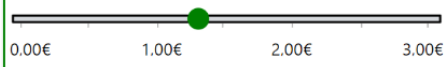
Question 1

Assuming you made the following decision:



If you are given Baked Feta Cheese with Rice: What is the maximum amount of your experiment payment you would be willing to give up in exchange for the cheese sandwich?

(Click on the gray bar to make the slider visible).



You would like to give up a maximum of **€1.30** of your experiment payment to receive the meal **Cheese sandwich** instead of the Baked Feta Cheese with Rice.

Here's what happens in this example (which you have no control over):

- The decision was carried out.
- You are first assigned your **less preferred meal**, the Baked Feta Cheese with Rice.
- A **surcharge of €0.70** is randomly determined.

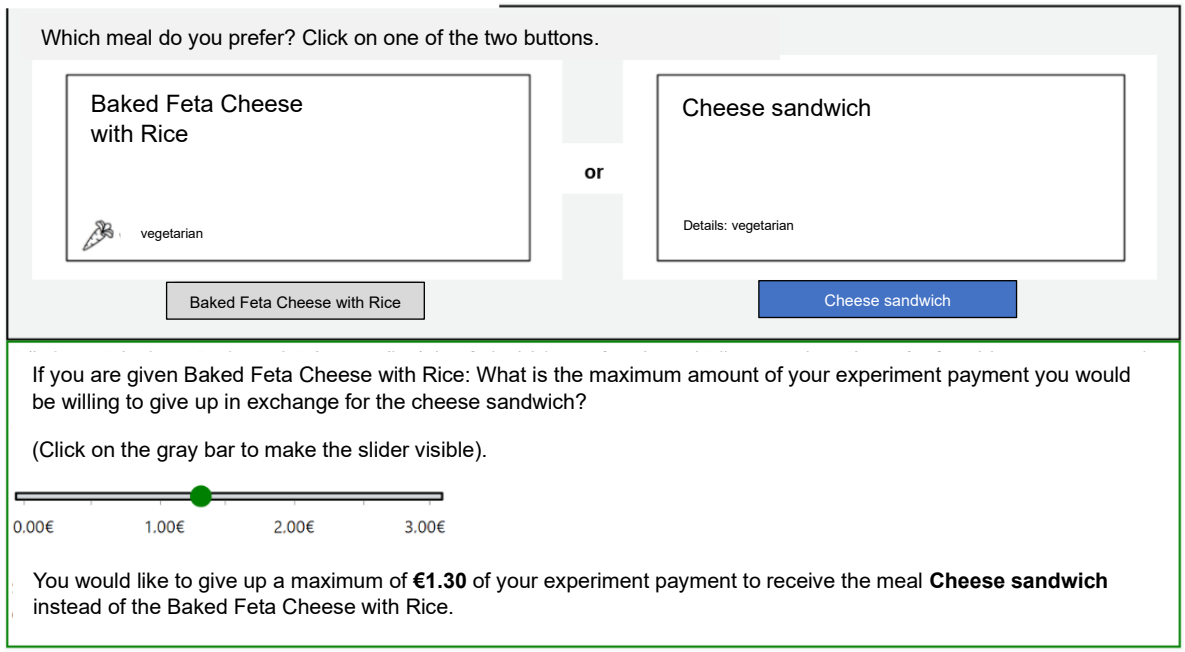
What do you receive?

- ○ The baked feta cheese with rice and your full experiment payment.
- ○ The baked feta cheese with rice and €0.70 will be deducted from your experiment payment.
- ○ The cheese sandwich and €0.70 will be deducted from your experiment payment.
- ○ The cheese sandwich and your full experiment payment.

Comprehension questions - Screenshot 2/3

Question 2

Assuming you made the following decision:



Here's what happens in this example (which you have no control over):

- The decision was carried out.
- You are assigned your **preferred meal**, the cheese sandwich.

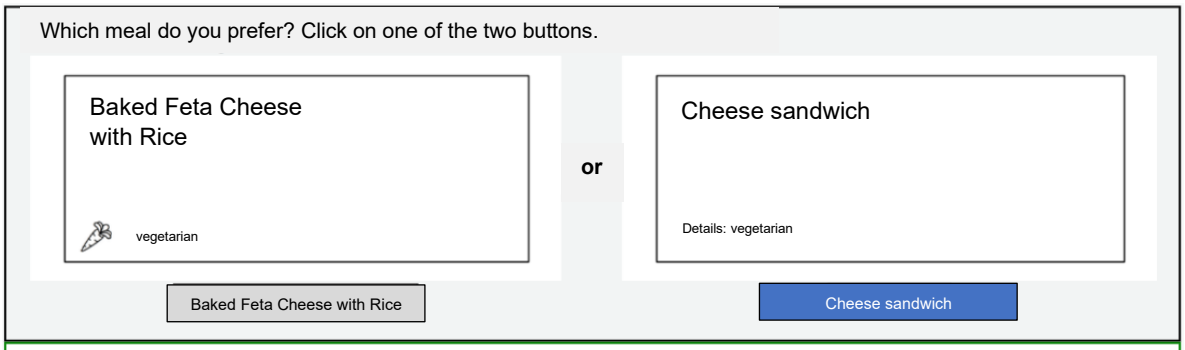
What do you receive?

- ○ The baked feta cheese with rice and your full experiment payment.
- ○ The baked feta cheese with rice and €0.70 will be deducted from your experiment payment.
- ○ The cheese sandwich and €0.70 will be deducted from your experiment payment.
- ○ The cheese sandwich and your full experiment payment.

Comprehension questions - Screenshot 3/3

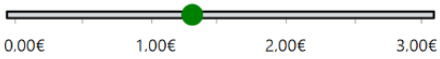
Question 3

Assuming you made the following decision:



If you are given Baked Feta Cheese with Rice: What is the maximum amount of your experiment payment you would be willing to give up in exchange for the cheese sandwich?

(Click on the gray bar to make the slider visible).



You would like to give up a maximum of **€1.30** of your experiment payment to receive the meal **Cheese sandwich** instead of the Baked Feta Cheese with Rice.

Here's what happens in this example (which you have no control over):

- The decision was carried out.
- You are first assigned your **less preferred meal**, the Baked Feta Cheese with Rice.
- A **surcharge of €2.70** is randomly determined.

What do you receive?

- ○ The baked feta cheese with rice and your full experiment payment.
- ○ The baked feta cheese with rice and €2.70 will be deducted from your experiment payment.
- ○ The cheese sandwich and €2.70 will be deducted from your experiment payment.
- ○ The cheese sandwich and your full experiment payment.

Question 4

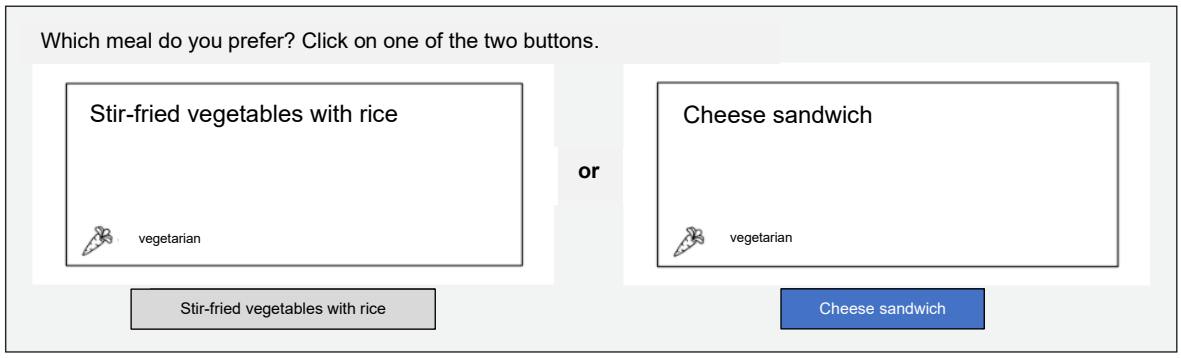
How many of the 15 decisions actually have an impact on the meal you are handed and your experiment payment?

- ○ All the 15 decisions have an impact.
- ○ Five of the 15 decisions have an impact.
- ○ One of the 15 decisions has an impact.
- ○ None of the 15 decisions has an influence.

Back to the explanation | Continue with the rest of the survey

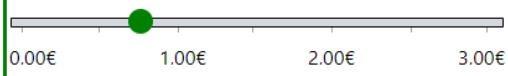
Example baseline decision

You can receive either a cheese sandwich or the meal 'Stir-fried vegetables with rice' with your payout.



If you are given the Stir-fried vegetables with rice: What is the **maximum** amount of your experiment payment you would be willing to give up in exchange for the cheese sandwich?

(Click on the gray bar to make the slider visible).



You would like to give up a maximum of **€0.75** of your experiment payment to receive the meal **cheese sandwich** instead of the Stir-fried vegetables with rice.

Continue

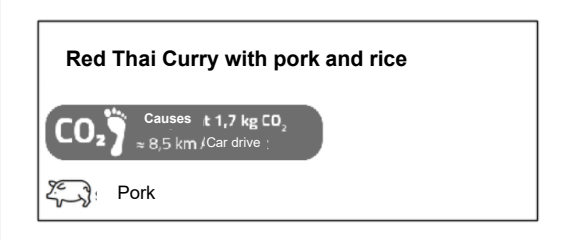
Experiment 1 emission-guessing instructions (11 questions)

You will now estimate the greenhouse gas emissions (in CO2 equivalents) caused by a total of eleven meals.

- You have **60 seconds** to answer each question.
- For each question in which your guess does not deviate from the correct value by more than 30%, **€0.10 is added to your payout**.

During each guessing question you will be shown the emissions caused by the meal “Red Thai Curry with pork and rice” **as a reference value**.

Your reference value:



What assumptions should be made for the guessing questions?

For the following questions you will not be shown any ingredient lists or a description of the origin of the ingredients. This is because we only want to give you the information which you would normally find in a restaurant. We would like to know how you, based only on the name of the meal on the menu, guess the magnitude of the emissions caused by a meal.

Of course, the emissions of a seemingly identical meal can differ, e.g., depending on the exact ingredients and depending on whether the ingredients were produced in an ecologically sustainable or in a conventional manner. Please assume a conventional production and a conventional meal preparation – just like you would expect it, if you are offered such a meal without any further information in a restaurant.

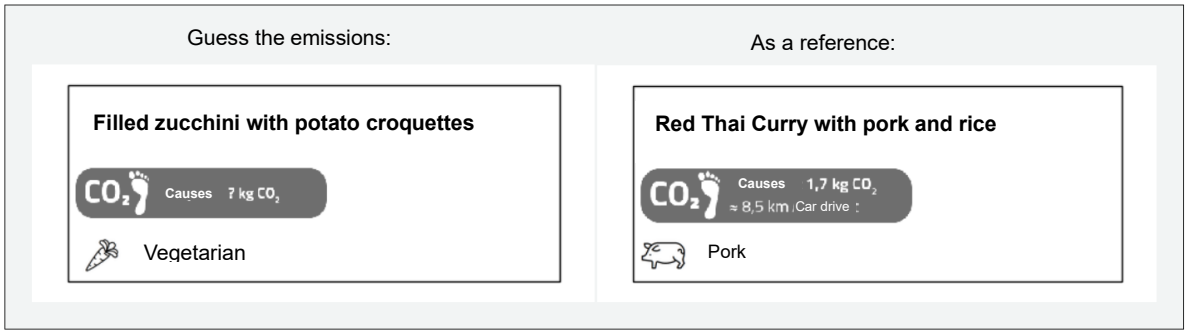
Please take into account all emissions caused in the agricultural production and in food processing, packaging, conservation and transport of ingredients, up until an ingredient can be purchased in the store. You do not need to take into account emissions which are caused by the transport of ingredients from store to restaurant

Continue

Experiment 1 example emission-guessing question

Remaining time on this page. **0:54**

What do you estimate: How high are the greenhouse gas emissions (in CO2-equivalents), which are caused by the meal “Filled zucchini with potato croquettes”?



I estimate that the meal “Filled zucchini with potato croquettes” causes emissions of

kg.

Continue

Introduction decisions with labels

You will now make four more of the 15 decisions. One of the 15 decisions will be implemented.

You will be shown the greenhouse gas emissions (in CO2 equivalents) of both meals for the upcoming decisions.

For those interested: More information on the calculation of greenhouse gas emissions:

What assumptions are made in the calculation?

In the calculation, the emissions attributable to a meal are calculated as the sum of the emissions generated in the production of the ingredients. The emissions of each ingredient are calculated "from farm to gate", i.e. all emissions are included that occur during agricultural production and during further processing, packaging, preservation and transport until the ingredient is available for purchase in shops. Not included are the transport from the shop to the restaurant or end consumer and the emissions that arise from any further refrigeration in the restaurant or at the end consumer, as well as the emissions that arise from cooking the meal.

When calculating the values, conventional (i.e. not specifically organically certified) agriculture is assumed. Otherwise, assumptions are made about production that reflect the production of the average product found on our supermarket shelves.

What data is the calculation based on?

The Eaternity database on which the calculations are based is currently the largest and most comprehensive database for calculating the climate-relevant emissions of meals and food products. It includes more than 550 ingredients and other parameters on organic and greenhouse production as well as production, processing, packaging and preservation. The Eaternity database is maintained by scientists from the Zurich University of Applied Sciences (ZHAW), the University of Zurich (UZH), the Swiss Federal Institute of Technology Zurich (ETH Zurich), the Research Institute of Organic Agriculture (FiBL), Quantis and other institutions.

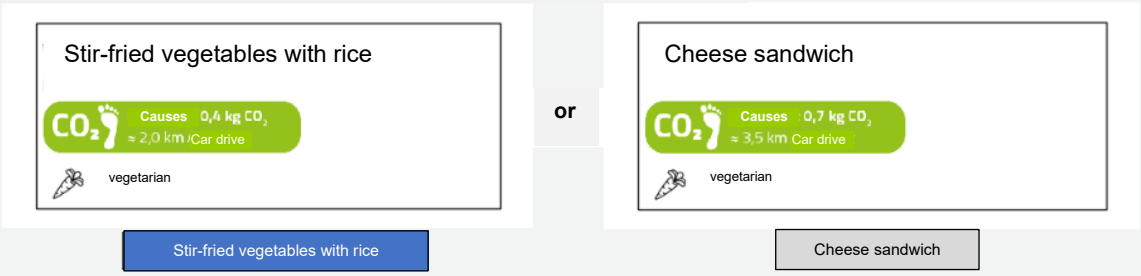
Source: Eaternity.

Continue

Example decision with labels

You can either get a cheese sandwich or the meal 'stir-fried vegetables with rice' with your payout.

Which meal do you prefer? Click on one of the two buttons.



Stir-fried vegetables with rice | Cheese sandwich

If you are given the cheese sandwich: What is the **maximum** amount of your experiment payment you would be willing to give up in exchange for stir-fried vegetables with rice?

(Click on the gray bar to make the slider visible).



You would like to give up a maximum of **€1.10** of your experiment payment to receive the meal **stir-fried vegetables with rice** instead of the cheese sandwich.

Continue

For those interested: More information on the calculation of greenhouse gas emissions:

What assumptions are made in the calculation?

In the calculation, the emissions attributable to a meal are calculated as the sum of the emissions generated in the production of the ingredients. The emissions of each ingredient are calculated "from farm to gate", i.e. all emissions are included that occur during agricultural production and during further processing, packaging, preservation and transport until the ingredient is available for purchase in shops. Not included are the transport from the shop to the restaurant or end consumer and the emissions that arise from any further refrigeration in the restaurant or at the end consumer, as well as the emissions that arise from cooking the meal.

When calculating the values, conventional (i.e. not specifically organically certified) agriculture is assumed. Otherwise, assumptions are made about production that reflect the production of the average product found on our supermarket shelves.

What data is the calculation based on?

The Eaternity database on which the calculations are based is currently the largest and most comprehensive database for calculating the climate-relevant emissions of meals and food products. It includes more than 550 ingredients and other parameters on organic and greenhouse production as well as production, processing, packaging and preservation. The Eaternity database is maintained by scientists from the Zurich University of Applied Sciences (ZHAW), the University of Zurich (UZH), the Swiss Federal Institute of Technology Zurich (ETH Zurich), the Research Institute of Organic Agriculture (FiBL), Quantis and other institutions.

Source: Eaternity.

Introduction decisions with offsetting

You will now make four more of the 15 decisions. One of the 15 decisions will actually be implemented.

If it is one of the now following four choices that is implemented, the greenhouse gas emissions of the meal you have been handed will be offset by a donation to the NGO Atmosfair. This happens regardless of whether the meal was originally assigned to you or whether you exchanged it for the other meal by paying a surcharge. Atmosfair uses the donation to support sustainable energy projects so that the emissions are saved elsewhere. In this way, the meal handed out to you becomes emission-neutral / carbon-neutral.

For those interested: Further information on carbon offsetting:

How does the carbon offset work?

The donation to Atmosfair is used to develop renewable energies in countries where they hardly exist yet, i.e. mainly in developing countries. In this way, Atmosfair saves CO2 that would otherwise have been produced by fossil energies in these countries.

Example projects

- Atmosfair uses donations to reduce the selling price of energy-efficient stoves in Nigeria. In Nigeria, 75% of families cook on open fires, and a family of 7 consumes 5 tonnes of wood per year. This enormous consumption of firewood has already led to almost total deforestation and the progressive spread of deserts, especially in the poor north of the country. Energy-efficient stoves use about 80% less wood.
- Atmosfair uses donations to make small-scale biogas plants more affordable in Nepal. This project targets families living in rural areas who previously used wood as an energy source for cooking. In this way, the increasing deforestation of Nepal's forests can be counteracted.
- Atmosfair uses donations to support a small hydropower plant in Honduras. In this way, four villages that previously used wood and diesel generators for energy supply could be connected to the electricity grid for the first time. In addition, electricity can be fed into the national grid, replacing electricity from gas-fired power plants.

Source: Atmosfair

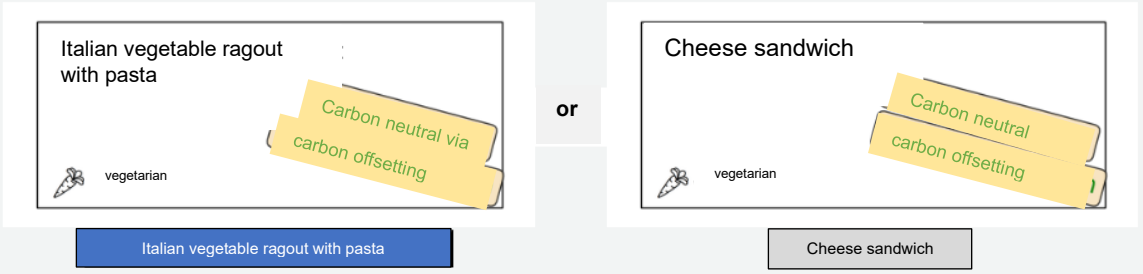
Continue

Example decision with offsetting

You can either receive a cheese sandwich or the meal 'Italian vegetable ragout with pasta' with your payout.

The emissions attributable to each meal are offset by a donation to the NGO Atmosfair. Atmosfair supports sustainable energy projects with the donation, so that the emissions are saved elsewhere.

Which meal do you prefer? Click on one of the two buttons.



Italian vegetable ragout with pasta | Cheese sandwich

If you are assigned the cheese sandwich: What is the **maximum** amount of your experiment payment that you would be willing to give up in exchange for Italian vegetable ragout with pasta? (Click on the gray bar to make the slider visible).



You would like to give up a maximum of **€0.75** of your experiment payment to receive the **Italian vegetable ragout with pasta** instead of the cheese sandwich.

Continue

For those interested: Further information on carbon offsetting:

How does the carbon offset work?

The donation to Atmosfair is used to develop renewable energies in countries where they hardly exist yet, i.e. mainly in developing countries. In this way, Atmosfair saves CO2 that would otherwise have been produced by fossil energies in these countries.

Example projects

- Atmosfair uses donations to reduce the selling price of energy-efficient stoves in Nigeria. In Nigeria, 75% of families cook on open fires, and a family of 7 consumes 5 tonnes of wood per year. This enormous consumption of firewood has already led to almost total deforestation and the progressive spread of deserts, especially in the poor north of the country. Energy-efficient stoves use about 80% less wood.
- Atmosfair uses donations to make small-scale biogas plants more affordable in Nepal. This project targets families living in rural areas who previously used wood as an energy source for cooking. In this way, the increasing deforestation of Nepal's forests can be counteracted.
- Atmosfair uses donations to support a small hydropower plant in Honduras. In this way, four villages that previously used wood and diesel generators for energy supply could be connected to the electricity grid for the first time. In addition, electricity can be fed into the national grid, replacing electricity from gas-fired power plants.

Source: Atmosfair

Introduction calorie estimation

You will now estimate the energy value of each meal in kilocalories (kcal) for a total of five meals. For each estimation question, the completion time is **limited to 60 seconds**. For each estimation question where your estimate does not deviate from the correct value by more than 30%, **your payout increases by €0.10**.

What assumptions should be made for the estimation?

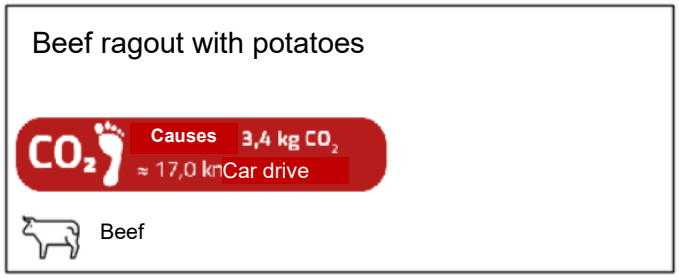
You will not be presented with ingredient lists for the following estimation questions. This is because we want to give you, as much as possible, only the information that you would find in the restaurant. We want to know how you estimate the energy value of a meal, based solely on the name of the meal in the menu.

Continue

Example calorie estimation

Remaining time on this page. **0:54**

What do you estimate: What is the energy value in kilocalories (kcal) of the meal ‘Beef ragout with potatoes’?



I estimate that the meal ‘Beef ragout with potatoes’ has

kcal.

Continue

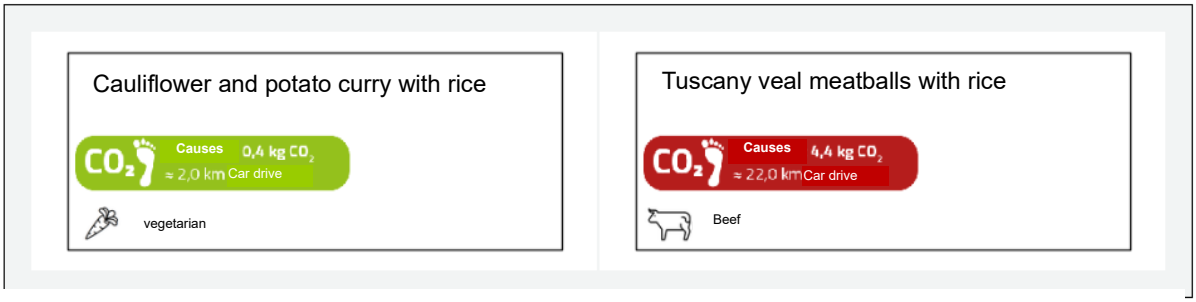
WTP for label presence

You are about to make the last three of the 15 decisions. One of the 15 decisions will actually be implemented.

But now there are two differences:

1. There are now three **new meals** that you have not seen in your previous decisions.
2. You can see **emissions labels** for these three meals. These labels show the greenhouse gas emissions of the meals in CO2 equivalents.

For example, two of the labels might look like this:



The display of the labels can either be preset so that:

- The labels are displayed to you, or
- The labels are not displayed to you.

Chance decides whether the display setting of the labels corresponds to your wishes without charge.

- Case 1 (probability 50%): We (do not) display the labels according to your wishes.
- Case 2 (probability 50%): The labels are initially preset so that it does not correspond to your wishes. For this case, you specify the **maximum** amount of your experiment payment you would like to give up in order to get your preferred display setting instead.

If **case 2** occurs, chance decides again:

- A **price** is determined randomly. Every value between €0 and €3 (in 5 cent steps) is equally probable.
- If the given amount is **higher** than the price, you will still get your **preferred display setting**. For this**, the charge will be deducted from your experiment payment**. However, this will only happen if one of the three meals shown actually determines your payout.
- If the specified amount is **less** than the price, you will receive your **non-preferred display setting for free**.

Which display settings do you prefer? Click on one of the two buttons.

Labels should be **shown** | Labels should **not be shown**

If the display of labels is **not** preset and one of the three decisions you are about to make actually determines your payout: What is the **maximum** amount of your experiment payment you would like to give up in order to have the labels displayed?

0.00€ 1.00€ 2.00€ 3.00€

You want to give up a maximum of **€1.70** of your experiment payment to **enable the display of labels**.

Continue

S.5 Experiment 2: Details on the experimental set-up

S.5.1 Pre-registration

I preregistered Experiment 2 on the 25th of April 2022 on aspredicted #95108.

S.5.2 Canteen set-up in Bonn

The natural field experiment was conducted in the student canteens of the University of Bonn from April 2022 to July 2022. The whole of April (four weeks) served as a pre-intervention phase in which baseline consumption decisions were observed. Emission labels were introduced in the treatment student canteen from the beginning of May to mid-June 2022 (seven weeks). From mid-June to mid-July 2022 (three weeks, which ended with the summer closing of the treated student canteen), I continue to observe consumption decisions to examine post-intervention behavior.

There are three student canteens in Bonn: The treatment student canteen, the first control restaurant (located 1.7 km from the treatment restaurant), and the second control restaurant (located 4.7 km from the treatment restaurant and frequented much less than the other two restaurants). Menu planning is centralized among the three student canteens, and there is thus a large overlap in the daily offering. All three student canteens offer two main meal components, which differ daily but are mostly the same across student canteens. In addition, each of the student canteens might offer additional options, which are student-restaurant-specific. The larger control restaurant sometimes offers pizza or pasta in addition, and all student canteens might serve leftover main meal components from the previous day, soup, and side dishes. The canteens did not adjust their centralized menu planning in response to the labeling intervention, and deviations from the common menu did not become more frequent in the treatment canteen during or after the intervention.

In the treatment restaurant, only the main meal components were equipped with carbon labels, and sides and leftover main meal components were not labeled. [82] Cor-

82. The main reason for this was that I wanted to test carbon labeling in a manner that was feasible for the student canteen to implement long-term. While main meal components are planned and known beforehand, sides and leftover dishes are decided spontaneously. Further, leftover main meal components only make up a smaller part of daily sales and the emissions caused by side dishes are relatively small compared to those of the main meal components, with main meal components accounting for about 70% of lunchtime emissions. Sales of all products are tracked, and label effects in the main sample are conservatively calculated over all main meal compo-

respondingly, the dependent variable in my main regression is whether the main meal component a restaurant guest chooses contains meat or is vegetarian.

The price of the main meat meal offered in the sample period ranged from €1.85 to €2.5, while that of the main vegetarian meal ranges from €1.35 to €2.4. The student canteen in Bonn is a non-profit organization aimed at serving the students. Prices are thus fixed to reflect the cost of the ingredients plus a part of the canteens' fixed costs (Studierendenwerk Bonn, 2025). The remaining fixed costs are financed through subsidies. In fast food restaurants located in the surrounding area, meals were priced at €4.00 upward during the sample period.

S.5.3 Canteen visiting patterns

An average student canteen guest visited the student canteen 10 times from April to mid-July. Around 34% visit 10 times or more, and around 15% visit 20 times or more. 92% of guests visited the same student canteen at least 80% of the time. In the pre-intervention survey (described in section S.5) over 40% of respondents report that they would have difficulty finding an affordable meal if the student canteens did not exist. This suggests that switching between student canteens and other gastronomic offers is not frequent. The labeling intervention does not seem to have led to a decrease in student canteen sales relative to the control restaurant (see Figure S.5). Section S.2.4 and Table S.47 examine whether the intervention affected regular canteen guests' likelihood of visiting a canteen. The estimated effect during the labeling period is negative but statistically insignificant.

Do canteen guests regularly frequent multiple canteens? Figure S.16 includes an analysis based on the trackable personal card payments. I classify restaurant guests as "Treatment" or "Control" visitors based on their consumption behavior in the first two weeks. 91% of those regularly frequenting canteens during these two weeks (i.e. at least twice) visit the same canteen at least 80% of the time. I classify guests as "Control" or "Treatment" guests based on these two weeks. Around 2% of purchases made by "Control" visitors are made in the treated restaurant in the remaining 12-week period, while around 4% of the canteen visits of those classified as "Treatment" guests are to one of the control canteens. Figure S.16 calculates weekly statistics on switching and shows

nents offered, i.e. including main meal components spontaneously added to the menu but not labeled.

time trends. The proportion of control-canteen purchases made by guests classified as "Treatment" fluctuates between 3% and 7% across weeks, while the proportion of treatment-canteen purchases made by guests classified as "Control" fluctuates between 2% and 3%. Thus, switching between canteens overall does not seem very prevalent. While switching seems to fluctuate a bit around weeks 5 and 6 of the intervention period, this pattern does not correlate with any increase or decrease in meat consumption when eating at the "non-home" canteen, as shown in Figure S.17. It thus seems unlikely that these small fluctuations are driven by guests wanting to consume higher-emission meals without seeing carbon labels.

Because treatment status remains defined by pre-intervention canteen affiliation, switching does not mechanically alter treatment assignment; any such switching is therefore incorporated in the ITT effect. While I base classification as a "Treatment" or "Control" guest on only two weeks of data in the descriptive analysis shown here (to be able to show a two-week pre-trend), I base classification on the entire 4-week pre-intervention period in the main text, reducing noise. In the specifications additionally including guest fixed effects, guest fixed effects absorb stable differences in meat-choice propensity, but changes in attendance can still alter guests' relative contribution to the transaction-weighted estimate.

The introduction of carbon labels in the treatment restaurant was displayed as a measure taken by the student canteens themselves, with no connection presented to the University of Bonn or me specifically as the researcher. The introduction of the emission labels was explained on billboards and leaflets available inside the student canteen, as shown in Figure S.18. I conducted two surveys accompanying the measure, one before the intervention period and one after the intervention period. The surveys and the labeling measures were advertised through different channels, and the survey was advertised as a chance to voice one's opinion on the offer of the student canteen. It is thus unlikely that restaurant guests drew a connection between the initiative and the survey.

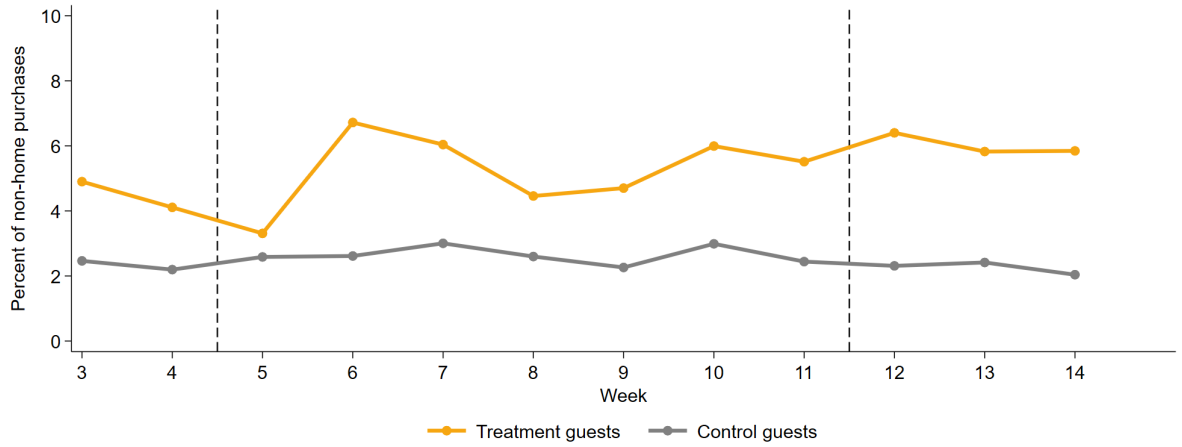


Figure S.16. Visits to the "non-home" canteen

Notes: In percent relative to total canteen visits. Classification as the "home" canteen is based on behavior in the first two weeks. The sample is similar to that in Table OA.17, Panel B, Column (1), but the classification as treatment is calculated based entirely on the first two weeks, based on a minimum of two visits during this period. Based on $N = 44,816$ observations in weeks 3–14.

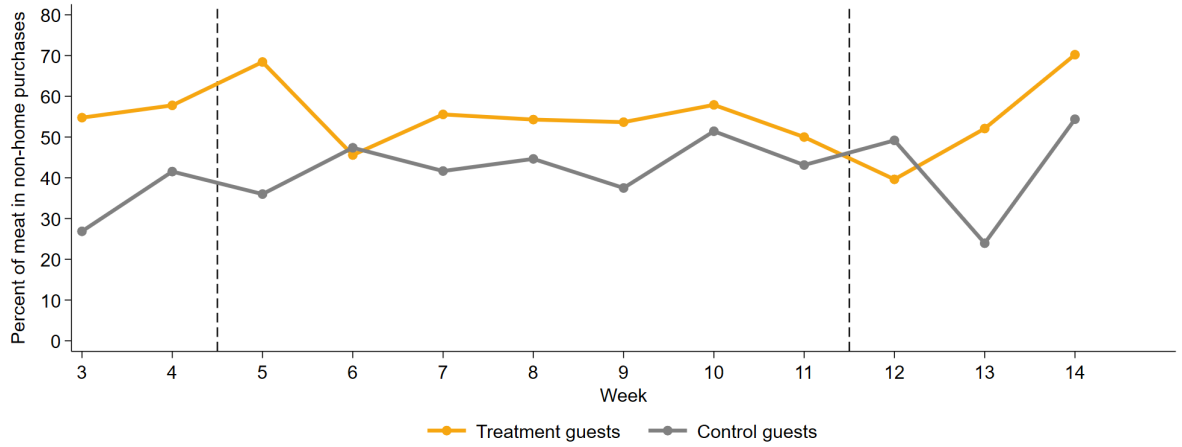


Figure S.17. Meat consumption in the "non-home" canteen

Notes: Graph shows average meat consumption at the non-home canteen. Classification as the "home" canteen is based on behavior in the first two weeks. The sample is similar to that in Table OA.17, Panel B, Column (1), but the classification as treatment is calculated based entirely on the first two weeks, based on a minimum of two visits during this period. Based on $N = 1,274$ observations of guests dining in a canteen not in line with their classification of treated or control in weeks 3–14.

S.5.4 Carbon label calculation

For the carbon labels, I calculated emission values with the application Eaternity Institute (2020), using ingredient lists provided by the student canteen. The design of the carbon labels was proposed by the student canteen, based on what is technically feasible and possibly implementable as a long-run measure. Examples are shown in Figure 4.

They were coded in a traffic-light system, with thresholds determined such that approximately a third of the main components offered by the student canteen during the study period would be classified as green, one-third as yellow, and one-third as red. This corresponded to thresholds of 0.5 kg and 1 kg. [83]

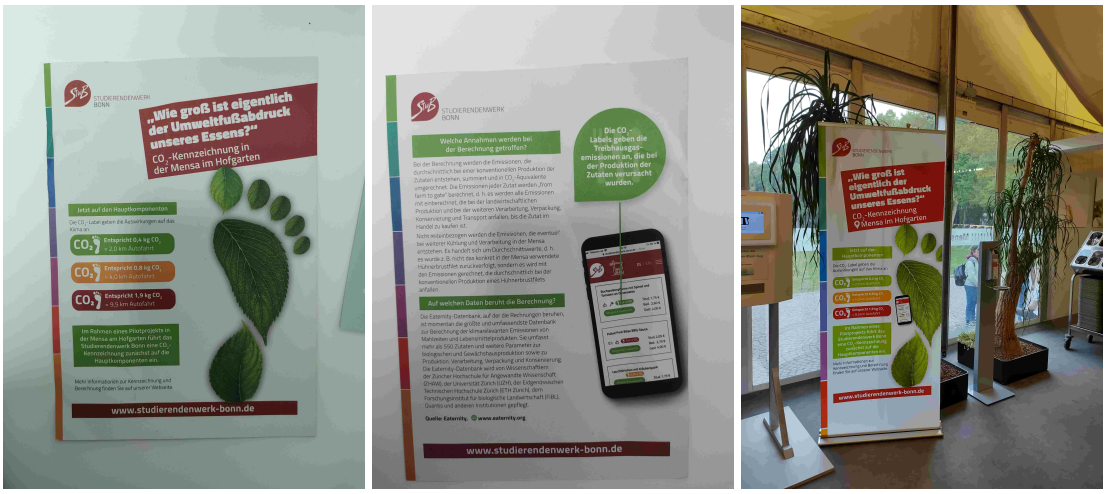


Figure S.18. Explanation of the carbon labeling initiative in the canteen
Notes: Leaflets (left and center) and billboards at the entrance of the student canteen (right).

S.5.5 Data set construction: Full sample

The raw data cover purchases from April 1 to July 8, 2022. April 1 is excluded before Specification (1) in Table S.45 because no meal-offer information is available for that date. Specification (1) otherwise uses weeks 1–14 before the sequential exclusions below. Table S.46 records the exact calendar construction. The preregistration envisaged six post-intervention weeks, but irregular canteen closures limited the comparable post-intervention window to three weeks. The main analysis further excludes the final of these three weeks, as described below.

- Starting from week 9 of the study period (May 30th to June 3rd), Ukrainian refugees received free meals in the treated student canteen and the larger control restaurant, using specific student canteen cards. I thus identify these sales and exclude them

83. Carbon emission labels for a given meal are calculated as the sum of the emissions caused by each of the ingredients. For each ingredient, emission values are calculated "from farm to gate". Hereby, it is assumed that the production process mirrors the average conventional production, e.g. I do not track the specific chicken breast bought by the student canteen but assume average conventional production. Emissions caused by the student canteen cooling, freezing, and cooking ingredients on-site are not included. These calculation details are explained to students on the student canteen website and on leaflets lying out on-site in the student canteen.

from all analyses. For the treated restaurant, they make up 12% of total sales in week 9,25% in week 10, and between 13% and 18% for the rest of the observation period. For the control restaurant, they make up between 2% and 4% of total sales during this time period. Spec. (2) in Table S.45 shows how this exclusion affects results.

- During the first week of the label period (May 2nd to May 6th), the display was irregular, as the student canteen needed some "trial and error" to get the system running. On some days, the labels were only displayed in the student canteen or online. Further, the student canteens (i.e. all three) had a special "Healthy Campus" week during the first week of May. The "Healthy Campus" week featured an additional, rarely chosen ($\approx$ 10% of sales) premium-priced meal and light promotional activities (e.g., nutrition quizzes). Since the event week affected all canteens, it should not confound the treatment effect unless its impact differed across locations. To be conservative, I omit this week in the main analysis. Spec. (3) in Table S.45 additionally excludes week 5 from the sample.
- I also exclude week 14 (July 4–8) from the main sample, since it is the last week before an irregular two-week closure of the treatment student canteen, followed by a two-week closure of the control canteens. The canteens usually do not close at this time of year, and the impending closures might have affected consumption behavior. Specification (4) in Table S.45 additionally excludes week 14, yielding the sample analyzed in the main paper.

For each purchase, I have data on the mode of purchase (student canteen card or debit card), meal category (combined with daily menus, this provides the specific meal name), student canteen card ID (if the purchase is made with the student canteen card), cash register number, date of purchase, time of purchase (exact to the minute), and purchase value.

Table S.45. Field estimates under sequential full-sample exclusions

	Meat meal chosen			
	(1) Full data	(2) Excl. Ukr.	(3) +Excl. W5	(4) +Excl. W14
Treatment × Label period	-0.017***	-0.030***	-0.022***	-0.022***
	(0.007)	(0.007)	(0.007)	(0.007)
Treatment × Post period	-0.011	-0.066***	-0.066***	-0.047***
	(0.007)	(0.008)	(0.008)	(0.009)
Treatment	-0.097***	-0.097***	-0.097***	-0.097***
	(0.005)	(0.005)	(0.005)	(0.005)
Date effects	No	No	No	No
Fixed effects	No	No	No	No
Observations	155,165	150,142	137,788	124,784

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: April 1 is omitted throughout because meal-offer information is unavailable. Specification (1) otherwise includes weeks 1–14. Specification (2) excludes purchases by Ukrainian refugees. Specification (3) additionally excludes week 5 (May 2–6), when label display was irregular and all canteens held the Healthy Campus event. Specification (4) additionally excludes week 14 (July 4–8), the final week before irregular canteen closures, and is the sample analyzed in the main paper. The columns vary only these sample exclusions and repeat the baseline difference-in-differences specification in Table OA.17, Panel A, Column (1), including treatment, label-period and post-period indicators and treatment–period interactions, without offer or week fixed effects. Standard errors are heteroskedasticity robust.

Table S.46. Exact calendar construction of the full field-regression sample

Dates	Study week(s)	Status	Reason
April 1, 2022	Day 1	Excluded	No meal-offer information
April 4–29, 2022	1–4	Included	Pre-intervention period
May 2–6, 2022	5	Excluded	Irregular display and Healthy Campus event
May 9–June 17, 2022	6–11	Included	Label period
June 20–July 1, 2022	12–13	Included	Post-intervention period
July 4–8, 2022	14	Excluded	Final week before irregular closures

Notes: The table lists every date range in the April 1–July 8 raw study window. Student canteens are ordinarily closed on weekends. Beginning May 30, purchases made with the cards issued for free meals to Ukrainian refugees are excluded within otherwise included dates; other eligible transactions on those dates remain in the sample.

S.5.6 Data set construction: Intention-to-treat sample

From the full sample data set detailed above, I construct the ITT sample data set:

- I restrict the sample to purchases made with a personal payment card (69% of purchases).

- Using the individual payment data, I can identify guests who purchased several meal components on a single day. These are 7% of the remaining sample. I drop these purchases since it is not clear whether the guest is purchasing the meal for themselves or for a friend. The ITT analysis therefore uses one main-meal purchase per guest-day.
- Further, I restrict the analysis to regular canteen guests, which I define as individuals who visited one of the student canteens more than four times during the pre-intervention period (51% of the remaining sample). Results are robust to different cut-off values, as Table S.48 shows.
- Finally, I restrict the sample to canteen guests visiting the same canteen in 80% of their pre-intervention visits (89% of the remaining sample). Results are robust to different percentage cutoff values, as Table S.49 shows.

Table S.47. Decision to visit one of the student canteens as outcome variable

	Likelihood of visiting canteen (and consuming...)		
	Visit	Visit+Meat	Visit+Veg
Treatment × Label period	-0.012	-0.016**	0.004
	(0.009)	(0.007)	(0.007)
Treatment × Post period	-0.047***	-0.045***	-0.002
	(0.013)	(0.009)	(0.010)
Offer controls	No	No	No
Week fixed effects	No	No	No
Guest fixed effects	Yes	Yes	Yes
Guests control	1,422	1,422	1,422
Guests treated	496	496	496
Observations	103,572	103,572	103,572

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: All columns are linear probability models using an extended guest-by-date panel that includes zero entries when a guest did not visit. Each guest is assigned to a home canteen from all pre-intervention visits using the same 80% rule as the main ITT analysis and appears once on each of 54 canteen dates. The panel contains all 1,918 guests in the ITT sample: 1,422 assigned to the control arm and 496 to the treated arm. Column (1) indicates a canteen visit on which the guest makes a card purchase satisfying the ITT sample criteria. Columns (2) and (3) indicate that such a visit includes a meat- or vegetarian-meal purchase, respectively. Coefficients are in probability units, so 0.01 corresponds to one percentage point. All specifications include guest fixed effects and label- and post-period indicators, which control for common changes relative to the pre-intervention period. They include no offer controls or week, calendar-date, or day-of-week fixed effects. Standard errors are clustered at the guest level.

Table S.48. ITT sensitivity to the minimum pre-intervention visit requirement

	Meat meal chosen				
	> 4 visits	> 2 visits	> 3 visits	> 5 visits	> 6 visits
Treatment × Label period	-0.022**	-0.016*	-0.017*	-0.020*	-0.024*
	(0.010)	(0.009)	(0.009)	(0.011)	(0.013)
Treatment × Post period	-0.041***	-0.030**	-0.038***	-0.035**	-0.039**
	(0.014)	(0.012)	(0.013)	(0.015)	(0.017)
Offer controls	Yes	Yes	Yes	Yes	Yes
Week fixed effects	Yes	Yes	Yes	Yes	Yes
Guest fixed effects	Yes	Yes	Yes	Yes	Yes
Guests control	1,422	2,414	1,847	1,079	792
Guests treated	496	858	653	369	261
Observations	37,284	50,516	43,686	30,833	24,584

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable equals one when a guest chooses a meat meal and zero otherwise; coefficients are in probability units, so 0.01 corresponds to one percentage point. Spec. (1) conducts the ITT analysis following the above described data preparation procedure, i.e. guests are classified as regular student canteen guests if they visit any student canteen more than four times during the pre-intervention period. Col. (2) instead requires more than two visits, Col. (3) more than three visits, Col. (4) more than five visits, and Col. (5) more than six visits. Each column repeats Table OA.17, Panel B, Column (3), changing only the indicated sample-selection threshold. Offer controls include indicators for the specific main vegetarian and meat meals and for the availability of a second vegetarian or meat option. All specifications include week, day-of-week, and guest fixed effects; standard errors are clustered at the guest level.

Table S.49. ITT sensitivity to the pre-intervention canteen-assignment threshold

	Meat meal chosen			
	80%	60%	70%	90%
Treatment × Label period	-0.022**	-0.020**	-0.023**	-0.027***
	(0.010)	(0.010)	(0.010)	(0.010)
Treatment × Post period	-0.041***	-0.038***	-0.040***	-0.045***
	(0.014)	(0.013)	(0.013)	(0.015)
Offer controls	Yes	Yes	Yes	Yes
Week fixed effects	Yes	Yes	Yes	Yes
Guest fixed effects	Yes	Yes	Yes	Yes
Guests control	1,422	1,505	1,455	1,281
Guests treated	496	541	513	438
Observations	37,284	39,884	38,476	33,858

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The dependent variable equals one when a guest chooses a meat meal and zero otherwise; coefficients are in probability units, so 0.01 corresponds to one percentage point. Guests are assigned to a home canteen if at least 80% of their pre-intervention visits are to that canteen. Col. (2) instead uses a 60% assignment rule, Col. (3) uses a 70% assignment rule, and Col. (4) uses a 90% assignment rule. Each column repeats Table OA.17, Panel B, Column (3), changing only the indicated sample-selection threshold. Offer controls include indicators for the specific main vegetarian and meat meals and for the availability of a second vegetarian or meat option. All specifications include week, day-of-week, and guest fixed effects; standard errors are clustered at the guest level.

S.5.7 Descriptive statistics on meat consumption and average emissions

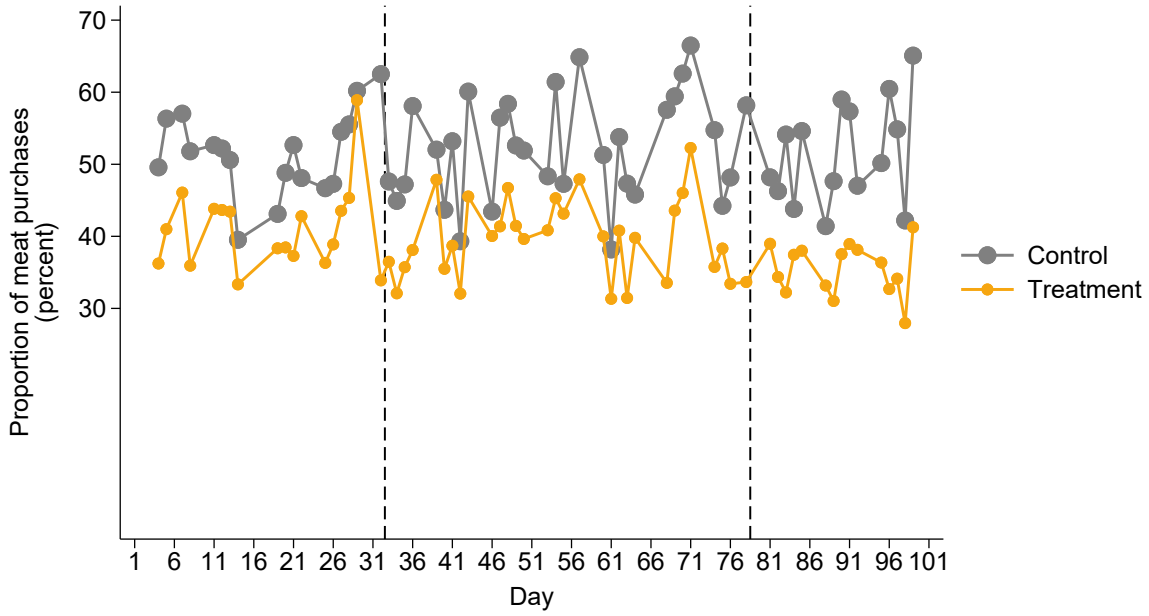


Figure S.19. Proportion of meat meals sold in the canteen

Notes: The sample excludes purchases by Ukrainian refugees but retains weeks 5 and 14. $N = 150{,}187$.

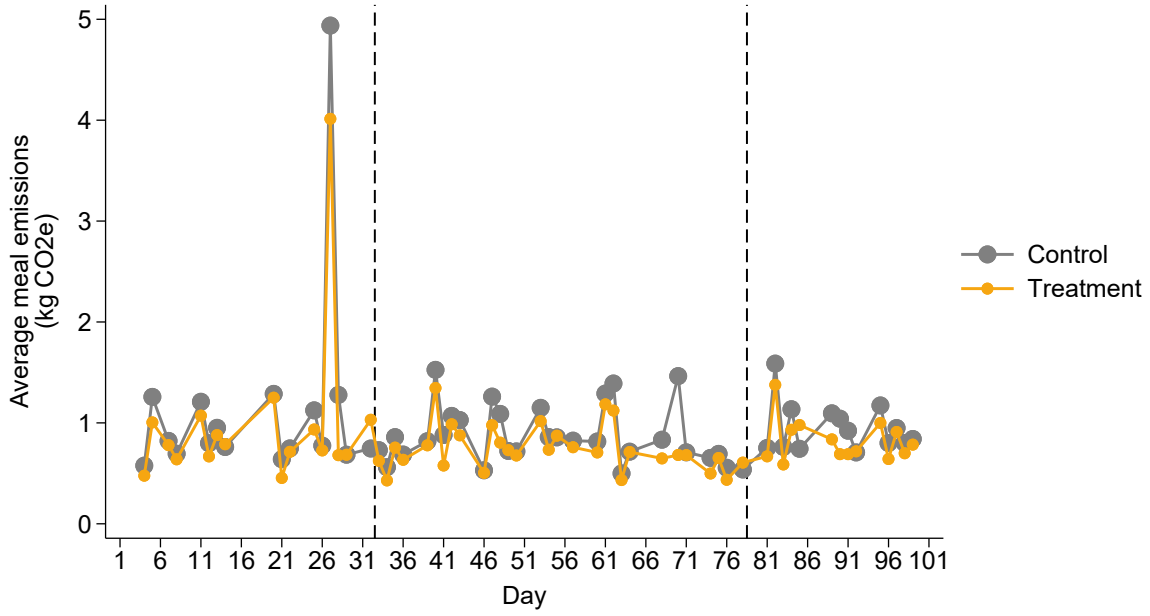


Figure S.20. Average emissions per meal sold in the canteen

Notes: The sample excludes purchases by Ukrainian refugees but retains weeks 5 and 14. It additionally excludes April 19, June 8, and June 27, when the maximum cross-canteen difference in main-meal offer emissions exceeded 1 kg CO_2e. $N = 143{,}518$.

S.5.8 Survey accompanying natural field experiment

Pre-intervention survey: Between April 18 and May 6, I conducted a survey among student canteen guests at the treatment student canteen and the first, larger, control restaurant. All active promotion of the survey was removed before May 2, and 14 dated responses among the 1,795 survey records were submitted on or after May 2.[84] Analyses treating first-survey responses as predetermined use only surveys completed before the labels were introduced on May 2. The survey was advertised as an opportunity to voice one's opinion on the offer of the student canteen, took participants around five minutes, and motivated potential participants with the chance to win one of ten €50 coupons for the student canteen. The survey was advertised through multiple channels. First, I put up posters advertising the survey in many faculties throughout the University of Bonn. Second, I distributed leaflets in front of the treatment restaurant and the larger control restaurant, together with research assistants (see Figure S.21). It is common for students and student groups to advertise surveys, projects, and events in this manner. Finally, the experimental lab at the University of Bonn sent out an e-mail to its entire participant pool advertising participation.

In the survey, respondents indicated their student canteen card number and consented to their survey responses being connected to their consumption decisions from April to July. They filled out questions on demographics, environmental attitudes, political preferences, and preferences towards the student canteen offer. Responses to the questions on student canteen offer and participant comments were analyzed, summarized, and presented to the gastronomic manager of the student canteens. 1,700 respondents participated in this first survey, 94% of these students.

Post-intervention survey: From the 22nd of June, I started sending out invitations to participate in a second survey. These were sent out by e-mail to participants in the first survey who provided an e-mail address and consented to be contacted again. Of the 1,558 people invited, 1,004 completed survey 2 and could be linked through their primary participant identifier. I invited participants in a staggered fashion over two weeks and sent a reminder on the 7th of July. Again, survey respondents had the opportunity to win one of ten €50 coupons for the student canteen.

84. These responses may come from people who acted on the promotion email several days after receiving it or scanned the QR code on a leaflet they had taken home.



Figure S.21. Leaflet advertising participation in the survey
Notes: Leaflet was distributed in front of the student canteen.

In survey 2, I repeated some of the questions from survey 1, to assess whether attitudes changed differentially in the treatment student canteen. The survey further included some questions of interest to the student canteen following the outcome of the first survey. At the end of the survey, participants could indicate whether and how they had perceived the emission labels, as well as voice their opinion on the initiative. Summary statistics for the linked survey and consumption samples are reported in Tables S.5 and S.6.

S.5.9 Post-experiment label installation

Building on the experience from this research project, the student canteens in Bonn initiated the development of a technical system enabling the permanent display of carbon labels on all canteen menus. This system was completed in early 2025. The installed labels resemble those tested in the project and are visible on the canteen menus (as of August 2025).

References for the Supplementary Material

Barahona, N., C. Otero, and S. Otero. 2023. "Equilibrium effects of food labeling policies." *Econometrica* 91 (3): 839–68 [S-47].

Brunner, F., V. Kurz, D. Bryngelsson, and F. Hedenus. 2018. "Carbon label at a university restaurant–label implementation and evaluation." *Ecological Economics* 146: 658–67 [S-45].

Bundesministerium für Umwelt, Klimaschutz, Naturschutz und nukleare Sicherheit. 2026. *KlimaTeller in der Mensa* [S-47].

DellaVigna, S., and E. Linos. 2022. "RCTs to Scale: Comprehensive Evidence From Two Nudge Units." *Econometrica* 90 (1): 81–116 [S-46].

Eaternity Institute. 2020. *Eaternity database* [S-91].

Haws, K. L., S. W. Davis, and U. M. Dholakia. 2016. "Control over what? Individual differences in general versus eating and spending self-control." *Journal of Public Policy & Marketing* 35 (1): 37–57 [S-26, S-57].

Ho, L., and L. Page. 2025. "Got Beef with Beef? Evidence from a Large-Scale Carbon Labeling Experiment." Working paper [S-46, S-55].

Imai, T., D. D. Pace, P. Schwardmann, and J. J. van der Weele. 2026. "Correcting consumer misperceptions about CO2 emissions." *Journal of the Association of Environmental and Resource Economists* 13 (5): 1399–435 [S-48].

Kehlbacher, A., R. Tiffin, A. Briggs, M. Berners-Lee, and P. Scarborough. 2016. "The Distributional and Nutritional Impacts and Mitigation Potential of Emission-Based Food Taxes in the UK." *Climatic Change* 137 (1–2): 121–41 [S-34].

Krupka, E. L., and R. A. Weber. 2013. "Identifying social norms using coordination games: Why does dictator game sharing vary?" *Journal of the European Economic Association* 11 (3): 495–524 [S-26, S-57].

List, J. A. 2020. "Non est Disputandum de Generalizability? A Glimpse into The External Validity Trial." Working Paper 27535. National Bureau of Economic Research [S-42, S-43].

Liu, Z., E. Ansink, and J. Bouma. 2025. "Plant-based meat complements." Discussion Paper 2024-046/VIII. Tinbergen Institute. Revised April 2025 [S-35].

Lohmann, P. M., E. Gsottbauer, A. Doherty, and A. Kontoleon. 2022. "Do carbon footprint labels promote climatarian diets? Evidence from a large-scale field experiment." *Journal of Environmental Economics and Management* 114: 102693 [S-45].

Lohmann, P. M., E. Gsottbauer, J. Farrington, S. Human, and L. A. Reisch. 2024. "Choice architecture promotes sustainable choices in online food-delivery apps." *PNAS Nexus* 3 (10): pgae422 [S-46].

Roosen, J., M. Staudigel, and S. Rahbauer. 2022. "Demand elasticities for fresh meat and welfare effects of meat taxes in Germany." *Food Policy* 106: 102194 [S-34, S-35].

Schulze Tilling, A. 2021. *Quantifying the role of greenhouse gas emissions in consumption choice* [S-13, S-14].

Studierendenwerk Bonn. 2025. *Essen & Trinken — „Unsere Preise" (Studierendenwerk Bonn)* [S-37, S-89].