

Firms and governments increasingly rely on environmental labels to guide consumers toward greener products. Governments regulate environmental claims, and firms use them to market their products. Both rest on the same premise: consumers value this information and reward it with their purchases. My research tests this premise. I ask how much consumers value carbon-neutral claims, whether more detailed information changes this value, and whether it is reflected in market prices. I combine preregistered choice experiments, quasi-experimental analysis of online marketplace data, and meta-analysis of existing literature.

My findings point to a gap between consumers' stated and revealed preferences for carbon neutrality. Stated preferences imply a positive price premium, but market prices show none. I also find no evidence that more detailed disclosure about how carbon neutrality is achieved increases stated willingness to pay.

My agenda builds on this gap and studies how sustainability labels work where consumers actually shop. It moves from stated to revealed preferences in two settings: products whose sustainability labels compete with other labels and commercial badges, and purchases guided by AI shopping assistants. For each setting, I plan to first use market data in a quasi-experimental design that needs no partner, and then to test the mechanisms behind the market patterns in a field experiment with a partner retailer. Both steps will cover countries where consumer behavior and the regulation of green claims differ, which allows a comparison across regulatory settings.

Current Research

Do consumers value carbon-neutral products?

Demand for Carbon-Neutral Products

With Stefano Carattini, Fabian Dvorak, and Ivana Logar. Revised and resubmitted, Journal of Environmental Economics and Management.

Many firms have made or pledged to make their operations carbon neutral due to pressure from investors and boards. Whether consumers reward these efforts, and thus privately contribute to a global public good, is an open empirical question. We combine two approaches. First, a meta-analysis of 37 studies on carbon-neutral and low-carbon products, with 29,666 participants and mostly based on stated preferences and controlled environments, finds an average willingness to pay of USD 1,993 per ton of CO₂ reduced, far above current carbon prices and many estimates of the social cost of carbon. Willingness to pay is positively associated with the amount of CO₂ reduced and the product price, and is higher in studies conducted in Europe. Second, we use weekly data on prices, product characteristics, and reviews for more than 120,000 products on Amazon in the US, the UK, and Germany. The staggered certification of products as carbon neutral creates a series of quasi-natural experiments, which we exploit with a hedonic difference-in-differences design. Unlike the meta-analysis, this approach finds no evidence that consumers value carbon neutrality: price effects are indistinguishable from zero in all three markets. Market prices thus suggest that private motivations are not strong enough to spur the adoption of carbon-neutral products.

Does transparency raise the value of carbon-neutral claims?

Transparency of Carbon-Neutral Labels: Evidence from a Choice Experiment

Job market paper. Single-authored. Revise and resubmit, Environmental and Resource Economics.

Carbon-neutral labels do not reveal how much of a claim rests on emission reductions and how much on offsets, which experts treat as a last resort. Regulators differ: the EU now prohibits claims based on offsetting, California requires firms to disclose the offsets behind them, the UK and Australia recommend disclosing the mix, and the US Federal Trade Commission has called for evidence on how consumers perceive such claims. Disclosure-based approaches assume that consumers care about the mix, an assumption that remains untested.

I test it with a preregistered discrete choice experiment with 1,339 UK tea consumers. Between subjects, a control group sees a standard carbon-neutral label, and two treatment groups see labels that disclose either 95% offsetting or equal shares of offsetting and reduction. Consumers state a premium of £0.55 for the standard label, about a fifth of the average tea price. Disclosure does not raise it, and consumers do not value a higher reduction share, even when they correctly distinguish offsets from reductions. Organic and ethical-trade labels substantially reduce the carbon-neutral premium and, combined, can reverse it. Disclosure alone is thus unlikely to reward firms that reduce rather than offset, and additional regulation may be needed. The crowd-out also points to label placement and prominence as policy levers.

Research Agenda

The question that motivates my agenda

Together, the two papers point to a puzzle. Consumers say they value carbon neutrality, yet this value neither shows up in prices nor responds to better information. Gaps between stated and revealed preferences are familiar, but why this one arises is an open question. One explanation is hypothetical bias, along with other survey biases such as social desirability: stated values may overstate what people would actually pay. Another is low salience: at the point of purchase, a claim competes with many other product attributes for limited attention, so even a genuine preference may not translate into purchases. The distinction matters for policy, because better presentation can address the second but not the first. At the same time, a growing share of purchases is filtered, recommended, or executed by AI algorithms. These observations motivate my agenda: studying sustainability labels in real purchases, where stated values meet actual choices, and asking whether these values survive when people delegate choices to AI.

The quasi-experimental studies build on the difference-in-differences methods of my current work. They rely on data I will collect myself from current and archived product pages or obtain from commercial trackers, so they do not depend on a partner. Both field experiments can run with the same partner, ideally an online retailer that already tests product badges and offers a shopping assistant, a partnership I intend to build through collaborators who run field experiments with firms. I plan to collect the market data in 2026–27 and to pilot both field designs with this partner in 2027–28.

Delegated virtue: do pro-social choices survive AI advice?

Delegated Virtue

Main project. Planned quasi-experiment and field experiment.

Consumers increasingly use AI assistants when they shop: Amazon reports that more than 250

million customers used its shopping assistant in 2025 (Amazon 2025). These assistants increasingly shape which products, labels, and prices shoppers see. This raises a broader question: when people delegate part of a choice to AI, what happens to preferences that benefit the environment and other people? Delegation can make people less honest or more selfish (Hamman et al. 2010; Köbis et al. 2025). Sustainable purchases are an ideal test case, because the buyer bears the cost and others receive the benefit.

An ongoing preregistered online study tests whether AI assistants prompted to promote sustainability shift shoppers toward green products (Karlsruhe Institute of Technology 2026). A more fundamental question is how AI assistants change green choices that people would make on their own. Theory gives opposite predictions. If shoppers do not mention sustainability when asking for help, it may drop out of the choice altogether, a form of attribute non-attendance. Yet an assistant that points to sustainable alternatives can make labels salient at the moment of choice, and an algorithm audit of AI hotel recommendations finds that assistants weigh eco-certification more heavily than human choices suggest (Baig et al. 2026). This project will test both predictions among real shoppers at a retailer, where purchases can differ from stated choices and from choices made in lab or online experiments.

The planned experiment has four arms. In the control arm, shoppers buy without an assistant. In the second, an AI assistant recommends products based only on what shoppers ask for. In the third, the assistant also points to a similar product with a sustainability label, such as an organic or carbon label, and asks whether the shopper wants it instead. The fourth arm adds the price difference to this question. The main outcome is whether a shopper buys a labeled product, and all shoppers count in the arm they were assigned to, whether or not they follow the assistant. Comparing the first two arms shows whether shopping with an assistant changes green purchases. Comparing the second and third shows whether pointing shoppers to a labeled alternative increases them. Comparing the third and fourth shows whether this increase survives once shoppers see the price difference, a direct link to the gap between stated and revealed preferences. A survey after purchase measures stated preferences, which, together with shoppers' requests to the assistant and their purchases, allows a comparison of stated, *delegated*, and revealed preferences for the same person. How assistants present sustainability and its cost is a concrete policy lever.

Before the field experiment, the staggered rollout of AI shopping assistants offers a quasi-experiment. Amazon introduced its assistant in the US in early 2024 and in other markets over the following months. Using historical price and sales-rank data that I will collect, I will compare labeled and unlabeled products in the same category before and after the assistant became available in each market. This design does not show what shoppers asked or saw. It does test whether the market share of labeled products changes once shoppers can delegate part of their search.

When do corporate environmental claims complement or crowd out one another?

Competing Environmental Claims

Additional project. Planned quasi-experiment and field experiment.

Green claims rarely reach consumers alone. A single product may carry a carbon label, an organic label, and a Fair Trade label, often next to commercial badges such as bestseller tags. These signals may reinforce one another, or they may crowd one another out: one sustainable attribute may be enough for shoppers to feel good about a purchase, and each additional label competes for limited attention. My job market paper finds such crowd-out in stated choices. Whether it also occurs in real purchases is an open question, because field experiments have so far studied a single sustainability label, alone or with price cues (Hainmueller et al. 2015; Feuß et al. 2022). The answer

matters for policy: when rules remove some claims, their effect depends on whether the remaining labels gain or lose value.

The planned experiment randomly shows or hides each label and badge on an online retailer’s product pages. A shopper may see none, some, or all of the labels a product holds, but never a label it does not hold. This design answers three questions with one sample. First, does the carbon label alone change real purchases? This directly tests the salience explanation in my current research. Second, do shoppers value environmental labels differently from social ones such as Fair Trade? Third, does a label lose value when others appear, and do commercial badges draw attention away from sustainability labels? Coffee, tea, and chocolate often carry several of these labels, and tea links the experiment to my job market paper. The main outcome is purchases. Clicks, time on page, and a short recall survey after purchase measure attention. Because every shopper contributes to every comparison, the design uses the sample efficiently, although tests of crowd-out still require a retailer with high traffic.

Before the field experiment, labels that appear or disappear at different times offer a quasi-experiment. Products gain certifications at different dates, and the EU ban on offset-based claims removes carbon-neutral claims in EU markets but not in the UK or the US. Using product-page data that I will collect from Amazon and other online retailers in these markets, I will test whether a carbon label loses value when a product gains a second label, and whether the remaining labels gain value when a claim is removed. This version measures prices and sales ranks rather than individual purchases, but it asks the same question in market data.

Outlook

My research follows one thread: from what consumers say they value, to what they actually buy, to what they buy when they delegate choices to AI. Regulators are setting new rules on green claims, and platforms are deciding how their AI assistants present them, with little evidence from real purchases on either. My agenda will provide this evidence by combining the valuation methods of my current work with quasi-experiments in market data and field experiments in real markets. Running these studies across countries with different green-claim rules will show which information policies change what people buy, and where. Alongside these projects, I plan a review of the economics of environmental labels, which brings together theory on labels as signals of credence attributes with evidence from stated, revealed, and field studies.

References

- Amazon (2025). Amazon Rufus: Amazon’s AI shopping assistant gets smarter and more personal. About Amazon, November 18. <https://www.aboutamazon.com/news/retail/amazon-rufus-ai-assistant-personalized-shopping-features>.
- Baig, M. S. A., S. A. Gillani, and A. Ali (2026). Whose hotel does the AI recommend? An algorithm audit of reputation signals in LLM-assisted hotel selection. arXiv preprint arXiv:2606.16344.
- Feuß, S., D. Fischer-Kreer, J. Majer, J. Kemper, and M. Brettel (2022). The interplay of eco-labels and price cues: Empirical evidence from a large-scale field experiment in an online fashion store. *Journal of Cleaner Production* 373, 133707.
- Hainmueller, J., M. J. Hiscox, and S. Sequeira (2015). Consumer demand for fair trade: Evidence from a multistore field experiment. *Review of Economics and Statistics* 97(2), 242–256.
- Hamman, J. R., G. Loewenstein, and R. A. Weber (2010). Self-interest through delegation: An additional rationale for the principal-agent relationship. *American Economic Review* 100(4), 1826–1846.
- Karlsruhe Institute of Technology (2026). ManyPrompts: A many design study on sustainable conversational shopping. Research project announcement, March 17. https://h-lab.win.kit.edu/news_3885.php.
- Köbis, N., Z. Rahwan, R. Rilla, B. I. Supriyatno, C. Bersch, T. Ajaj, J.-F. Bonnefon, and I. Rahwan (2025). Delegation to artificial intelligence can increase dishonest behaviour. *Nature* 646(8083), 126–134.