

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Lanz, Bruno; Wurlod, Jules-Daniel; Panzone, Luca; Swanson, Timothy

Working Paper The Behavioral Effect of Pigovian Regulation: Evidence from a Field Experiment

IRENE Working Paper, No. 17-01

Provided in Cooperation with: Institute of Economic Research (IRENE), University of Neuchâtel

Suggested Citation: Lanz, Bruno; Wurlod, Jules-Daniel; Panzone, Luca; Swanson, Timothy (2017) : The Behavioral Effect of Pigovian Regulation: Evidence from a Field Experiment, IRENE Working Paper, No. 17-01, University of Neuchâtel, Institute of Economic Research (IRENE), Neuchâtel

This Version is available at: https://hdl.handle.net/10419/191496

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

University of Neuchatel Institute of Economic Research IRENE Working paper 17-01

The Behavioral Effect of Pigovian Regulation: Evidence from a Field Experiment

Bruno Lanz Jules-Daniel Wurlod Luca Panzone Timothy Swanson



la etitut de

Institut de recherches économiques

The behavioral effect of Pigovian regulation: Evidence from a field experiment*

Bruno Lanz^{\dagger} Jules-Daniel Wurlod^{\ddagger}

Luca Panzone[§]

Timothy Swanson[¶]

This version: January 2017

Abstract

Pigovian regulation provides monetary penalties/rewards to incentivize prosocial behavior, and may thereby trigger behavioral effects beyond a more standard response associated with a change in relative prices. This paper quantifies the magnitude of these behavioral effects using data from an experiment on real product choices together with a structural model of consumer behavior. First, we show that information about external effects (products' embodied carbon emissions) triggers voluntary substitution towards cleaner alternatives, and we estimate that this effect is equivalent to a change in relative prices of GBP30.69-165.15/tCO₂. Second, comparing a Pigovian intervention (GBP19/tCO₂) with a neutrallyframed price change of the same magnitude, we find a negative behavioral effect associated with regulation. Compensating this bias would require increasing the Pigovian price signal by up to 48.06/tCO₂. Finally, based on a cross-product comparison, we show that the magnitude of behavioral effects declines with substitutability between clean and dirty product alternatives, a measure of effort to reduce emissions.

Keywords: Externalities; Pigovian regulation; Consumer behavior; Information; Field experiments; Environmental policy.

JEL Codes: C93; D03; D12; H23; Q58.

^{*}We would like to thank Valentina Bosetti, Mehdi Farsi, Nick Hanley, Guy Meunier, Emi Uchida, and seminar participants at Neuchâtel, WCERE, Bioecon conference, and the EAAE/AAEA conference for useful comments and discussions. We also thank Grischa Perino and Denise Leung for their involvement in this work, and the supermarket chain Sainsbury's for permission to run the experiment in their stores. Excellent research assistance was provided by Ghislain Lang. Funding from EU-POPP is gratefully acknowledged. Any remaining errors are ours.

[†]Corresponding author. University of Neuchâtel, Department of Economics and Business; ETH Zurich, Chair for Integrative Risk Management and Economics; Massachusetts Institute of Technology, Joint Program on the Science and Policy of Global Change. Mail: A.-L. Breguet 2, CH-2000 Neuchâtel, Switzerland. email: bruno.lanz@unine.ch.

[‡]Boston Consulting Group, Geneva, Switzerland.

[§]Newcastle University, School of Agriculture, Food and Rural Development, UK.

[¶]Graduate Institute of International and Development Studies, Centre for International Environmental Studies, Switzerland.

1 Introduction

In a traditional framework, Pigovian regulation sets up a corrective tax/subsidy to make agents internalize external effects associated with consumption or production decisions (Pigou, 1920). This classic approach to regulation is, however, increasingly being refined by a literature that considers how behavioral agents process information and, in turn, how behavioral traits affect optimal policy design (e.g. Allcott et al., 2014; Farhi and Gabaix, 2015).¹ In particular, evidence from economics and experimental psychology suggests that when agents are willing to voluntarily exert an effort, for example because of (self-)image concerns, explicit financial incentives may have a detrimental impact on effort provision (see Gneezy et al., 2011; Bowles and Polanía-Reyes, 2012, for a review). As Pigovian interventions associate an external monetary reward/penalty with prosocial behavior, this type of regulation may attenuate the effectiveness of a change in relative prices (see also Frey and Oberholzer-Gee, 1997; Bénabou and Tirole, 2003).

While theoretical properties of these mechanisms have been studied extensively (see e.g. Brekke et al., 2003; Bénabou and Tirole, 2006), empirical evidence as to their relevance for the design of public policies, and in particular those regulating externalities, remains scant. Moreover, given the growing use of market-based instruments for environmental policy, the in-expensive nature of non-price interventions (Bertrand et al., 2010; Allcott, 2011), as well as the emerging literature on behavioral public finance (e.g. Chetty et al., 2009; Mullainathan et al., 2012), quantifying the policy implications of behavioral effects associated with Pigovian interventions is important. This paper provides a first step in that direction. We exploit data from a field experiment in which subjects make real consumption decisions about ordinary grocery products, with a set of clean and dirty alternatives as determined by their embodied carbon emissions.² After an initial product choice, we randomly assign subjects to three treatments manipulating attributes of the choice between dirty and clean alternatives: (i) information about

¹ A parallel literature focuses on the role of technology and innovation in the transition from processes or products that generate relatively large external costs (i.e. *dirty*) to others with relatively low external costs (i.e. *clean*). See for example Acemoglu et al. (2012), Aghion et al. (2016) and Acemoglu et al. (2016).

² One interesting aspect of carbon emissions reduction, as a global public good, is that personal contributions have negligible direct private benefits (for example in terms of climate), so that choices solely reflect prosocial motivations. In other cases, contributions may have direct private benefits. For example, information about energy use of durable products may affect perceptions about both private energy expenditures and external effects associated with energy use. Thus focusing on climate change allows us to net-out consumption-related personal benefit from individual choices.

embodied emissions, revealing the propensity to voluntarily contribute to the emission reduction effort; (ii) a Pigovian price change, combining a change in relative prices in proportion to external costs and information about the regulatory nature of the price signal; and (iii) a neutrally framed change in relative prices, which mimics market-driven price variations and has the exact same magnitude as the Pigovian intervention.

Based on a descriptive analysis of the same experimental data, Perino et al. (2014) show that, as expected, all treatments increase the *aggregate* choice probability of the clean alternatives (i.e. considering all grocery products together). Perino et al. (2014) further find that the proportion of respondents who switch from dirty to clean alternatives is smaller for the Pigovian subsidy than for both the information treatment and the neutral price change. This suggests that a change in relative prices framed as an explicit external policy intervention induces a negative behavioral effect, as suggested by evidence reviewed in Gneezy et al. (2011). However, while Perino et al. (2014) focus on differences in aggregate choice frequencies pre- and postinterventions, in the present paper we exploit the discrete nature of product choices to estimate a structural demand model for differentiated products.³ In turn, the main novelty of our work is to use this model to do welfare analysis, and to derive quantitative implications for the design of public policies.

Another departure from the work of Perino et al. (2014) is that we consider each product category included in the experiment separately (with dirty and clean alternatives in brackets here): cola-type sodas (in aluminum cans and in plastic bottles), spreads (margarine and butter), milk (skimmed, semi-skimmed and whole), and meat (chicken and beef). This is important because relative carbon content differs across product categories, and therefore experimental treatments applied to different products are heterogeneous. Consequently, incentives associated with the treatments differ across product categories. Moreover, carbon emissions are tied to product characteristics, so that substitutability between clean and dirty alternatives also varies across products. Concretely, substituting among cola products in cans or in plastic bottles is not the same as switching from beef to chicken products in the meat category. Preferences for these product characteristics are associated with the effort of switching from dirty to clean alterna-

³ Formally, our estimation strategy identifies the structural parameters of the underlying consumer decision problem using Lancaster's (1966) multi-attribute utility theory and McFadden's (1974) random utility model. In this framework, conditional average treatment effects (CATE) are estimated controlling for preferences over observed product characteristics, and can be given a welfare-theoretic interpretation.

tives, and estimated treatment effects can therefore be expected to vary across products.

Our product-level structural estimation strategy allows us to make two key contributions. First, we exploit the neutrally-framed price change to quantify substitutability of clean and dirty alternatives across product categories. This allows us to assess the effectiveness of alternative interventions as a function of product substitutability. The second and more interesting contribution afforded by our approach is the ability to estimate money-metric welfare measures associated with the treatments. Specifically, we estimate an equivalent price metric (EPM, as per Chetty et al., 2009; Allcott and Taubinsky, 2015) of the information intervention, which measures the change in relative prices that would yield the same change in choice probabilities as providing information about carbon emissions. In our setting, the EPM for information provides an estimate of consumers' valuation of a reduction of carbon emissions, which can be compared against estimates of the social cost of carbon (which determine the level of the Pigovian price signal). Similarly, we obtain an EPM quantifying the behavioral effects of the Pigovian intervention aside from the induced change in relative prices. Intuitively, this EPM quantifies, in monetary terms, the differential impact of a change in relative prices when it is framed as a Pigovian intervention and when it reflects market-induced variations. As we detail below, it provides a measure of the increase of Pigovian prices needed to compensate the negative behavioral effects associated with the regulatory intervention. Our structural approach thereby provides novel quantitative evidence on the behavioral effect associated with the regulatory dimension of a Pigovian price change, and how these behavioral traits should be factored in for policy design.

Our key results are as follows.⁴ First, we find that providing information about the carbon content induces a voluntary transition towards cleaner products, which signals the existence of some form of preferences for the public good embedded in the marketed goods (Harsanyi, 1955; Margolis, 1982; Kahneman and Knetsch, 1992; Nyborg, 2000).⁵ The EPM estimates associated

⁴ As noted by Perino et al. (2014), external validity of the results may suffer from the fact that participants know they participate in an experiment. We note, however, that the experiment reproduces an online shopping environment in which respondents completed the choice tasks individually and anonymously through a computer, mitigating any experimenter demand effects. Moreover, as we discuss below, the changes in choice probabilities we measure are similar in magnitude to those estimated using actual transactions (e.g. Teisl et al., 2002; Bjorner et al., 2004).

⁵ A related literature on altruism and private provision public goods identifies at least three different motivation for voluntary contributions. First, agents may derive utility from the (shared) private benefits from the public good (Kotchen, 2005). Second, voluntary provision of public goods can originate from pure altruism (see Becker, 1974; Kotchen, 2006). Third, agents might also derive utility from their own contribution through a warm-glow effect (Andreoni, 1990).

with information range from GBP0.03 to GBP1.28 depending on the product category, which corresponds to GBP30.69-165.15/tCO₂. An implication is that consumers' implicit valuation of emissions is above most estimates of the social cost of carbon (for reference, we employ a value of GBP19/tCO₂ to compute the Pigovian subsidy, taken from DEFRA, 2002).⁶

Second, we find evidence of a significant negative behavioral effect associated with Pigovian regulation. Quantitatively, our EPM estimates amount to GBP48.06/tCO₂ for cola products, GBP37.46/tCO₂ for milk, while for spread and meat products EPM estimates are not statistically significantly different from zero (our point estimates are GBP6.07/tCO₂ and GBP0.22/tCO₂, respectively). In the context of global public goods, the literature suggests two behavioral traits that could explain our findings. The first relies on self-image motivation (Ariely et al., 2009), which could be negatively affected by the presence of monetary rewards, for example through 'moral licensing' (Schotter et al., 1996; Jacobsen et al., 2012).⁷ A second mechanism works through consumers' prior beliefs about the importance of external effects associated with their choices. As initially suggested by Gneezy and Rustichini (2000a,b), when these beliefs differ from reality, providing information about the momentary value of behavior through a tax/subsidy may induce consumers to update their beliefs. In cases where they over-estimate the impact of external effects associated with their choices, a monetary incentive may lead to reduced effort or provision.

Third, we exploit the neutrally framed price change to show that substitutability between clean and dirty alternatives varies substantially across product categories. Further, a crossproduct comparison reveals that policy interventions are more effective for products with high substitutability across alternatives. This finding is intuitive, and it is already recognized in the literature (e.g. Bjorner et al., 2004). However, the ability to assess, in a controlled environment, how substitutability affects the behavioral impact of price and non-price interventions is novel. We also observe that the extent of behavioral effects associated with Pigovian regulation also varies with substitutability: for products with close substitutes (cola and milk in our setting) we observe very substantial negative behavioral effects, while for products where substitution

⁶ There is a lot of uncertainty regarding the value of external costs associated with carbon emissions, and DEFRA (2002) was the relevant source at the time the experiment was designed. More recent estimates used by the UK government for the appraisal of public projects in 2016 tend to be slightly lower, with a central value of $5.94/tCO_2$ and an upper bound of $23.40/tCO_2$ (see DECC, 2013).

⁷ Moral licensing theory rationalizes findings from psychology of how individuals self-justify "bad" behavior by performing "good" actions. See Monin and Miller (2001) for example.

requires more effort (spreads and meat) behavioral effects are small and statistically indistinguishable from zero. Intuitively, if prosocial behavior requires more effort, self-image benefits will have relatively less weight in the decision to switch towards cleaner alternatives. Therefore, we can conjecture that when self-image concerns are given relatively less weight in consumption decisions, the detrimental impact of external monetary interventions declines.

Aside from quantifying the extent of behavioral effects associated with Pigovian regulation, our results also contribute to the growing empirical literature on how information about external effects affect consumption decisions. In particular, several studies have shown that voluntary public good provision can be triggered by information, and that this can have a significant effect on market outcomes. For example, Teisl et al. (2002) use data from the U.S. to show that information labels led to an increase in the market share of "Dolphin-friendly" canned tuna. Similarly, Bjorner et al. (2004) use a large sample of Danish consumer from 1997 to 2001 to identify a positive marginal willingness to pay for the label "Nordic Swan." These studies have established the role of information provision using day-to-day transactions, and our work provides an interesting complement in which the choice set and substitution patterns are controlled experimentally. Importantly, the magnitude of our findings for the information treatment, with changes in choice probabilities ranging from 10 to 30 percent, is similar to the studies using market share observations, suggesting that our results are reflective of (non-experimental) market behavior.

Our work is also related to a number of recent papers studying the behavioral effect of taxrelated prices changes, as compared to less salient variation in market prices (e.g. the seminal paper by Chetty et al., 2009; on the demand for gasoline products, see Davis and Kilian, 2011, Li et al., 2014, and Rivers and Schaufele, 2015). One particularly relevant paper is Rivers and Schaufele (2015), who estimate that the impact of a carbon tax targeting explicitly a reduction of gasoline demand is about four times larger than the effect induced by price fluctuations unrelated to environmental policy. While at first sight this evidence may seem to contradict our results, a key difference compared to our analysis is the role of salience. In observational data, market-driven changes in prices are less salient as compared to policy-driven interventions, which would tend to exacerbate the impact of the latter. By contrast, in our experiment, salience of Pigovian and neutrally-framed price changes is held constant. Therefore, relative to Rivers and Schaufele (2015), our contribution is to quantify the effect of information included in Pigovian prices, controlling for the magnitude of the change in relative prices as well as its salience.

The remainder of this paper is organized as follows. In Section 2, we describe the experimental setting, including the four different consumption goods we consider and the three treatment interventions. In Section 3, we present our empirical strategy, including identification of the EPM for information and the behavioral effect of Pigovian regulation. Section 4 presents our results. Section 5 concludes.

2 Experimental Design

Data on consumer choices are collected in seven supermarkets in the greater London area.⁸ Consumers entering the supermarket are offered to participate in a "university-sponsored grocery shopping study" with a GBP5 voucher compensation. The experiment is described as neutrally as possible, "studying how people make REAL LIFE grocery shopping decisions." No other information on the purpose of the experiment is provided to avoid self-selection of respondents with a specific interest in environmental issues. Participants in the experiment are selected by identifying those who intended to purchase a product in one out of four categories: cola-type sodas, milk, spreads (margarine and butter) and meat (chicken and beef).⁹

The experiment consists of two steps. First, recruited subjects have to complete an initial product choice for the product categories in which they intended to make a purchase during their shopping trip. The product choice set, which is a key building block of our identification strategy, is discussed in details below. This part of the experiment takes place at the entrance of the supermarket, using a laptop, and closely replicates an online shopping environment. To make the choice environment as realistic as possible, the presentation of products is very close to the actual online shopping platform of the store, and product prices match those in the store on the day of the experiment.

⁸ Perino et al. (2014) provides information about the sampling of the locations and how the experiment was setup within each supermarket. Here we focus on the aspects of the experiment that are most relevant for our structural identification strategy.

⁹ Potential participants are turned down if younger than 21 years old, if they cannot speak or read English, or if they had participated in the experiment previously. While the selection mechanism means that our sample is non-random, as we essentially focus on sub-populations purchasing the products we selected, it includes the relevant population of customers purchasing the products of interest in this geographical area. Note that we do not, however, observe the characteristics for this population, and it is not necessarily representative of the population in other locations. Nevertheless Appendix A shows that our sample includes a diverse socio-economic background.

Following their initial product choice, participants are then offered to receive general information about either nutritional matters, carbon emissions and climate change, or both.¹⁰ Subjects are then randomly assigned to one out of three treatments (described in more details in the subsections below): (i) an information label showing embodied carbon emissions associated with clean vs. dirty alternatives; (ii) a Pigovian subsidy on the clean alternatives; and (iii) a neutrally framed price reduction of the clean alternatives of the same magnitude as the Pigovian subsidy. After treatment, subjects are allowed to revise their initial choice. In the final part of the computerized step, socio-demographic data are collected.

In the second step of the experiment, subjects complete their planned grocery trip in the main shopping area of the supermarket, before coming back to the entry of the store to collect the GBP5 voucher. This part allows us to enforce truthful preference revelation by making payment of the voucher conditional on the purchase of the good selected in the post-treatment consumption choice (subjects were informed about this condition beforehand).¹¹ The compliance rate is 96 percent, and non-compliers are dropped from the sample.¹² This yields a total of 854 shoppers who completed the task independently, complied with all terms and conditions of the experiment, and are thus included in the sample. Table B1 in Appendix B summarizes randomized treatment assignment, illustrating that subsamples are balanced on observables (t-tests of difference in means are insignificant for all but one variable, gender).

In the following, we further discuss the product choice set, and then describe each treatment intervention.

2.1 Product alternatives

Each product category includes a fixed and finite number of alternatives, reported in Table 1, which we classify as 'clean' or 'dirty' according to relative embodied carbon emissions (measured in CO_2 equivalent). Carbon emissions are associated with a particular feature of the product

¹⁰ These information screens are the same for all participants, and explain in general terms the meaning of food labels and implications of carbon emissions. Importantly, this information is presented in a neutral manner, does not relate in any way to the products on offer, and it has no impact on the results, as estimated treatment effects for subjects who receive this additional information and those who did not are statistically indistinguishable. The same finding is also reported in Perino et al. (2014) when all products are considered together.

¹¹ For the subsidy and neutral price change treatments, participants who selected the option with an experimentally altered price also received the difference between the experimental price (reduced by the subsidy) and the instore price.

¹² We acknowledge that dropping these observation may introduce some systematic selection, although given the high compliance rate this is unlikely to affect our results significantly.

Products	Clean alternati	ves		Dirty alternativ	ves	
	Options	Price (GBP)	Emissions (kgCO ₂)	Options	Price (GBP)	Emissions (kgCO ₂)
Cola	Coca Cola in PET bottle (2l)	1.69	0.50	Coca Cola in ALU cans (2l)	2.85	1.02
	Coca Cola Diet in PET bottle (2l)	1.69	0.50	Coca Cola Diet in ALU cans (2l)	2.85	1.02
	Coca Cola Zero in PET bottle (21)	1.69	0.50	Coca Cola Zero in ALU cans (2l)	2.85	1.02
	Pepsi Regular in PET bottle (21)	1.00-1.69	0.50	Pepsi Regular in ALU cans (21)	2.75	1.02
	Pepsi Diet in PET bottle (2l)	1.00-1.69	0.50	Pepsi Diet in ALU cans (2l)	2.75	1.02
	Pepsi Max in PET bottle (2l)	1.00-1.69	0.50	Pepsi Max in ALU cans (2l)	2.75	1.02
Milk	Skimmed milk (2 pints)	0.86	1.40	Whole milk (2 pints)	0.86	1.80
				Semiskimmed milk (2 pints)	0.86	1.60
Spread	Lurpak Spread (500g)	2.58	0.68	Lurpak butter (500g)	2.76	11.90
	Sainsbury's spread (500g)	1.00	0.68	Sainsbury's Basics butter (500g)	1.76	11.90
	Anchor Spreadable (500g)	2.18	0.68	Anchor butter (500g)	2.40	11.90
	Flora Original spread (500g)	1.18	0.68	Country life butter (500g)	2.36	11.90
	Clover (500g)	1.49	0.68	Kerrygold butter (500g)	1.90	11.90
Meat	Chicken breast (300g)	2.39	1.50	Beef braising steak (440g)	3.49	7.04
	Chicken fillet (500g)	2.18 - 4.00	2.50	Beef mince (500g)	2.20	8.00
	Chicken thighs & drumsticks (721g)	2.37 - 3.00	3.61	Diced casserole steak (440g)	2.50	7.04

Table 1: Product categories and clean / dirty alternatives

Notes: Table displays the exhaustive list of options available to consumers in each product category. For some alternatives in the cola and meat product categories the supermarket modified its price over the course of the experiment, and for consistency it was reflected in the experiment.

category. For cola products low-emissions alternatives are sold in a 2L PET bottle, whereas the high emissions alternatives are sold in aluminum cans. For milk, carbon emissions are proportional to the fat content, for spreads the carbon content is higher for butter relative to margarine (produced mainly from vegetable oil), and for meat it is higher for beef products. Preferences over product characteristics related to emissions (such as plastic vs. aluminum packaging or the type of meat) are a key determinant of substitutability, and therefore capture the effort associated with switching from dirty to clean alternatives. Thus for example, if the type of meat matters to consumers, they will be more resistant to substitute away from the beef alternatives.

2.2 Emissions information treatment

The information treatment consists in a carbon "footprint" label in the form of a stylized footprint and shows the amount of CO_2 (in grams) emitted over the product's production process (i.e. embodied emissions). As shown in Table 1, the difference in carbon emissions between clean and dirty alternatives varies significantly across product categories.

In order to avoid overemphasizing the importance of the information on emissions, which would allow respondent to easily guess the theme of the experiment, we also provided nutritional information. Because this information is readily provided on product packages, consumers who have preferences for these characteristics of the products would already be aware of them and hence it should not overly influence choices.¹³

2.3 Pigovian subsidy to the clean alternative

The Pigovian subsidy treatment decreases the price of the clean alternatives in proportion to embodied carbon emissions. For example, in the case of cola products, respondents are told that "There has been a price change. Products in plastic bottles have a 5p discount due to a GOVERNMENT SUBSIDY received on account of its low carbon footprint." This provides information about differences in relative emissions between alternatives, and makes clear that the change in price is associated with an external government intervention as a way to reduce carbon emissions associated with consumption.

The value of the Pigovian subsidy is determined by the externality created by the consumption of different alternatives. More specifically, starting from an estimate for the social cost of carbon of GBP19/tCO₂ taken from DEFRA (2002), the subsidy is calculated by using the difference in embodied CO_2 emissions between clean and dirty alternatives. The final values of the subsidies are: GBP0.05 for cola products in PET bottles; GBP0.03 for semi-skimmed milk, or GBP0.06 for skimmed milk; GBP0.43 for margarine; and GBP0.21 per kg of chicken.

2.4 Neutrally framed price reduction of the clean alternatives

The change in price in this treatment is equivalent to the subsidy, but the justification is framed in a neutral manner. For example the neutral price change for cola products is presented as follows: "There has been a price change. Products in plastic bottles have a 5p discount because of a change in the price of materials." The change in relative prices is thus caused by market conditions unrelated to the regulation of externalities.

This treatment allows us to quantify how an exogenous price change induces consumers to substitute towards the clean alternative without reference to an external intervention targeting a reduction of carbon emissions. Moreover, this treatment has several advantages. First, it yields an internally consistent estimate of price responsiveness, capturing the willingness (or effort)

¹³ Note that it could potentially be the case that consumers factor in carbon emissions in their initial choices. In our analysis, this would be captured by preferences for product characteristics in each product category as estimated from the baseline product choice.

to substitute between clean and dirty alternatives. Second, this allows calculating monetary equivalents for differences in effectiveness across interventions. Finally, as we show below, it can be used to identify the magnitude of the behavioral effect associated with Pigovian regulation aside from the change in relative prices.

3 Estimation Strategy

In each of the two sequential choice occasions, consumers select one product from a finite set of alternatives, and a natural estimation framework is McFadden's (1974) model for differentiated products. Specifically, in the initial choice, consumers reveal their preferences for the attributes of each product by selecting their preferred alternative in the absence of any interventions. In the second choice, product characteristics are manipulated by the treatments, altering the public good attributes (information and Pigovian subsidy treatments) and relative prices (Pigovian subsidy and neutral price change treatments). By using a structural representation of the choice process we are able to quantify the CATE controlling for preferences over observed product characteristics, and derive money-metric welfare measures associated with the treatments.

In the following we first describe the conceptual framework and proceed by describing our maximum likelihood estimation procedure. Finally we explain how we quantify the EPM for information and for the behavioral effect of Pigovian regulation.

3.1 Conceptual framework

Denote the utility that consumer n derives from alternative j by U_n^j , the price of j by p^j and the utility of all observed and unobserved (non-price) characteristics of that alternative by u_n^j , so that: $U_n^j = u_n^j - p^j$. Further denote relative utility of dirty and clean options as $u_n = u_n^{\text{dirty}} - u_n^{\text{clean}}$ and relative prices as $p = p^{\text{dirty}} - p^{\text{clean}}$. Consumer n will select a dirty alternative if:

$$U_n^{\text{dirty}} > U_n^{\text{clean}} \Leftrightarrow u_n > p.$$
(1)

After observing an initial choice, experimental treatments manipulate both the relative utility from consuming each good and the relative prices.

When clean and dirty alternatives are good substitutes, a small change in relative prices will dominate the difference in utility derived from consuming the two goods. Identification of substitutability is achieved by the neutrally framed price treatment in the form of a price elasticity. Let β_n^{price} denote the change in relative utility induced by neutral price change Δp . The neutral price change will induce consumer *n* to switch to the clean alternative if:

$$\beta_n^{\text{price}} \Delta p > u_n - p.$$
⁽²⁾

In words, the utility weight associated with a reduced price for the clean alternative has to outweigh the surplus derived from consuming the dirty instead of the clean alternative. In turn, price responsiveness of consumers provides a measure of how close or substitutable the two alternatives are.

Turning to the information treatment, denote embodied CO_2 emissions of alternative j by e^j . Providing information reveals individuals' preferences for the public good component of each product, denoted by β_n^{info} . Thus under the information treatment a consumer will switch to the clean alternative when:

$$\beta_n^{\text{info}} \Delta e > u_n - p \,, \tag{3}$$

where $\Delta e = e^{\text{dirty}} - e^{\text{clean}}$. When clean and dirty alternatives are perceived to be good substitutes, the right hand side will be relatively small, and information about the extent of external effects may significantly increase the probability of choosing one of the clean alternatives.

The final treatment is the Pigovian subsidy. This treatment changes relative prices in the same way as the neutral price change does, but it also frames the monetary change as a regulatory intervention targeting relative carbon emissions. We can thus write that a consumer initially choosing the dirty alternative will switch to one of the clean alternatives provided that:

$$\beta_n^{\text{pigou}} \Delta s = \beta_n^{\text{price}} \Delta p + \beta_n^{\text{regul}} > u_n - p , \qquad (4)$$

where Δs is subsidy amount, Δp is the monetary price signal defined above, and β_n^{regul} measures the behavioral impact of the Pigovian regulatory intervention over and above the impact of the change in relative prices. In particular, β_n^{regul} captures the effect of framing the price change as an explicit government intervention related to the external effects associated with each product.

3.2 Structural estimation: Multinomial choice

For each category of product, there is a finite number of alternatives *J* from which the consumers can choose from (see Table 1), and each alternative is described by a set of characteristics or attributes. These are summarized in Table 2. For instance, in the case of cola products, characteristics are packaging (2L PET bottle or aluminum cans), price (in cent), brand (Coca-Cola or Pepsi), and 'Light' or 'Zero/Max' versions.¹⁴ For cola, spread and meat product categories the first attribute (attribute 1) is an indicator variable equal to one if a particular product is one of the dirty alternatives, zero otherwise. In the case of milk products there are 3 different alternatives that vary only in terms of the amount of fat (in grams), and is perfectly collinear with embodied carbon emissions. Thus preferences for clean and dirty versions of milk products are given by preferences for the fat content.

Assuming that individual n chooses alternative j if the utility of j is greater than any other alternatives i in the choice set, the probability that option j is selected by individual n is:

$$P_n^j = \operatorname{Prob}(U_n^j > U_n^i), \quad \forall i \neq j.$$
(5)

Following McFadden (1974), we decompose the utility of product j into a deterministic part observed by the researcher, denoted by V_n^j , and an unobserved part denoted ε_n^j , so that: $U_n^j = V_n^j + \varepsilon_n^j$. Given the notation developed above, we specify the observed part of utility as:

$$V_n^j = \gamma_n' Z^j + I_{\text{info}}^{\text{clean}} \cdot \beta_n^{\text{info}} \Delta e + I_{\text{pigou}}^{\text{clean}} \cdot \beta_n^{\text{pigou}} \Delta s + I_{\text{price}}^{\text{clean}} \cdot \beta_n^{\text{price}} \Delta p$$
(6)

where Z^{j} is a vector of observed product attributes as defined in Table 2, I_{info}^{clean} , I_{pigou}^{clean} and I_{price}^{clean} are indicator variables equal to one if a particular choice is done under a given treatment *and* option j is one of the clean alternatives, and γ, β are parameters to be estimated from the data.

The unobserved part of the utility ε_n^j is assumed to be identically and independently distributed according to an extreme value type 1 distribution, so that choice probabilities take the

¹⁴ These attributes represent an exhaustive list of observed dimensions across which product alternatives differ. Preferences for attributes that do not vary across alternatives, such as for example the country of origin, are not identified. Obviously, there can be other factors that influence choices, and as we show below the importance of unobservable characteristics will be reflected in the size of the structural error term.

	Cola	Milk	Spread	Meat		
Nr. of alternatives ^a	12	3	10	6		
Product attributes						
Attribute 1 ^b Attribute 2 Attribute 3 Attribute 4 Attribute 5 Attribute 6 Attribute 7 Attribute 8	ALU cans (=1) Price (GBP cent) Coca-Cola brand (=1) Light (=1) Zero/Max (=1) - -	Fat cont. (g) - - - - - - - -	Butter (=1) Price (GBP cent) Lurpak brand (=1) Sainsbury brand (=1) Anchor brand (=1) Proteins (g) Fat (g) Salt (g)	Beef (=1) Price (GBP cent) Protein (g) Salt (g) Fat (g) Weight (g) -		
Policy treatmentsInformation label Pigovian subsidyDifference in embodied carbon emissions between clean and dirty alternatives (kgCO2) Pigovian subsidy to the price of the clean options (GBP cent)Neutral price changeNeutrally framed decrease in price of clean options (GBP cent)						

Table 2: Choice set, product attributes and policy treatments

Notes: Table lays out the data structure underlying estimation of a discrete choice demand model. ^{*a*}Each product category includes the number of product alternatives reported in Table 1. ^{*b*}Attribute 1 determines whether a product alternative belongs to the set of dirty alternatives, and thus captures preferences for the dirty version of each product. For milk, in which the semi-skimmed alternative has a carbon footprint in between that of whole milk and skimmed milk alternatives, we use a continuous measure for the fat content.

convenient logit form:

$$P_n^j = \operatorname{Prob}(V_n^j - V_n^i > \varepsilon_n^j - \varepsilon_n^i) = \frac{e^{V_n^j}}{e^{V_n^j} + \sum_i e^{V_n^i}}, \quad \forall i \neq j$$
(7)

and the log-likelihood function writes:

$$\log L = \sum_{n=1}^{N} \sum_{j=1}^{J} \sum_{t=1,2} d_{nt}^{j} \log P_{nt}^{j}$$
(8)

where d_{nt}^{j} is an indicator function equal to 1 if alternative j is selected in choice t, zero otherwise. When preference parameters are the same for each individuals, the model reduces to the standard multinomial logit (MNL) framework, which makes maximum likelihood estimation of the structural parameters straightforward. However, the MNL model implies restrictive substitution patterns, the so-called independence of irrelevant alternatives (IIA) property. To exploit the panel structure of the data and allow the error term to be correlated across alternatives, and thereby relax the IIA requirement, we account for unobserved preference heterogeneity using a random parameter or mixed logit (MXL) specification (Revelt and Train, 1998; McFadden and Train, 2000). The MXL model is estimated via simulated maximum likelihood, where unob-

served preference parameters are assumed to be normally distributed in the population, and we approximate the integral of the unconditional probability of each panel choice using 200 Halton draws.¹⁵

In a discrete choice demand model a change in one of the attributes affects the choice probabilities (or market shares) of all options, so that the vector of estimated parameters is not directly tied to marginal effects on choice probabilities. In addition, because the estimated coefficients are not separately identified from the variance of the error term (or scale parameter), coefficients cannot be directly compared across estimated models. Thus in order to compare results across product categories we use the estimated structural model to simulate the impact of the treatments on the choice probability of clean alternatives.¹⁶

3.3 Equivalent price metric for policy interventions

One benefit of estimating a structural utility-maximization model is the ability to conduct welfare analysis, and to more broadly inform policy design. Here we build on the work of Chetty et al. (2009) who compare a *tax-demand* curve and a *price-demand* curve, and Allcott and Taubinsky (2015) who similarly define the EPM as follows:

$$EPM^{treatment} = \frac{D^{treatment}(p) - D(p)}{D'(p)}$$
(9)

where $D^{\text{treatment}}(p) - D(p)$ is the change in demand of the clean alternative induced by the treatment, and D'(p) is the price responsiveness of demand. In our experiment, the neutral price change treatment provides a relevant measure of D'(p), as it directly manipulates relative prices of clean and dirty alternatives.¹⁷

Given the notation developed above, the EPM of information is equal to the ratio between

¹⁵ Note that the taste normality assumption mainly serves tractability of the simulation process and that the preference parameters measuring treatment effects are held fixed. We also considered specifications with random coefficients for the treatment effects but encountered numerical convergence issues likely caused by the fact that we only observe one choice per respondent in the presence of a treatment.

¹⁶ For MXL models choice probabilities have no closed-form expressions, and we rely on bootstrapping to obtain standard errors. As the simulation-based estimation procedure for MXL specifications is computationally intensive we rely on 1000 replications. Although this number is relatively small, the ensuing inference yields similar conclusions to the closed-form results drawn from MNL specifications, which suggests that this is appropriate.

¹⁷ Note that the structural model effectively provides a framework to extrapolate the treatment effect in an internally consistent manner. This is a valid approach as long as the model of behavior is appropriate and the social value of carbon (and hence the Pigovian price) is fixed. As we discuss further in the concluding section, if there is uncertainty in the correct Pigovian price, our research design could be extended to estimate the average marginal bias along the demand curve (Allcott and Taubinsky, 2015).

Statistic	Definition		Units
EPM ^{info}	$\frac{\Delta e \beta^{\text{info}}}{\beta^{\text{price}}}$	EPM for the information label	GBP cent
	$\frac{\beta^{\text{info}}}{\beta^{\text{price}}}$	EPM for the information label per unit of emissions	GBP per tCO ₂
EPM ^{regul}	$\frac{\beta^{\text{regul}}}{\beta^{\text{price}}}$	EPM for the behavioral effect of Pigovian regulation	GBP cent
	$\frac{\beta^{\rm regul}}{\beta^{\rm price}} \frac{1}{\Delta e}$	EPM for the behavioral effect of Pigovian regulation per unit of emissions	GBP per tCO_2

Table 3: Definition of equivalent price metric (EPM) statistics

the utility weight associated with the information treatment and that of the neutral price change treatment:

$$EPM^{info} = \frac{\beta^{info}}{\beta^{price}}.$$
 (10)

Intuitively, EPM^{info} measures the change in relative prices that would generate a behavioral change of the same magnitude to that of the information treatment, capturing consumers' valuation of relative carbon emissions embodied in the products.

Similarly, our framework can be used to quantify the behavioral effect associated with Pigovian regulation, and given the notation developed above the associated EPM is given by:

$$\text{EPM}^{\text{regul}} = \frac{\beta^{\text{regul}}}{\beta^{\text{price}}} = \frac{\beta^{\text{pigou}}\Delta s - \beta^{\text{price}}\Delta p}{\beta^{\text{price}}}.$$
 (11)

In words, EPM^{regul} reflects the behavioral effect (in monetary terms) of information provided by a Pigovian intervention netting out the change in relative prices. If an explicit external intervention attenuates incentives to behave prosocially and reduce emissions, this measure quantifies the change in relative prices that would compensate the negative behavioral effects. For example, as initially put forward by Gneezy and Rustichini (2000b) in a different context, if consumers' valuation of emissions reduction is higher than the Pigovian price signal, a Pigovian intervention may crowd out intrinsic motivation to switch towards a cleaner alternative. The Pigovian price would therefore need to be set higher in order to compensate this behavioral trait (Allcott et al., 2014; Farhi and Gabaix, 2015).

Table 3 summarizes the statistics used to quantify the monetary equivalent of the information treatment and the behavioral effect associated with Pigovian regulation. Note that these quan-

tities are free of the scale parameter and are thus directly comparable across models. Moreover, to compare EPM estimates across product categories, we control for the fact that products differ with respect to the level of embodied emissions (Δe) and estimate EPM per unit of emissions.

4 Data and Results

4.1 Descriptive statistics: Choice frequencies before and after treatments

Table 4 shows choice frequencies for the clean alternatives across product categories, before and after each treatment. The initial share of clean alternatives ranges from around 10 percent for milk products to 50 percent for spreads. These differences suggest that controlling for preferences over observed product characteristics is important in order to appropriately identify the treatment effects.

Descriptive statistics further show that all the treatments induced significant increases in the choice probability of clean alternatives. There is, however, ample variation both across treatments and across product categories. Comparing the impact of treatments within product categories, the proportion of consumers who switched towards clean alternatives is somewhat larger with an information label as compared to a Pigovian subsidy. Moreover, the neutrally framed price change also has a larger impact as compared to the Pigovian subsidy, suggesting that the regulatory intervention has a negative behavioral effect. These observations and their statistical significance (in terms of changes in choice frequencies) are discussed in details by Perino et al. (2014), although they aggregate all product categories together.

Comparing the impact of treatments across product categories, the largest percentage change is generally observed for cola products (substitution towards products in plastic packaging). However both the size of the treatment and the initial choice frequency differs, rendering comparisons difficult.

4.2 Econometric results

We now turn to the estimation results for the structural model which provides evidence on: (i) the CATE for each product controlling for preferences over product characteristics and embodied carbon emissions; (ii) substitutability across clean and dirty alternatives; (iii) an EPM for the information treatment and behavioral bias associated with the Pigovian intervention.

		Information label	Pigovian subsidy	Neutral price change
Cola	Before treatment	47.6	31.3	27.0
	After treatment	66.7	50.0	69.8
	Difference	19.1	18.7	42.8
Milk	Before treatment	12.3	6.1	8.3
	After treatment	19.3	10.9	14.3
	Difference	7.0	4.8	6.0
Spread	Before treatment	55.2	56.3	39.0
-	After treatment	82.8	63.4	51.2
	Difference	27.6	7.1	12.2
Meat	Before treatment	12.5	20.6	21.7
	After treatment	32.1	30.2	33.3
	Difference	19.6	9.6	11.6

Table 4: Observed choice frequencies of clean alternatives by product category (percentage)

Notes: This table reports the percentage of respondents who selected one of the clean alternative before and after each treatments.

Estimation results from MNL and MXL models are reported in Table 5. Recall that coefficients on attribute 1 provide evidence about preferences for the dirty alternatives. Except for spread products, these estimates have a positive sign and are highly statistically significant, which is consistent with the relatively large initial choice frequency for dirty alternatives reported in Table 4. Other variables capturing preferences for product attributes are mostly statistically significant at conventional levels, suggesting that the structural model provides a good account of observed choices. This is confirmed by comparing simulated choice probabilities of clean alternatives with the actual choice frequencies observed in our sample (Figure 1). The MXL specification further suggests the presence of preference heterogeneity, as demonstrated by the statistically significant standard deviation estimates. The goodness-of-fit statistics generally favor the MXL models and, since it provides a more flexible representation of behavior, in the rest of the paper we consider only results from the MXL specifications.

The main estimates of interest are those associated with the three treatments, as they quan-

		Cola			Milk			Spread			Meat	
	MNL (1)	IXM	(2)	(E) MNI	IXW	L (4)	MNL (5)	IXIN	(9)	MNL (7)	TXW	(8)
		Mean	Std-dev.		Mean	Std-dev.		Mean	Std-dev.		Mean	Std-dev.
Information label	3.16^{***}	10.12^{*}	I	4.17***	4.30^{***}	I	0.15***	1.05^{***}	I	0.03	5.80^{***}	I
	(0.85)	(5.60)		(0.61)	(0.60)	I	(0.04)	(0.31)		(0.05)	(0.52)	
Pigovian subsidy	0.19^{***}	0.86^{***}	I	0.18^{***}	0.20^{***}	I	0.01^{***}	0.08^{**}	I	0.03^{**}	1.87^{***}	I
	(0.05)	(0.23)		(0.02)	(0.02)	I	(0.005)	(0.04)		(0.01)	(0.18)	
Neutral price change	0.36^{***}	1.71^{***}	I	0.25^{***}	0.26^{***}	I	0.003	0.09**	I	0.04***	1.89^{***}	I
	(0.06)	(0.52)		(0.02)	(0.02)	I	(0.005)	(0.04)		(0.01)	(0.27)	
Attribute 1	0.82**	3.42^{**}	7.80***	0.24^{***}	0.25^{***}	0.11^{***}	0.43	2.27	12.49^{***}	2.68***	110.18^{***}	115.98
	(0.39)	(1.34)	(2.52)	(0.01)	(0.02)	(0.04)	(0.36)	(1.71)	(4.33)	(0.55)	(4.09)	(00.0)
Attribute 2	0.001	0.004	0.001	I	I	I	-0.02***	-0.04***	0.07**	0.002	0.00	0.03^{***}
Attribute 2	(0.002) 1 E4***	(0.004) 2 24***	(0.004) 2 27***				(0.005) 2 EA***	(0.02) 1 2 1	(0.03) 15 05**	(0.004) 0.01***	(0.01)	(10.0)
	(0.21)	(0.77)	0.80) (0.80)	I	I	l	(0.52)	(4.45)	(7.02)	0.01	(0.01)	(00.0)
Attribute 4	-0.64***	-2.67***	5.62^{***}	I	I	I	-1.79***	-9.30**	6.12^{**}	0.24**	0.47***	0.02
	(0.18)	(0.97)	(1.39)				(0.36)	(4.07)	(2.93)	(0.10)	(0.17)	(0.02)
Attribute 5	-1.63***	-7.38***	6.69***	I	I	I	0.23	-4.20*	8.34***	-19.07***	-45.72***	3.39
	(0.25)	(1.96)	(1.54)				(0.29)	(2.38)	(2.96)	(5.50)	(16.28)	(7.32)
Attribute 6	I	I	I	I	I	I	-3.29***	-11.92^{**}	24.06^{***}	0.29^{***}	0.73***	0.53^{***}
							(0.70)	(5.41)	(8.95)	(0.10)	(0.24)	(0.19)
Attribute 7	I	I	I	I	I	I	0.11^{***}	0.66***	0.17^{**}	I	I	I
							(0.02)	(0.23)	(0.08)			
Attribute 8	I	I	I	I	I	I	1.28^{***}	5.38^{**}	8.43**	ļ	I	I
							(0.25)	(2.22)	(3.36)			
Respondents	148	1	48	372	.03	72	182	18	32	152	152	
Log Pseudo-LL	-611.3	-45	6.5	-882.5	-88	80.9	-753.5	-53	8.9	-453.9	-359	6
AIC	1239.5	101	19.0	1773.0	177	71.8	1529.0	111	15.7	925.9	747.	2
BIC	1262.6	105	58.0	1788.7	179	91.4	1564.3	117	76.6	953.1	790	2
Pseudo \mathbb{R}^2	0.195	0.9	346	0.144	0.1	146	0.137	0.3	383	0.233	0.39	2
<u>Notes: This table report</u>	s preference	parameters	estimates fo	r MNL and MX	L specificati	ions. For MX	L models we e	stimate both	the mean an	d standard-devi	iation (std-dev	:) of
normally distributed pr	eference para	meters. Sai	mple size for	each product	category is 1	twice the nu	mber of respor	ndent (every	respondent n	nakes two choic	tes). Standard	errors
are clustered at the res	pondent level	and report	ed in parentl	hesis. ***, **, *	: statisticall	ly significant	at 1, 5 and 10	percent resp	pectively. The	list of attribute	s for each pro	duct is
in Table 2.												

Table 5: Estimation of product-specific multinomial choice models



Figure 1: Initial market shares of clean alternatives across products(%)

Notes: Figure 1 plots the observed market share of clean alternatives before the treatments against that predicted by MNL and MXL models reported in Table 5.

tify the impact of each treatment on choice probabilities.¹⁸ We find that the estimated utility weights associated with the information label, Pigovian subsidy and neutral price change all have the expected sign, having a positive impact on the choice probability of the clean alternatives. Moreover, the treatments variables in MXL models are all statistically significant at conventional levels.

As mentioned previously, estimation results cannot be directly compared across product categories, and in the following we study the impact of treatments on simulated choice probabilities (or market shares). We start with substitutability between clean and dirty and alternatives, and then quantify the effectiveness of policy instruments. Finally we report estimates of the EPM for the information treatment and quantifying the behavioral effect associated with the Pigovian intervention.

4.2.1 Measures of substitutability for each product category

Evidence about substitutability between clean and dirty alternatives is based on the responsiveness to a neutral price change. In Table 6 we report changes in simulated choice probabilities for the clean alternatives associated with the neutral price change treatment, derived from MXL model. The impact of the neutral price change specified in the experiment (i.e. the CATE) refers

¹⁸ As Table 1 reports, the information treatment is coded as the difference (in kg of CO₂) between the clean and dirty alternatives (for milk product we take whole milk as the baseline). The neutrally framed price change and the Pigovian subsidy are coded in GBP cent.

	Cola	Milk	Spread	Meat
Neutral price change: CATE	42.10***	11.27***	11.07***	8.52***
	(2.52)	(0.61)	(3.04)	(3.43)
10% neutral price change	70.06***	14.96***	4.46***	10.92***
	(4.13)	(0.63)	(1.23)	(4.52)
Neutral price change (GBP cent)	9.02***	1.84***	0.25***	0.45***
	(0.75)	(0.13)	(0.09)	(0.16)
Respondents	148	372	182	152

Table 6: Substitutability between clean and dirty alternatives

Notes: This table displays the marginal impact of the neutrally framed price change on simulated choice probabilities (or market shares) of clean alternatives, reported in percentage points difference. We report changes in simulated choice probabilities corresponding to the conditional average treatment effect (CATE), a 10% price reduction, and a 1 GBP cent price reduction. Simulated choice probabilities are derived from the MXL specifications reported in Table 5. Bootstrapped standard errors clustered at the respondent level reported in parenthesis. ***, **, *: statistically significant at 1, 5 and 10 percent respectively.

to different price changes for each product. To make comparison across products possible, we employ the estimated model to simulated changes in probabilities for a normalized treatment effect. More specifically, we consider the effect of a relative (10%) and absolute (1 GBP cent) neutrally framed price change.

Results indicate that the CATE is positive and highly statistically significant for all the products, with a 42 percent increase in the choice probability of cola products in plastic bottles, and around 10 percent for milk, spread and meat product categories. These figures are close to the descriptive statistics reported in Table 4, which again suggests that the model fits the data well.

Comparing results across products based on normalized treatment sizes, results show that changes in simulated choice probabilities is highest for cola products, followed by milk, meat and spread products. The ranking of products in terms of substitutability between clean and dirty alternatives is similar for absolute and proportional treatment sizes, and it is also confirmed if we consider the change in choice probabilities for clean products relative to a situation without the treatment. Specifically, for a unit change in relative prices, the choice probability of clean cola products increases from around 32% to 41%, or a 35% increase. The corresponding increase is 17% for milk, 1.2% for meat and 0.5% for spread.

4.2.2 Policy instruments: Comparison across products categories

Having established that substitutability between clean and dirty alternatives is highest for cola products, followed by milk, meat and spread products, we now study the impact of information and Pigovian treatments. Results derived from the MXL models are reported in Table 7. Panel A shows the CATE measured by the change in choice probabilities of clean alternatives. Because embodied carbon emissions of clean and dirty alternatives differ across product categories, these results can only be compared within products. To compare treatments across products, Panel B reports simulated changes in clean choice probabilities for normalized treatment sizes. For the information treatment, we consider a proportional 10% difference in emissions from clean and dirty alternatives as well as a unit difference measured in kg of CO₂. Similarly, for the Pigovian subsidy we consider both a 10% and a one GBP cent subsidy.

We find that the CATE is economically and statistically significant, ranging from 12 to 28 percent for the information label, and from 5.5 to 22 percent for the Pigovian subsidy. Comparing the impact of the treatments within products, the CATE of information is larger than that of the Pigovian subsidy for all product categories. This result is in line with aggregate results reported in Perino et al. (2014), and is also reasonably close to observational studies such as Teisl et al. (2002) and Bjorner et al. (2004).

Turning to evidence from normalized treatment effects, we find important differences across products. In particular, there is clear evidence that the effectiveness of policy interventions is related to substitutability between clean and dirty alternatives. For the information treatment, the choice probability of clean alternatives is most responsive for cola and milk products, while spread and meat are significantly less sensitive. Results for the Pigovian treatment similarly shows that cola products are very responsive to the intervention, whereas meat and in particular spread products are not. The ranking of products is again similar for absolute and proportional treatment sizes.

While these results accord with expectations, they should be contrasted with those reported in Panel A (and in Table 4), which suggest sizable impacts of the information and Pigovian subsidy treatments on both spread and meat products, and a small impact for milk products. Controlling for the initial market share of the dirty products (and underlying preferences for observed product characteristics) and for variation in embodied carbon across products (and associated size of the treatment effect) thus highlights the role of substitutability between clean

	Cola	Milk	Spread	Meat
Panel A: Within product comparison				
Information label: CATE	27.00***	12.23***	29.45***	13.78*
	(1.52)	(1.03)	(9.17)	(6.99)
Pigovian subsidy: CATE	21.82***	8.63***	8.47***	8.59***
	(1.24)	(0.72)	(2.48)	(2.89)
Panel B: Across product comparison				
10% information label	5.43***	5.68***	3.19**	1.52*
	(0.74)	(0.63)	(1.34)	(0.82)
Information label (kgCO ₂)	48.23***	19.70***	2.69***	0.93***
	(2.04)	(0.54)	(0.89)	(0.34)
10% Pigovian subsidy	59.96***	12.01***	3.47***	10.47***
	(1.90)	(0.82)	(0.87)	(2.34)
Pigovian subsidy (GBP cent)	4.46***	1.37***	0.20***	0.45***
	(0.18)	(0.13)	(0.04)	(0.03)
Respondents	148	372	182	152

Table 7: Effectiveness of policy instruments within and across products

Notes: Panel A displays the marginal impact of the information treatment and Pigovian subsidy on simulated choice probabilities (or market shares) of clean alternatives, reported in percentage points differences. We report changes in simulated choice probabilities corresponding to the conditional average treatment effect (CATE). Panel B displays the same but for a normalized treatment size representing a 10% difference (in relative emissions or relative prices) or a unit difference (in kgCO₂ and GBP cent for the information label and Pigovian subsidy respectively). Simulated choice probabilities are derived from the MXL specifications reported in Table 5. Bootstrapped standard errors reported in parenthesis. ***, **, *: statistically significant at 1, 5 and 10 percent respectively.

and dirty alternatives.

4.2.3 EPM for information and the behavioral effect of Pigovian regulation

This section concludes the comparison of regulatory interventions by reporting EPM estimates for information and the behavioral effect of Pigovian regulation, using the statistics laid out in Table 3. Results derived from MXL models are reported in Table 8. Panel A reports EPM estimates for the information treatment both for the specified size of the treatment effect (i.e. the EPM of the CATE measured in GBP cent) and per unit of emission (measured in GBP per tCO₂). The latter estimates provide a basis for a comparison of results across products. Panel B

	Cola	Milk	Spread	Meat			
Panel A: Equivalent price metric f	or informatio	n (EPM ^{info})					
$\Delta e eta^{ ext{info}} / eta^{ ext{price}}$ (GBP cent)	3.07**	6.61***	128.25**	33.76***			
	(1.47)	(1.12)	(53.8)	(3.19)			
$\beta^{ m info}/\beta^{ m price}$ (GBP per tCO ₂)	59.08**	165.15***	114.26**	30.69***			
	(28.35)	(28.07)	(47.93)	(11.19)			
Panel B: Equivalent price metric for behavioral impact of Pigovian regulation (EPM ^{regul})							
$eta^{\mathrm{regul}}/eta^{\mathrm{price}}$ (GBP cent)	-2.50***	-1.50**	-6.82	-0.25			
	(0.68)	(0.72)	(15.77)	(1.93)			
$eta^{ ext{regul}}/eta^{ ext{price}}rac{1}{\Delta e}$ (GBP per tCO ₂)	-48.06***	-37.46**	-6.07	-0.22			
	(13.14)	(17.92)	(14.05)	(1.76)			
Respondents	148	372	182	152			

Table 8: EPM for information and the behavioral effect of Pigovian regulation

Notes: Panel A displays the equivalent price metric (EPM) for the information label measured in GBP cent (i.e. the EPM of the CATE) and in GBP per tCO₂ (allowing a comparison across product categories). Panel B displays the EPM for the behavioral effect of Pigovian regulation measured in GBP cent (referring to the CATE) and in GBP per tCO₂ (allowing a comparison across product categories). See Table 3 for a definition of the EPM statistics. All estimates are derived from the MXL specifications reported in Table 5. Standard errors clustered at the respondent level obtained via the delta methods reported in parenthesis. ***, **, *: statistically significant at 1, 5 and 10 percent respectively.

reports the EPM associated with the behavioral effect of Pigovian regulation, both for the CATE and per unit of emissions.

For the information treatment, the EPM of the CATE ranges from GBP0.03 for cola products to GBP1.28 for spread products, and all estimates are statistically significantly different from zero at conventional levels. Recall that for both cola and spread products the CATE of information is relatively large (Table 7, Panel A), but the neutrally framed price change has a much larger impact on choice probabilities for cola products relative to spread products (Table 6).

However, once we control for variation in products' embodied emissions, EPM estimates are more similar across product categories, ranging from around GBP31/tCO₂ for meat products to around GBP165/tCO₂ for milk. Because $\beta^{info}/\beta^{price}$ can be interpreted as a marginal rate of substitution between embodied carbon emissions and money, it also provides an estimate for consumers' monetary valuation of embodied carbon emissions. For all products, these numbers are larger than the Pigovian price used in the experiment (GBP19/tCO₂, see DEFRA, 2002) and most other estimates of the social cost of carbon.

Reported in Panel B, the EPM associated with the behavioral effect of Pigovian regulation is negative, reflecting the fact that for all products the Pigovian subsidy treatment is less effective than a neutrally framed price change (see equation 11). When measured per tCO_2 , this measure of the behavioral bias also has interesting implications for the setting of the Pigovian tax rate, as it measures the change in relative prices that would be required to compensate the negative behavioral effects associated with Pigovian regulation. We find that the Pigovian tax rate should be increased by around GBP48/tCO₂ and GBP37.5/tCO₂ for for cola and milk products respectively, while our results indicate no statistically significant behavioral bias for spread and meat products.

Therefore, an important conclusion from the exercise is that the negative behavioral impacts are related to the effort of behaving prosocially, here switch to one of the less preferred cleaner alternatives, as measured by the substitutability between clean and dirty alternatives. More specifically, the EPM capturing the behavioral effects of Pigovian regulation is significantly larger when consumers perceive clean alternatives to be close substitutes to the dirty ones, and for these products (cola and milk here) adjustments to the Pigovian tax rate are substantial. For the other products we consider, meat and spread, substitutability is measured to be lower, and in turn our EPM estimates are much smaller and not statistically different from zero.

5 Discussion and Conclusion

Market-based instruments, and in particular Pigovian regulation, have the potential to make consumers internalize external effects associated with their choices. However, in the presence of behavioral agents framing a change in relative prices as an explicit intervention to encourage or reward the provision of a global public good may backfire. This could notably be due to the fact that intrinsic motivation to behave prosocially can decline in response to an external intervention to promote effort in this direction Gneezy et al. (2011). From this perspective, regulatory interventions ought to be adjusted to account for these behavioral traits (Allcott et al., 2014; Farhi and Gabaix, 2015).

In this paper we have used data on consumption behavior in a controlled supermarket shopping environment to shed light on the magnitude and policy relevance of these behavioral effects. In our experiment, we found that consumers responded to information by exerting an effort in the form of a substitution from a dirty product alternative to cleaner one. Estimates from our structural model indicate that the implied value of carbon emissions, as revealed by choices once information about emission was revealed, was significantly above most estimates of Pigovian tax rates. Moreover, experimental results revealed that a monetary incentive explicitly motivated by the internalization of carbon emissions was less effective as compared to a neutrally framed change in relative prices of the same magnitude. An implication is that the price signal of Pigovian regulation would need to be set above its socially efficient level (i.e. the marginal damages) in order to compensate the negative behavioral effect associated with the external regulatory intervention.

While, to the best of our knowledge, this paper constitutes the first attempt to draw policy design implications based on the behavioral effect of Pigovian regulation, our results are inevitably incomplete and open the door to further research on a number of fronts. First, we find evidence that the negative behavioral effect of Pigovian regulation is significantly larger for products with close substitutes available, which suggests that environmental taxes ought to be set higher for products with price elastic demand. While this finding can be related to variations in self-image benefits of prosocial behavior (Bénabou and Tirole, 2003, 2006), additional work on the role of effort in relation to financial incentives is warranted. Second, our investigation has focused on a point-estimate of the EPM, referring to a single estimate of the Pigovian price signal (or social cost of carbon). In principle, however, our design could be extended to evaluate the behavioral impact of environmental taxes along the entire demand curve. In fact, evidence reported in Allcott and Taubinsky (2015) suggests that the EPM can vary at different relative prices, so that a broader examination of the behavioral effect of environmental taxes is potentially important. Third, in our experiment we consider a subsidy to the clean alternatives. Further research should consider the behavioral effects associated with a Pigovian tax instead. Fourth, in actual applications of externality-correcting price signals, behavioral effects associated with Pigovian regulation may be a function of salience. In cases where salience of the policy intervention declines with time, the behavioral effects we measure could potentially be attenuated. On the other hand, empirical evidence reported in Rivers and Schaufele (2015) suggests that consumers' response to environmental taxes differ from market-induced price variations. Further research could therefore investigate, over time, how salience and the behavioral

effects we measure interact.

Finally, in light of consumers' implicit valuation of carbon emissions, which is significantly larger than a Pigovian subsidy level, increasing the Pigovian subsidy rate may reduce the behavioral effects by providing information that is in line with consumer's prior valuation (Gneezy and Rustichini, 2000b). Indeed it may be that consumers inferred from the relatively low Pigovian price signal that climate change is not as problematic as they thought it might be. However, if negative behavioral effects are driven by moral licensing, so that paying for emissions relieves the moral cost of socially harmful behavior, it is conceivable that increasing the Pigovian price signal would further erode the effectiveness of regulation, as it also increases the ability to alleviate guilt. Discriminating among these two sources of behavioral bias thus appears to be an important research avenue.

Appendix A Sample composition

	Mean	Std. Dev.	Min	Max
Cola subsample (N=148)				
Male indicator	0.43	0.49	0	1
Age (in years)	34.8	11.87	21	72
Education ^{<i>a</i>}	1.80	0.81	1	3
Income ^b	3.86	2.70	1	9
Children in the household c	0.69	1.09	0	6
Non-white indicator	0.44	0.50	0	1
Milk subsample (N=372)				
Male indicator	0.35	0.48	0	1
Age (in years)	37.51	11.98	18	80
Education ^a	1.79	0.74	1	3
Income ^b	4.01	2.73	1	9
Children in the household c	0.63	0.98	0	6
Non-white indicator	0.33	0.47	0	1
Spread subsample (N=182)				
Male indicator	0.29	0.46	0	1
Age (in years)	38.55	12.06	21	79
Education ^{<i>a</i>}	1.81	0.76	1	3
Income ^b	3.63	2.50	1	9
Children in the household c	0.65	1.01	0	6
Non-white indicator	0.38	0.49	0	1
Meat subsample (N=152)				
Male indicator	0.35	0.48	0	1
Age (in years)	38.20	12.14	21	79
Education ^a	1.75	0.71	1	3
Income ^b	4.26	2.83	1	9
Children in the household c	0.62	1.09	0	6
Non-white indicator	0.28	0.45	0	1

Table A1: Demographic variables by product category

Notes: ^{*a*}Education is coded as: 1 - Non-university education or equivalent; 2 - University education (includes current undergraduate students); and 3 - Postgraduate level (includes current postgraduate students). ^{*b*}In GBP thousand per year. ^{*c*} Number of children in the household.

Appendix B Treatment randomization

	Information	Pigov	vian	Pri	ce	Prod	luct
	label	subs	idy	chai	1ge	remo	oval
	Mean	Mean	Diff.	Mean	Diff.	Mean	Diff.
Male	0.35	0.40	-0.05	0.30	0.05	0.42	-0.07*
indicator	(0.48)	(0.49)	(0.04)	(0.46)	(0.04)	(0.49)	(0.04)
Age (in	36.92	37.54	-0.62	37.21	-0.30	35.72	1.20
years)	(11.97)	(12.42)	(1.10)	(12.17)	(1.08)	(12.40)	(1.08)
Education ^a	1.78	1.79	-0.01	1.83	-0.05	1.75	0.03
	(0.75)	(0.74)	(0.07)	(0.05)	(0.07)	(0.73)	(0.07)
Income ^b	30.61	32.05	-1.45	32.45	-1.84	29.54	1.07
	(20.77)	(20.81)	(2.04)	(20.53)	(2.03)	(19.92)	(1.96)
Children in	0.58	0.57	0.01	0.56	0.02	0.65	-0.07
household ^c	(1.01)	(0.96)	(0.09)	(0.91)	(0.09)	(0.97)	(0.09)
Non-white	0.39	0.37	0.02	0.34	0.04	0.38	0.001
indicator	(0.49)	(0.48)	(0.04)	(0.48)	(0.04)	(0.49)	(0.04)

Table B1: Differences in means across treatments

Notes: Means are reported by sub-samples with standard deviations in parenthesis below. The column with 'differences' reports differences in means between the respective sub-samples and the information label treatment, with t-statistics reported in parenthesis below. ***, **, *: statistically significant at 1, 5 and 10 percent respectively. ^aEducation is coded as: 1 – Non-university education or equivalent; 2 – University education (includes current undergraduate students); and 3 – Postgraduate level (includes current postgraduate students). ^bIn GBP thousand per year. ^c Number of children in the household.

References

- Acemoglu, D., P. Aghion, L. Bursztyn, and D. Hemous (2012) "The environment and directed technical change," *American Economic Review*, 102, pp. 131–166.
- Acemoglu, D., U. Akcigit, D. Hanley, and W. Kerr (2016) "Transition to clean technology," *Journal* of *Political Economy*, 124 (1), pp. 52 104.
- Aghion, P., A. Dechezlepretre, D. Hemous, R. Martin, and J. Van Reenen (2016) "Carbon taxes, path dependency and directed technical change: Evidence from the auto industry," *Journal of Political Economy*, 124 (1), pp. 1 51.
- Allcott, H. (2011) "Social norms and energy conservation," *Journal of Public Economics*, 95, pp. 1082 1095.
- Allcott, H., S. Mullainathan, and D. Taubinsky (2014) "Energy policy with externalities and internalities," *Journal of Public Economics*, 112, pp. 72 88.
- Allcott, H. and D. Taubinsky (2015) "Evaluating behaviorally-motivated policy: Experimental evidence from the lightbulb market," *American Economic Review*, 105 (8), pp. 2501 2538.
- Andreoni, J. (1990) "Impure altruism and donations to public goods: A theory of warm-glow giving," *The Economic Journal*, 100, pp. 464 477.
- Ariely, D., A. Bracha, and S. Meier (2009) "Doing good or doing well? Image motivation and monetary incentives in behaving prosocially," *American Economic Review*, 99 (1), pp. 544–555.
- Becker, G. S. (1974) "A theory of social interactions," *Journal of Political Economy*, 82, pp. 1063 1093.
- Bertrand, M., D. Karlan, S. Mullainathan, E. Shafir, and J. Zinman (2010) "What's advertising content worth? Evidence from a consumer credit marketing field experiment," *Quarterly Journal of Economics*, 125, pp. 263–306.
- Bjorner, T. B., L. Hansen, and C. S. Russell (2004) "Environmental labelling and consumers' choices an empirical analysis of the effect of the Nordic Swan," *Journal of Environmental Economics and Management*, 47 (3), pp. 411 434.
- Bénabou, R. and J. Tirole (2003) "Intrinsic and extrinsic motivation," *Review of Economic Studies*, 70, pp. 489 520.
- (2006) "Incentives and prosocial behavior," *American Economic Review*, 96, pp. 1652 1678.
- Bowles, S. and S. Polanía-Reyes (2012) "Economic incentives and social preferences: Substitutes or complements?" *Journal of Economic Literature*, 50(2), pp. 368–425.
- Brekke, K. A., S. Kverndokk, and K. Nyborg (2003) "An economic model of moral motivation," *Journal of Public Economics*, 87, pp. 1967 1983.
- Chetty, R., A. Looney, and K. Kroft (2009) "Salience and taxation: Theory and evidence," *American Economic Review*, 99 (4), pp. 1145–1177.
- Davis, L. W. and L. Kilian (2011) "Estimating the effect of a gasoline tax on carbon emissions," *Journal of Applied Econometrics*, 26, 7, pp. 1187–1214.

- DECC (2013) "Updated short-term traded carbon values used for UK public policy appraisal." Department of Energy and Climate Change, London UK.
- DEFRA (2002) "Valuing the social cost of carbon emissions: Defra guidance." London: DEFRA.
- Farhi, E. and X. Gabaix (2015) "Optimal taxation with behavioral agents." Working paper.
- Frey, B. S. and F. Oberholzer-Gee (1997) "The cost of price incentives: An empirical analysis of motivation crowding," *American Economic Review*, 87, pp. 746 755.
- Gneezy, U., S. Meier, and P. Rey-Biel (2011) "When and why incentives (don't) work to modify behavior," *Journal of Economic Perspectives*, 25 (4), pp. 191 210.
- Gneezy, U. and A. Rustichini (2000a) "Pay enough or don't pay at all," *The Quarterly Journal of Economics*, 115 (3), pp. 791–810.
- (2000b) "A fine is a price," *Journal of Legal Studies*, 29 (1), pp. 1–18.
- Harsanyi, J. (1955) "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility," *Journal of Political Economy*, 63, pp. 309–321.
- Jacobsen, G., M. Kotchen, and M. Vandenbergh (2012) "The behavioral response to voluntary provision of an environmental public good: Evidence from residential electricity demand," *European Economic Review*, 56, pp. 946 960.
- Kahneman, D. and J. Knetsch (1992) "Valuing public goods: The purchase of moral satisfaction," *Journal of Environmental Economics and Management*, 22, pp. 57 70.
- Kotchen, M. (2005) "Impure public goods and the comparative statics of environmentally friendly consumption," *Journal of Environmental Economics and Management*, 49, pp. 281–300.
- (2006) "Green markets and private provision of public goods," *Journal of Political Economy*, 114, pp. 816–845.
- Lancaster, K. J. (1966) "A new approach to consumer theory," *Journal of Political Economy*, 74, pp. 132 157.
- Li, S., J. Linn, and p. . E. Muehlegger Am. Econ. J. Econ. Policy, 6 (4) (2014) (2014) "Gasoline taxes and consumer behavior," *American Economic Journal: Economic Policy*, 6 (4), pp. 302 342.
- Margolis, H. (1982) Selfishness, Altruism, and Rationality: University of Chicago Press, Chicago.
- McFadden, D. (1974) "Conditional logit analysis of qualitative choice behavior," in P. Zarembka ed. *Frontiers in econometrics*, pp. 105 142: Academic Press, New York.
- McFadden, D. and K. Train (2000) "Mixed mnl models of discrete response," *Journal of Applied Econometrics*, 15, pp. 447 470.
- Monin, B. and D. T. Miller (2001) "Moral credentials and the expression of prejudice," *Journal* of *Personality and Social Psychology*, 81, pp. 33 43.
- Mullainathan, S., J. Schwartzstein, and W. J. Congdon (2012) "A reduced-form approach to behavioral public finance," *Annual Review of Economics*, 4, pp. 511 540.
- Nyborg, K. (2000) "Homo economicus and homo politicus: Interpretation and aggregation of environmental values," *Journal of Economic Behavior and Organization*, 42, pp. 305 322.

- Perino, G., L. A. Panzone, and T. Swanson (2014) "Motivation crowding in real consumption decisions: Who is messing with my groceries?" *Economic Inquiry*, 52 (2), pp. 592 607.
- Pigou, A. (1920) The Economics of Welfare: London: Macmillan and Co.
- Revelt, D. and K. Train (1998) "Mixed logit with repeated choices: Households' choices of appliance efficiency level," *Review of Economics and Statistics*, 80, pp. 647–657.
- Rivers, N. and B. Schaufele (2015) "Salience of carbon taxes in the gasoline market," *Journal of Environmental Economics and Management*, 74, pp. 23 36.
- Schotter, A., A. Weiss, and I. Zapater (1996) "Fairness and survival in ultimatum and dictatorship games," *Journal of Economic Behavior and Organization*, 31, pp. 37 56.
- Teisl, M. F., B. Roe, and R. L. Hicks (2002) "Can eco-labels tune a market? Evidence from dolphin-safe labeling," *Journal of Environmental Economics and Management*, 43, pp. 339 359.