Research Paper 32 | **2014**

THE BEHAVIOURAL EFFECT OF PIGOVIAN REGULATION: EVIDENCE FROM A FIELD EXPERIMENT

B. LANZ, J.-D. WURLOD, L. PANZONE, AND T. SWANSON

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

Bruno Lanz† Jules-Daniel Wurlod‡ Luca Panzone§ Timothy Swanson¶

This version: April 2016

Abstract

Monetary incentives associated with Pigovian regulation may crowd-out intrinsic motivation to mitigate an externality. We quantify this behavioral trait using data from an experiment with real product choices together with a structural model of consumer behavior. First we show that information about products' embodied carbon induces a voluntary shift to cleaner products, and we estimate that intrinsic motivation is equivalent to a change in relative prices of GBP25-167/tCO_2. Second, the impact of a Pigovian intervention in proportion to embodied carbon (GBP19/tCO_2) suggests significant motivation crowding, and compensating this behavioral bias would require increasing the Pigovian price signal by GBP0.60-170/tCO_2. Finally, based on a cross-product comparison, we find that motivation crowding declines with the cost (or effort) of provision. This suggests that, on account of motivation crowding, environmental taxes ought to be set higher for products with price elastic demand.

Keywords: Motivation crowding; Consumer behavior; Pigovian regulation; Information; Field experiments.

JEL Codes: C91; D03; D12; L15; L50; Q58.

---

*We would like to thank Valentina Bosetti, Nick Hanley, Guy Meunier, Emi Uchida, as well as seminar participants at the 2014 WCERE, the 2014 Bioecon conference, and the 143rd 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. Funding from EU-POPP is gratefully acknowledged. Any remaining errors are ours.

†Corresponding author. Graduate Institute of International and Development Studies, Department of Economics and Centre for International Environmental Studies, Eugène Rigot 2, 1202 Geneva, Switzerland; ETH Zurich, Chair for Integrative Risk Management and Economics; Massachusetts Institute of Technology, Joint Program on the Science and Policy of Global Change. email: bruno.lanz@graduateinstitute.ch.

‡Graduate Institute of International and Development Studies, Department of Economics and Centre for International Environmental Studies, Switzerland.

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

¶Graduate Institute of International and Development Studies, Department of Economics and 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 interventions (e.g. Allcott et al., 2014; Farhi and Gabaix, 2015).[1] The objective of this paper is to provide quantitative evidence about the magnitude of such behavioral traits for the design of policies.

We consider two main types of policy interventions affecting internalization of externalities (here carbon emissions embodied in consumer products). The first supplies information about external effects associated with consumption decisions and relies on intrinsic motivation of consumers to behave prosocially. In principle, conditional on the existence of some form of preferences for the public good embedded in marketed products, the provision of information may induce a voluntarily transition towards cleaner alternatives (Harsanyi, 1955; Margolis, 1982; Kahneman and Knetsch, 1992; Nyborg, 2000).[2] The second intervention, a Pigovian price signal, modifies relative prices in proportion to the public good component of the products. By relating explicitly a change in relative prices to a regulatory intervention targeting the ranking of products in terms of their public good contents, substitution away from dirty alternatives may be reinforced relative to a basic change in relative prices. However, in instances where consumers are willing to voluntary internalize the externality (once given relevant information), explicitly associating monetary incentives to the decision to behave prosocially may crowd-out intrinsic motivation to do so (see Gneezy et al., 2011; Bowles and Polanía-Reyes, 2012, for a review). As a consequence, Pigovian interventions that change relative prices so as to reflect the socially optimal value of emissions may attenuate the effectiveness of the change in relative prices (e.g. Frey and Oberholzer-Gee, 1997; Bénabou and Tirole, 2003).

In the context of global public goods, the literature suggests two main mechanisms that could

[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] This is also related to the literature on altruism and the private provision public goods, which 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).

induce motivation crowding. The first relies on 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 monetary value of behavior through a tax/subsidy may induce consumers to update their beliefs. In cases where they over-estimate the value of external effects associated with their choices, so that the Pigovian price is lower than the subjective value they are willing to give up for their contribution, monetary incentives may reduce effort or provision. A second mechanism work through image motivation (Ariely et al., 2009). In particular, monetary rewards could affect self-image benefits through 'moral licensing' (Schotter et al., 1996; Jacobsen et al., 2012),[3] and thus the perceived ability to compensate external effects by a monetary payment may reduce incentives to behave prosocially.

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 empirical 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, inexpensive 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; Farhi and Gabaix, 2015), quantifying the behavioral impact of alternative externality-correcting interventions is important. This paper provides an attempt in that direction. We use data from a field experiment in which respondents make real consumption decisions about four commonly-purchased product categories (with clean and dirty alternatives in brackets):[4] 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). As described in details below, each product category offers an exhaustive number of options to consumers, with clean and dirty alternatives determined by their associated embodied carbon emissions. Emissions are tied to product characteristics such as packaging, fat content, or the type of raw material. Preferences for these private product characteristics will affect substitutability between clean and dirty alternatives alternatives, and thus the willingness (or effort) of supplying the associated

[3] 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.

[4] The experimental design has features of both a framed and a natural field experiment (Harrison and List, 2004), as it takes place in the field but subject are aware that they participate in an experiment. In particular, we carry out a computerized experiment replicating an online shopping experience, and do so within a supermarket.

public good.[5]

After a first baseline choice in one of the four product categories, subjects are randomly assigned to one of four treatments: (i) an information label showing carbon emissions embodied in different product alternatives; (ii) a Pigovian subsidy to the clean alternatives; (iii) a neutrally framed price reduction of the clean alternatives of the same magnitude as the Pigovian subsidy (i.e. the change in relative prices is justified by market reasons unrelated to environmental outcomes); and (iv) a neutrally framed removal of the dirty alternatives, giving participants the option to either choose one of the remaining clean alternatives or purchase nothing. Subjects are then allowed to reconsider their initial choice. On the one hand, observing two choices for each participant (within subject variations) permits us to separately identify inherent preferences for the characteristics of each product from the impact of the treatments. On the other hand, random allocation of subjects to different treatments after their initial choice (between subject variations) allows us to quantify how different interventions affect behavior.

A descriptive analysis of the same experiment by Perino et al. (2014) shows that all treatments increase the *aggregate* market share of the clean alternatives (i.e. for all products together). They further find that the change in clean products market shares is smaller for the Pigovian subsidy than for both the information treatment and the neutrally framed price change, and interpret this as evidence that the impact of a monetary incentive with information on relative emissions is sub-additive.[6] A first contribution of the present paper is to recognize that treatments applied to different products are in fact heterogeneous. The relative carbon content of dirty and clean alternatives differs across product categories, so that information about embodied emissions and the level of the Pigovian subsidy differs. Moreover, as carbon emissions are tied to product characteristics, substitutability between clean and dirty alternative differs across products. Concretely, substituting among cola products in cans or in plastic bottles is not

[5] One interesting feature of carbon emissions (and associated climate change) as a global public good is that personal contributions have negligible direct private benefits, so that choices solely reflect prosocial motivations. By contrast, in other settings 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.

[6] 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 further below, the market share impacts we measure are similar in magnitude to those estimated using actual transactions (e.g. Teisl et al., 2002; Bjorner et al., 2004).

the same as choosing from beef or chicken products in the meat category.

In order to account for observed product characteristics and control for the size of the treatment across products, we treat each alternative in the choice set as a differentiated product and estimate the structural parameters of the underlying consumer decision problem, using Lancaster's (1966) multi-attribute utility theory and McFadden's (1974) random utility model. Given the utility-maximizing framework, preferences over observed product characteristics are identified from the baseline product choice, which allows us to estimate conditional average treatment effects (CATE) as structural parameters of the decision problem controlling for these preferences.

As a first benefit of this framework, we exploit responses to the two neutrally-framed treatments to quantify substitutability of clean and dirty alternatives in each product category. This enables us to assess the effectiveness of alternative interventions as a function of product substitutability. But the second and more interesting contribution afforded by our approach is that we can obtain money-metric welfare measures for both the information label and motivation crowding. 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 behavioral impact as providing information about carbon emissions. In our setting the EPM for information provides an estimate of consumers' valuation of a voluntary reduction of carbon emissions, and can thus be compared against estimates of the social cost of carbon (which determines the level of the Pigovian price signal). Similarly, we obtain an EPM quantifying the extent of motivation crowding by comparing the behavioral response to the change in relative prices when it is framed as a Pigovian intervention and when it is framed neutrally. This provides direct evidence on how the regulatory dimension of a Pigovian price signal affects behavior, and how much Pigovian prices ought to be augmented in order to overcome the negative behavioral effect associated with monetary incentives and motivation crowding.

Our key results are as follows. First, using evidence from the two neutrally framed treatments, we show that substitutability between clean and dirty alternatives varies substantially across product categories, and that policy interventions are more effective for products with high substitutability. This finding is intuitive, and it is already recognized in the literature (e.g. Bjorner et al., 2004), although the ability to assess, in a controlled environment, how substitutability affects the impact of price and non-price interventions is novel. Second, we estimate that the EPM associated with information ranges from GBP0.02 to GBP1.66 depending on the product category, which corresponds to GBP25-167/tCO_2. An implication is that consumers's valuation of emissions is significantly above most estimates of the social cost of carbon (we employ a value of GBP19/tCO_2 to compute the Pigovian subsidy). Third, we find evidence of motivation crowding, and observe that its extent is increasing with the effort to behave prosocially: for products with close substitutes (cola and milk in our setting) we observe very substantial motivation crowding, while for products where substitution requires more effort (spreads and meat) it is significantly smaller. Quantitatively, in order to compensate the negative effect of a monetary price signal, the Pigovian subsidy rate would need to be augmented by approximately GBP0.60/tCO_2 for meat products, GBP5/tCO_2 for spreads, GBP54/tCO_2 for cola products, and GBP170/tCO_2 for milk.

Aside from quantifying the extent of motivation crowding, our results also contribute to the growing empirical literature quantifying the effect of information about the public good content of products, and several studies have shown that voluntary contributions triggered by information can have a significant effect on market outcomes. For example, Teisl et al. (2002) uses 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 market shares ranging from 10 to 30 percent, is similar to the studies using market 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 tax-related prices changes, as compared to less salient variation in market prices (see Davis and Kilian, 2011; Li et al., 2014; Rivers and Schaufele, 2015, on the demand for gasoline products). For example, Rivers and Schaufele (2015) 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, an effect attributed to the salience of the policy intervention. By contrast, in our experiment salience of Pigovian and neutrally-framed price changes are held constant, and our contribution is to quantify the effect of information included in Pigovian prices, controlling for the magnitude of the price change and 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 four treatment interventions. In Section 3, we present our empirical strategy, including identification of the EPM for information and motivation crowding. 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.[7] Consumers entering the supermarket are offered to participate voluntarily in a "university-sponsored grocery shopping study." 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 environmentally friendly respondents. Respondents have to complete the tasks independently on a computer at the entrance of the supermarket, closely replicating an online shopping environment.

The main task consists of two choices. First, subjects make an initial purchasing decisions, and participants in the experiment are selected by identifying those who intend to purchase a product in one of four categories: cola-type sodas, milk, spreads (margarine and butter) and meat (chicken and beef).[8] Each product category includes a fixed and finite number of alternatives, reported in Table 1. Participants are offered a GBP5 voucher provided that they actually purchase the goods they select during the experiment.[9] The enforcement by making payment conditional on the actual purchase of goods selected is meant to induce truthful preference rev-

[7] 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.

[8] Potential participants were turned down if younger than 21 years old 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, Appendix A shows that our sample includes diverse socio-economic background.

[9] Note also that 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 in-store price.

Table 1: Product categories and clean / dirty alternatives

Products	Clean alternatives			Dirty alternatives		
	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 (2l)	1.69	0.50	Coca Cola Zero in ALU cans (2l)	2.85	1.02
	Pepsi Regular in PET bottle (2l)	1.00-1.69	0.50	Pepsi Regular in ALU cans (2l)	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

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.

elation. The compliance rate is 96 percent, and non-compliers are dropped from the sample.[10]

Participants are then randomly assigned to one of four treatments (described in more details in the following subsections): (i) an information label showing embodied carbon emissions associated with each alternatives; (ii) a Pigovian subsidy on the clean (low-emissions) alternatives; (iii) a neutrally framed price reduction of the clean alternatives of the same amount as the subsidy; and (iv) a neutrally framed removal of the dirty alternatives. After being subject to one of the treatments, respondents are allowed to revise their initial choice.

In each product category, we classify alternatives 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 category. For cola products low-emissions alternatives are sold in 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 private attributes related to emissions (such as plastic vs. aluminum packaging or the type of meat) will be a key determinant of substitutability, and introspection suggests that it should vary across products. For example, if the type of meat

[10] 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.

matters for consumers, they will be more resistant to substitute away from beef alternatives.

After the second choice, socio-demographic data on the respondents are collected. A total of 993 shoppers 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 treatment assignment, showing that subsamples are balanced across randomly assigned treatments (t-tests of difference in means being insignificant for all but one variable, gender).

2.1 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.[11]

2.2 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 a 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/tCO_2 taken from DEFRA (2002), the subsidy is calculated by using the differ-

[11] 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 choice.

ence in embodied CO_2 emissions between clean and dirty alternatives.[12] The final values of the subsidies are: GBP0.05 for cola products in aluminum cans; GBP0.03 for semi-skimmed milk, or GBP0.06 for skimmed milk; GBP0.43 for margarine; and GBP0.21 per kg of chicken.

2.3 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 providing explicit information about the public good dimension of each alternative. Moreover, this treatment has several advantages. First, it yields an internally consistent estimate of price responsiveness, capturing the willingness to substitute between clean and dirty alternatives. Second, this allows calculating monetary equivalents for differences in effectiveness across interventions.

2.4 Neutrally framed removal of dirty alternatives

In this treatment all dirty alternatives are removed from the choice set, leaving consumers to choose between the remaining clean alternatives or not purchasing anything at all. The neutral product removal is introduced with the following statement (in the case of cola products): "There has been a change in product availability. Products are not supplied in cans on account of the lack of availability of the necessary materials."

This treatment provides a different measure of the substitutability between clean and dirty alternatives, as it also captures the 'essentiality' of the product category. More specifically, if a given product category is not important in the consumption basket, consumers will be more

[12] There is a lot of uncertainty regarding the estimation of a social cost of carbon, and was 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 are in fact slightly lower. For 2014, estimates from DECC (2013) suggest an upper bound of GBP16.73/tCO_2.

likely to opt out of the market rather than select one of the less preferred clean alternatives.[13]

3 Estimation Strategy

In the two sequential choice occasions, consumers select one product alternative from a finite set of discrete 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), relative prices (Pigovian subsidy and neutral price change treatments), or product availability (product removal treatment). 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 of information and motivation crowding.

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:

$$U_n^{\text{dirty}} > U_n^{\text{clean}} \Leftrightarrow u_n > p\,. \tag{1}$$

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

[13] Note that the removal treatment is, in a sense, hypothetical in that subjects may still purchase one of the dirty alternatives after completing the study (e.g. in a different supermarket). Nevertheless, it provides an interesting complement to the neutral price change treatment to identify substitutability.

will dominate the difference in utility derived from consuming the two goods. Identification of the 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:

$$\beta_n^{\text{price}} \Delta p > u_n - p \,. \tag{2}$$

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.

A second measure of substitutability between clean and dirty alternatives is afforded by the exogenous removal treatment. In this treatment, we observe how consumers behave when the (preferred) dirty version is removed from the choice set. Denoting the utility of the outside option $U_n^{\text{outside}} = \beta_n^{\text{outside}}$, a consumer initially selecting the dirty alternative and assigned to the removal treatment will select the outside option when $U_n^{\text{dirty}} > U_n^{\text{outside}} > U_n^{\text{clean}}$. Thus when a consumer selects one of the clean alternatives, he reveals that the difference in utility between clean and dirty alternatives is larger than that of not buying anything at all. Furthermore, for product categories in which clean options are perceived to be relatively good substitutes (as measured by β_n^{price}), a relatively large increase in the market share of the clean product would also capture the importance or 'essential nature' of the product category.

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 a 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 small, and information about the public good component may significantly increase the market share 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 the clean alternative provided that:

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

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 aside from the price signal. In particular, β_n^{regul} captures the effect of framing the intervention as an explicit government intervention related to the public good content of the products.

3.2 Structural estimation: Multinomial choice

For each category of product, there is a finite set 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. 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 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:

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

Following McFadden (1974), we decompose the utility of product j into a deterministic part observed by the researcher, denoted by V_n^j, and a unobserved part denoted ε_n^j, so that: $U_n^j =$

[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.

Table 2: Choice set, product attributes and policy treatments

	Cola	Milk	Spread	Meat
Nr. of alternatives[a]	13	4	11	7
Product attributes				
Attribute 1[b]	ALU cans (=1)	Fat cont. (g)	Butter (=1)	Beef (=1)
Attribute 2	Price (GBP cent)	–	Price (GBP cent)	Price (GBP cent)
Attribute 3	Coca-Cola brand (=1)	–	Lurpak brand (=1)	Protein (g)
Attribute 4	Light (=1)	–	Sainsbury brand (=1)	Salt (g)
Attribute 5	Zero/Max (=1)	–	Anchor brand (=1)	Fat (g)
Attribute 6	–	–	Proteins (g)	Weight (g)
Attribute 7	–	–	Fat (g)	–
Attribute 8	–	–	Salt (g)	–
Policy treatments				
Information label	Difference in embodied carbon emissions between clean and dirty alternatives ($kgCO_2$)			
Pigovian subsidy	Pigovian subsidy to the price of the clean options (GBP cent)			
Neutral price change	Neutrally framed decrease in price of clean options (GBP cent)			
Product removal	Removal of the dirty options (=1)			

Notes: Table lays out the data structure underlying estimation of a discrete choice model. [a]Each product category includes the number of product alternatives reported in Table 1 plus the outside (or opt-out) option. [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.

$V_n^j + \varepsilon_n^j$. Given the notation developed above, we specify the observed part of utility as:

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

where Z^j is a vector of observed product attributes as defined in Table 2, $I_{\text{info}}^{\text{clean}}$, $I_{\text{pigou}}^{\text{clean}}$ and $I_{\text{price}}^{\text{clean}}$ 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, $I_{\text{remove}}^{\text{dirty}}$ is an indicator variable equal to one if j is a dirty alternative removed from the choice set, and γ, β are parameters to be estimated from the data. In addition, the utility of the outside option is specified as $V_n^{\text{outside}} = I_{\text{outside}} \cdot \beta_n^{\text{outside}}$, where β_n^{outside} measures preferences for the outside option.

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 convenient logit form:

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

and the log-likelihood function writes:

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

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 assumes that observations are independent and implies restrictive substitution patterns, the so-called independence of irrelevant alternatives (IIA) property. To account for the panel structure of the data and allow choices made by each respondent to be correlated, 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 unobserved 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 500 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 market share of clean alternatives.[16]

[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 500 replications. Although this number is relatively small, the ensuing inference yields similar conclusions to the closed-form results drawn from MNL specifications, which gives us confidence that this is appropriate.

3.3 Equivalent price metric for information and motivation crowding

Following Allcott and Taubinsky (2015), we define the EPM as follows:

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

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.

Given the notation developed above, the EPM of information is equal to the ratio between the utility weight associated with the information treatment and that of the neutral price change treatment:

$$\text{EPM}^{\text{info}} = \frac{\beta^{\text{info}}}{\beta^{\text{price}}}. \tag{10}$$

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.

As initially put forward by Gneezy and Rustichini (2000b) in a different context, if consumers' valuation of emissions is higher than the Pigovian price signal, a Pigovian intervention may crowd out intrinsic motivation to switch towards one of the cleaner alternatives. An implication is that, in the presence of motivation crowding, the Pigovian price would need to be set higher to compensate this behavioral trait (Allcott et al., 2014; Farhi and Gabaix, 2015). Given the notation developed above, we have that $\beta^{\text{regul}} = \beta^{\text{pigou}}\Delta s - \beta^{\text{price}}\Delta p$ and the EPM reflecting the behavioral component of the Pigovian regulation is given by:

$$\text{EPM}^{\text{regul}} = \frac{\beta^{\text{regul}}}{\beta^{\text{price}}}. \tag{11}$$

In words, $\text{EPM}^{\text{regul}}$ reflects the behavioral effect (in monetary terms) of information provided by a Pigovian intervention over and above the change in relative prices, and thus quantifies the extent of motivation crowding associated with the regulation. In turn, this provides direct evidence about the change in relative prices that would compensate motivation crowding.

Table 3 summarizes the statistics used to quantify the behavioral effect of information and

Table 3: Definition of equivalent price metric statistics

Statistic	Definition	Units
$\frac{\Delta e \beta^{\text{info}}}{\beta^{\text{price}}}$	Equivalent price metric of the information label	GBP cent
$\frac{\beta^{\text{info}}}{\beta^{\text{price}}}$	Equivalent price metric of the information label per unit of emissions	GBP per tCO_2
$\frac{\beta^{\text{regul}}}{\beta^{\text{price}}}$	Equivalent price metric of motivation crowding	GBP cent
$\frac{\beta^{\text{regul}}}{\beta^{\text{price}}} \frac{1}{\Delta e}$	Equivalent price metric of motivation crowding per unit of emissions	GBP per tCO_2

motivation crowding. Note that these measures are free of the scale parameter and are thus directly comparable across models. However, to compare EPM estimates for different product categories, we need to control for the fact that products differ with respect to the level of embodied emissions (Δe). This is achieved by estimating an EPM per unit of emissions.

4 Data and Results

4.1 Descriptive statistics: Market share before and after treatments

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

Descriptive statistics further suggest that all the treatments induced significant increases in the market share of clean alternatives. There is, however, ample variations 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 motivation crowding has a detrimental impact on the effectiveness of the regulatory intervention. These observations and their statistical significance (in terms of changes in market shares) are discussed in details by Perino et al. (2014), although they aggregate all product categories together.

Table 4: Observed market share of clean alternatives by product category (percentage)

		Information label	Pigovian subsidy	Neutral price change	Product removal
Cola	Before treatment	47.6	31.3	27.0	46.1
	After treatment	66.7	50.0	69.8	86.8
	Difference	19.1	18.7	42.8	40.7
Milk	Before treatment	12.3	6.1	8.3	13.7
	After treatment	19.3	10.9	14.3	54.3
	Difference	7.0	4.8	6.0	40.6
Spread	Before treatment	55.2	56.3	39.0	51.7
	After treatment	82.8	63.4	51.2	93.1
	Difference	27.6	7.1	12.2	41.4
Meat	Before treatment	12.5	20.6	21.7	25.0
	After treatment	32.1	30.2	33.3	78.6
	Difference	19.6	9.6	11.6	53.6

Notes: For the product removal treatment respondents who did not choose one of the remaining clean alternatives exited the market by choosing not to purchase anything.

Comparing the impact of treatments across product categories, the largest percentage change in market share is generally observed for cola products (substitution towards products in plastic packaging). However both the size of the treatment and the initial market share 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 information and motivation crowding.

Estimation results from the 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, estimates have a positive sign and are highly statistically significant, which is consistent with the relatively large initial market shares for dirty alternatives reported in Table 4. Other variables capturing preferences for product attributes are mostly statistically significant

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

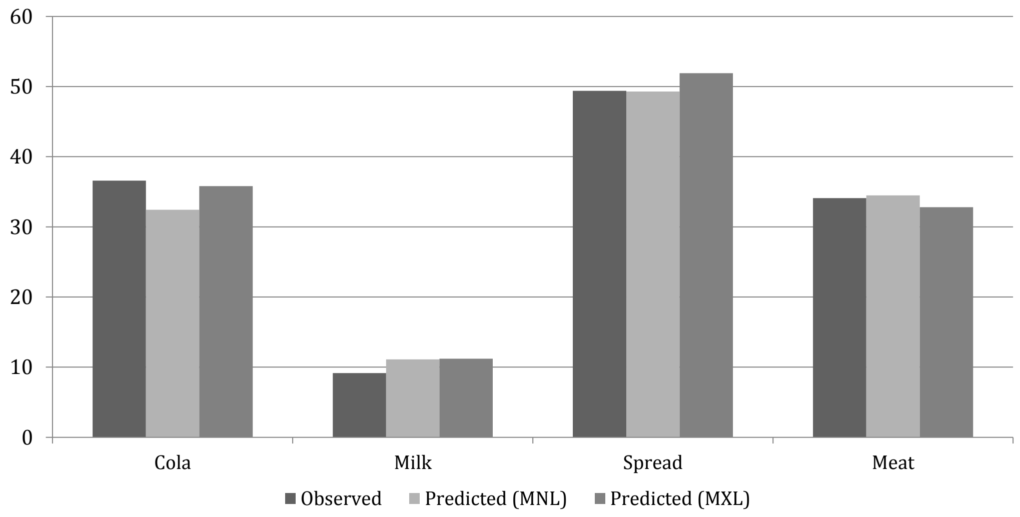

Notes: Figure 1 plots the observed market share of clean alternatives before the treatments against that predicted by MNL and MXL models reported in Table 5.

at conventional levels, suggesting that the structural model provides a good account of observed choices. This is confirmed by comparing simulated market shares of clean alternatives with the actual market shares 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 model.

The main estimates of interest are those associated with the four treatments, as they quantify the impact of each treatment on market shares.[17] We find that the estimated utility weights associated with the treatments are mostly highly statistically significant, confirming that the treatments affect choice probabilities, and all have the expected sign. The information label, Pigovian subsidy and neutral price change increase the market share of the clean alternatives, while the product removal treatment is negative and very large, reflecting the fact that under that treatment the market share of the dirty alternatives drops to zero.

[17] As Table 1 reports, the information treatment is coded as the difference (in kg of CO_2) 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. Finally, the removal treatment is coded as a categorical variable, and the utility of the outside option is captured by an additional indicator variable ('remove outside').

Table 5: Estimation of product-specific multinomial choice models

	Cola			Milk			Spread			Meat		
	MNL (1)	MXL (2)		MNL (3)	MXL (4)		MNL (5)	MXL (6)		MNL (7)	MXL (8)	
		Mean	Std-dev.		Mean	Std-dev.		Mean	Std-dev.		Mean	Std-dev.
Information label	2.74***	17.12*	–	3.93***	3.96***	–	0.14***	1.75***	–	0.02	5.43***	–
	(0.87)	(9.76)		(0.61)	(0.62)		(0.04)	(0.48)		(0.05)	(1.92)	
Pigovian subsidy	0.15***	1.24**	–	0.17***	0.17***	–	0.01***	0.10**	–	0.02*	2.04***	–
	(0.05)	(0.63)		(0.02)	(0.02)		(0.01)	(0.05)		(0.01)	(0.67)	
Neutral price change	0.31***	2.82**	–	0.24***	0.24***	–	0.002	0.12***	–	0.03***	2.11***	–
	(0.06)	(1.38)		(0.02)	(0.02)		(0.005)	(0.04)		(0.01)	(0.69)	
Product removal	-16.81***	-64.76**	–	-20.34***	-24.42***	–	-17.19***	-90.45***	–	-18.89***	-573.64***	–
	(0.17)	(26.10)		(0.11)	(0.17)		(0.16)	(22.07)		(0.21)	(189.72)	
Attribute 1	0.67**	4.15*	14.24**	0.23***	0.23***	0.04	0.26	-0.38	19.23***	1.96***	116.98***	132.67***
	(0.29)	(2.45)	(7.05)	(0.01)	(0.01)	(0.07)	(0.24)	(1.93)	(5.63)	(0.38)	(39.46)	(45.09)
Attribute 2	0.0001	0.002	0.001	–	–	–	-0.01***	-0.03***	0.03**	-0.0001	-0.0001	0.04***
	(0.002)	(0.003)	(0.001)				(0.003)	(0.01)	(0.01)	(0.003)	(0.01)	(0.01)
Attribute 3	1.63***	3.12***	3.05***	–	–	–	1.83***	-0.67	5.80***	0.005**	0.02***	0.01***
	(0.18)	(0.45)	(0.57)				(0.35)	(2.38)	(2.09)	(0.002)	(0.01)	(0.003)
Attribute 4	-0.60***	-3.14***	6.43***	–	–	–	-1.32***	-5.37**	4.63**	0.19**	0.73***	0.25**
	(0.14)	(0.96)	(1.25)				(0.25)	(2.54)	(1.96)	(0.07)	(0.21)	(0.10)
Attribute 5	-1.73***	-9.76***	8.28***	–	–	–	0.33	-3.58	8.62***	-13.52***	-67.61***	18.22**
	(0.22)	(2.95)	(2.10)				(0.23)	(3.05)	(3.30)	(3.98)	(20.02)	(7.46)
Attribute 6	–	–	–	–	–	–	-1.99***	-2.87	11.75***	0.25***	1.02***	0.52***
							(0.39)	(2.08)	(2.70)	(0.08)	(0.28)	(0.14)
Attribute 7	–	–	–	–	–	–	0.07***	0.22***	0.01	–	–	–
							(0.01)	(0.07)	(0.02)			
Attribute 8	–	–	–	–	–	–	1.03***	2.85***	3.18***	–	–	–
							(0.16)	(0.76)	(1.12)			
Outside option	0.49	1.22**	–	-0.13	-0.13	–	2.58***	6.90***	–	6.43**	14.92***	–
	(0.40)	(0.54)		(0.15)	(0.15)		(0.72)	(2.35)	–	(3.23)	(4.70)	
Respondents	224	224		547	547		269	269		208	208	
Log Pseudo-LL	-887.2	-733.0		-1220.2	-1220.1		-1101.3	-866.5		-617.8	-518.5	
AIC	1794.3	1496.0		2452.3	2454.2		2228.5	1775.0		1257.6	1071.0	
BIC	1861.4	1547.2		2478.1	2484.3		2275.2	1850.5		1294.3	1127.7	
Pseudo R^2	0.228	0.362		0.196	0.196		0.146	0.328		0.237	0.359	

Notes: This table reports preference parameters estimates for MNL and MXL specifications. For MXL models we estimate both the mean and standard-deviation (std-dev.) of normally distributed preference parameters. Sample size for each product category is twice the number of respondent (every respondent makes two choices). Standard errors are clustered at the respondent level and reported in parenthesis. ***, **, *: statistically significant at 1, 5 and 10 percent respectively. The list of attributes for each product can be found in Table 2.

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 to quantify motivation crowding.

4.2.1 Measures of substitutability for each product category

Evidence about substitutability between clean and dirty alternatives is based on two different measures, namely the responsiveness to a neutral price change and the impact of the product removal treatment. Results derived from the MXL models are reported in Table 6. In Panel A, we report changes in simulated market shares for the clean alternatives associated with the neutral price change treatment. The impact of the neutral price change specified in the experiment (i.e. the CATE) refers to different price changes for each product. In addition, 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. In Panel B, we report simulated changes in the market share of the clean alternative in response to the removal of dirty alternatives.

Results indicate that the CATE is positive and statistically significant for all the products, with a 40 percent increase in the market share 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 market shares 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 confirmed if we consider how the market share of clean products relative to a situation without the treatment. Specifically, for a unit change in relative prices, the market share of clean cola products increases from around 35% to 45%, or a 30% increases. The corresponding increase is 20% for milk, 1.5% for meat and 0.5% for spread.

Turning to the product removal treatment, we find that changes in market shares are very similar for all products we consider. However, since the clean alternative of milk and meat

Table 6: Measures of substitutability between clean and dirty alternatives

	Cola	Milk	Spread	Meat
Panel A: Impact of the neutral price change on the market share of clean alternatives[a]				
Neutral price change: CATE	39.48***	11.05***	10.83***	8.77**
	(1.47)	(0.59)	(3.25)	(3.75)
10% neutral price change	64.11***	14.82***	4.29***	11.14**
	(3.41)	(2.56)	(1.45)	(4.82)
Neutral price change (GBP cent)	9.77***	1.79***	0.26***	0.51***
	(0.59)	(0.13)	(0.10)	(0.18)
Panel B: Impact of product removal on the market share of clean alternatives[b]				
Product removal: CATE	55.83***	52.65***	54.51*	52.24***
	(2.64)	(2.56)	(1.32)	(4.89)
Respondents	224	547	266	208

Notes: Panel A 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 market shares corresponding to the conditional average treatment effect (CATE), a 10% price reduction, and a 1 GBP cent price reduction. Panel B shows the marginal impact of the product removal treatment on the simulated market share of clean alternatives, reported in percentage points difference. Simulated market shares 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.

products have a relatively small initial market share, the proportional increase is much larger: following the product removal treatment, the market share of the clean milk alternatives increase by a factor of four, and three for clean meat alternatives. This suggests that milk and meat products are more essential to the consumers, which is in line with expectations.

4.2.2 Policy instruments: Comparison across products categories

We now study the impact of information and Pigovian treatments and results derived from the MXL models are reported in Table 7. Panel A shows the CATE measured by the change in the market share 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 market shares 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

Table 7: Effectiveness of policy instruments within and across products

	Cola	Milk	Spread	Meat
Panel A: Within product comparison				
Information label: CATE	26.72***	12.11***	33.17***	12.88*
	(1.41)	(1.05)	(8.26)	(7.66)
Pigovian subsidy: CATE	19.47***	8.10***	9.26**	8.51**
	(0.88)	(0.74)	(3.61)	(3.44)
Panel B: Across product comparison				
10% information label	6.47***	5.58***	4.67**	1.91*
	(0.65)	(0.63)	(1.80)	(1.07)
Information label (kgCO_2)	45.64***	20.14***	3.96**	1.36**
	(1.99)	(0.52)	(1.53)	(0.69)
10% Pigovian subsidy	51.51***	11.38***	3.74**	10.49**
	(1.79)	(0.85)	(1.52)	(4.21)
Pigovian subsidy (GBP cent)	4.76***	1.28***	0.23**	0.45***
	(0.17)	(0.14)	(0.10)	(0.16)
Respondents	224	547	266	208

Notes: Panel A displays the marginal impact of the information label and Pigovian subsidy treatments on simulated choice probabilities (or market shares) of clean alternatives, reported in percentage points differences. We report changes in simulated market shares corresponding to the conditional average treatment effect (CATE), and allow for a within product comparison. 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 kgCO_2 and GBP cent for the information label and Pigovian subsidy respectively). Normalized treatments allow for comparisons across product categories. Simulated market shares 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.

difference measured in kg of CO_2. 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 33 percent for the information label, and from 8 to 20 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).

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 clean market share of cola and milk products is most responsive, while spread and meat are far less sensitive. Similarly, cola products are very responsive to Pigovian subsidies, whereas meat and in particular spread products are not. The ranking of products is 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 and dirty alternatives.

4.2.3 Equivalent price metric of information and motivation crowding

This section concludes the comparison of regulatory interventions by reporting EPM estimates for information and motivation crowding, 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 tCO_2). The latter estimates provide a basis for a comparison of results across products. Panel B reports the EPM associated with motivation crowding, both for the CATE and per unit of emissions.

For the information treatment, the EPM of the CATE ranges from GBP0.02 for milk (not statistically significantly different from zero) to GBP1.66 for spread products. Recall that for milk products the impact of information is relatively small (Table 7, Panel A) while the neutrally framed price change has a significant impact on choice probabilities (Table 6, panel A). By contrast, spread is most responsive to the information treatment but displays the lowest responsiveness to a neutral price change.

However, once we control variations in products' embodied emissions, EPM estimates are more similar across product categories, ranging from around GBP26/tCO_2 for meat products to around GBP167/tCO_2 for milk. Because $\beta^{\text{info}}/\beta^{\text{price}}$ can be interpreted as a marginal rate of substitution between embodied carbon emissions and money, it also provides an estimate for

Table 8: Equivalent price metric for information and motivation crowding

	Cola	Milk	Spread	Meat
Panel A: Equivalent price metric for information				
$\Delta e \beta^{\text{info}} / \beta^{\text{price}}$ (GBP cent)	3.16***	2.41	166.39***	28.29***
	(0.88)	(1.63)	(43.95)	(1.97)
$\beta^{\text{info}} / \beta^{\text{price}}$ (GBP per tCO_2)	60.69***	166.63***	148.23***	25.72***
	(17.00)	(31.45)	(39.15)	(1.79)
Panel B: Equivalent price metric associated with motivation crowding				
$\beta^{\text{regul}} / \beta^{\text{price}}$ (GBP cent)	-2.80***	-6.79***	-5.75	-0.68
	(0.56)	(2.06)	(16.00)	(1.49)
$\beta^{\text{regul}} / \beta^{\text{price}} \frac{1}{\Delta e}$ (GBP per tCO_2)	-53.76***	-169.68***	-5.12	-0.62
	(10.73)	(30.97)	(14.25)	(1.35)
Respondents	224	547	266	208

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 tCO_2 (allowing a comparison across product categories). Panel B displays the EPM for motivation crowding measured in GBP cent (referring to the CATE) and in GBP per tCO_2 (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.

consumers' monetary valuation of embodied carbon emissions. For all products, these numbers are significantly larger than the Pigovian price used in the experiment (GBP19/tCO_2, see DEFRA, 2002) and most other estimates of the social cost of carbon.

Reported in Panel B, the EPM associated with motivating crowding is negative, reflecting the fact that for all products the Pigovian subsidy treatment is less effective than a neutrally framed price change. The EPM for motivation crowding is, however, only statistically significant for cola and milk products.

When measured per tCO_2, the EPM associated with motivation crowding 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. In particular, we find that the Pigovian tax rate should be increased by around GBP0.6/tCO_2 for meat products and up to GBP170/tCO_2 for milk products, in order to compensate for motivation crowding.

An other important conclusion from the exercise is that motivation crowding is 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. In particular, motivation crowding is much more significant 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 that display intrinsic motivation to behave prosocially, framing a change in relative prices as an explicit intervention to encourage or reward the provision of a global public good may backfire. 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 this issue. In our experiment, we found that consumers responded to information by behaving prosocially, and that the implied value of carbon emissions was significantly above most estimates of Pigovian tax rates. Moreover, monetary interventions explicitly motivated by the internalization of carbon emissions were less effective that neutrally framed changes in relative prices. An implication is that the price signal of Pigovian regulation would need to be set above its socially efficient level in order to compensate the negative behavioral effect associated with motivation crowding.

In practice however our results open the door to further research on at least two fronts. First, we find evidence that motivation crowding is related to the effort of behaving prosocially, and the monetary equivalents of motivation crowding are significantly larger for products with close substitutes available. This suggests that environmental taxes ought to be set higher than the Pigovian level for products with price elastic demand. Second, 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 motivation crowding 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 motivation crowding is 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 motivation crowding thus appears to be an important research avenue.

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 anough 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.

Harrison, G. and J. List (2004) “Field experiments,” *Journal of Economic Literature*, 42, pp. 1009–1055.

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.

Appendix A Sample composition

Table A1: Demographic variables by product category

	Mean	Std. Dev.	Min	Max
Cola subsample (N=346)				
Male indicator	0.44	0.50	0	1
Age (in years)	33.62	11.61	18	72
Education[a]	1.77	0.77	1	3
Income[b]	3.92	2.76	1	9
Children in the household[c]	0.63	1.04	0	6
Non-white indicator	0.45	0.50	0	1
Milk subsample (N=825)				
Male indicator	0.36	0.48	0	1
Age (in years)	37.07	12.03	18	80
Education[a]	1.80	0.73	1	3
Income[b]	4.02	2.79	1	9
Children in the household[c]	0.63	1.00	0	6
Non-white indicator	0.36	0.48	0	1
Spread subsample (N=431)				
Male indicator	0.34	0.47	0	1
Age (in years)	38.26	12.24	18	79
Education[a]	1.79	0.75	1	3
Income[b]	3.81	2.68	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=322)				
Male indicator	0.37	0.48	0	1
Age (in years)	38.39	12.22	18	79
Education[a]	1.75	0.73	1	3
Income[b]	3.90	2.64	1	9
Children in the household[c]	0.54	1.01	0	6
Non-white indicator	0.27	0.45	0	1

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

Table B1: Differences in means across treatments

	Information label	Pigovian subsidy		Price change		Product removal	
	Mean	Mean	Diff.	Mean	Diff.	Mean	Diff.
Male indicator	0.35 (0.48)	0.40 (0.49)	-0.05 (0.04)	0.30 (0.46)	0.05 (0.04)	0.42 (0.49)	-0.07* (0.04)
Age (in years)	36.92 (11.97)	37.54 (12.42)	-0.62 (1.10)	37.21 (12.17)	-0.30 (1.08)	35.72 (12.40)	1.20 (1.08)
Education[a]	1.78 (0.75)	1.79 (0.74)	-0.01 (0.07)	1.83 (0.05)	-0.05 (0.07)	1.75 (0.73)	0.03 (0.07)
Income[b]	30.61 (20.77)	32.05 (20.81)	-1.45 (2.04)	32.45 (20.53)	-1.84 (2.03)	29.54 (19.92)	1.07 (1.96)
Children in household[c]	0.58 (1.01)	0.57 (0.96)	0.01 (0.09)	0.56 (0.91)	0.02 (0.09)	0.65 (0.97)	-0.07 (0.09)
Non-white indicator	0.39 (0.49)	0.37 (0.48)	0.02 (0.04)	0.34 (0.48)	0.04 (0.04)	0.38 (0.49)	0.001 (0.04)

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. [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.