Understanding Markets with Socially Responsible Consumers*

Marc Kaufmann
Central European University

Botond Kőszegi
briq

March 10, 2023

Abstract

Many consumers care about climate change and other broad externalities. We model and analyze the market behavior of such "socially responsible consumers," derive properties of the resulting competitive equilibria, and study the effectiveness of different policies. In violation of price taking, a vanishingly small consumer cares about her impact on the behavior of the rest of the market to a non-vanishing extent. That impact on others endogenously dampens the consumer's direct effect on the externality, undermining responsible behavior. Dampening implies that even if all consumers value the externality like the social planner, they mitigate too little in any equilibrium, and may coordinate on the worst of multiple equilibria. To motivate consumers to lower the externality in a closed economy, a unit tax is superior to a cap-and-trade system, but there are policies that are better than a tax. Furthermore, under trade with a large or very polluting partner, a cap is better than a tax. When there are two products that are perfect substitutes in consumption but generate different externalities, there is always an equilibrium in which the products have the same price and consumers are indifferent between them. Under conditions we identify, this selfish equilibrium is the unique equilibrium. In a selfish equilibrium, a cap and a unit tax on the dirty product can achieve the same outcomes. In non-selfish equilibria, a proportional subsidy on the clean product dominates both a unit tax and a cap.

*Wenjun Zheng provided fantastic research assistance. We thank Peter Andre, Simon Gächter, Hendrik Hakenes, Mats Köster, Mechthild von Knobelsdorff, Stephan Lauermann, Tímea Molnár, Christoph Semken, Philipp Strack, and Paul Voss for insightful discussions. Kőszegi thanks the European Research Council for financial support under Grant #788918.

1 Introduction

When investigating the implications of externalities from market trade, economists overwhelmingly assume that individuals are selfish, maximizing their private consumption utilities. But many consumers care about the consequences of their marketplace behavior: they are willing to reduce or modify their consumption to alleviate the associated climate change (Andre et al., 2022, Imai et al., 2022) and animal and human suffering (Auger et al., 2003, Hertel et al., 2009, Hainmueller et al., 2015). The goal of this paper is to investigate equilibrium behavior, welfare, and the effectiveness of interventions in a classical market when such “socially responsible consumers” are present. We focus in particular on rational consumers.

Our starting point is an observation that departs from previous research investigating social preferences in large product markets. We demonstrate that price taking is violated: a socially responsible consumer who is a vanishingly small part of the market cares about her price-mediated impact on the behavior of the rest of the market to a non-vanishing extent. Specifically, if she consumes more, the resulting infinitesimal increase in the price induces others to consume less, dampening her impact on the externality.

We show that due to dampening, markets are typically bad at coordinating the behavior of socially responsible consumers. Even if all consumers “internalize the externality” — they put the same weight on it as the social planner — they consume too much in any competitive equilibrium, and may coordinate on the worst of multiple equilibria. Furthermore, consumers may in equilibrium be indifferent between two products that are perfect substitutes in consumption but generate different externalities. Turning to policies to lower the externality, we establish that a unit tax is better than a cap in a closed economy, but under trade with a large or very polluting partner, a cap is better. We also identify policies that are better than both a tax and a cap.

We begin in Section 2 by analyzing how a single socially responsible consumer thinks about her impact in a single-good market. Her utility function is $u(c) - pc - kg$, where $c > 0$ is her consumption, $u(c)$ is her consumption utility, $p > 0$ is the price, g is an externality equal to total consumption in the market, and $k > 0$ is a constant. The last term captures the consumer's concern regarding the externality. There are I identical other consumers in the market, as well as

I identical suppliers, who have exogenously given demand and supply curves. Whatever demand c the consumer submits, the price p is chosen to clear the entire market.

In contrast with the typical view, we demonstrate that the consumer needs to keep thinking about her impact on the market price even as $I \to \infty$. If she lowers her demand to reduce the externality, the price drops and induces others to consume more and consequently generate more of the externality. Furthermore, while the price effect is proportional to $1/I$, there are I other consumers who react, so their dampening effect does not vanish. This lowers the consumer's motive to mitigate, and means that the market mechanism erodes her moral behavior: she consumes more than she would under home production.

Based on the above analysis, we introduce our framework for markets with many socially responsible consumers. We define a competitive equilibrium as a situation in which the market clears, each consumer maximizes her utility given the dampening generated by the market demand curve, and that demand curve is consistent with consumer behavior.[1] In addition, we define social welfare as the sum of the private utilities of consumers and the producer surplus, minus K times the externality g.[2]

In Section 3, we show that markets fail in the presence of externalities not only when individuals do not care about it, but also when individuals do care about it. As an illustrative if extreme case, suppose that all consumers have $k = K$ — they assign the same weight to the externality as the social planner. This means that a consumer and the planner are willing to pay the same amount out of the consumer's funds to reduce CO2 in the atmosphere by a ton. Nevertheless, dampening implies that any equilibrium features overconsumption, and multiple inferior equilibria may arise. Further, due to path dependence it is plausible that society converges to the highest-consumption, and hence worst, equilibrium. These are *market* failures: for consumption not traded in the market (e.g., burning garden waste) fully socially responsible consumers produce the optimal outcome.

The above market failures imply that dealing with externalities cannot be based solely on consumer responsibility. In Section 4, therefore, we analyze the effectiveness of different policies

[1] In Appendix A, we provide game-theoretic microfoundations for our notion of competitive equilibrium.

[2] As we argue in more detail below, with socially responsible consumers it would be inappropriate to define welfare as a (weighted) sum of individual utilities.

in improving outcomes. We consider market-based interventions in which a linear permit fee τ for production arises, and define a policy as superior to another if it can achieve the same decrease in the externality produced at a lower fee or consumer price. For selfish consumers ($k = 0$), different policies are equivalent in this sense (as in Weitzman, 1974). For socially responsible consumers, however, differences arise because policies can differ greatly in the dampening they induce. A policy that induces lower dampening is better at motivating consumers, so it is superior.

As a starting point, we confirm in our setting the main insight of Herweg and Schmidt (2022), that a unit tax is better than a cap. Under a (binding) cap, there is a fixed quantity of permits, and the fee τ for using one is determined in equilibrium. Then, dampening is full, as the consumer realizes that she does not affect the quantity traded. Under a tax, in contrast, the permit fee τ is fixed, and the quantity is determined in equilibrium. Then, dampening is not full.

Going further, we show that the planner can improve on a tax. In particular, a permit fee that is decreasing in the consumer price induces less dampening and hence provides a stronger motivation for socially responsible consumers. Intuitively, when a consumer lowers her consumption, the price falls. This raises the permit fee, hindering others from consuming more.

In addition, we identify conditions under which a cap is better than a tax when there is trade, and the policymaker can only control domestic policy. As is well-understood by policymakers and researchers, trade facilitates "carbon leakage," whereby production moves abroad to avoid regulation. If the foreign source is large and/or much more polluting than domestic producers, carbon leakage implies that socially responsible consumers are more motivated to cut consumption under a cap. When a consumer reduces her consumption in this situation, she realizes that the dirty foreign producers will respond, which she greatly cares about.

In Section 5, we consider the scope for mitigating the externality through product selection rather than consumption reduction. We assume that there are two products, a clean one and a dirty one. They are perfect substitutes in consumption utility, but the clean one generates less of an externality. Consumers all have unit demand, and are heterogeneous in the weight k they attach to the externality. Then, there is always an equilibrium in which the two products sell at the same price, yet consumers are indifferent between them. This means that socially responsible consumers behave as if they were selfish, and the clean product enjoys no advantage in the market. Intuitively, if a consumer expects others to make their decisions based just on the price, then she expects dampening to be full: if she switches to the clean product, its price rises, so someone else switches to the dirty product to re-equilibrate the market at the previous price. Hence, the consumer does not affect the externality, and makes her decision based just on the price too.

Worse than a possibility, the above selfish equilibrium is unique if the clean product is not much less polluting than the dirty one. For an intuition, suppose as a start that consumers expect dampening to be weak. Still, since the two products generate similar externalities, consumers with different k behave quite similarly. This means that the price sensitivity of demand is quite high, so dampening must be non-trivial. As a result, consumers behave even more similarly, so the price sensitivity of demand is even higher. In this fashion, the equilibrium unravels. These results on the prominence of the selfish equilibrium are an extreme manifestation of the market's failure to facilitate outcomes reflecting the preferences of socially responsible consumers. As such, they reconcile extensive survey evidence that consumers have ethical concerns (Loureiro and Lotade, 2005, Auger et al., 2008, in addition papers cited above) with our casual observation that this preference is not translated into more ethical market outcomes.

When the dirty good is sufficiently dirtier and other conditions hold, then there is at least one equilibrium in which the clean product is sold at a premium, more responsible consumers buy the clean product, and less responsible consumers buy the dirty product. Even in such a non-selfish equilibrium, however, dampening operates, so the clean product's market share remains below the socially optimal level.

Given the possibility of massive market failure, we revisit the role of policy. We establish that with either a cap or a unit tax on the dirty product, the selfish equilibrium continues to exist. Furthermore, in a selfish equilibrium, the logic of Weitzman (1974) applies, so the two policies are equivalent. In a non-selfish equilibrium, in contrast, the insight of Herweg and Schmidt (2022) operates, and a tax is better than a cap. Again, however, the policymaker can do better than with a unit tax. Specifically, a proportional subsidy on the price of the clean product dominates a unit tax on the dirty product and a unit subsidy on the clean product, which both dominate a proportional tax on the dirty product. Intuitively, when a consumer switches to the clean product, she raises its price, which under a proportional subsidy encourages an especially large response in supply. As a result, dampening is not as strong, motivating the consumer to act responsibly.

As we have mentioned, our theory assumes that consumers are rational; in particular, they correctly understand dampening. This approach follows the common methodology in behavioral economics of introducing one new assumption at a time (Rabin, 1998, 2013). Roughly consistent with our premise, several observations suggest that many individuals understand dampening. Notably, dampening is a variant of the "replacement logic," a common justification for immoral behavior under which a person claims that if she did not do it, someone else would. Experimental research by Falk et al. (2020) and especially Ziegler et al. (2022) demonstrates that individuals understand, and their behavior responds to, the replacement logic. Similarly, people worry about the issue of "additionality" in CO2 reductions — asking whether an action creates additional reductions or just replaces others — and in social impact investment — asking whether a project would have been funded anyway (Perino, 2015, Green and Roth, 2021, Krahnen et al., 2021, Dietz and Grabs, 2022, Fraser and Fiedler, 2022, Oehmke and Opp, 2022). Our paper studies the equilibrium implications of such understanding. This is especially relevant since a full understanding by all consumers is not necessary for a competitive equilibrium to obtain.

Nevertheless, there are also consumers who are not fully rational. Indeed, we translate some of our results into insights that may be helpful for consequentialist consumers who do not fully understand their market impacts. For instance, we have often heard the claim that if a person consumes a polluting good sold under a cap, she does not affect pollution levels. As we have explained above, under trade this logic can be majorly wrong, with a consumer having a large impact exactly because of the cap.

We conclude in Section 6 with potential topics for future research. Since our equilibrium framework can be adapted in a formulaic way to other situations, it opens the possibility for studying the behavior of consequentialist socially responsible consumers — and the question of how to motivate them — in myriad other settings with externalities. At the same time, incorporating into our framework consumers who are not fully rational or have non-consequentialist motivations is a natural next step.

Related Literature Our research relates to several literatures. No previous paper, however, demonstrates that dampening operates for infinitesimally small consumers in a product market, and hence no other paper incorporates such dampening into a model of competitive equilibrium or studies its implications.

There is a large literature on how to think about equilibria in markets with small consumers. In this literature, price taking is a typical axiom, which for self-interested preferences has broadly applicable game-theoretic foundations (Mas-Colell, 1983). Because behavior in a market depends only on prices, an intuitive implication of taking prices as fixed is that a person takes others' behavior as fixed as well. Existing research on the market consequences of social preferences has overwhelmingly — and often implicitly — adopted such an interpretation (Sobel, 2007, Dufwenberg et al., 2011, Hakenes and Schliephake, 2021, Pástor et al., 2021, Dewatripont and Tirole, 2022, Herweg and Schmidt, 2022, Piccolo et al., 2022, Aghion et al., 2023, Arnold, 2023). The starting insight of our paper is that this version of price taking does not apply to rational consequentialist consumers who care about a global externality. Furthermore, we define a variant of competitive equilibrium that appropriately accounts for socially responsible consumers' incentives.

Although they do not frame it in these terms, the model of Broccardo et al. (2022) also features a violation of price taking. The authors show that because exit (i.e., divestment or boycotts) has a non-trivial private cost, it is worse at aligning the individual and social incentives of socially responsible investors than shareholder voting. A consumer's price impact acts through managers' investment decisions, and while it alters the impact of exit, it is not crucial for the paper's main result. In contrast, our dampening acts through the product market, and our primary interest is in studying its implications, leading to different results.[3,4]

[3] At a more technical level, Broccardo et al. assume that agents commit to their actions and firms can observe these actions before making their own decisions. Our equilibrium and its microfoundations do not require such arguably unrealistic public commitment.

[4] A more distantly related literature studies voluntary donations by small agents to fund public goods, which is not a market interaction. One strand, notably Sugden (1982) and Sugden (1985), highlights a strategic form of dampening that can occur in these non-market settings, whereby one donor's action may be (partially or fully) offset by other donors. Another strand does not investigate dampening, but the complementary concern of free-riding (see for instance Andreoni, 1988, Bergstrom et al., 1986, Gaube, 2006).

Recent research explores how the problem of additionality affects the funding decisions of socially responsible investors with market power. For instance, Green and Roth (2021) show that investors may invest in less green projects to raise the returns on — and thereby attract less responsible investors to — profitable greener projects (see also Moisson, 2020, Hakenes and Schliephake, 2021, Pástor et al., 2021, Oehmke and Opp, 2022). These indirect effects are related to our dampening, but because we consider the consumer side, our results are completely different.

Many researchers have explored the effectiveness of different policies to regulate externalities. We contribute to this literature by studying policies in markets with socially responsible consumers. Most related to our policy analysis, Herweg and Schmidt (2022) show that a tax dominates a cap in a closed economy. As we have mentioned, the difference arises due to a difference in (what we call) dampening under the two policies. Herweg and Schmidt capture the difference through an exogenous parameter (β^R), whereas we derive dampening from economic primitives. This allows us to consider other policy questions, such as identifying better policies or comparing a cap and a tax under trade, without making additional exogenous assumptions.

2 Large Markets with Socially Responsible Consumers

2.1 A Single Consumer's Perspective

As a first step in developing our framework, we analyze how a single socially responsible consumer thinks about her behavior in a large market. Our analysis yields some basic points and serves as an input into the definition of competitive equilibrium below.

Setup The market mechanism, consistent with the notion of a Walrasian auctioneer, is the following. First, the consumer submits her demand $c \geq 0$. Then, the price $p \geq 0$ is chosen so that the market clears: total demand, which consists of c and others' demand at p, equals supply at p. (Others' demand and supply will be described below.) Finally, outcomes are determined.

The consumer correctly predicts the above outcomes, and maximizes

$$\underbrace{u(c)}_{\substack{\text{consumption}\\ \text{utility}}} - pc - \underbrace{kg}_{\substack{\text{social}\\ \text{concern}}}, \tag{1}$$

where $u(\cdot)$ is a thrice differentiable strictly concave utility function, g is an externality equal to the total quantity q of the good produced and consumed in the market, and $k \geq 0$ is a constant.[5] The term $u(c)$ is the person's consumption utility, $u(c) - pc$ is her private utility, and the last term $-kg$ captures her concern regarding the externality. To ensure that the equilibrium we define below features positive consumption, we impose that $u'(0) > k$ and $\lim_{c\to\infty} u'(c) \leq 0$.

In the above formulation, the consumer internalizes the harm g caused by everyone. An equivalent assumption is that she cares about her own impact on the externality. Formally, let $g(c)$ be the externality that results if the consumer consumes c. Then, her utility function is $u(c) - pc - k\left(g(c) - g(0)\right)$. Since $g(0)$ is a constant, this yields the same behavior as the utility function above.

Our interest is in the behavior of a consumer, such as an individual, household, or small organization, who is tiny relative to the market. A long-standing axiom regarding such consumers is that they think of the market price as fixed (Debreu, 1959, Arrow and Hahn, 1971). Indeed, extensive research studying self-interested consumers establishes that such price taking is the limit of optimal strategic behavior as the consumer becomes a negligible part of the economy (e.g., Mas-Colell, 1980, 1983, and others in the former paper's issue). We reconsider this question for socially responsible consumers by adapting an existing methodology, a "replicator economy" (Shubik, 1973, Roberts and Postlewaite, 1976). This involves introducing identical copies of the rest of the economy, and seeing how the consumer behaves as there are more and more copies.

Suppose, then, that there are I other consumers and I suppliers. The consumers all have the same continuously differentiable demand curve $D(p)$ with $D'(p) < 0$ everywhere, and the suppliers all have the same continuously differentiable supply curve $S(p)$ with $S'(p) > 0$ everywhere. There

[5] Hence, in this basic model the notation g is redundant — we could have defined the consumer's utility in terms of just q. We distinguish g and q because below, we will extend our model to markets with multiple products that generate different externalities. By assuming that the consumer cares about g, we can introduce such modifications without changing the consumer's utility function.

is a price $p^* > 0$ for which $S(p^*) = D(p^*)$, i.e., for which the rest of the market would clear, and $\lim_{p\to\infty} S(p) - D(p) = \infty$. For simplicity, we assume that as $I \to \infty$, the weight k the consumer places on the externality remains constant. Our points remain unchanged as long as k does not vanish. A non-vanishing k, in turn, follows from our definition of a socially responsible consumer: that she is willing to modify her consumption to mitigate the externality associated with it.

Analysis For any $c \geq 0$, there is a unique market-clearing price $p(c) \geq 0$, which solves $c + ID(p(c)) = IS(p(c))$. Taking the total derivative with respect to c and rearranging gives

$$p'(c) = \frac{1}{I\left(S'(p(c)) - D'(p(c))\right)}.$$

The total quantity produced in the market is $q(c) = IS(p(c))$. Hence, the effect of the consumer's demand on the market quantity, which we call the market responsiveness, is

$$q'(c) = IS'(p(c))p'(c) = \frac{S'(p(c))}{S'(p(c)) - D'(p(c))}.$$

Taking limits gives:

Proposition 1 (Violation of Price Taking)**.** *For any $D(\cdot)$ and $S(\cdot)$, a vanishingly small consumer:*

I. has a negligible impact on the price: for any c, $\lim_{I\to\infty} p(c) = p^*$ and $\lim_{I\to\infty} p'(c) = 0$.

II. has a non-negligible impact on others through the price: for any c, market responsiveness is

$$\lim_{I\to\infty} q'(c) = \frac{S'(p^*)}{S'(p^*) - D'(p^*)} < 1. \tag{2}$$

Part I replicates the insight of Roberts and Postlewaite (1976) and others that in a large economy, a consumer has a negligible price impact. Even so, Part II says that the consumer's impact on the market quantity is less than one-to-one, which means that she impacts others' consumption to a non-vanishing extent. Indeed, others have a dampening effect of size

$$\lim_{I\to\infty} \left| \frac{\mathrm{d}[ID(p(c))]}{\mathrm{d}c} \right| = \lim_{I\to\infty} \left[-ID'(p(c))p'(c)\right] = \frac{-D'(p^*)}{S'(p^*) - D'(p^*)}, \tag{3}$$

exactly $\lim_{I\to\infty}(1 - q'(c))$. Intuitively, the consumer's demand c raises the price, leading others to consume less. Although the price impact is proportional to $1/I$, there are I other consumers, so their total consumption response is comparable to c. Because a socially responsible consumer by definition cares about the impact of c on the externality, she must care about others' response as well. In this sense, she is not a price taker.

It is now easy to determine how an infinitesimally small consumer chooses c. Namely, for any $p^* \geq 0$ her maximization problem has a unique solution satisfying the first-order condition[6]

$$\lim_{I\to\infty} \frac{\mathrm{d}}{\mathrm{d}c}\left[u(c) - cp(c) - k \cdot q(c)\right] = \lim_{I\to\infty} \left[u'(c) - p(c) - cp'(c) - k \cdot q'(c)\right]$$
$$= u'(c) - p^* - k \cdot \frac{S'(p^*)}{S'(p^*) - D'(p^*)} \leq 0, \text{ with equality if } c > 0. \tag{4}$$

Two immediate points follow. First, because the consumer's disutility from consumption includes not only the price p^* she pays, but also her net impact on the externality from Proposition 1, she consumes less than a selfish consumer. Second, however, dampening implies that she consumes more than she would if c came from home production with the same private marginal cost, p^*. With home production, she would not affect others' consumption, so dampening would be zero. Consistent with an experimental literature starting from Falk and Szech (2013) that people behave more selfishly in a market than in an individual-decisionmaking setting, therefore, the market erodes the consumer's moral behavior. In these experiments, the erosion of moral behavior is arguably due to the replacement logic built into the strategic setting by the researchers. Proposition 1 shows that the erosion is a fundamental property of the standard price-based market mechanism.

Equation (3) implies that dampening, and therefore how much the market erodes moral behavior, depends on a single sufficient statistic: the elasticity of demand relative to the elasticity of supply $(-D'(p)/S'(p))$. This determines the extent to which demand rather than supply responds to the price increase resulting from the consumer's demand. If the relative elasticity is low, it is mostly supply that responds, so the dampening effect of other consumers is small. An example is when there is an abundance of suppliers with the same marginal cost. But if the relative elasticity is high, it is mostly demand from other consumers that responds, so their dampening effect is large. An example is when there is an abundance of consumers with the same marginal utility.

[6] Since $u(\cdot)$ is strictly concave and $\lim_{c\to\infty} u'(c) \leq 0$, Equation (4) has a unique solution for any $p^* \geq 0$. The strict concavity of $u(\cdot)$ also ensures that the solution satisfies the second-order condition.

2.2 Competitive Equilibrium

We now build on the above analysis to define equilibrium in markets with many socially responsible consumers. Because the demand curve both affects and is determined by consumer behavior, it must be endogenously defined.

We assume that there is a mass 1 of homogeneous consumers, who have the utility function in the previous subsection. (Incorporating consumers who are heterogeneous in their consumption utilities u or weights k is straightforward, but requires additional notation.) Because our insights derive from demand-side considerations, for the rest of the paper we make the simplifying assumption that the supply curve is exogenous and linear. Here, we let $S(p) = sp$, with a slope $s > 0$.

Definition 1. A competitive equilibrium is a quantity $q^* > 0$, price $p^* > 0$, consumer price responsiveness $q_p^* \in \mathbb{R}$, and market responsiveness $q_c^* \in \mathbb{R}$ such that

1. Supply equals q^*: $q^* = S(p^*)$.
2. Demand equals q^*: $u'(q^*) = p^* + k \cdot q_c^*$.
3. Market responsiveness is consistent with consumer price responsiveness: $q_c^* = s/(s - q_p^*)$.
4. Consumer price responsiveness is consistent with optimization: $q_p^* = 1/u''(q^*)$.

We define a competitive equilibrium based on two standard ingredients, the quantity q^* and price p^*, as well as two new ingredients, the price responsiveness of demand q_p^* and the market responsiveness q_c^*. These must satisfy four conditions. Condition 1 is standard: supply equals the equilibrium quantity. Condition 2 says that demand equals the equilibrium quantity, where we use Equation (4) and the market responsiveness to account for a consumer's mindfulness of her net effect on the externality. Condition 3 states how the market responsiveness is determined. Applying Equation (2), it depends on the slopes of the supply and demand curves. Finally, Condition 4 requires that the slope of the demand curve be consistent with consumer optimization. To obtain the condition, we totally differentiate Condition 2 and impose that the consumer treats the market responsiveness as constant within the range of an infinitesimal price change.

In Appendix A, we provide game-theoretic microfoundations for our equilibrium concept, showing that it arises as a limit case when each consumer becomes infinitesimally small in the economy. Our analysis uses methods starting from Kyle (1989) that have been developed to model the behavior of financial-market participants with non-trivial price impacts (see Rostek and Yoon, 2020, for a review). In these market games, players submit not only a demand to the market, but also how their demand responds to price changes. Analogously to Section 2.1, others' price responsivenesses determine how an individual player's consumption affects the market quantity. To give each player an incentive to truthfully report her price responsiveness, we introduce noise into supply. This results in price shocks, so that the submitted price responsiveness becomes relevant on the equilibrium path. We develop methods for solving such games for a finite number of players. The notion of competitive equilibrium in Definition 1 describes the limit of equilibrium outcomes as the number of players approaches infinity and the noise in supply vanishes.

Within this framework, we consider two different games. In the first, we allow each player to submit an arbitrary demand schedule, and look for equilibrium in linear schedules. Much like in the finance literature, this approach is only tractable for quadratic consumption utility (Rostek and Yoon, 2020). Hence, this microfoundation applies for a quadratic approximation of $u(\cdot)$ around the equilibrium point. In the second game, we restrict strategies to be linear: each player submits a quantity and a scalar price responsiveness. This leads to a microfoundation for any $u(\cdot)$.

Our notion of a competitive equilibrium collapses to the standard notion when $k = 0$. As above, Conditions 1 and 2 then say that supply and demand must both equal q^*. But for $k = 0$, the market responsiveness q_c^* plays no role in a consumer's optimization, so it, and Condition 3 on it, are redundant. Furthermore, since the consumer price responsiveness q_p^* is only used to define q_c^*, it, and Condition 4 on it, are also redundant.

As a basic point, we note:

Observation 1. *A competitive equilibrium exists.*

To complete our setup, we define social welfare when everyone consumes an amount q as

$$u(q) - \int_0^q S^{-1}(x)\mathrm{d}x - Kq.$$

One part of social welfare is total consumption utility net of the costs of production. This is the sum of consumers' private utilities $u(c_i) - pc_i$ and the producer surplus. In addition, the social planner puts an exogenously given weight $K > 0$ on the externality. First-pass conventional logic might dictate that we include the sum of individuals' utilities here too. But the weight a consumer puts on the externality already incorporates a concern for society, so including each such term in the social welfare function amounts to multiple-counting the same concern.

A natural assumption is $k \leq K$. The reason that a socially responsible consumer puts a weight on the externality is that she cares about others in addition to herself. But in such a tradeoff, most people put a larger weight on themselves, whereas a social planner uses a more equal weight.

3 Failures of Socially Responsible Consumerism

In this section, we demonstrate ways in which the market's ability to coordinate socially responsible behavior is limited. Our first result is a basic market failure:

Proposition 2 (Overconsumption)**.** *There is a unique socially optimal quantity q^{FB}. For any $k \leq K$, any competitive-equilibrium quantity q^* is strictly greater than q^{FB}.*

Proposition 2 adds to our understanding of the basic economics of markets. It is well-known that when there are no externalities or other frictions, markets perform well despite everyone favoring their own private consumption. It is also well-known that when each person's consumption creates an externality she does not care about, markets perform poorly. Going beyond these insights, the proposition says that markets perform poorly even when each person's consumption contributes to an externality she does care about. Specifically, dampening reduces each consumer's incentive to cut back, leading to overconsumption.

As an extreme but illustrative special case, suppose that $k = K$ — i.e., consumers put the same weight on the externality as the social planner. Imagine, for example, that the planner values €200 of a citizen's consumption the same as reducing CO2 in the atmosphere by a ton. Then, the assumption $k = K$ means that the citizen has the same valuation, willing to lower her consumption by €200 to reduce CO2 in the atmosphere by a ton. As everyone's preferences internalize the externality, the only disagreement between individuals regards their private consumption, so one would think that — as without externalities — the social optimum again obtains.[7] Proposition 2 says that it does not.

Like the erosion of moral behavior in the previous section, overconsumption is due to the market. In particular, it is not just the presence of externalities from a good that causes overconsumption; it is that the good is obtained from the market. To see this formally, consider the following modification of our model. We let the number of individuals I be finite. We assume that c_i is a good, such as how much of her garden waste consumer i burns, that is produced at home. Consumer i's utility is $u_i(c_i) - kg$, where $u_i(c_i)$ is her benefit from consumption and g is the externality equal to the total quantity consumed by everyone. The utility function $u_i(\cdot)$ satisfies the same conditions as $u(\cdot)$ above. The social welfare function is $[\sum_i u_i(c_i)] - Kg$. Then, the following is obvious:

Observation 2. *With non-market consumption, there is a unique socially optimal consumption profile* $(c_1^{\text{FB}}, \ldots, c_I^{\text{FB}})$*. If* $k = K$*, then in the unique Nash equilibrium, consumer* i *chooses* c_i^{FB}*.*

With the market not eroding her behavior, a fully responsible consumer ($k = K$) trades off the private benefit of consumption with the full externality generated. Hence, she chooses the socially optimal level of consumption.

The market failure we have identified is relevant for society's vision for dealing with externalities. The observation that many consumers care about social problems raises the hope that we can use a decentralized approach, allowing consumers to optimally balance their private needs with social concerns. Indeed, Giesler and Veresiu (2014) and Chater and Loewenstein (forthcoming) argue that there has been a shift toward emphasizing consumer responsibility rather than systemic reform as a solution to social problems. The founder of the World Economic Forum (WEF) in Davos, Klaus Schwab, summarizes the case succinctly. He argues that neither "shareholder capitalism," based on pure profit maximization, nor "state capitalism," where states take an active role, are the way forward. Instead, he argues for individuals as consumers and investors — rather than citizens

[7] This logic is reinforced by the rationale economists typically give for Pigouvian taxation. The rationale is that a tax equal to a person's externality generates optimal outcomes because it leads individuals to internalize the externality they would otherwise neglect (e.g., Gruber, 2005). Hence, the observation that consumers' preferences already internalize the externality would appear to obviate the need for intervention.

— to affect change on social and environmental problems through "stakeholder capitalism."[8] A manifestation of this push is the individual carbon footprint calculator that consumers are supposed to use to guide their choices. An obvious flaw with this approach is that consumers will free-ride because they do not care sufficiently. But our result that the social optimum does not obtain even with $k = K$ shows that stakeholder capitalism fails on its own terms: dealing with externalities cannot be successfully outsourced even to extremely responsible individuals.

We now show that beyond overconsumption in any equilibrium, multiple equilibria can arise.

Proposition 3 (Multiple Equilibria). *Fix any $u(\cdot)$, k, and s, and take a corresponding competitive equilibrium (q^*, p^*, q_p^*, q_c^*). If $u'''(q^*)$ is sufficiently high, then there is (i) a competitive equilibrium q^+, p^+, q_p^+, and q_c^+ with $q^+ > q^*$, $p^+ > p^*$, $|q_p^+| > |q_p^*|$, $q_c^+ < q_c^*$; and (ii) a competitive equilibrium with q^-, p^-, q_p^-, and q_c^- with $q^- < q^*$, $p^- < p^*$, $|q_p^-| < |q_p^*|$, $q_c^- > q_c^*$. Among multiple equilibria, social welfare is strictly decreasing in the equilibrium quantity.*

For a simple intuition, suppose that consumers get a large utility boost from taking one flight per year, but do not care as much for further flights. Then, if consumers expect everyone to take one flight, they expect demand to be price-insensitive, so by Equation (3) they conclude that dampening is low. Consistent with equilibrium, therefore, consumers are motivated to mitigate and they take one flight. In contrast, if consumers expect everyone to take multiple flights, they expect demand to be price-sensitive and hence dampening to be high. Again consistent with equilibrium, therefore, consumers are not motivated to mitigate and take multiple flights.[9]

By Proposition 3, multiple equilibria exist if $u'''(q^*)$ is sufficiently high. This condition means that consumer price responsiveness ($1/u''(c)$) changes sufficiently fast for the above self-reinforcing

[8] See, e.g., "Why we need the 'Davos Manifesto' for a better kind of capitalism" (Schwab 2019, https://www.weforum.org/agenda/2019/12/why-we-need-the-davos-manifesto-for-better-kind-of-capitalism/). This sentiment is echoed by managers at the WEF interviewed by Giesler and Veresiu (2014). A health insurance manager said: "Taxes and bureaucratic initiatives do one thing and one thing only, they prevent partnership and collaboration because they create constraints rather than opportunities. Instead be enablers! Create a climate of partnership and collaboration where we can work together to engage consumers." Similarly, a sustainability manager said during his interview: "We're all passionately committed to implementing absolute sustainability. Absolute sustainability in the sense of nurturing an entirely new set of water consumption habits and attitudes in people, in short, creating a fundamentally new generation of water consumers who readily connect environmental issues with their personal water lifestyles and what they aspire to as individuals."

[9] The prediction of our model that a person's pro-social behavior depends on what she thinks others will do, and the possibility of multiple equilibria, are also defining features of conditional cooperation (Fischbacher et al., 2001). Hence, our model can be seen as identifying a novel reason for conditionally cooperative behavior in markets.

logic to take hold. Existing evidence shows that the third derivative of the consumption-utility function is substantially positive (Parker and Preston, 2005, Ventura and Eisenhauer, 2006, Deck and Schlesinger, 2014, Noussair et al., 2014), indicating that the condition is plausible. But we are not aware of evidence that would allow us to perform a precise calibration.

Given the possibility of multiple equilibria, it is natural to ask which one is selected. Informal arguments based on path dependence suggest that the worst equilibrium is the most likely outcome. The detrimental effect of many global externalities, such as global warming or air pollution, has not been understood or appreciated until recently. In such situations, the market has historically been in a high-consumption equilibrium (that with $k = 0$). Even as consumers find out about the high social cost of consumption, they also realize that the current equilibrium is one in which everyone consumes a lot. Then, it is plausible that society equilibrates at the highest-consumption equilibrium. Reinforcing this miscoordination issue is that different consumers may recognize the problem at different points in time, creating no obvious focal point for switching to a better equilibrium.

4 Policy

We have shown that even extremely responsible consumers overconsume in equilibrium. We now analyze the effectiveness of different policies in mitigating this market failure.

4.1 Intervention and Definition of Competitive Equilibrium

We first consider policies under which the planner acts as a supplier of pollution permits producers must obtain to sell units of the good. Such a "permit-supply policy" is described by the curve $\pi g - (1 - \pi)\tau + \pi_0 = 0$, where g is the amount of permits the planner makes available, τ is the fee for a permit, and $\pi \in [0, 1]$ and $\pi_0 \in \mathbb{R}$ are constants that parameterize the policy. Notice that the lower is π, the more permit supply responds to changes in the permit fee. Two commonly analyzed types of policies arise as special cases. A binding cap obtains when $\pi = 1$; then, $-\pi_0$ is the cap. This creates a completely inelastic supply of permits. And a fixed unit tax obtains when $\pi = 0$; then, π_0 equals the tax. This creates an infinitely elastic supply of permits.

To define competitive equilibrium, notice that if the consumer price is p, then the producer price is $p-\tau$, so supply is $S(p-\tau)$. Furthermore, since the available number of permits must equal supply ($g = S(p-\tau)$), we must have

$$\pi S(p-\tau) - (1-\pi)\tau + \pi_0 = 0. \tag{5}$$

There is a unique fee τ that solves (5). We use the solution to define net supply $S_{\text{net}}(p) = S(p-\tau)$. With supply and demand both functions of the consumer price p, we can use the same logic as in Sections 2.1 and 2.2 to define market responsiveness. This requires the price responsiveness of net supply, $S'_{\text{net}}(p)$. Totally differentiating (5) with respect to p yields $\mathrm{d}\tau/\mathrm{d}p = \pi s/((1-\pi)+\pi s)$, so

$$S'_{\text{net}}(p) = S'(p-\tau) \cdot \left(1 - \frac{\mathrm{d}\tau}{\mathrm{d}p}\right) = \frac{(1-\pi)s}{(1-\pi)+\pi s} \equiv s_{\text{net}}. \tag{6}$$

We are now ready to modify Definition 1 to account for intervention:

Definition 2. A competitive equilibrium with policy is a quantity $q^* > 0$, consumer price $p^* > 0$, permit fee τ^*, consumer price responsiveness $q_p^* \in \mathbb{R}$, and market responsiveness $q_c^* \in \mathbb{R}$ such that

1. a. Supply equals q^*: $q^* = S(p^* - \tau^*)$.

   b. The market for permits clears: $\pi S(p^* - \tau^*) - (1-\pi)\tau^* + \pi_0 = 0$.

2. Demand equals q^*: $u'(q^*) = p^* + k \cdot q_c^*$.

3. Market responsiveness is consistent with the responsiveness of consumers and net supply: $q_c^* = s_{\text{net}}/(s_{\text{net}} - q_p^*)$, where s_{net} is given by Equation (6).

4. Consumer price responsiveness is consistent with optimization: $q_p^* = 1/u''(q^*)$.

Definition 2 differs from Definition 1 in three respects. In Part 1.a, supply in the product market accounts for the permit price τ^*. In Part 1.b, there is the new condition that given the policy curve, the market for permits clears. And in Part 3, the market responsiveness accounts for intervention as well. Here, the condition in the definition plugs Equation (6) into Equation (2).

We suppose that in choosing its policy, the social planner is looking to implement a given feasible externality level g^* — or, equivalently, quantity level $q^* = g^*$ — that improves on the no-intervention outcome. Further, we assume that less intervention is better than more intervention.

More precisely, we say that policy type A (e.g., a unit tax) is strictly better than policy type B (e.g., a cap) if for any $q^*, \tau_B^* > 0$ that is part of an equilibrium with a B-type policy, there is an A-type policy and a corresponding equilibrium with quantity q^* and fee $\tau_A^* < \tau_B^*$. Since in equilibrium the market must clear (i.e., $q^* = S(p^* - \tau^*)$), this also means that the consumer price p^* is lower under the A-type policy. Lower permit prices might be socially preferable because they have lower administrative, enforcement, or political costs. In addition, lower consumer prices might be socially preferable because consumer surplus is more important or politically expedient than producer surplus.

4.2 Comparison of Policies

Principle In our setting with no uncertainty, a version of the classical result by Weitzman (1974) holds: with selfish consumers ($k = 0$), different policies are equivalent in the above sense. The same is not true with socially responsible consumers due to a general principle. Namely, a policy that generates less dampening is better at motivating consumers to mitigate, and hence can achieve the same outcome with less intervention. To understand how policies rank, therefore, it is sufficient to understand how they impact dampening. Hence, we explain our results in terms of this principle.

Results Our benchmark result is:

Proposition 4 (Permit-Supply Policies)**.** *Policies with parameter π are strictly better than policies with parameter $\pi' > \pi$.*

Among permit-supply policies, more responsive policies are better. Intuitively, a more elastic supply of permits to product suppliers translates into a higher price responsiveness of product supply, which by Equation (3) implies lower dampening. An implication is that a tax ($\tau = 0$) is the best permit-supply policy, while a cap ($\tau = 1$) is the worst. Replicating Herweg and Schmidt's (2022) insight, therefore, a tax is better than a cap. Herweg and Schmidt, however, reach that conclusion by exogenously imposing that consumers are more willing to mitigate under a tax. Our framework endogenizes this willingness, and thereby allows us to study other policy questions without new exogenous assumptions about how consumers think.

We consider two such further questions.[10] First, we show that a regulator can do better than with a unit tax. A conceptually interesting example is a fixed-price policy: the planner fixes the consumer price p^*, and chooses the fee τ^* to clear the market. Because dampening operates through changes in the price, it now equals zero. Intuitively, when a consumer cuts her consumption, the fee responds to offset the resulting decrease in the price. Hence, the behavior of other consumers does not change. Therefore:

Observation 3. *A fixed-price policy is strictly better than a tax.*

While a fixed-price policy is not realistic in practice, its logic suggests more plausible alternatives. We analyze the class of tax policies that depend on the consumer price in a linear way: $\tau = \tau_0 + \tau_1 p$. The case $\tau_1 = 0$ corresponds to a unit tax we have considered above. We get:

Proposition 5 (Price-Dependent Policies)**.** *Taxes that are decreasing in the consumer price ($\tau_1 < 0$) are better than fixed taxes ($\tau_1 = 0$).*

The intuition is related to that under the fixed-price policy above. If a consumer cuts her consumption, the resulting decrease in the price raises the tax. This attenuates the price drop, so the response of other consumers is lower. Hence, dampening is weaker than under a fixed tax.

As a second further question, we compare cap and tax policies when there is trade, and the planner can only control domestic policy. In this case, it is generally understood that "carbon leakage" can occur: the price on pollution may lead supply to shift toward foreign sources. We analyze how this affects a socially responsible consumer's thinking and the optimal policy.

We assume that there is a foreign net supply curve $S_f(p) = s_f p$, and a home supply curve $S_h(p - \tau) = s_h(p - \tau)$. Foreign supply is a function of the consumer price p because foreign production is not subject to regulation. In addition, foreign production is more polluting: a unit of home and foreign production cause externalities $e_h \geq 0$ and $e_f > e_h$, respectively.[11] Hence, $g = e_h q_h + e_f q_f$, where q_h and q_f are the quantities produced at home and abroad, respectively. A consumer cannot distinguish the products; for instance, she does not know where her electricity

[10] Here and in our policy analysis in Section 5, adapting the definition of competitive equilibrium to the particular environments is straightforward, and hence omitted.

[11] For instance, foreign suppliers may use a different technology due to the differences in environmental regulation.

is coming from.[12] Finally, for simplicity in stating our results, we assume that u is quadratic, and denote its second derivative by u_{cc}.

To adapt Definition 2 for this situation, we distinguish between the market responsivenesses of home- and foreign-supplied quantities, $q_{h,c}$ and $q_{f,c}$. These can be calculated similarly to q_c in the definitions above. Denoting by q_p the price responsiveness of demand and letting $s_{\text{net},h} = (1-\pi)s_h/[(1-\pi)+\pi s_h]$, we get

$$q_{h,c} = \frac{s_{\text{net},h}}{s_{\text{net},h} + s_f - q_p} \qquad \text{and} \qquad q_{f,c} = \frac{s_f}{s_{\text{net},h} + s_f - q_p}. \tag{7}$$

We give the full definition of competitive equilibrium in the appendix.

Proposition 6 (Cap versus Tax Under Trade)**.**

I. If

$$-1/u_{cc} < \frac{e_f - e_h}{e_h} \cdot s_f, \tag{8}$$

then a cap is strictly better than a tax.

II. If Inequality (8) *goes strictly the other way, then a tax is strictly better than a cap.*

To understand the proposition, compare the responsivenesses (7) under a cap — where $\pi = 1$ and hence $s_{\text{net},h} = 0$ — and a tax — where $\pi = 0$ and hence $s_{\text{net},h} = s_h$. Both policies generate dampening, but there are two differences in how dampening operates in the two cases. On the one hand, home supply is more elastic under a tax than under a cap, so total supply is more elastic too. This means that the dampening effect on the overall quantity is smaller under a tax. On the other hand, the consumer has a greater impact on the more polluting foreign supply — the source she cares more about — under a cap, since this is then the only source that can respond to her behavior. Comparing the strengths of these two effects yields the two parts of the proposition.

Specifically, whether a cap or a tax is better depends on how much dirtier foreign production is, and how price responsive foreign supply is relative to demand. If foreign supply is much dirtier, then a cap is better. Intuitively, in this case a consumer cares mostly about lowering foreign production, and a cap allows her to do that. In addition, if foreign supply is quite elastic, than the

[12] The results are identical if consumers can distinguish but are in equilibrium indifferent between the two products. Analogously to Section 5 below, such an equilibrium always exists.

dampening effect on total consumption is small. Hence, both under a cap and under a tax, total consumption responds about one-to-one to a consumer's own consumption. But this all comes from foreign supply with the cap, and only partially from foreign supply with the tax.

The above logic yields a comparative static with respect to the importance of foreign supply in the domestic market. Proposition 6 reduces to a variant of Proposition 4 if the foreign partner is small in terms of production. Then, we can expect s_f to be small, so that a tax dominates a cap. If there are many foreign suppliers, in contrast, a small increase in the price can create excess supply that is large relative to the domestic market. Hence, s_f is likely to be large. Then, a cap is better than a tax.

The logic of Proposition 6 clarifies a misperception some observers seem to have regarding a consumer's impact on emissions. We have often heard that for consumption under a cap-and-trade system (e.g., flying within the EU), that impact is zero. As has already been understood under the rubric of carbon leakage, this argument is incorrect with international trade. Our analysis adds that in Case I, the argument is incorrect in a major way: it is *exactly because* of the cap that a consumer's impact on pollution is large.

5 Substitute Products

In our basic model, there is a single good traded in the market. We now analyze a variant in which there are substitute products with different externalities competing against each other.

We assume that there are two products, a clean one and a dirty one. They are perfect substitutes in consumption utility, and generate externalities $e_c \geq 0$ and $e_d > e_c$, respectively. For example, a consumer can power her appliances equally well with green and brown electricity, but the former is more environmentally friendly to produce. Similarly, fair-trade and non-fair-trade products are for many people functionally equivalent, and humanely sourced meat tastes similar to the non-humane kind. Denoting the market quantities of the two products by q_c and q_d, the total externality is $g = e_c q_c + e_d q_d$. Consumers are heterogeneous in the weight k they put on the externality, with k distributed on $[\underline{k}, \overline{k}] \subset (0, K]$ according to the cumulative distribution function F with density f bounded away from zero. To isolate the new issue, which of two products rather than how much of a product consumers choose, we impose that consumers have unit demand.[13] Suppliers provide the clean and dirty products according to the exogenous supply curves $S_c(p_c) = s_c p_c$ and $S_d(p_d) = s_d p_d$, respectively, where $s_c, s_d > 0$.

We define an equilibrium by adapting Definition 1. Crucially, there is a market responsiveness that determines how much the equilibrium quantities of the two products respond if a consumer moves her consumption from the dirty to the clean market. Each consumer's behavior, then, must be optimal given the equilibrium market responsiveness and her k. This generates a (local) demand curve for the two products as a function of their price difference. Finally, the market responsiveness must be consistent with the resulting demand curve and the exogenously given supply curves. We provide a formal definition in the appendix.

5.1 Laissez-Faire Equilibria

Proposition 7 characterizes the key features of competitive equilibria.

Proposition 7 (Substitute Products)**.** *Take any F.*

I. There is a competitive equilibrium in which the two products have the same price ($p_c^* = p_d^*$), and all consumers are indifferent between them.

II. If $e_d - e_c$, s_c, or s_d is sufficiently small, then there is no other competitive equilibrium.

III. If $e_d - e_c$ is sufficiently large, then there is a competitive equilibrium in which $p_c^* > p_d^*$, and some consumers strictly prefer the clean product.

IV. Among multiple competitive equilibria, the higher is $p_c^* - p_d^*$, the greater is social welfare.

Part I says that there is always an equilibrium in which the two products sell at the same price, and consumers are indifferent between them. This outcome is identical to that when all consumers are selfish. To see the intuition, suppose that each consumer believes that everyone else behaves selfishly, always choosing the cheaper product. Then, a consumer believes that the two prices must always equalize, and therefore dampening is full: if she switches from the dirty to the clean product, the resulting price difference will lead someone else to switch the other way,

[13] I.e., each consumer is looking to buy exactly one unit of the good. This is a limiting case of the previous model where the utility function becomes infinitely concave at a point.

re-equilibrating the markets at the previous price. She concludes that total quantities in the two markets, and therefore total pollution, do not depend on which product she chooses. Consistent with equilibrium, therefore, she chooses purely based on which product is cheaper.

Worse, Part II implies that if the products are not too different in the externalities they generate ($e_d - e_c$ is sufficiently small), then the above selfish equilibrium is the only equilibrium. The intuition derives from an unravelling of the dampening effect. Suppose, to start, that consumers expect the dampening effect to be small. Still, since the two products generate similar externalities, consumers with different k behave quite similarly. This means that the price sensitivity of demand is quite high, so the dampening effect must be non-trivial. As a result, consumers behave even more similarly, so the price sensitivity of demand is even higher, etc.

A similar unraveling logic holds when s_c or s_d is low, but the starting point acts in part through supply. A low s_c or s_d implies that reallocating production between the dirty and clean sectors requires a large change in the relative price. By moving her consumption to the clean sector, therefore, a consumer induces a large price change, so that many consumers move in the opposite direction. This means that the dampening effect is non-trivial, kickstarting the unraveling.

By Part III, at least one other equilibrium exists if the products' pollution levels differ sufficiently. In such a non-selfish equilibrium, the clean product is more expensive, and consumers decide which product to get according to the relative price as well as how much they care about the externality. Since consumers do not care only about prices, dampening is not full, motivating a share of the population to mitigate. Part IV says that this results in higher social welfare.

The fact that the selfish equilibrium exists, and may be the only one, is an extreme manifestation of the market's inability to motivate socially responsible behavior by socially responsible consumers. Although each consumer is willing to pay to mitigate the externality stemming from her own consumption, the equilibrium is identical to that when all consumers are selfish. Hence, consumers' social preferences are not reflected in their behavior or the market outcomes at all, and in particular the clean good enjoys no advantage in the market.[14]

[14] This insight is superficially related to what happens in previous models where consumers with social preferences act selfishly in a competitive equilibrium (Dufwenberg et al., 2011, Arnold, 2023). But previous theories assume a priori that each consumer is unwilling to pay to mitigate the effect of her consumption, so that she is not socially responsible by our definition. Similarly, although Fehr and Schmidt's (1999) inequity-averse agents act selfishly in

For situations in which the selfish equilibrium is played, two additional observations immediately follow. First, since socially responsible consumers do not induce a price premium for the clean product, their presence provides no incentive to develop cleaner technologies. Second, when observers see that the clean product is no more successful than the dirty one, they may naturally — but falsely — conclude that consumers are selfish. They may then, for example, underestimate support for policies to mitigate externalities. Furthermore, all these observations apply in weaker form to other equilibria as well. Even in non-selfish equilibria, dampening implies that individuals' choice between the products may provide a lower bound on how much they care. This reduces firms' incentive to innovate, and implies that observers may underestimate consumers' social concerns.

We now turn from the logic of selfish equilibria to the possibility of good equilibria.

Proposition 8 (Consensus Boost)**.** *For any* $\underline{f} > 0$*, there is a* $\overline{MS} < 1$ *such that for any* F *with* $\inf_{k \in [\underline{k}, \overline{k}]} f(k) \geq \underline{f}$*, the following holds.*

I. If $\underline{k}(e_d - e_c) > 1/s_c$*, then a competitive equilibrium with* $q_c^* = 1$ *exists.*

II. If $\underline{k}(e_d - e_c) < 1/s_c$*, then any competitive equilibrium has* $q_c^* \leq \overline{MS}$*.*

Proposition 8 identifies a condition under which a competitive equilibrium with everyone buying the clean product exists; but it also says that if the condition is not satisfied, the market share of the clean product is discontinuously lower. To understand the statements more precisely, note that $1/s_c$ is the (relative) price at which clean producers supply the entire market. Hence, $\underline{k}(e_d - e_c) > 1/s_c$ means that everyone is willing to pay this necessary price to do their part in lowering the externality. Part I says that when such a consensus is in place, an equilibrium implementing it exists. Intuitively, if a consumer assumes that dampening is zero, then she prefers the clean product for prices near $1/s_c$. Hence, she cannot be induced to buy the dirty product, and dampening is indeed zero. On the other hand, Part II says that if the consensus is violated to an arbitrarily small extent, then a non-trivially worse equilibrium obtains; indeed, this could be the selfish equilibrium. Without a consensus, dampening necessarily comes into play, lowering consumers' motivation to act responsibly. In this sense, a full consensus in all or part of the population that investing in

"proposer competition," this is because they are in a disadvantageous position where they are unwilling to sacrifice for others.

mitigation is privately worth it can provide a discrete boost to mitigation efforts.[15]

Three factors determine whether the beneficial consensus in the proposition obtains. The first is the responsibility of the population, specifically that of the least responsible consumer. Hence, slight increases in responsibility can substantially improve outcomes, but slight decreases can lead to an unraveling of the good equilibrium. The latter could occur, for instance, due to a disinformation campaign that the costs of climate change are not high, an opening of the product market to the global market with a different population, or a recession that raises the value of money relative to that of the externality. The second factor is the externality advantage of the clean product ($e_d - e_c$). Hence, increasing the environmental friendliness of the clean product can make a big difference, but doing the same for the dirty product can backfire. And the third factor is the production efficiency of the clean product, as measured by s_c. If suppliers can easily expand production, then the price premium necessary for the clean product to dominate the market is low, so consumers can more easily agree that the clean product is worth it.

As in the previous section, our results help clarify a potential misperception regarding the environmental impact of consumer choices. Casual observation may suggest that when the price difference between substitute products with different impacts is small or zero, a consumer can do good at little or no cost to herself. Unfortunately, a consumer is not off the hook so easily. Our model says that if the products are equally good deals, then the market is in a selfish equilibrium. Then, buying either product causes the same pollution — a weighted average of the pollution levels of the two products. In the extreme, therefore, taking the train may pollute as much as flying, and green electricity may pollute as much as brown electricity. In that case, a consumer can reduce pollution only by consuming less, which is personally costly. In a non-selfish equilibrium, a socially responsible consumer can do good by choosing the cleaner product. Again, however, this is costly for her, now because the clean product is more expensive. Indeed, the price difference is a market signal indicating that consuming the clean product is effective in reducing pollution.

Our model in this section is related to the model of consumer boycotts by Broccardo et al. (2022).

[15] Although the condition in Proposition 8 is stated in terms of a full consensus, this is not necessary for the logic to operate. For instance, if there is a positive share of selfish consumers, and the rest is described by the distribution F, a similar discontinuity result holds.

In both settings, socially responsible consumers choose between a dirty and a clean product, and the main question is how the two products fare. In their model, however, the price difference between the products is fixed. This implies that if consumers are not sufficiently responsible, a boycott cannot be sustained. But for the same reason, there is no analogue to our main results, all of which depend on the endogeneity of prices. In particular, there is no analogue to our finding that despite all consumers being socially responsible, the clean product may in the unique equilibrium not enjoy any price advantage.

5.2 Policy

Given the possibility of massive market failure, we return to analyzing the effects of policy. We start with the usual candidates. Under a tax, there is a unit tax τ on the dirty good. Under a cap, there is a fixed number of permits, and to sell a unit of the dirty good, a supplier must purchase a permit. The permit fee τ is determined in equilibrium. Analogously to Section 4, we consider a type of policy to be strictly better than another if it can achieve the same feasible externality — or, equivalently, the same market share for the dirty product — with less intervention (lower τ^*).

Proposition 9 (Cap versus Tax — Substitute Products)**.**

I. Under both a tax and a cap, a selfish equilibrium (Part I of Proposition 7) exists. In such an equilibrium, the two policies are equivalent: they implement a given externality g with the same τ^*.

II. Suppose that the best competitive equilibrium is always selected, and that with no intervention, there is a non-selfish equilibrium. To achieve an externality g that is strictly lower than without intervention, a tax is strictly better than a cap.

Part I considers selfish equilibria. An intervention does not change the logic of this equilibrium — if a consumer expects others to choose the cheaper product, dampening is full, so she herself wants to choose the cheaper product — so the equilibrium still exists. Furthermore, with all consumers acting selfishly, the logic of Weitzman (1974) applies: a planner can achieve the same outcomes with a tax and a cap.

Part II analyzes policies when with no intervention a non-selfish equilibrium exists, and such equilibria tend to be selected by the market. For presentational simplicity, we assume here that the best equilibrium is selected. Then, the logic of Herweg and Schmidt (2022) holds, so that a unit tax is better than a cap. Under a binding cap, consumers realize that they have no effect on the externality, and hence are indifferent between the products. As a result, the equilibrium collapses to a selfish equilibrium. With a tax, in contrast, a non-selfish equilibrium still exists, so consumers are willing to reduce consumption to mitigate.

But there are better policies. We compare a unit tax on the dirty product to a unit subsidy on the clean product, as well as a proportional tax on the dirty product and a proportional subsidy on the clean product. Under a proportional policy, a tax or subsidy is levied in proportion to the price: if the consumer price is p, the price the producer is paid equals $p(1-\tau_1)$. A policy is strictly better than another if it achieves the same market shares with a lower fee (τ^* or $|\tau_1^* p^*|$).

Proposition 10 (Better Policies — Substitute Products)**.** *Suppose that the best competitive equilibrium is always selected, and that with no intervention, there is a non-selfish equilibrium. To achieve an externality g that is strictly lower than without intervention, a proportional subsidy on the clean product is strictly better than other unit or proportional policies.*

A proportional subsidy on the clean product is the best policy. Intuitively, if a consumer switches to the clean product, she raises its price. Because the product enjoys a proportional subsidy, this increases the subsidy producers receive, boosting supply. As a result, supply plays a greater role in responding to the consumer's demand, lowering dampening.

6 Conclusion

Our paper introduces a novel framework for thinking about behavior, equilibrium outcomes, welfare, and policy in large markets with rational consequentialist socially responsible consumers. The key behind our notion of competitive equilibrium is that a consumer recognizes and takes into account the degree of dampening in the market. Beyond the questions we have analyzed, our equilibrium framework is portable to any of the myriad other issues that arise with externalities. To analyze a new situation, one can follow a recipe: (1) use standard price-theoretic (supply-demand) analysis to derive the degree of dampening from properties of the situation; and (2) incorporate the degree of dampening into consumer optimization when defining competitive equilibrium.

We mention two simple potential questions of interest. In Section 4, we have studied the effects of policy under free trade. But the European Union recently enacted the Carbon Border Adjustment Mechanism (or CBAM, colloquially known as a carbon tariff). Under this system, an importer wishing to sell in the domestic market must buy a "CBAM certificate" for the pollution caused abroad, paying a price equal to that of a domestic permit. Incorporating this consideration into our model affects the thinking of socially responsible consumers, and we conjecture that the implications now depend on other aspects of the economy. And in Section 5, we have studied a situation with two substitute products. There may, however, be many other alternatives in-between that are also substitutes in consumption utility, but generate a range of externalities. We conjecture that all products must then trade at the same price, as a price difference between products with slightly different pollution levels cannot be supported.

As a methodologically natural and empirically plausible first pass, our paper assumes that consumers are rational — they understand the net effects of their consumption choices. This approach is supported by evidence that many consumers understand the main mechanism (the dampening effect of markets) driving our results. Nevertheless, many consumers may not fully appreciate how the effects of their actions wind through the economy. For example, some consumers may evaluate their choices by (explicitly or implicitly) assuming a dampening of zero or one. How consumers think about dampening is a natural question for future empirical research, and how potential bounded rationality affects behavior and market outcomes is a natural question for future theoretical research.

Finally, while we have assumed that consumers are consequentialists, they may also have other types of social motivations. In a type of "warm glow" (Andreoni, 1990), for example, a consumer may care about her direct contribution to reducing the externality. In our framework, this is equivalent to a consequentialist consumer who assumes that dampening is zero. Alternatively, a consumer may attempt to adhere to social norms, a consideration that does not readily fit in the consequentialist framework. Furthermore, different motives can interact in non-trivial ways. For a rational consequentialist consumer, for example, the presence of others who adhere to social norms may lead to a multiplier effect rather than a dampening effect. If she does more to mitigate the externality, she might change the social norm, encouraging others to do more as well.

References

Aghion, Philippe, Roland Bénabou, Ralf Martin, and Alexandra Roulet, "Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?," *American Economic Review: Insights*, 2023, *5* (1), 1–20.

Andre, Peter, Teodora Boneva, Felix Chopra, and Armin Falk, "Misperceived Social Norms and Willingness to Act Against Climate Change," 2022. Working Paper.

Andreoni, James, "Privately provided public goods in a large economy: the limits of altruism," *Journal of public Economics*, 1988, *35* (1), 57–73. Publisher: Elsevier.

_ , "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving," *Economic Journal*, 1990, *100* (401), 464–477.

Arnold, Lutz G., "On the Neutrality of Socially Responsible Investing," *Theoretical Economics*, 2023, *18* (1), 65–95.

Arrow, Kenneth J. and Frank Hahn, *General Competitive Analysis*, Holden-Day, 1971.

Auger, Pat, Paul Burke, Timothy M. Devinney, and Jordan J. Louviere, "What Will Consumers Pay for Social Product Features?," *Journal of Business Ethics*, February 2003, *42* (3), 281–304.

_ , Timothy M. Devinney, Jordan J. Louviere, and Paul F. Burke, "Do social product features have value to consumers?," *International Journal of Research in Marketing*, September 2008, *25* (3), 183–191.

Bergstrom, Theodore, Lawrence Blume, and Hal Varian, "On the private provision of public goods," *Journal of Public Economics*, February 1986, *29* (1), 25–49.

Broccardo, Eleonora, Oliver Hart, and Luigi Zingales, "Exit versus Voice," *Journal of Political Economy*, December 2022, *130* (12), 3101–3145.

Chater, Nick and George Loewenstein, "The I-Frame and the S-Frame: How Focusing on Individual-Level Solutions Has Led Behavioral Public Policy Astray," *Behavioral and Brain Sciences*, forthcoming.

Debreu, Gerard, *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*, Yale University Press, 1959.

Deck, Cary and Harris Schlesinger, "Consistency of Higher Order Risk Preferences," *Econometrica*, 2014, *82* (5), 1913–1943.

Dewatripont, Mathias and Jean Tirole, "The Morality of Markets," 2022. Working Paper.

Dietz, Thomas and Janina Grabs, "Additionality and Implementation Gaps in Voluntary Sustainability Standards," *New Political Economy*, March 2022, *27* (2), 203–224.

Dufwenberg, Martin, Paul Heidhues, Georg Kirchsteiger, Frank Riedel, and Joel Sobel, "Other-Regarding Preferences in General Equilibrium," *Review of Economic Studies*, 2011, *78* (2), 613–639.

Falk, Armin and Nora Szech, "Morals and Markets," *Science*, 2013, *340* (6133), 707–711.

_ , Thomas Neuber, and Nora Szech, "Diffusion of Being Pivotal and Immoral Outcomes," *Review of Economic Studies*, 2020, *87* (5), 2205–2229.

Fehr, Ernst and Klaus M. Schmidt, "A Theory of Fairness, Competition, and Cooperation," *Quarterly Journal of Economics*, 1999, *114* (3), 817–868.

Fischbacher, Urs, Simon Gächter, and Ernst Fehr, "Are people conditionally cooperative? Evidence from a public goods experiment," *Economics Letters*, June 2001, *71* (3), 397–404.

Fraser, Alastair and Tanya Fiedler, "Net-zero targets for investment portfolios: an analysis of financed emissions metrics," Technical Report 2022.

Gaube, Thomas, "Altruism and charitable giving in a fully replicated economy," *Journal of Public Economics*, September 2006, *90* (8), 1649–1667.

Giesler, Markus and Ela Veresiu, "Creating the Responsible Consumer: Moralistic Governance Regimes and Consumer Subjectivity," *Journal of Consumer Research*, October 2014, *41* (3), 840–857.

Green, Daniel and Benjamin Roth, "The Allocation of Socially Responsible Capital," 2021. Working Paper.

Gruber, Jonathan, *Public Finance and Public Policy*, first ed., Worth Publishers, 2005.

Hainmueller, Jens, Michael J. Hiscox, and Sandra Sequeira, "Consumer Demand for Fair Trade: Evidence from a Multistore Field Experiment," *Review of Economics and Statistics*, 2015, *97* (2), 242–256.

Hakenes, Hendrik and Eva Schliephake, "Responsible Investment and Responsible Consumption," 2021. Working Paper.

Hertel, Shareen, Lyle Scruggs, and C. Patrick Heidkamp, "Human rights and public opinion: From attitudes to action," *Political Science Quarterly*, 2009, *124* (3), 443–459. Publisher: JSTOR.

Herweg, Fabian and Klaus M. Schmidt, "How to Regulate Carbon Emissions with Climate-Conscious Consumers," *Economic Journal*, 2022, *132* (648), 2992–3019.

Imai, Taisuke, Davide D. Pace, Peter Schwardmann, and Joel J. van der Weele, "Correcting Consumer Misperceptions About CO $_2$ Emissions," Technical Report 2022.

Krahnen, Jan Pieter, Jörg Rocholl, and Marcel Thum, "A primer on green finance: From wishful thinking to marginal impact," Working Paper 87, SAFE White Paper 2021.

Kyle, Albert S., "Informed Speculation with Imperfect Competition," *Review of Economic Studies*, 1989, *56* (3), 317–355.

Loureiro, Maria L. and Justus Lotade, "Do fair trade and eco-labels in coffee wake up the consumer conscience?," *Ecological Economics*, April 2005, *53* (1), 129–138.

Mas-Colell, Andreu, "Noncooperative Approaches to the Theory of Perfect Competition: Presentation," *Journal of Economic Theory*, 1980, *22* (2), 121–135.

_ , "The Cournotian Foundations of Walrasian Equilibrium Theory: An Exposition of Recent Theory," in Werner Hildenbrand, ed., *Advances in Economic Theory*, Cambridge University Press, 1983, pp. 183–224.

Moisson, Paul-Henri, "Ethics and impact investment," *Toulouse School of Economics Working Paper*, 2020.

Noussair, Charles N., Stefan T. Trautmann, and Gijs van de Kuilen, "Higher Order Risk Attitudes, Demographics, and Financial Decisions," *The Review of Economic Studies*, 2014, *81* (1), 325–355.

Oehmke, Martin and Marcus M. Opp, "A Theory of Socially Responsible Investment," 2022. Working Paper.

Parker, Jonathan A. and Bruce Preston, "Precautionary Saving and Consumption Fluctuations," *American Economic Review*, 2005, *95* (4), 1119–1143.

Perino, Grischa, "Climate Campaigns, Cap and Trade, and Carbon Leakage: Why Trying to Reduce Your Carbon Footprint Can Harm the Climate," *Journal of the Association of Environmental and Resource Economists*, September 2015, *2* (3), 469–495. Publisher: The University of Chicago Press.

Piccolo, Alessio, Jan Schneemeier, and Michele Bisceglia, "Externalities of Responsible Investments," August 2022.

Rabin, Matthew, "Psychology and Economics," *Journal of Economic Literature*, 1998, *36* (1), 11–46.

_ , "An Approach to Incorporating Psychology into Economics," *The American Economic Review*, May 2013, *103* (3), 617–622.

Roberts, Donald John and Andrew Postlewaite, "The Incentives for Price-Taking Behavior in Large Exchange Economies," *Econometrica*, 1976, *44* (1), 115–127.

Rostek, Marzena and Ji Hee Yoon, "Equilibrium Theory of Financial Markets: Recent Developments," 2020. Working Paper.

Shubik, Martin, "Commodity Money, Oligopoly, Credit and Bankruptcy in a General Equilibrium Model," *Economic Inquiry*, 1973, *11* (1), 24–38.

Sobel, Joel, "Do markets make people selfish," Technical Report 2007. Publisher: Citeseer.

Sugden, Robert, "On the Economics of Philanthropy," *The Economic Journal*, June 1982, *92* (366), 341–350.

_ , "Consistent conjectures and voluntary contributions to public goods: why the conventional theory does not work," *Journal of Public Economics*, June 1985, *27* (1), 117–124.

Ventura, Luigi and Joseph G. Eisenhauer, "Prudence and Precautionary Saving," *Journal of Economics and Finance*, 2006, *30* (2), 155–168.

Weitzman, Martin L., "Prices vs. Quantities," *Review of Economic Studies*, 1974, *41* (4), 477–491.

Ziegler, Andreas, Giorgia Romagnoli, and Theo Offerman, "Morals in multi-unit markets," Technical Report 2022.

Ľuboš Pástor, Robert F. Stambaugh, and Lucian A. Taylor, "Sustainable investing in equilibrium," *Journal of Financial Economics*, November 2021, *142* (2), 550–571.