Understanding Markets with Socially Responsible Consumers*

Marc Kaufmann
Central European University

Botond Kőszegi
briq

May 13, 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, Oliver Hart, Mats Köster, Mechthild von Knobelsdorff, Andreas Lange, Stephan Lauermann, Tímea Molnár, Klaus Schmidt, Fred Schroyen, Christoph Semken, Philipp Strack, Paul Voss, as well as seminar audiences 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 do not care about the externalities they cause. But many consumers do care: 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 markets with such "socially responsible consumers." We focus in particular on rational consumers.

Departing from previous research investigating social preferences in large product markets, we first 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 fully "internalize the externality" — they put the same weight on it as the social planner — they consume too much in any competitive equilibrium, and they 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 output in the market, and $k \geq 0$ is a constant. To understand how the consumer behaves when she is small relative to the market, we introduce identical copies of the other participants. There are I other consumers and I 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.

Contradicting 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 thus 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.

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, social welfare equals total consumption utility minus production costs and K times the externality g.

In Section 3, we show that the market fails not only when consumers do not care about the externality they cause, but also when they do care. As an illustrative if extreme case, suppose that all consumers have $k = K$. This means that a consumer and the planner assign the same weight to the externality; e.g., they are willing to pay the same amount out of her 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., campfires with wood from one's garden), fully socially responsible consumers bring forth the socially optimal outcome.

The above market failures imply that dealing with externalities cannot be based solely on consumer responsibility. In Section 4, therefore, we analyze interventions that improve outcomes through the introduction of a permit fee τ for production. We define a policy as superior to another if it can achieve a greater decrease in the externality with the same τ; incidentally, a better policy also leads to a lower consumer price. While for selfish consumers ($k = 0$) different policies are equivalent (as in Weitzman, 1974), the same is not true for socially responsible consumers. A policy that induces lower dampening is better at motivating consumers, so it is superior.

As a starting point, we confirm and generalize in our setting the main insight of Herweg and

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

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, 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 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, trade facilitates leakage, whereby production moves abroad to avoid regulation. If the foreign source is large or much more polluting than domestic producers, leakage implies that socially responsible consumers are more motivated to cut consumption under a cap. When a consumer reduces her consumption in this situation, the dirty foreign producers respond, making her sacrifice effective. This insight is also relevant for policy debates regarding the effectiveness of abatement efforts under a cap.

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 always choose the cheaper product, 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 chooses the cheaper product herself.

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 for a moment that consumers expect dampening to be weak, and consider the consumer who is indifferent between the two products. Since the products generate similar externalities, a small price change implies a large change in the indifferent consumer's social concern k. This implies that the price sensitivity of demand is quite high, so dampening must be non-trivial. Net of dampening, then, the products generate even more similar externalities. As a result, a small price change implies an even greater change in the indifferent consumer's k. In this fashion, the equilibrium unravels.

The existence and potential uniqueness of the selfish equilibrium are dramatic manifestations 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 to papers cited above) with our casual observation that these concerns are not translated into more ethical market outcomes.

If the dirty good is sufficiently dirtier and other conditions hold, 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. Unless the market for the dirty product shuts down completely, however, dampening operates even in such a non-selfish equilibrium, so the clean product's market share still does not reflect consumers' preferences.

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) operates, so the two policies are equivalent. But in a non-selfish equilibrium, the insight of Herweg and Schmidt (2022) applies, so 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 or proportional tax on the dirty product as well as a unit subsidy on the clean product. Intuitively, when a consumer switches to the clean product, she raises its price, which under a proportional subsidy encourages an especially large supply response. Consequently, dampening is not as strong.

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), and is roughly consistent several observations. 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 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 and social impact investment, asking whether one's action creates additional benefits or just replaces others' actions (Perino, 2015, Green and Roth, 2021, Krahnen et al., 2021, Dietz and Grabs, 2022, Fraser and Fiedler, 2022, Oehmke and Opp, 2022). Studying the equilibrium implications of such considerations 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, some of our results may be helpful in correcting misperceptions about the environmental impact of consumption choices. 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 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 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. 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 work, however, recognizes the violation of price taking for socially responsible consumers, and no other work incorporates the resulting dampening into a model of consumer behavior and competitive equilibrium in a standard product market.

There is a large literature on how to think about equilibria in markets with many consumers. A long-standing axiom is that small consumers take the market price as fixed (Debreu, 1959, Arrow and Hahn, 1971), and extensive research establishes conditions under which this is approximately optimal (e.g., Mas-Colell, 1980). Because market behavior 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 sometimes 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). We show that price taking does not hold for our socially responsible consumers, and we define a variant of competitive equilibrium that appropriately accounts for such consumers' incentives.

Starting from Heinkel et al. (2001), a few researchers have pointed out that price changes induced by unselfish behavior can alter the benefits of that behavior. Our dampening effect derives from the same mechanism — which, unlike most previous research, we establish applies to vanishingly small consumers. Although they do not frame it as a violation of price taking and show dampening in a different form, this starting point parallels analyses of small agents by Broccardo et al. (2022) and especially Norwood and Lusk (2011, Chapter 8). Norwood and Lusk heuristically derive the dampening effect in the context of advice about animal-product purchases to consumers who care about animal welfare. Broccardo et al. study the equilibrium behavior of small investors who take into account their price-mediated impacts. The authors show that because exit (i.e., divestment or boycotts) has a non-trivial private cost, it is worse at aligning the private and social incentives of socially responsible investors than shareholder voting. A boycotting 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.[2]

Other research exploring the above types of price-mediated effects focuses on 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). Because we consider the consumer side and analyze competitive equilibria, our results are completely different.

Many researchers have investigated the effectiveness of different policies to regulate externalities. We contribute to this literature by studying policies in markets with socially responsible consumers. Most related, Herweg and Schmidt (2022) show that a tax dominates a cap in a closed economy with

[2] At a 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 are based on standard notions of a Walrasian auctioneer, who asks for demand schedules to find the equilibrium price.

such consumers. In our terminology, this is because a tax generates a lower degree of dampening. Herweg and Schmidt capture the difference through an exogenous parameter (β^R) that measures consumers' motives to mitigate, 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

We develop our equilibrium framework in two steps. First, we analyze the behavior of a single socially responsible consumer. Then, we lay out what this implies for competitive equilibrium.

2.1 A Single Consumer's Perspective

Setup We consider a single-good market, and study a consumer who is tiny relative to the market. To do so, we use a "replicator economy" (Shubik, 1973, Roberts and Postlewaite, 1976): we introduce identical copies of the other participants, and let the number of copies approach infinity. We describe this rest of the market first.

There are I other consumers and I suppliers. The other 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 is a price $p^* > 0$ for which $S(p^*) = D(p^*)$, and $\lim_{p\to\infty} S(p) - D(p) = \infty$.

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 $p(c) > 0$ is chosen to clear the market: $c + ID(p(c)) = IS(p(c))$; $p(c)$ clearly exists and is unique. Finally, outcomes are realized.

The consumer correctly predicts the above outcomes, and maximizes

$$u(c) - p(c)c - kg(c), \tag{1}$$

where $u(\cdot)$ is a thrice differentiable strictly concave utility function, $g(c)$ is an externality equal to the total quantity $q(c)$ of the good produced and consumed in the market, and $k \geq 0$ is a constant.[3] The term $u(c)$ is the person's consumption utility, $u(c) - p(c)c$ is her private utility, and the term $-kg(c)$ captures her concern regarding the externality. While this means that the consumer internalizes the harm caused by everyone, an equivalent formulation is that she cares about her own impact on the externality. Indeed, since $g(0)$ is a constant, the utility function as $u(c) - p(c)c - k\big(g(c) - g(0)\big)$ yields the same behavior as (1). To ensure that the equilibrium we define below features positive consumption, we impose that $u'(0) > k$ and $\lim_{c\to\infty} u'(c) \leq 0$. For simplicity, we also 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.[4]

Analysis Differentiating the market-clearing condition with respect to c and rearranging gives

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

The total quantity produced is $q(c) = IS(p(c))$, on which the consumer's demand has an effect of

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

Taking limits yields:

Proposition 1 (Violation of Price Taking)**.** *Take any $D(\cdot)$ and $S(\cdot)$.*

I. A vanishingly small consumer has a negligible impact on the price: for any c,

$$\lim_{I\to\infty} p(c) = p^* \quad \text{and} \quad \lim_{I\to\infty} p'(c) = 0.$$

II. The same consumer has a non-negligible impact on others' consumption: for any c,

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

[3] Hence, in this model the notation g is redundant. 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.

[4] To simplify matters, we assume that the consumer's utility is linear in both money and the externality. This could lead to the non-existence of a utility-maximizing choice if the consumer could exchange directly between the two, e.g., if she had access to offsets. To analyze such situations, it is necessary to adjust the utility function slightly, for instance by assuming non-linear utility in money.

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 "market responsiveness" $q_c \equiv \lim_{I\to\infty} q'(c)$ to a vanishingly small consumer's demand is less than one, which means that she impacts others' consumption to a non-vanishing extent. In particular, others have a dampening effect of

$$\lim_{I\to\infty}\left|\frac{\mathrm{d}[ID(p(c))]}{\mathrm{d}c}\right| = \lim_{I\to\infty}\left[-ID'(p(c))p'(c)\right] = \frac{-D'(p^*)}{S'(p^*)-D'(p^*)} = 1-q_c. \tag{3}$$

Intuitively, the consumer's demand c raises the price, which leads others to consume less. Although the price impact vanishes at rate $1/I$, there are I other consumers, so their total response is comparable to c and hence relevant for a socially responsible consumer. In the sense of having to consider the implications of her effect on the price, therefore, the consumer is not a price taker.

By Equation (3), the degree of dampening is an increasing function of the relative responsivenesses or elasticities of demand and supply, $-D'(p^*)/S'(p^*) = (-D'(p^*)/D(p^*))/(S'(p^*)/S(p^*))$. This statistic determines the extent to which each side responds to the price increase resulting from an increase in the consumer's demand. If the relative elasticity is low, then it is mostly supply that responds, so the dampening effect of other consumers is small. An example is when suppliers can easily increase production because their technology approximates constant returns to scale. But if the relative elasticity is high, then it is mostly demand from other consumers that responds to price increases, so the dampening effect is large. An example is when producers cannot flexibly expand production due to capacity or input constraints.

Based on the above, an infinitesimally small consumer chooses c to solve[5]

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

Two immediate points follow. First, because the consumer has an impact on the externality ($q_c > 0$) and she derives disutility from this, she consumes less than a selfish consumer ($k = 0$). In contrast, some previous models predict that in the market, an agent with social preferences chooses

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

her consumption selfishly (Dufwenberg et al., 2011, Arnold, 2023). The difference arises because these models impose price taking, and assume that an agent is unwilling to pay to mitigate the external effect of her consumption, i.e., she is not socially responsible by our definition.

Second, however, dampening ($q_c < 1$) implies that the consumer consumes more than she would if c came from home production with the same private marginal cost, p^*. With home production, the consumer would not affect others' consumption, so her marginal utility would be $u'(c) - p^* - k$. 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 existing 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.

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 behave like the consumer 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, we make the simplifying assumption that the supply curve $S(\cdot)$ is exogenous and linear: $S(p) = sp$, with $s > 0$.

Definition 1. A competitive equilibrium is a quantity $q^* > 0$, price $p^* > 0$, consumer price responsiveness $q_p^* < 0$, and market responsiveness $q_c^* > 0$ 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: $q_p^* = 1/u''(q^*)$.

Our notion of competitive equilibrium is based on two classical ingredients, the quantity q^* and the 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 that supply equals the equilibrium quantity. Condition 2 says that demand, determined by plugging the equilibrium market responsiveness into a consumer's first-order condition (4), equals the equilibrium quantity. Condition 3 applies Equation (2) to formulate how the market responsiveness is determined, taking into account dampening. Finally, Condition 4 requires that consumers' price responsiveness is consistent with optimization. To obtain the condition, we totally differentiate Condition 2 and impose that a consumer treats the market responsiveness as constant within the range of an infinitesimal price change.[6]

In Appendix A, we provide some foundational analysis for our equilibrium concept. First, we outline a way to think about equilibrium determination graphically. Second, more importantly, we develop game-theoretic microfoundations for our concept, showing that it arises as a limit case when each consumer becomes infinitesimally small. Our analysis uses methods Kyle (1989) and the subsequent literature have developed to model the behavior of financial-market participants with non-trivial price impacts (see Rostek and Yoon, 2020, for a review). In these games, each player submits not only her demand, but also her price sensitivity of demand, to the market. Analogously to Section 2.1, others' price responsivenesses determine how a 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 solve the 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 requires quadratic utility (Rostek and Yoon, 2020), so it

[6] Our concept of competitive equilibrium collapses to the standard concept 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 unnecessary. Furthermore, since the consumer price responsiveness q_p^* is only used to define q_c^*, it, and Condition 4 on it, are also unnecessary.

yields a microfoundation 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 $u(\cdot)$.

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

$$u(q) - \int_0^q S^{-1}(x)\mathrm{d}x - Kq. \tag{5}$$

One part of social welfare is total consumption utility net of the costs of production. This is the sum of consumers' private utilities $u(q) - pq$ and the producer surplus. In addition, the social planner puts an exogenously given weight $K > 0$ on the externality. A natural assumption is $k \leq K$: although socially responsible consumers care about the disutility others suffer from the externality, most do not fully internalize this. In Appendix A.4, we motivate these assumptions in more detail, and clarify the specification of the consumer's utility function (1) as well.

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. A fundamental Econ-1 lesson is that when there are no externalities or other frictions, markets perform efficiently despite everyone favoring their own private consumption. Another fundamental Econ-1 lesson is 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. One would then 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 not only to the presence of externalities, but also to the market. To see this formally, consider the following modification of our model. We let the number of individuals $I + 1$ be finite. We assume that there are no suppliers, and c_i is a good, such as how much of the wood growing in her garden consumer i burns, that is produced and consumed at home. Consumer i's utility is $u_i(c_i) - kg$, where $u_i(\cdot)$ satisfies the same conditions as $u(\cdot)$ above, and $g = \sum_{i'} c_{i'}$ as before. Social welfare equals $[\sum_i u_i(c_i)] - Kg$. Then, the following is obvious:

Observation 2. *With non-market consumption, there is a unique socially optimal consumption profile* $(c_1^{\mathrm{FB}}, \ldots, c_{I+1}^{\mathrm{FB}})$. *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.

In addition to its message for economics, 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, voluntary approach. 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.

[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.

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 to affect change as consumers and investors 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 such an approach is that many consumers do not care sufficiently and will therefore free-ride. 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 resulting competitive equilibrium (q^*, p^*, q_p^*, q_c^*). If $u'''(q^*)$ is sufficiently high, then there are competitive equilibria (i) q^+, p^+, q_p^+, and q_c^+ with $q^+ > q^*$, $p^+ > p^*$, $|q_p^+| > |q_p^*|$, $q_c^+ < q_c^*$; and (ii) 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.*

The condition for multiple equilibria to arise is that $u'''(c)$ is high over a range, so that consumer price responsiveness $(-1/u''(c))$ increases fast. This means that an individual consumer's demand curve is quite convex. As an intuitive example, the consumer's utility function for air travel may decrease sharply below, but flatten out quickly above, one flight. Then, if a consumer expects everyone to take one flight, she expects demand to be price-insensitive, so by Equation (3) she concludes that dampening is low. Consistent with equilibrium, therefore, the consumer is motivated to mitigate and takes one flight. But if the consumer expects everyone to take multiple flights, she expects demand to be price-sensitive and hence dampening to be high. Again consistent with equilibrium, therefore, she is not motivated to mitigate and takes multiple flights.[9]

[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). For instance, 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."

[9] The prediction of our model that a person's pro-social behavior depends on what she thinks others will do, and

While we have not developed formal selection criteria, informal arguments based on path dependence suggest that in stable supply conditions, the worst equilibrium is the most likely outcome. The detrimental effects of many externalities, such as global warming or air pollution, have not been appreciated until recently. Therefore, the market has 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.

At the same time, temporary shifts in supply may permanently shift the equilibrium. Suppose, for example, that an oil shock increases the price of air travel to a level where even selfish consumers would take just one flight. This reduces dampening, so as the shock dissipates and prices drop, consumers may naturally stay in the low-consumption equilibrium.

4 Policy

We have shown that even extremely responsible consumers overconsume in equilibrium. We therefore analyze the effectiveness of different policies in mitigating this market failure. We focus on market-based approaches in which producers must purchase a permit to manufacture a unit of the good, and these permits are available at a fee τ. Proceeds from the permits are lump-sum-redistributed to consumers.

4.1 Intervention and Definition of Competitive Equilibrium

We first consider policies under which the planner acts as a supplier of permits. Such a "permit-supply policy" is described by the curve $\pi g - (1 - \pi)\tau + \pi_0 = 0$, where $g \geq 0$ is the amount of available permits, $\tau \geq 0$ is the fee for a permit, and $\pi \in [0, 1]$ and $\pi_0 \in \mathbb{R}$ are constants that parameterize the policy. The lower is π, the more permit supply responds to changes in the permit

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 reason for conditionally cooperative behavior in markets.

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. With the planner willing to supply any number of permits at the same price, 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

$$\pi S(p-\tau) - (1-\pi)\tau + \pi_0 = 0. \tag{6}$$

Consider a $p^* > 0$ and $\tau^* > 0$ with $p^* - \tau^* > 0$ that solve (6). Then, for any p in the neighborhood of p^*, there is a unique τ that solves (6). 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 (6) with respect to p and using that $S(p-\tau) = s(p-\tau)$ yields $\mathrm{d}\tau/\mathrm{d}p = \pi s/((1-\pi)+\pi s)$, so

$$S'_{\text{net}}(p) = S'(p-\tau) \cdot \left(1 - \frac{\mathrm{d}\tau}{\mathrm{d}p}\right) = \frac{(1-\pi)s}{(1-\pi)+\pi s} \equiv s_{\text{net}}. \tag{7}$$

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

Definition 2. A competitive equilibrium with a permit-supply policy is a quantity $q^* > 0$, consumer price $p^* > 0$, permit fee $\tau^* > 0$, 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 (7).

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 (7) into Equation (2).

We consider situations in which the policymaker is constrained from implementing the socially optimal outcome, so that the equilibrium quantity necessarily remains too high. Such a constraint could, for instance, come from prohibitive enforcement or political costs associated with very high permit fees. In this constrained environment, we think of a policy type as superior to another if it can achieve a lower equilibrium externality level $g^* = q^*$ with the same degree of intervention. More precisely, we say that policy type A (e.g., a tax) is strictly better than policy type B (e.g., a cap) if for any $q_B^*, \tau^* > 0$ with $u'(0) > k + \tau^*$ that is part of an equilibrium with a B-type policy, there is an A-type policy and a corresponding equilibrium with permit fee τ^* and quantity $q_A^* < q_B^*$.[10] Since in equilibrium the market must clear (i.e., $q_i^* = S(p_i^* - \tau^*)$ for $i = A, B$), this also means that the A-type policy leads to a lower consumer price ($p_A^* < p_B^*$). Beyond having a lower externality, a lower consumer price might be socially preferable because consumer surplus is more important than producer surplus. Analogously, we say that the two types of policy are equivalent if they can achieve the exact same outcomes (i.e., the above definition holds with $q_A^* = q_B^*$ rather than $q_A^* < q_B^*$).

4.2 Comparison of Policies

Principle As in the classical result by Weitzman (1974), different policies are equivalent with selfish consumers ($k = 0$).[11] 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 better outcomes. To understand how policies rank, therefore, it is sufficient to understand how they impact dampening or the complementary notion of market responsiveness. We explain our results in terms of this principle.

[10] We impose the condition $u'(0) > k + \tau^*$ to guarantee that with both types of policies, an equilibrium (defined to feature positive consumption) exists with fee τ^*.

[11] To see this, consider any demand curve $D(p)$ and supply curve $S(p)$ satisfying the conditions in Section 2. In equilibrium $q^* = S(p^* - \tau^*) = D(p^*)$. Hence, there is a unique pair p^*, q^* that is consistent with a given τ^*.

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 ($\pi = 0$) is the best permit-supply policy, while a cap ($\pi = 1$) 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. First, we show that a regulator can do better than with a unit tax.[12] 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 ($\tau_1 = 0$), which are better than taxes that are increasing in the consumer price ($\tau_1 > 0$).*

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

While most existing tax regimes feature constant or increasing taxes, our model says that a tax that is decreasing in the consumer price is better at motivating socially responsible consumers. 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 policy applies only at home. It is generally understood that leakage results: to avoid the fee on pollution, more production locates abroad than without intervention. We analyze how this affects a socially responsible consumer's thinking and the optimal policy.

We assume that there is a single product with home and foreign producers, and (for simplicity) only home consumers. A consumer cannot distinguish the products; for instance, she does not know where her electricity is coming from.[13] Accordingly, there is a single consumer price p at which both home and foreign products are sold. The foreign supply curve, which is a function of p because there is free trade and foreign producers are not subject to regulation, equals $S^f(p) = s^f p$ with $s^f > 0$. The home supply curve, which instead is a function of $p - \tau$ because home producers must pay the permit fee τ to supply a unit of the good, equals $S^h(p-\tau) = s^h(p-\tau)$ with $s^h > 0$. Market activities generate a total externality of $g = e^h q^h + e^f q^f$, where q^h and q^f are the quantities produced at home and abroad, respectively, and $e^f > e^h > 0$. This means that foreign production is more polluting, for instance because foreign suppliers use a different technology due to the differences in environmental regulation. For simplicity in stating our results, we assume that u is quadratic, and denote its second derivative by u_{cc}. As our proof establishes, this implies that there is a unique equilibrium. Then, we say that policy type A (a cap or tax) is strictly better than policy type B (a tax or cap) if for any equilibrium g_B^*, τ^* with $u'(S^f(\tau^*)) > \tau^* + ke^f$ that obtains under a B-type policy, there is an A-type policy that in equilibrium generates the permit fee τ^* and an externality $g_A^* < g_B^*$.[14]

To adapt Definition 2 for this situation, we distinguish between the market responsivenesses of

[13] 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.

[14] We impose the condition on u' to guarantee that under both policies, home production remains positive.

home- and foreign-supplied quantities, q_c^h and q_c^f. These can be calculated similarly to q_c in the definitions above. Denoting the price responsiveness of demand by q_p, they are

	q_c^h	q_c^f
tax	$\dfrac{s^h}{s^h + s^f - q_p}$	$\dfrac{s^f}{s^h + s^f - q_p}$
cap	0	$\dfrac{s^f}{s^f - q_p}$

(8)

A consumer's effect on the externality is then $q_c^h e^h + q_c^f e^f$, and this is what she takes into account when choosing her consumption. We give the full definition of competitive equilibrium in the appendix.

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

I. If

$$-1/u_{cc} < \frac{e^f - e^h}{e^h} \cdot s^f, \tag{9}$$

then a cap is strictly better than a tax.

II. If Inequality (9) *goes strictly the other way, then a tax is strictly better than a cap.*

The market's response to a consumer's demand — and hence her motive to mitigate — differs in two ways under a tax and a cap. On the one hand, with home supply being flexible, the responsiveness of total production $(q_c^h + q_c^f)$ is greater under a tax. On the other hand, with home supply unable to respond, the responsiveness of foreign supply — the source the consumer cares more about — is greater under a cap. As a result of this second effect, a cap is better than a tax if foreign supply is sufficiently dirty, or sufficiently responsive relative to demand. The latter is especially likely to be the case when foreign supply dominates production in the market.

The logic of Proposition 6 clarifies a misperception some observers seem to have regarding a consumer's or even country's impact on emissions. We have often heard that for consumption under a cap-and-trade system, for instance buying steel in the EU, the environmental impact is zero because emissions will be at the cap anyway. As has already been understood at least for large buyers under the rubric of carbon leakage, this argument is incorrect with international trade.

Nevertheless, there is a continuing policy debate regarding the argument (e.g., Intergovernmental Panel on Climate Change, 2022, page 1396). Our analysis says 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. Furthermore, the logic extends to non-traded products, such as flying within the EU, covered by the same cap. This is because a decrease in the permit price due to the reduction in demand for the non-traded good lowers the consumer prices of traded goods, leading foreign suppliers to decrease production.

5 Substitute Products

In our basic model, a consumer can reduce the externality by consuming less. We now consider the scope for mitigating the externality through the selection of a less polluting product instead.

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 continuous density f. For presentational simplicity, we assume that f is bounded away from zero on $[\underline{k}, \overline{k}]$.[15] To isolate the new issue — which of two products to select rather than how much to consume — we impose that consumers have unit demand.[16] 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 $p^c, p^d \geq 0$ are the product prices and $s^c, s^d > 0$. Typically, one would expect $s^c > s^d$ because the clean product is more difficult to produce, but we do not impose this.

[15] The key possibilities we identify below can also occur when f approaches zero at $\underline{k}$ and $\overline{k}$, but the conditions for them are then more difficult to state. This applies to both the potential uniqueness of the selfish equilibrium (Proposition 7) and the discontinuity in outcomes at full consensus (Proposition 8).

[16] 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.

We define an equilibrium by adapting Definition 1. Crucially, there is a market responsiveness that determines how much the total quantities of the two products respond if a consumer moves her consumption from the dirty to the clean market. Each consumer's behavior must be optimal given the market responsiveness, prices, and her k. This consideration yields 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. For the formal definition, see the appendix.

To evaluate outcomes, we define social welfare as the negative of suppliers' production costs minus K times the externality g. Since consumers have unit demand and the products are perfect consumption substitutes, consumption utility is always the same and hence can be ignored when making welfare comparisons.

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 $e^d - e^c$ or s^c 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 the clean product is more expensive ($p^{c*} > p^{d*}$), and some consumers strictly prefer the clean product.

IV. Among multiple competitive equilibria, the greater is the clean product's price premium ($p^{c*} - p^{d*}$), 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 expects everyone else to behave selfishly, i.e., to always choose the cheaper product. Then, a consumer believes that the two prices will 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, re-equilibrating the markets at the previous price. The consumer concludes that the quantities in the two markets, and therefore total pollution, do not depend on which product she chooses. Consistent with equilibrium, therefore, she chooses the cheaper product herself.

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 due to a mutually reinforcing interaction with consumers' price sensitivity. Suppose, to start, that consumers expect dampening to be zero, and consider the social concern k that makes a consumer indifferent between the two products. Since the products generate similar externalities, a small change in their price difference implies a large change in this cutoff k. Hence, the price sensitivity of demand is quite high, so dampening must be non-trivial. Taking dampening into account, then, purchasing the two products generates even more similar externalities. As a result, a small price change implies an even greater change in the cutoff k. Hence, the price sensitivity of demand is even higher, etc.

A similar unraveling logic holds when s^c is low, but in that case the first step acts in part through supply. Analogously to Equation (3), a low s^c implies that when a consumer moves her demand to the clean sector, it is mostly other consumers rather than suppliers who respond to the price increase. This means that dampening 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, but dampening is not full, so consumers with a sufficiently high concern k are willing to pay the higher price to mitigate. Part IV says that this results in higher social welfare.

The existence and potential uniqueness of the selfish equilibrium are drastic manifestations 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 the clean good enjoys no advantage in the market.[17]

[17] 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 as we have mentioned, in

For situations in which the selfish equilibrium is played, two additional observations follow. First, since socially responsible consumers do not induce a price premium for the clean product, they provide 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 conclude that consumers are selfish. They may then, for example, underestimate support for policies to mitigate externalities.

Furthermore, observations about selfish equilibria apply in weaker form to non-selfish equilibria as well. Even in such equilibria, dampening implies that individuals' choice between the products may provide a lower bound on how much they care about the externality. Hence, the market partially fails in aggregating consumers' social preferences. This reduces firms' incentive to innovate, and means 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 an* $\overline{MS} < 1$ *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 there is a competitive equilibrium in which all consumers buy the clean product* ($q^{c*} = 1$)*.*

II. If $\underline{k}(e^d - e^c) < 1/s^c$*, then in any competitive equilibrium, the share of the clean product is at most* $\overline{MS}$ ($q^{c*} \leq \overline{MS}$)*.*

Proposition 8 identifies a condition under which a fully clean equilibrium exists; but it also warns that if the condition is not satisfied, then the equilibrium is discontinuously worse. To understand the statements more precisely, notice 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, then, says that when such a consensus is in place, an equilibrium implementing it exists. Intuitively, if a consumer expects dampening to be zero, then she prefers the clean product for prices near $1/s^c$, so she cannot be induced to buy the dirty product. With all consumers thinking the same, dampening is indeed zero. On the other hand, Part II says that if the consensus is violated to an arbitrarily small

previous theories consumers are not socially responsible by our definition. Similarly, Fehr and Schmidt's (1999) inequity-averse agents act selfishly in "proposer competition" because they are in a disadvantageous position where they are unwilling to sacrifice for others.

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 mitigation is privately worth it can provide a discrete boost to mitigation efforts.[18]

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 the unraveling of a good equilibrium. The latter could occur 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 insights 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 pollution levels 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, so buying either product causes the same pollution. That pollution equals 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 this case, a consumer can reduce pollution only by consuming less, which is personally costly.

In a non-selfish equilibrium, a consumer can do good by choosing the cleaner, more expensive

[18] 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 of the population is described by the distribution F, then a similar discontinuity result holds for the latter group.

product. 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). In both settings, socially responsible consumers choose between a dirty product and a clean product, and the main question is how the two products fare. In Broccardo et al.'s model, however, the price difference between the products is fixed by the positive cost firms must pay to be clean. This implies that if consumers are not sufficiently responsible, they are unwilling to pay the premium, so a boycott cannot be sustained. But the fixed price premium also means that there are no analogues to our main results, all of which depend on how demand affects 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 begin with the usual candidates, the tax and the cap. Under a tax, there is a unit tax τ^* on the dirty good. Under a cap, there is a fixed number of permits to sell the dirty good, and the permit fee τ^* is determined in equilibrium. We restrict attention to binding, strictly positive caps and taxes $\tau^* < 1/s^c$. For $\tau^* \geq 1/s^c$, the entire market would be supplied by the clean product even if consumers were selfish, so motivating socially responsible consumers would not be relevant.

We begin by noting:

Observation 4. *Under both a tax and a cap, a selfish equilibrium (Part I of Proposition 7) exists.*

Neither intervention changes the logic of the selfish equilibrium: if a consumer expects others to choose the cheaper product, she expects dampening to be full, so she herself chooses the cheaper product. Consequently, the selfish equilibrium still exists.

Because in the current model the comparison of policies depends on the equilibrium being played, we consider two opposite equilibrium-selection assumptions. In one, the selfish equilibrium is always selected, and in the other, the best equilibrium is always selected. Under either assumption, our model predicts a unique outcome both with and without intervention. To focus on situations in which policy is necessary to improve outcomes, we assume that without intervention a positive amount of the dirty good is produced ($q^{d*} > 0$). Analogously to Section 4, we consider a type of policy to be strictly better than another if it can achieve a greater improvement on the no-intervention externality level — or, equivalently, a greater decrease in the share of the dirty product — with the same intervention (same τ^*).[19] Two types of policy are equivalent if they can achieve the same improvements in externality levels with the same interventions.

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

I. If the selfish equilibrium is always selected, then a cap and a tax are equivalent.

II. Suppose that the best competitive equilibrium is always selected, and that with no intervention, there is a non-selfish equilibrium. Then, a tax is strictly better than a cap.

Part I considers selfish equilibria. 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. 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, instead, 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 a strictly lower market share for the dirty product with the same wedge between producer and consumer prices (i.e., the same 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. Then, a*

[19] That is, policy type A is strictly better than policy type B if for any dirty quantity $q_B^{d*} > 0$ and fee $\tau^* > 0$ that is part of an equilibrium with a B-type policy, there is an A-type policy that yields permit fee τ^* and quantity $q_A^{d*} < q_B^{d*}$.

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 under that policy, she raises its price and therefore also the subsidy on it. This boosts the market for the clean product, reducing dampening.

By analogous logic, a proportional tax on the dirty product is the worst policy. When a consumer switches away from the dirty product under that policy, she lowers its price and thereby also lowers the tax on it. This boosts the market for the dirty product, raising dampening. In some situations, the increase in dampening can be sufficiently large for the non-selfish equilibrium to be eliminated. Then, the policy has the unintended effect of inducing a selfish equilibrium.

6 Conclusion

Our paper introduces a novel framework for thinking about behavior, equilibrium outcomes, welfare, and policy in markets with rational consequentialist socially responsible consumers. Beyond the questions we have studied, our framework is portable to any of the myriad other issues that arise with externalities. To study 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; (2) incorporate the degree of dampening into consumer optimization when defining competitive equilibrium; and (3) investigate one's research question using the new notion of 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.

Our framework also suggests questions that are not about externalities, but require similar principles to analyze. As an example, suppose that a consumer is "socially sensitive" in that she is willing to modify her consumption to reduce inequality between two strangers, and her disutility from inequality in society is linear. This consumer violates price taking because her vanishingly small price impact can alter inequality between many others. Hence, for instance, she may go out of her way to patronize small farmers rather than large food producers, if that tilts prices in favor of the former. This is true even if a small farmer is a price taker, and hence — equating the marginal cost of expanding production with the market price of her produce — she does not benefit from an increase in purchases holding prices fixed. To analyze how such social-sensitivity motives play out in equilibrium, it is necessary to incorporate the relevant price effects into a small consumer's maximization problem.

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. But many consumers may not fully appreciate how the effects of their actions wind through the economy.[20] 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 forms of social motivations. In a type of "warm glow" (Andreoni, 1990), for example, a consumer may care about her direct contribution to reducing an externality. In our framework, this is equivalent to a consequentialist consumer who assumes that dampening is zero. Quite in contrast, a consumer may resent that others dampen her efforts to mitigate, and therefore put a greater weight on the dampening effect than on her direct effect. Alternatively, a consumer may attempt to

[20] Albeit in a completely different setting, Dal Bó et al. (2018) find that many subjects in an experimental game fail to understand (even simple) equilibrium feedback effects.

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, "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*, 2008, *25* (3), 183–191.

Bergstrom, Theodore C., "Benefit-Cost in a Benevolent Society," *American Economic Review*, 2006, *96* (1), 339–351.

Bó, Ernesto Dal, Pedro Dal Bó, and Erik Eyster, "The Demand for Bad Policy when Voters Underappreciate Equilibrium Effects," *Review of Economic Studies*, 2018, *85* (2), 964–998.

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.

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*, 2001, *71* (3), 397–404.

Fraser, Alastair and Tanya Fiedler, "Net-Zero Targets for Investment Portfolios: An Analysis of Financed Emissions Metrics," 2022. Working Paper.

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.

Heinkel, Robert, Alan Kraus, and Josef Zechner, "The Effect of Green Investment on Corporate Behavior," *Journal of Financial and Quantitative Analysis*, 2001, *36* (4), 431–449.

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.

Intergovernmental Panel on Climate Change, "Climate Change 2022: Mitigation of Climate Change," Technical Report, World Meteorological Organization and United Nations Environment Programme 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*, 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.

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

Norwood, F. Bailey and Jayson L. Lusk, *Compassion, by the Pound: The Economics of Farm Animal Welfare*, Oxford University Press, 2011.

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

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

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," 2007. Working Paper.

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," 2022. Working Paper.

A Background on Competitive Equilibrium

A.1 Graphical Illustration

Figure 1 shows how to think about equilibrium determination using modifications of a standard supply-demand diagram, and illustrates the possibility of multiple equilibria. With slight abuse of notation, we write demand and supply curves as functions of the quantity rather than of the price (as we do in the rest of the paper). We can start by drawing the exogenous and linear supply curve $S(q)$. To proceed, we start from the demand curve $D_0(q)$ that would obtain if consumers were selfish ($k = 0$). This is a standard demand curve given by $D_0(q) = u'(q)$. For simplicity in making our points, we assume that this demand curve is two-piece linear and kinked. As indicated on the figure, the intersection of the supply curve and the selfish demand curve gives the unique selfish competitive equilibrium.

Now assume $k > 0$. Based on Definition 1, the equilibrium quantity q and price p must satisfy $u'(q) - p - kq_c = 0$, or $p = u'(q) - ks/(s - 1/u''(q))$. Using that $u'(q) = D_0(q)$ and therefore $u''(q) = D_0'(q)$, we define a "virtual demand curve"

$$D_k(q) = D_0(q) - k \cdot \frac{s}{s - 1/D_0'(q)}.$$

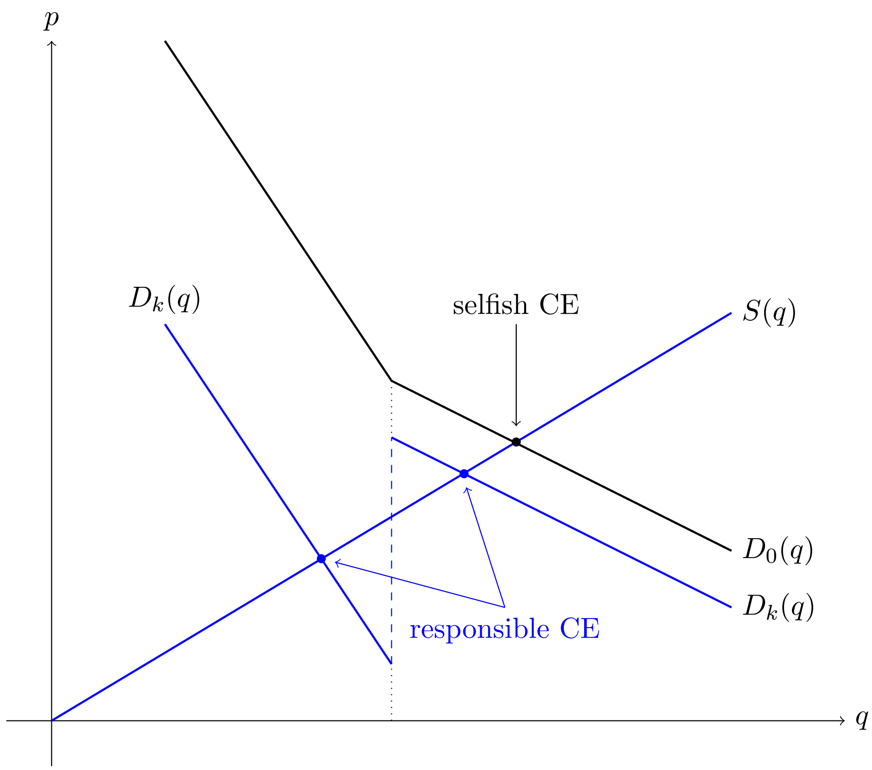

Figure 1: Illustration of Competitive Equilibrium and Multiple Equilibria

It is easy to see that intersections of this virtual demand curve with the supply curve correspond to competitive equilibria. We refer to $D_k(q)$ as virtual because it cannot be easily observed in its entirety. For instance, a common way of identifying the demand curve, shifts in supply, do not in general trace out the virtual demand curve as the curve depends on the supply curve.

Multiple competitive equilibria can occur because even for a downward-sloping selfish demand curve $D_0(q)$, the virtual demand curve $D_k(q)$ can increase. Indeed, for our kinked $D_0(q)$, $D_k(q)$ jumps up. Intuitively, at the kink the price sensitivity of demand and hence dampening jump up, so a consumer's willingness to mitigate jumps down. This results in the two competitive equilibria identified in the figure. Both feature lower consumption than the selfish competitive equilibrium, but the consumption levels are quite different.

A.2 Microfoundation I

We consider an economy with $I+1$ identical consumers and I identical suppliers. The following assumptions are made:

1. Consumer i's utility is
$$U = Ac_i - \frac{1}{2}Bc_i^2 - pc_i - k\sum_j c_j,$$
where c_i denotes consumer i's consumption and $A, B > 0$.

2. The supply of each seller is given by
$$S(p) = sp, \quad \text{s} > 0$$

3. Total supply is
$$s_0 + Isp$$
where s_0 is a non-degenerate random variable whose support includes 0.

The timing of the game is as follows. First, each consumer submits a (weakly) decreasing demand schedule $c_i(p)$. Then, given all submitted schedules, the market-clearing price is determined. Since supply is linear with a strictly positive slope and demand is weakly decreasing, the market-clearing price exists and is unique. Finally, outcomes are determined and utilities are realized.

We look for symmetric linear Bayesian Nash equilibria. In such an equilibrium, all consumers submit the same linear demand schedule $c_i(p) = a - bp$ with $b > 0$.

Proposition 11. *There is a unique symmetric linear Bayesian Nash equilibrium. In this equilibrium, each consumer chooses the schedule $c_i(p) = a - bp$, where*
$$b = \frac{-(IBs - I + 1) + \sqrt{\Delta}}{2IB}, \ \textit{with } \Delta = (IBs - I + 1)^2 + 4BI^2 s, \textit{ and}$$
$$a = \left(A - \frac{ks}{b+s}\right) \cdot \frac{1}{B + \frac{1}{I(b+s)}}.$$

We now use the analysis above to provide a microfoundation for Definition 1. Notice that in Proposition 11, the parameters describing the equilibrium, a and b, are functions of I, so that we can write them as $a(I)$ and $b(I)$. We define:

Definition 3. A limiting Kyle equilibrium is a strategy profile in which a consumer chooses $c_i(p) = a_\infty - b_\infty p$, where $a_\infty = \lim_{I\to\infty} a(I)$ and $b_\infty = \lim_{I\to\infty} b(I)$.

We also say that A and B quadratically approximate the utility function u around c_i if $B = -u''(c_i)$ and $A - Bc_i = u'(c_i)$.

Proposition 12. *The following are equivalent.*

I. The quantity $q^* > 0$, price $p^* > 0$, consumer price responsiveness $q_p^* \in \mathbb{R}$, and market responsiveness $q_c^* \in \mathbb{R}$ constitute a competitive equilibrium.

II. The pair q^*, p^* constitutes the realized outcome in the limiting Kyle equilibrium of the economy in which consumer utility is quadratically approximated around q^*, and $s_0 = 0$ (i.e., when $s_0 = 0$, the realized quantity and price are q^* and p^*). In this limiting Kyle equilibrium, consumers' price responsiveness equals $c_i'(p) = -b_\infty = 1/u''(q^*)$.

A.3 Microfoundation II

We consider the game as in Section A.2 with the following modifications. Consumer i's utility is $u(c_i)$ — as in the text, i.e., not necessarily quadratic. At the same time, a consumer can only submit a linear demand schedule, i.e., she submits $c_i(p) = a_i - b_i p$, where she can specify $a_i \in \mathbb{R}$ and $b_i > 0$. We consider symmetric pure-strategy equilibria when $s_0 = 0$, but within that class impose a kind of robustness requirement with respect to shocks to supply. This is defined in the following way. Let $\bar{s}_0$ be a random variable that is continuously distributed with support $[-1, 1]$. When optimizing with respect to the strategies of other consumers, consumer i assumes that $s_0 = \epsilon \bar{s}_0$, where $\epsilon > 0$. Consumer i's strategy must be the limit of optimal responses as $\epsilon \to 0$.

Proposition 13. *The pair a, b, price p, and resulting consumption level $q = a - bp$ constitute a robust equilibrium if and only if market clearing is satisfied with q and p, and*

$$0 = u'(q) - p - \frac{1}{I(b+s)}q - k\frac{s}{b+s}, \text{ and}$$
$$b = \frac{1}{(-u''(q)) + 1/(I(b+s))}.$$

We now use the above characterization to provide a foundation for competitive equilibrium.

Proposition 14. *The following are equivalent.*

I. The quantity $q^* > 0$*, price* $p^* > 0$*, consumer price responsiveness* $q_p^* \in \mathbb{R}$*, and market responsiveness* $q_c^* \in \mathbb{R}$ *constitute a competitive equilibrium.*

II. There is a sequence $I_n \to \infty$ *and a sequence of robust equilibria* $a(I_n)$*,* $b(I_n)$*,* $p(I_n)$*, such that* $p^* = \lim_{n\to\infty} p(I_n)$*,* $q^* = \lim_{n\to\infty} a(I_n) - b(I_n)p(I_n)$*,* $q_p^* = \lim_{n\to\infty}(-b(I_n))$*, and* $q_c^* = s/(s - q_p^*)$*.*

A.4 Private Utility, Social Concerns, and Social Welfare

In this subsection, we motivate the social welfare function (5) in the text. The calculation also clarifies the utility function (1) we have assumed for an individual consumer.

As above, there are $I+1$ identical consumers. Consumer i's private utility is

$$u(c_i) - pc_i - K \cdot \frac{\sum_{i'=1}^{I+1} c_{i'}}{I+1}.$$

Hence, each consumer's utility is decreasing in the average consumption in society, for instance because of the health consequences of pollution or social consequences of global warming. Consumer i realizes that her consumption contributes to the externality, which affects others negatively. The total disutility others suffer from the externality is

$$K \cdot \frac{I}{I+1} \cdot \sum_{i'=1}^{I+1} c_{i'}.$$

Being socially responsible, the consumer partially or fully takes into account this disutility, or her effect on it. Including this social concern, then, her utility is

$$u(c_i) - pc_i - K \cdot \frac{\sum_{i'=1}^{I+1} c_{i'}}{I+1} - k \cdot \frac{I}{I+1} \cdot \sum_{i'=1}^{I+1} c_{i'} = u(c_i) - pc_i - \left(k \cdot \frac{I}{I+1} + K \cdot \frac{1}{I+1}\right) \cdot \sum_{i'=1}^{I+1} c_{i'},$$

where $k \le K$. As $I \to \infty$, the weight the consumer attaches to the total market quantity $\sum_{i'=1}^{I+1} c_{i'}$ approaches k, as we have assumed in the text. Note also that if $k = 0$, in the limit the consumer ignores her effect on the externality, as in the classical case of a selfish consumer.

Our formulation assumes that in choosing her consumption, a socially responsible consumer cares about her effect on the externality, but not about her effect on others' private utilities. Hence, for instance, in choosing how much to fly the consumer thinks about global warming and its effect on humanity, but does not internalize others' enjoyment of flying. Beyond realism, this assumption is helpful in connecting our basic market-failure results to previous ones. In classical settings, the efficiency of markets does not require consumers to internalize others' private utilities.

Crucially, we assume that the social-welfare function equals the average of individuals' *private* utilities. First-pass conventional logic might dictate that the social welfare function equals the average of individuals' total utilities. 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. To illustrate this in another way, consider the following example, adapted to our setting from Bergstrom (2006). Suppose that each consumer is consuming $c_i = 1$, generating an externality of $g = 1$. How much should society be willing to pay per person to eliminate the harm from the externality? The natural answer is K. If we were willing to pay more than K, then we would be willing to impose a tax greater than K and use it to eliminate the harm. But this intervention would strictly lower all consumers' private utilities. It would be exceedingly odd to use consumers' concerns for others' disutility from the externality to justify a policy that makes everyone privately worse off.

Hence, the social welfare function is

$$\frac{\sum_{i=1}^{I+1}[u(c_i) - pc_i - Kc_i]}{I+1}.$$

When all consumers choose the same c_i, the above reduces to the social welfare function in the text.