



The Morality of Markets

Mathias Dewatripont
ECARES, I3h & TIMES², Université libre de Bruxelles

Jean Tirole
TSE & IAS, University of Toulouse

September 2023

ECARES working paper 2023-23

The Morality of Markets*

Mathias Dewatripont[†] Jean Tirole[‡]

September 20, 2023

Abstract: Scholars and civil society have argued that competition erodes supplier morality. This paper establishes a robust irrelevance result, whereby intense market competition does not crowd out consequentialist ethics; it thereby issues a strong warning against the wholesale moral condemnation of markets and pro-competitive institutions. Intense competition, while not altering the behavior of profitable suppliers, however may reduce the standards of highly ethical suppliers or not-for-profits, raising the potential need to protect the latter in the marketplace.

Keywords: Competition, consequentialism, replacement logic, non-profits, corporate social responsibility, race to the ethical bottom.

JEL numbers: D21, D43, D6, I11, L21.

*This project received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 669217 - ERC MARK-LIM). Jean Tirole acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program), and financial support from the Chaire SCOR "Market risks and Value creation". We gratefully acknowledge the comments of participants at conferences at Bonn, Chicago-LSE, CRESSE, Foundation Getulio Vargas, MIT, SIOE, and TSE, as well as seminars at Université Libre de Bruxelles, Reggio de Calabria, Harvard, and Western Ontario Law, and of Tim Besley, Emir Kamenica, Louis Kaplow, Paul-Henri Moisson, Charles Pébereau, François Salanié, Joel Sobel, Nicolas Treich and four referees.

[†]Université libre de Bruxelles (Solvay Brussels School of Economics and Management, ECARES, I3h and TIMES²), mathias.dewatripont@ulb.be.

[‡]Toulouse School of Economics (TSE) and Institute for Advanced Study in Toulouse (IAST), jean.tirole@tse-fr.eu.

1 Introduction

Whether markets impede ethical behavior has always engendered a variety of views. Many 18th century thinkers believed that such concerns are either irrelevant or mistaken. Adam Smith stressed that self-interest could lead to socially optimal outcomes. Condorcet, Hume, Montesquieu (with his “doux commerce”), and Turgot viewed market institutions as creating trust among otherwise unrelated individuals; see Hirschman (1977) and the economic history work of McCloskey (2006) and Mokyr (2016).¹ A different tradition, dating back to Karl Marx and popular in today’s public opinion and among social scientists, politicians and religious leaders, argues in contrast that markets promote unethical² behavior. For instance, numerous prominent contemporary philosophers have warned against the religion of the marketplace, with a variety of viewpoints from the necessity to ban repugnant markets to the stance that a market economy is an unlikely path to a harmonious society (see Anderson, 1993, Sandel 2012, Satz, 2010 and Walzer, 2008). The critique that market competition obliterates our moral compass is the focus of this paper.³

The morality-based criticism of competition often builds on the “replacement logic” narrative, the idea that if a supplier refuses to engage in an immoral trade, “someone else will”.⁴ In that, the criticism echoes widespread narratives used by firms and countries selling weapons to dictatorships or bribing officials to win a contract, by banks selling toxic products or providing short-term incentives to talents they want to attract, by employees ingratiating themselves to their superiors in order to be promoted, by doctors overprescribing opioids, antibiotics and sick leave, by professional athletes taking illicit drugs to defeat their competitors, by farmers exploiting animals,⁵ or by companies whitewashing their products’ potential shortcomings (their brittleness or high fat and sugar content).

Deconstructing the informal argument behind the replacement logic, suppliers must balance a profit motive with social concerns (in possibly heterogenous fashion), goods

¹In conformity with this view of markets, Dufwenberg et al. (2022) find experimental support for individuals having reciprocal preferences and for successful market interactions (interpreted as the efficient equilibrium outcome in a cooperative coordination game) triggering generosity in a dictator game. Our perspective is different in that we focus on how the nature of market interactions themselves affect players’ own tradeoffs between profits and ethical concerns. This being said, like Dufwenberg et al., we assume that market interactions do not change intrinsic preferences.

²We will use “ethics” and “morals” indifferently in this paper. For our purpose, it does not matter whether the social preferences of suppliers or stakeholders refer to rules provided by an external source, or reflect an individual’s own principles regarding right and wrong.

³To be certain, some of the critiques reflect a desire to move all the way to an economy consisting only of state-owned or democratically-run firms. However, doubts about the morality of competitive markets are much more pervasive in society and often do not reflect such ulterior motives.

⁴In the policy debate, the “replacement logic” is sometimes called “first-mover disadvantage”: “If I reduce my carbon footprint, I will lose market share”.

⁵Animal exploitation induces an externality/harm on other sentient beings and is considered by philosophers as morally problematic. The rhetoric of animal farmers is often based on the replacement effect: “We like animals but if we did not put animals in cages, we would import cheaper and less humane meat from competitors.”

must be substitutes (the more so, the stronger the argument), and purchasers must benefit from suppliers behaving unethically (as in the examples above). To contrast the latter with ethical (socially responsible) consumers, whose demand increases with the moral content of the product (a much-studied case in the economics literature), we will call such consumers “unethical”. The replacement logic also embodies a consequentialist overtone, as the supplier worries about the ultimate consequences of their behavior. In that, it is a special case of a large class of preferences that will play a key role in our analysis. We will say that the supplier is consequentialist if it internalizes the impact of a change in its own moral choice in proportion of the quantity it sells (say, it internalizes the social cost of 2 tons of carbon emissions as twice that of a single ton).

It is therefore important to stress right away that our paper is not specifically about the replacement logic. We put forward this logic, as it has been a key narrative for critiques of the market economy and also is often used as an excuse by suppliers to “justify” their unethical behavior. We give a formal content to the replacement logic through the internalization by suppliers of what we label “broad bracketing” or “ethical welfare” (say, the total amount of pollution or opioids). Our insights and results however apply to the much larger class of consequentialist social preferences. Other examples of the latter include narrow bracketing (say, the supplier internalizes only its own pollution or prescription of opioids), as well as other notions of social welfare that further account for consumer net surplus or the misallocation of consumers to products (social preferences then reflect total welfare).

Our results relative to unethical purchasers (say, officials asking for bribes or dictators buying weapons) will apply more generally to any environment in which unethical behavior by the supplier leads to an increase in the firm’s demand. Such is also the case when the consumer suffers from present bias, as is illustrated by the opioid case. Overprescriptions of opioids raise sales, but not consumer welfare. A third illustration of demand-increasing unethical supplier behavior occurs when the consumers are ill-informed about what they purchase. Whether they are rational (and just uninformed), naïve or gullible, the failure to disclose flaws or the use of misleading advertising boosts demand, at a cost to society whenever the consumer would not have purchased, absent the misrepresentation. For this reason, we will consider the broader class of “UPI (*unethical/present-biased/influenceable*)” consumers. In contrast, the demand of “ethical” (“*indifferent*”) consumers increases (remains constant) when the supplier’s offering is more ethical.

The irrelevance result. We ask: does the combination of unethical (or more generally UPI) consumers and of suppliers with consequentialist social preferences imply that moral behavior deteriorates under more intense competition? Our answer to this question is “no”. Indeed, under weak assumptions the degree of competitive pressure is irrelevant to ethical behavior (moral choices are independent of demand functions) if prices are flexible.

The intuition behind the irrelevance result goes as follows: When a supplier faces more intense competition (a more elastic demand), raising ethical behavior has a bigger negative impact on the supplier’s market share and is therefore costlier for the supplier; *ceteris*

paribus this makes suppliers cut ethical corners in reaction to the increase in competition, as indicated in the conventional wisdom. However, there is a second, reduced-stakes effect: A more intense competition reduces prices and markups, making supplier ethical concerns loom larger relative to material ones. We show that a sufficient condition for these two effects to exactly offset each other is that suppliers have consequentialist preferences and returns to scale are constant.

The irrelevance result, which applies as well to ethical or indifferent consumers, is important not only because it sheds light on the validity of the widespread concern about markets expressed by the public opinion, social scientists, politicians and religious leaders, but also because it affects our stance vis-a-vis key competition-enhancing public policies such as the opening of borders to free trade, competition policy and the deregulation of industries. The irrelevance result is also in stark contrast with earlier theoretical results on the irrelevance of social preferences in highly competitive environments, in particular with Dufwenberg et al (2011) and Sobel (2015): In our case the social preferences of suppliers and of consumers matter regardless of the competitive pressure, and it is the intensity of competition that is irrelevant. The difference is driven in particular by the fact that, in their settings, one can only affect others' utilities through one's impact on their quantities traded or the market price, an impact which vanishes under perfect competition. In our setting, an individual may want to change her action just because it is objectionable to herself or others, even if this does not affect their ability to trade, a feature which is widespread in the real world. See the literature review for a detailed comparison.

We then show that the irrelevance result is robust to various forms of competition. In particular it holds under strategic substitutes as well, i.e., when firms compete in capacities. There is then no replacement effect as a supplier's increase in moral content does not affect their rivals' output; this result confirms that the irrelevance result is consistent with the replacement logic, but by no means hinges on its existence. The result also accommodates a wide range of consequentialist preferences, from the case in which the suppliers care only about the moral consequences (e.g., the emissions) of their own production to that in which they care about overall welfare. Finally, the irrelevance result extends to imperfect consumer information, to some forms of non-linear price discrimination, and to some environments with non-constant returns to scale.

When does the irrelevance property fail? As the reduced-stakes effect suggests, the clue lies in the rigidity of prices. Prices may be rigid for one of two reasons. First, prices may be exogenously set by either a regulator (taxis, notaries, doctors in some health systems) or by a private party (apps and franchising environments). Keeping the regulated prices fixed, a more intense competition impedes moral behavior under UPI consumers (validating the common criticism of "markets" when competitors are constrained in their ability to lower price to gain market share), and fosters moral behavior when consumers are ethical.

Second, prices may be endogenously downward-constrained by limited liability. Firms with different corporate forms, for-profits and not-for-profits, may co-exist. Indeed, it is often suggested that in industries with strong moral overtones (health, education), the

profit motive should be eliminated. The not-for-profits must align revenue with cost and so their prices, while endogenous, are not fully flexible. Alternatively, when suppliers are all for-profits but differ in their social preferences, the more ethical suppliers' preferred policy may put them in the red when competition is sufficiently intense; this implies that they are de facto, although not de jure, not-for-profits. We show that with UPI consumers moral choices of not-for-profits or highly ethical suppliers mimick those of less-ethical ones; and so, to make a difference, the former must be insulated from an intense competitive pressure from for-profits. This strategy is a better response than a weakening of competition policy enforcement to concerns about insufficient ethics in our market economy.

Roadmap

This paper is organized as follows: Section 2 develops the baseline model. Suppliers operate in an imperfectly competitive industry and select two actions: a price and a moral action. The moral action affects demand and/or production cost. The product's consumers are defined as ethical, UPI or indifferent depending on whether a more moral action increases, decreases or does not affect the firm's demand.

Even though our focus is on suppliers, who operate the ethical choices and whose social preferences are the more novel part of the paper, we consider a general model in which stakeholders, namely consumers but workers and investors as well, are also driven by both a material motive and social preferences. We require stakeholders to be consequentialists as well, which accommodates, like for suppliers, a wide range of social preferences. For instance, the warm glow experienced by some investors when trading an ESG security for a brown one is a form of narrow bracketing; in contrast, other investors exhibit broader bracketing when using the replacement excuse that "If I divest oil companies from my portfolio, someone else will buy the shares anyway". Similarly, accounting for the fact that my consuming green electricity from a hydroelectric dam with a fixed capacity displaces other consumers' purchases toward brown power implies broad bracketing, while my ignoring equilibrium (or leakage) effects reflects narrow bracketing.

A supplier's social preferences are most simply interpreted as either those of the manager in the case of an owner-managed firm (entrepreneur, doctor) or those of shareholders under shareholder value. Alternatively, they might reflect a mixture of the two, with different weights depending on the extent of agency. "Shareholders" stand for active investors, who exert voice to impact the firm's choice. In contrast, passive investors have no such impact but may accept a lower return when investing in an ethical firm (their influence will then be reflected in the cost function).

Section 2 develops the framework and discusses the three assumptions that are key to the irrelevance result: Consequentialism, price flexibility, and "constant returns to scale" (returns need not be constant in output, rather a supplier's marginal cost of raising the morality of her production is proportional to its output). Section 3.1 derives the basic irrelevance result, and Section 3.2 performs the various extensions discussed above.

Section 4 considers limits to price flexibility. Section 4.1 and 4.2 study rigid prices. Section 4.1 first shows that for given prices, moral choices are strategic complements under UPI consumers. Two reasons underly this strategic complementarity: an “elasticity effect” and a “social responsibility effect”. Section 4.2 then demonstrates that, as announced above, a more intense competition impedes (fosters) moral behavior when consumers are UPI (ethical). Finally, Section 4.3 analyzes competition between suppliers when corporate form or social preferences heterogeneity leads to a break-even concern; it demonstrates the interdependence among policies adopted by rival corporate forms and derives some policy implications.

Section 5 demonstrates the relevance of the analysis to shed light on current debates and the real world, and, while emphasizing the need for more empirical work, discusses various forms of evidence supporting the theory. Section 6 relates the paper to the existing literature. Section 7 summarizes the main insights. Omitted proofs and more specific material are relegated to the Online Appendix.

2 Framework

Our baseline model is one of differentiated Bertrand competition. There are n suppliers, $i \in \{1, \dots, n\}$ and a mass 1 of unit-demand, price-taking consumers. The outside option is indexed by 0; it can be interpreted as the absence of consumption, or the consumption of a substitute good with fixed ethical implications. Suppliers compete in price and non-price dimensions. Supplier i selects its price p_i as well a moral or ethical choice a_i , both in \mathbb{R}^+ . We will use formal assumptions to flag those driving the irrelevance result.

Impact of ethical choice. Besides price p_i , supplier i picks a level of morality $a_i \in [0, \bar{a}_i]$ with $\bar{a}_i \leq +\infty$. Choice a_i has per-unit-of-output direct welfare impact $W_i(a_i)$, say (minus) a per-unit externality cost. For example, a_i might be a choice of technology; a CO₂ emission rate of $\psi_i(a_i)$ per unit of output yields welfare $W_i(a_i) = -\psi_i(a_i)e$, where e is the social cost of carbon. A higher value of a_i indexes a more moral choice: $W'_i(a_i) > 0$ on $[0, \bar{a}_i)$ and $W'_i(\bar{a}_i) = 0$.⁶ We assume that $W''_i(a_i) < 0$ for all a_i and $W'_i(0) = +\infty$. The outside option, “good 0”, generates exogenous welfare impact $w_0 \equiv W_0(a_0)$. For example, the absence of purchase of conventional electricity generation might involve no pollution (energy sobriety) or else be highly polluting (return to coal or wood-burning). Let $\mathbf{a} \equiv (a_1, \dots, a_n)$ denote the vector of ethical choices. Ethical choices are observable unless otherwise indicated.⁷ Looking ahead, the moral choice a_i will matter not only to

⁶ \bar{a}_i is finite in all examples provided as microfoundations in Online Appendix A.

⁷The context may influence the stakeholders’ social preferences. Suppose that the supplier first produces and then brings its production to the market; broad bracketing (the replacement excuse) then seems logical for consumers: “the animal was raised cruelly, whether the meat is eaten or not: my buying the meat does not alter this”.

In contrast, under “production only upon order”, the animal is not raised cruelly if I choose not to consume meat. Then a_i does damage/benefit only when supplier i ’s good is purchased, and so narrow bracketing seems more logical.

suppliers, who will value both their material payoff and the morality of their behavior, and to consumers, whose demand will depend on a_i and p_i , but also to workers and investors, who will accept a lower wage or return for being associated with a more moral firm (the impact of supplier moral choices on workers and investors will operate on the cost, rather than the demand side).

Consumer attitudes. Incentives for suppliers to choose a given action a_i will depend on consumer attitudes towards a_i . “Consumers” are parties who impact the demand side. These may be ordinary consumers of goods and services or an “agent” or “purchaser” selecting on behalf of them (officials selecting a contractor, current incarnation); we will generically use the term “consumer” to encompass each acceptance. The consumers’ cost or benefit of the moral action is captured through its monetary equivalent $\phi_i(a_i)$ with $\phi_i'' \geq 0$ (to ensure the concavity of optimization programs), such that the consumers’ demand for product i depends only on its net price \hat{p}_i (and on the net prices charged by other suppliers)

$$\hat{p}_i \equiv p_i + \phi_i(a_i). \quad (1)$$

So we assume that the extent to which consumers care about a_i is independent of the price, in the same way we model the impact of a sales tax in econ 101. Note also that the consumers’ cost or benefit of the moral action could be heterogenous. The function $\phi_i(a_i)$ would then stand for the average cost or benefit (there is a formal equivalence for a linear demand system). The separability assumption seems reasonable provided that the consumer’s utility is separable in disposable income and accomplishment of one’s moral duty.⁸

Definition (*social responsibility*). Consumers are

- (i) UPI (unethical/present biased/influenceable) when $\phi_i'(a_i) > 0$ (their demand decreases with the morality of the firm’s offer);
- (ii) ethical when $\phi_i'(a_i) < 0$ (their demand increases with the morality of the firm’s offer);

The contrast between the two is clear in a static model. Things would be more complex in a dynamic one: My buying a steak encourages future factory farming, as it displaces the suppliers’ perception of the demand curve. So even in the more common case of production prior to sales, narrow bracketing might be more logical.

Importantly, the exact nature of stakeholders’ social preferences is not essential to our analysis as long as they reflect the scale of the trade, and so we do not need to question whether they are logical. Put differently, while we find such consumer reasoning appealing, we are agnostic as to how it is formed as long as it is consequentialist. The results carry over to when it is not applied, for example when consumers remain guided by the impact of their own purchase when the supplier produces and then brings its production to the market.

⁸It can be motivated by assuming that there are many goods and taking a linear approximation. Namely, the consumers could be consuming many such goods, indexed by x and have utility $\xi(y - \int p_x dx, \int -\phi_x(a_x) dx, -\int \varepsilon_x dx)$ where ξ is increasing in the three arguments, y is the endowment, (p_x, a_x) are the price and morality of his choices of sub-brand of good x , and ε_x is the hedonic benefit of his choice of sub-brand to his preferred specification (ε_x is consumer idiosyncratic). Taking a linear approximation yields the model studied in this paper.

(iii) indifferent when $\phi'_i(a_i) = 0$.

All cases are relevant, even though they typically depict different contexts. Modeling ethical consumers' concern is straightforward: $-\phi_i(a_i)$ may be equated to $\alpha_C W_i(a_i)$, the benefit from feeling one is doing the right thing (α_C is the consumer's internalization coefficient), thereby boosting demand. Ethical consumers derive a psychological benefit from consuming carbon-free or fair-trade products.⁹

More interesting for this paper are UPI consumers. As the terminology suggests, we provide three distinct rationales for the disconnect between the low moral standards desired by the purchasing agent and what is good for society (these are sketched here and detailed in Online Appendix A). Two of them create a private benefit for the purchasing agent, that is decreasing in the morality of the action. The first possible wedge (motivating the “ U ”) may stem from an externality (as when doctors deliver fake medical certificates to allow their client not to be vaccinated or to take sick leave, or when a firm bribes an official who awards a government contract or supplies weapons to a dictator; the client cynically benefits from the supplier's immoral behavior). The second wedge (motivating the “ P ”) may be traced to an internality (a doctor over-prescribes opioids, which are attractive to the client's “current self” but being addictive, detrimental to her “long-term self”, who is then the victim). A third possibility (motivating the “ I ”) arises when the moral action refers to the truthfulness of product disclosure. In this case, the consumers are victims when the supplier behaves less morally. For example, misleading advertising - the absence of disclosure of the product's flaws or limitations (a low a_i) - raises demand. A more complex case (but one covered by our framework) arises when the flaw is the necessity for the consumer to later purchase an unforeseen add-on from the supplier (as in the “shrouded attributes” literature initiated by Gabaix and Laibson 2006). The non-disclosure then not only increases demand, but also generates for the supplier deferred profits, whose expectation is akin to a reduction in the supplier's marginal cost.

Finally, indifferent consumers are either of the homo-economicus type (their preferences are purely material) or, more interestingly, they have social preferences but cannot express them in the marketplace as they do not observe the suppliers' moral choices prior to their purchase and furthermore the realized moral choice does not affect their demand.¹⁰

Net prices and demands. The vector (p_i, a_i) determines the net price \hat{p}_i perceived by the consumers (\hat{p}_0 is the net price for the exogenous outside option). Supplier i faces demand function $q_i = D_i(\hat{\mathbf{p}})$, where $\hat{\mathbf{p}} \equiv (\hat{p}_1, \dots, \hat{p}_n)$ denotes the vector of supplier net prices. We

⁹That the consumer internalizes the welfare associated with her choice does not imply a “narrow internalization”. To be certain, an alternative choice (say, firm j) would have welfare consequences that depend on a_j . However, this is taken into consideration by the consumer when selecting a supplier. Letting ε_{hk} denote the valuation of consumer h for good k , the consumer compares $\varepsilon_{hi} - [p_i - \alpha_C W_i(a_i)]$ with $\varepsilon_{hj} - [p_j - \alpha_C W_j(a_j)]$. That is, the consumer accounts for the welfare impact of alternative choices.

¹⁰In contrast with the disclosure examples just mentioned in which the consumer also does not observe the realization of the moral action, here the actual choice of this action does not affect demand (think of unobserved use of child labor or of pollution: the consumer's demand does not react to the realized choice of moral action).

will also write firm i 's demand as $D_i(\hat{p}_i, \hat{\mathbf{p}}_{-i})$, where $\hat{\mathbf{p}}_{-i}$ denotes the vector of net prices charged by supplier i 's rivals. Firm i 's demand is decreasing in its own (net) price. [In Sections 4.2 and 4.3, we will specialize to the case of a fixed total demand (everyone needs a doctor or a school, say), with a mass 1 of unit-demand consumers and $\sum_{i=1}^n D_i(\hat{\mathbf{p}}) = 1$ in the relevant range of net prices; we will then say that the market is “covered”.]

For instance, the demand function $D_i(\hat{\mathbf{p}})$ stems from a consumer discrete choice model: Consumers have unit demands with valuations $\{\varepsilon_{hi}\}_{i \in \{0, \dots, n\}}$ drawn from some smooth joint distribution. Consumer h therefore buys from supplier i if $\varepsilon_{hi} - \hat{p}_i > \max_{\substack{j \neq i \\ j \geq 0}} \{\varepsilon_{hj} - \hat{p}_j\}$ and does not if the inequality is in the opposite direction. As we will later show that the irrelevance result extends to the Cournot model, we should note that the perfect-substitutes demand function is a special case of the discrete choice model, with perfect correlation of the differential between the oligopolists' products and the outside option: $\varepsilon_{hi} - \varepsilon_{h0} = \varepsilon_{hj} - \varepsilon_{h0} = v$, where v is the valuation, distributed according to some c.d.f. $F(v)$.

Demand elasticity. The suppliers are substitutes ($\partial D_i / \partial \hat{p}_i < 0 < \partial D_i / \partial \hat{p}_j$), and the profit function satisfies the standard assumptions. Supplier i 's marginal revenue is decreasing in price, keeping the ethical action constant ($(p_i - c_i)D_i(\hat{\mathbf{p}})$ is concave in p_i). We will let $\eta_i(\hat{\mathbf{p}}) \equiv (-\partial D_i / \partial \hat{p}_i) / (D_i / p_i)$ denote the price elasticity of demand for supplier i 's services (note that $\partial D_i / \partial \hat{p}_i$ is the price sensitivity of demand from (1)).¹¹ We assume that the goods are (local) strategic complements: Supplier i 's elasticity of demand increases with competitive pressure:

$$\frac{\partial \eta_i}{\partial \hat{p}_j} < 0.$$

Costs. We make the following assumption on the possible dependence of supplier i 's cost on the ethical choice a_i :

Assumption 1 (“constant returns to scale”). *A supplier's marginal cost of raising the morality of her production is proportional to her output: Her cost as a function of her output q_i and her moral choice a_i can be written as $C_i(q_i, a_i) = c_i(a_i)q_i + d_i(q_i)$, where $c'_i(a_i) \geq 0$ and $c''_i(a_i) > 0$ for all a_i .*

Note that returns need not be constant with respect to output; instead, the marginal cost of moral behavior scales up with this output. The firm may use child labor or fossil fuel sources of energy in order to keep its cost low, in which case $c'_i(a_i) > 0$. Alternatively, the ethics-dependent cost function captures investor and worker social responsibility. The latter may be willing to forego some return or some wage to be associated with a more ethical enterprise. Suppose, for the sake of illustration, that investors (resp. workers) are willing to accept a reduction in their return equal to $\alpha_I W_i(a_i)$ (resp. in their wage equal to $\alpha_W W_i(a_i)$) to be associated with firm i . Assuming that 1 unit of output requires 1 unit

¹¹In Section 4, we will index η_i by a parameter $\sigma \in \mathbb{R}^+$ of intensity of competition. For instance, σ might be the inverse transportation cost in the Hotelling model but there are many alternative interpretations.

of labor and 1 unit of investment, and letting $\gamma_i(a_i)$ denote firm i 's operating cost (where γ_i' is typically weakly positive), then

$$c'_i(a_i) = \gamma'_i(a_i) - \alpha_I W'_i(a_i) - \alpha_W W'_i(a_i),$$

and the analysis carries over with $c'_i < 0$ if γ_i is constant. A moral action then reduces the cost of doing business. The sign of $c'_i(a_i)$ thus hinges on the context.

Under ethical or indifferent consumers, if $c'_i(a_i) \leq 0$, then a more moral behavior does not reduce demand, morally pleases the supplier, and does not increase cost: the optimal choice of firm i is a no-brainer.¹² To avoid mentioning such trivial corner solutions, we require that $c'_i(a_i) > 0$ when consumers are ethical and indifferent. More generally we will rule out corner solutions for expositional simplicity.

Suppliers' objective functions. Suppliers care about profit, but have social preferences, as reflected in their internalization of welfare.¹³ Supplier i 's internalization of social welfare, $\mathcal{W}_i(\hat{\mathbf{p}}, \mathbf{a})$, depends on net prices and ethical choices. Let $\alpha_i \geq 0$ denote supplier i 's (common knowledge) intrinsic ethics, that is the weight on welfare relative to that on profit. Supplier i maximizes the sum of profit and internalized perceived social welfare; letting $\alpha_i \geq 0$ denote the intensity of her social preferences,¹⁴ her utility function is:

$$V_i \equiv [p_i - c_i(a_i)]D_i(\hat{\mathbf{p}}) + \alpha_i \mathcal{W}_i(\hat{\mathbf{p}}, \mathbf{a}). \quad (2)$$

Note that in corporations that are run by managers the relative weight suppliers put on profit and social welfare hinges on their compensation scheme. For example, the behavior of a supplier i who is an agent with social preferences α_i and receives a fraction ξ_i of the profit associated with their activity is indistinguishable from that of a residual claimant for the firm's profit with social preference parameter $\hat{\alpha}_i \equiv \alpha_i/\xi_i$. For example, private equity and LBOs are usually characterized by high-powered incentives (high ξ_i).¹⁵

We assume that suppliers care about the social impact of the industry's aggregate activity ("broad bracketing"), for example the resulting total pollution or opioid overuse, and so we can drop the subscript i for $\mathcal{W}_i(\hat{\mathbf{p}}, \mathbf{a})$. We define this "ethical welfare" as:

$$\mathcal{W}(\hat{\mathbf{p}}, \mathbf{a}) = \mathcal{E}(\hat{\mathbf{p}}, \mathbf{a}) \equiv \sum_{j=0}^n W_j(a_j)D_j(\hat{\mathbf{p}}).$$

The frequent appeal to the replacement excuse in the policy debate justifies for expositional purposes this particular choice of internalization (which in Proposition 3 will imply

¹²For example, with $c'_i(a_i) = -(\alpha_I + \alpha_W)W'_i(a_i)$, then supplier i , when endowed with social preferences as described shortly, chooses $a_i = \bar{a}_i$.

¹³See Section 3.2.4 for alternative moral imperatives.

¹⁴Were social preferences not common knowledge, suppliers might be reputation-conscious, in which case the objective function below would have to be augmented with an image term as in, e.g., Bénabou-Tirole (2006).

¹⁵Such reinterpretations must be kept in mind when thinking about the opioid scandal, as Purdue pharma had access to the doctors' prescription data and could (and did) provide high-powered incentives to its sales representatives (see US General Accounting Office, 2003, also discussed in Section 5.2).

that for given prices, unethical behavior by other suppliers “vindicates” one’s own unethical choice). However, Section 3.2.1 will show that consequentialism accommodates a large class of internalizations by the suppliers (namely, those for which $\partial \mathcal{W}_i / \partial a_i = \Gamma_i(a_i) D_i(\hat{\mathbf{p}})$ for some non-increasing, non-negative function $\Gamma_i(a_i)$; $\Gamma_i(a_i) = W'_i(a_i)$ in the special case of ethical welfare internalization). Namely, we make the more general assumption:

Assumption 2 (*consequentialism*). *A supplier internalizes the impact of a change in its own moral choice in proportion of the quantity it sells. Stakeholders (consumers, workers, investors)’ perception of the social impact of their trade is proportional to the size of this trade.*

Consequentialist preferences have been explicitly assumed for suppliers. They have been assumed more implicitly and mechanically for other players as they transact only one unit of good, labor or savings. However the theory carries over to arbitrary trade sizes, as long as the internalized welfare impact of ethical choice scales with quantity (for example, a consumer consuming q units from supplier i at tariff $T_i(q)$ internalizes net tariff $T_i(q) + \phi_i(a_i)q$).

Strategies and equilibrium. We look at Nash equilibria of the industry game in which the suppliers select simultaneously their price and their ethical action and then consumers select their supplier or choose the outside option. An important assumption is:

Assumption 3 (*flexible prices*). *Prices are (locally) flexible at equilibrium price configuration $\hat{\mathbf{p}}$. Namely, for equilibrium choices $(p_j, a_j)_{j=1, \dots, n}$, any local change in ethical behavior δa_i can be offset by a price change $\delta p_i = -\phi'_i(a_i)\delta a_i$ so as to keep supplier i ’s net price \hat{p}_i and therefore demand $D_i(\hat{\mathbf{p}})$ constant.*

Flexible pricing is a central assumption in much of economics. While it is a natural leading assumption, it does not apply to every context. Prices are flexible at some price configuration $\hat{\mathbf{p}}$ if (i) price p_i is not locally constrained by a public or private regulation, and (ii) supplier i ’s corporate charter or limited liability constraint does not preclude it from increasing or decreasing its price. The second condition is violated if the supplier is not-for-profit, even though its price is then endogenous. It also fails to hold if supplier i ’s would lose money at its optimal choice. Therefore, Assumption 3 will be made (and verified as it is an endogenous assumption) in Section 3, but not in Section 4.

3 The irrelevance result

3.1 Derivation

We first derive the paper’s main result.

Proposition 1 (*irrelevance*). *Suppose that prices are flexible at an equilibrium $(\hat{\mathbf{p}}, \mathbf{a})$. Supplier i 's ethical behavior a_i^\dagger is then uniquely defined by*

$$\alpha_i W'_i(a_i^\dagger) = c'_i(a_i^\dagger) + \phi'_i(a_i^\dagger) \quad (3)$$

It is therefore independent of the demand curve D_i faced by firm i , and thus of the intensity of competition.

Proof of Proposition 1. Using the definition of net prices, we can rewrite supplier i 's objective function as

$$V_i = [\hat{p}_i - \phi_i(a_i) - c_i(a_i)]D_i(\hat{\mathbf{p}}) + \alpha_i \mathcal{W}_i(\hat{\mathbf{p}}, \mathbf{a}). \quad (4)$$

Because $\partial \mathcal{W}_i / \partial a_i = W'_i(a_i)D_i(\hat{\mathbf{p}})$ and (from the envelope theorem) $\partial V_i / \partial \hat{p}_i = 0$, supplier i 's optimal ethical choice satisfies

$$\frac{\partial V_i}{\partial a_i} = 0 = [-\phi'_i(a_i) - c'_i(a_i) + \alpha_i W'_i(a_i)]D_i(\hat{\mathbf{p}}). \quad (5)$$

Hence, supplier i 's ethical behavior a_i is independent of the demand function, and so of the intensity of competition. The first-order condition w.r.t. a_i yields condition (3).¹⁶

The left-hand side of Condition (3) (the supplier's marginal demand for ethical behavior) is decreasing in a_i , while the right-hand side (the generalized marginal cost) is increasing; furthermore, as a_i tends to \bar{a}_i (resp. 0), the left-hand side goes to 0 (resp. $+\infty$). So even though the sign of c'_i and ϕ'_i can be positive or negative as we discussed, given that $\alpha_i W_i - c_i - \phi_i$ is strictly concave, Condition (3) defines a unique level of ethics a_i^\dagger .¹⁷ ■

The simple, but striking irrelevance result runs counter the conventional wisdom that competition erodes firms' moral compass. It calls for four comments:

(a) *Intuition.* When facing UPI consumers, say, a more elastic demand increases the market share loss from ethical behavior and makes the supplier cut ethical corners, as suggested by the conventional wisdom. However, there is a second, reduced-stakes effect: A more intense competition reduces prices and markups, making ethical concerns loom larger relative to material ones. These two effects exactly offset each other when suppliers have consequentialist preferences and returns to scale are constant.

Why the offset is complete can be grasped from the following cost-minimization reinterpretation, in which the supplier wants to transfer utility to consumers as efficiently as possible. Rewrite the objective function V_i as:

$$V_i = [\hat{p}_i - c_i(a_i) - \phi_i(a_i) + \alpha_i W_i(a_i)]D_i(\hat{\mathbf{p}}) + K(\hat{\mathbf{p}}, \hat{\mathbf{a}}_{-i}),$$

¹⁶See Online Appendix B for the verification of the global second-order condition in the case of a covered-market, symmetric equilibrium.

¹⁷As stated earlier, we ignore corner solutions at $a_i = \bar{a}_i$. If $\alpha_i W'_i(\bar{a}_i) - \phi'_i(\bar{a}_i) - c'_i(\bar{a}_i) > 0$, then a_i^\dagger is still unique and equal to \bar{a}_i . If overly pro-social actions ($W'_i(a_i) < 0$) were allowed, the optimum might again be interior. In any case, the equilibrium moral action remains (i) unique and (ii) independent of the demand curve.

and so supplier i 's unit cost is $c_i(a_i) + \phi_i(a_i) - \alpha_i W_i(a_i)$. Price flexibility, together with consequentialism and constant returns to scale, enables a decoupling between cost minimization and the choice of net price.¹⁸

(b) *Role of price flexibility.* The proof of Proposition 1 relies solely on the first-order condition with respect to a_i when supplier i also picks the net price \hat{p}_i ; the flexible-price assumption then implies that the choice of a_i can be performed keeping the net price \hat{p}_i and therefore demand D_i constant. Consider the first-order condition with respect to the ethical choice ($\partial V_i / \partial a_i = 0$), for *given* prices $\mathbf{p} = (p_1, \dots, p_n)$. When prices are not flexible, there is no such possible adjustment in the net price, and so a change $\delta a_i = \varepsilon$ is accompanied with a change in the net price $\delta \hat{p}_i = \phi'_i(a_i)\varepsilon$. Thus condition (5) now involves a total derivative, with $\frac{dV_i}{da_i} = \frac{\partial V_i}{\partial a_i} + \phi'_i \frac{\partial V_i}{\partial \hat{p}_i}$. Behaving more ethically (increasing a_i) has now three effects on supplier i 's payoff function $V_i = [p_i - c_i(a_i)]D_i(\hat{\mathbf{p}}) + \alpha_i \mathcal{W}_i(\hat{\mathbf{p}}, \mathbf{a})$:

$$\frac{\partial V_i}{\partial a_i} = \underbrace{(p_i - c_i)\phi'_i \frac{\partial D_i}{\partial \hat{p}_i} - c'_i D_i}_{\text{impact on profit}} + \underbrace{\alpha_i W'_i D_i}_{\substack{\text{ethical} \\ \text{impact on} \\ \text{supplier } i\text{'s} \\ \text{inframarginal} \\ \text{consumers}}} + \underbrace{\alpha_i \phi'_i \frac{\partial \mathcal{W}_i}{\partial \hat{p}_i}}_{\substack{\text{ethical} \\ \text{impact} \\ \text{of gain/} \\ \text{loss in} \\ \text{market} \\ \text{share}}} = 0. \quad (6)$$

using

$$\frac{\partial D_i}{\partial a_i} = \phi'_i \frac{\partial D_i}{\partial \hat{p}_i} = \phi'_i \frac{\partial D_i}{\partial p_i}. \quad (7)$$

Sections 4.1 and 4.2 will study the case of fixed prices in detail. For example, in the case of a symmetric oligopoly with a covered market and uniform regulated price p , we will show that, dropping subscripts, the equilibrium morality, a^* , is given by

$$\frac{\alpha W'(a^*) - c'(a^*)}{\phi'(a^*)} = \eta L$$

where $L = \frac{p-c}{p}$, the Lerner index, is exogenously given for the firms. Moral choices then depend on the elasticity of demand η , in a way that hinges on whether the consumers are UPI or ethical ($\phi' \geq 0$). [Price flexibility ensures that $L = 1/\eta$.]

(c) *When are prices flexible?* Consequentialism and constant returns to scale are embodied into the model. In contrast, the third key assumption, price flexibility, is an endogenous assumption, to be verified ex post by looking at the putative equilibrium. Suppose that prices are unregulated but that the suppliers cannot lose money (supplier

¹⁸The same reasoning holds under Cournot competition, replacing the vector of net prices $\hat{\mathbf{p}}$ by that of quantities \mathbf{q} (Section 3.2.2). It also holds for multi-unit demand consumers when each supplier i offers a (possibly non-linear) tariff $T_i(q_i)$, and demand is $D_i(\mathbf{T})$ where $\mathbf{T} \equiv (T_1(\cdot), \dots, T_n(\cdot))$ (Online Appendix C). Finally, regardless of whether competition is in price or quantity, one could add suppliers' choices of a dimension of quality that is devoid of moral connotation, again yielding the same demonstration of the irrelevance result.

i 's choices must satisfy: $p_i - c_i(a_i) \geq 0$). Prices are indeed flexible if this break-even constraint is non-binding.

Definition (*social responsibility index*). The “social responsibility index” S_i is defined as:

$$S_i \equiv \sum_{\substack{j \neq i \\ j \geq 0}} \sigma_{ij}(\hat{\mathbf{p}})[W_i(a_i) - W_j(a_j)],$$

where $\sigma_{ij}(\hat{\mathbf{p}}) \equiv [\partial D_j / \partial p_i] / [-\partial D_i / \partial p_i]$ (so $\sum_{\substack{j \neq i \\ j \geq 0}} \sigma_{ij} = 1$) measures the fraction of the market share gain by supplier i that comes from supplier j 's customers when supplier i lowers her price by one unit.

Note that $\partial S_i / \partial a_j < 0$ and that at a symmetric, covered-market equilibrium $S_i = 0$.

Proposition 2 (*flexible prices*). Conditions that are individually sufficient for prices to be flexible at the putative equilibrium, include:

- Either the equilibrium is symmetric and covered.
- Or the equilibrium is symmetric and $a^\dagger \leq a_0$, where a_0 is the morality of the outside option.
- Or, *ceteris paribus*, ethical concerns α_i are small enough.
- Or else competition, as measured by the semi-elasticities of individual demands, is not too strong.

Proof of Proposition 2. The FOC corresponding to the maximization of V_i w.r.t. \hat{p}_i with respect to prices yields a generalized Lerner formula:

$$\frac{p_i - c_i(a_i^\dagger) + \alpha_i S_i(\mathbf{a}^\dagger)}{p_i} = \frac{1}{\eta_i}. \quad (8)$$

So, prices are locally flexible if and only if $p_i \geq c_i(a_i^\dagger)$, or

$$\alpha_i S_i(\mathbf{a}^\dagger) \leq p_i / \eta_i. \quad (9)$$

Consider a symmetric equilibrium. Either the market is covered (all consumers purchase) and then $S_i(\mathbf{a}^\dagger)$ is equal to 0. Or the market is not covered, and in a symmetric equilibrium a^\dagger , $S_i(\mathbf{a}^\dagger) \leq 0$ if and only if $a^\dagger \leq a_0$. This case arises if the absence of trading by the industry yields a virtuous outcome (e.g. no pollution or no corruption).

Finally, consider a family $\{\alpha_i = \lambda \alpha_i^1\}_i$; then one can show that for $\lambda \leq \bar{\lambda}$ for some $\bar{\lambda} > 0$, equilibrium prices exceed unit costs. As λ become small, a_i^\dagger converges to the level that obtains for $\alpha_i = 0$ and $\alpha_i S_i(\mathbf{a}^\dagger)$ tends to 0. Moreover, when the semi-elasticity η_i / p_i is small, the RHS of (9) is large. ■

(d) *What are the drivers of ethics under price flexibility?* From Equation (3), equilibrium ethics under flexible prices is independent of the degree of competition but is influenced by

the ethical urges of the suppliers and stakeholders. Indeed, if we come back to our earlier characterization of $c'_i(a_i)$ as consisting of an operating cost $\gamma_i(a_i)$ minus a “discount” reflecting workers’ and investors’ social concerns, condition (3) can be rewritten as:

$$(\alpha_i + \alpha_W + \alpha_I)W_i(a_i) - \phi'_i(a_i) = \gamma'_i(a_i). \quad (10)$$

When consumers are ethical, $-\phi'_i(a_i)$ becomes $+\alpha_C W_i(a_i)$ and equilibrium ethics then increases in $\alpha_i + \alpha_W + \alpha_I + \alpha_C$, the sum of the ethical urges of all stakeholders and supplier i . When consumers are indifferent, $-\phi'_i(a_i)$ disappears and equilibrium ethics increases in $\alpha_i + \alpha_W + \alpha_I$. And when they are UPI, equilibrium ethics still increases in $\alpha_i + \alpha_W + \alpha_I$, the sum of the ethical urges of suppliers, workers and investors, but decreases when ϕ'_i increases.

3.2 Robustness

This section performs a few robustness checks. It focuses in turn on alternative forms of consequentialism, on other forms of competition (Cournot competition, tacit collusion) and on two of the three key assumptions: constant returns to scale and consequentialist preferences (the relaxation of price flexibility is found in Section 4). Online Appendix C further shows that the irrelevance result remains valid under imperfect consumer information and under volume-based price discrimination. In contrast, the intensity of competition has an ambiguous impact on ethics under ethics-based price discrimination. The punchline is that the irrelevance result is pretty general, the key exceptions being in Section 4. A reader wishing to move on to the study of non-flexible prices can skip this robustness section without loss of understanding.

3.2.1 Alternative forms of consequentialism

Suppliers. We assumed that suppliers internalize ethical welfare $\mathcal{E} = \sum_j W_j(a_j)D_j(\hat{\mathbf{p}})$. More generally, the proof of Proposition 1 shows that it still holds as long as marginal internalized welfare impacts scale with actual impacts, i.e. are proportional to demands: There exists a non-negative, non-increasing function $\Gamma_i(a_i)$ such that $\lim_{a_i \rightarrow 0} \Gamma_i(a_i) = +\infty$ and $\lim_{a_i \rightarrow \bar{a}_i} \Gamma_i(a_i) = 0$, and

$$\frac{\partial \mathcal{W}_i}{\partial a_i} = \Gamma_i(a_i)D_i(\hat{\mathbf{p}}). \quad (11)$$

That $\partial \mathcal{W}_i / \partial a_i$ is proportional to demand D_i is required by consequentialism: Ethical choices are uniform over supplier i ’s customers and so their impact on well-being is proportional to demand. The condition that Γ_i be non-increasing simply expresses the idea that returns to the ethical choice are non-increasing. This consequentialist internalization admits a wide variety of perceptions of social well-being.¹⁹ Besides ethical welfare, condition (11) is in particular satisfied by :

¹⁹Similarly, we assumed that workers and investors (partly) internalize $W_i(a_i)$. While this is natural, the irrelevance result does not hinge on this assumption. As for suppliers, one could assume that they

Narrow ethical welfare. Suppliers sometimes take a narrower view of ethical welfare, associated with the impact of their own production on well-being:²⁰

$$\mathcal{W}_i(\hat{\mathbf{p}}, \mathbf{a}) = \mathcal{E}_i^n(\hat{\mathbf{p}}, \mathbf{a}) \equiv W_i(a_i)D_i(\hat{\mathbf{p}}).$$

Note that such “narrow bracketing” is not specific to suppliers and is also relevant for stakeholders when they experience a warm glow. As we noted in the introduction, our analysis just assumes consequentialist preferences for suppliers and stakeholders.

Broader internalization. Conversely, what economists would call “welfare” usually encompasses other inefficiencies than those channeled through the choice of a_i . Online Appendix C shows that the analysis remains unchanged if suppliers internalize, on top of ethical welfare,

- *consumer surplus* (a drop in price reduces consumption distortions if the market is not covered)
- *product misallocation* (under asymmetric oligopoly or when net prices differ in a symmetric oligopoly, consumers are misallocated to products, and thus do not necessarily consume their preferred product).

In both cases, suppliers’ social preferences reflect inefficiencies that depend only on the vector of net prices. The efficiency-based foundation of irrelevance still holds.

3.2.2 Other forms of competition

(i) *Strategic complements vs. strategic substitutes (prices vs. quantities)*

Is the irrelevance property specific to the differentiated-products Bertrand model? Consider the Cournot model with perfect substitutes.²¹ Each supplier i first picks (q_i, a_i) . The suppliers then bring their production to the market. Finally, a Walrasian auctioneer clears the market (which then has a fixed supply). Under Cournot competition and for total quantity $Q = \sum_{j=1}^n q_j$, all net prices are equalized in the market clearing process:²²

$$P(Q) = p_i + \phi_i(a_i)$$

There is no replacement effect here, as an increase in a_i under UPI consumers does not induce an increase in q_j . Formally, $\sigma_{ij} = 0$ for $j \neq i, 0$ and $\sigma_{i0} = 1$. Supplier i ’s internalize an arbitrary $\Lambda_i(a_i)$ per unit (with $\Lambda_i' > 0 > \Lambda_i''$; they could also have internalizations that differ across stakeholders).

²⁰The distinction between ethical and narrow ethical welfares is reminiscent of Oehmke-Opp (2023)’s distinction between broad and narrow mandates and Green and Roth (2023)’s contrast between sophisticated and naive social investors (or “impact” and “value” investors).

²¹This is only for conciseness. The following reasoning applies also to the differentiated-products Cournot model.

²²If $F(v)$ is the distribution of valuations, then the inverse demand function $P(Q)$ is given by $Q = 1 - F(P(Q) - \phi_0(a_0))$ (or $1 - F(P(Q) - p_0 - \phi_0(a_0))$ if the outside option has a non-zero price).

social responsibility index is therefore (up to a term that does not depend on $\{q_i, a_i\}$) $S_i(a_i) = W_i(a_i) - W_0(a_0)$. Supplier i solves:

$$\max_{(q_i, a_i)} [P(\sum_{j=1}^n q_j) - c_i(a_i) - \phi_i(a_i)] q_i + \alpha_i S_i(a_i) q_i.$$

The FOC w.r.t. a_i yields the irrelevance result for strategic substitutes for the non-moral choice:

$$a_i = a_i^\dagger \quad \text{where} \quad \alpha_i W_i'(a_i^\dagger) = c_i'(a_i^\dagger) + \phi_i'(a_i^\dagger).$$

The outcome in quantities is given by the Cournot outcome with unit cost:

$$\hat{c}_i \equiv c_i(a_i^\dagger) + \phi_i(a_i^\dagger) - \alpha_i [W_i(a_i^\dagger) - W_0(a_0)].$$

(ii) *Tacit collusion*

Consider a symmetric, perfect-substitutes, n -firm oligopoly with total demand function $D(\hat{p})$ (where \hat{p} is the lowest price - the common price in a symmetric equilibrium), unit cost $c(a)$, and consumer internalization $\phi(a)$, with the assumptions made earlier. Suppliers' social preferences exhibit narrow bracketing (broad bracketing raises difficult conceptual issues for cartels and tacit collusion, which are discussed in Online Appendix C). Suppose that suppliers can collude tacitly, with reversal to Nash (the static outcome) in case of collusion breakdown; let V^{Nash} denote the per-period Nash payoff. Time is discrete: $t \in \{0, 1, 2, \dots\}$. The discount factor is δ . The *per-firm* payoff in a symmetric collusive outcome $\{\hat{p}, a\}$ is:

$$V \equiv \max_{\{\hat{p}, a\}} \left\{ [\hat{p} - \phi(a) - c(a)] \frac{D(\hat{p})}{n} + \alpha W(a) \frac{D(\hat{p})}{n} \right\}.$$

Sustainability requires that a supplier does not benefit from undercutting rivals:

$$\frac{V}{1 - \delta} \geq \max_{\{a'\}} \left\{ [\hat{p} - \phi(a') - c(a') + \alpha W(a')] D(\hat{p}) + \delta \frac{V^{\text{Nash}}}{1 - \delta} \right\}.$$

Thus both the cartel's optimal policy²³ and the optimal deviation adopt the cost-minimizing moral behavior (given by $\alpha W'(a^\dagger) = \phi'(a^\dagger) + c'(a^\dagger)$); and so the irrelevance property holds.²⁴

²³Note that it is not worth distorting the cartel's policy to deter deviations: Undercutting always takes the form of a better offer $\hat{p} - \varepsilon$ to consumers, and given this better offer, cost minimization is optimal for the deviator.

²⁴We do not need to derive V^{Nash} . But under Bertrand competition with perfect substitutes, then $p^* = c(a^\dagger) - \alpha W(a^\dagger)$, a limit case of condition (8) (with say, $W(a^\dagger) = -\psi(a^\dagger)e$ in the externality interpretation). To see this, consider a supplier's deviation to $\{p, a\}$. To benefit the supplier, the deviation must attract consumers ($p^* + \phi(a^\dagger) > p + \phi(a)$) and benefit the supplier ($p - c(a) + \alpha W(a) > p^* - c(a^\dagger) + \alpha W(a^\dagger)$). This however contradicts the fact that a^\dagger maximizes $\alpha W(a) - c(a) - \phi(a)$. And so, $V^{\text{Nash}} = 0$. So the Nash payoff is also the minmax payoff.

3.2.3 Non-constant returns to scale

We listed constant returns to scale as a key assumption for the irrelevance result. To see why, consider an arbitrary cost function $C_i(q_i, a_i)$. The generalization of condition (3) is then:

$$\frac{\partial C_i(q_i, a_i)/\partial a_i}{q_i} + \phi'_i(a_i) = \alpha_i W'_i(a_i) \quad (12)$$

There are interesting cases in which returns are not constant in *quantity*, and yet competition is irrelevant for moral choices:

(a) *Separability*. Suppose (as we did earlier) that the moral action impacts cost proportionally to output while returns need not be constant: $C_i(q_i, a_i) = c_i(a_i)q_i + d_i(q_i)$. Condition (12) then implies the irrelevance property.²⁵

An important example of separability arises when the moral incentive does not reside on the cost side (C_i depends only on q_i) but on the demand side, as is the case in the examples with UPI consumers mentioned in the introduction. Then (12) boils down to $\phi'_i(a_i) = \alpha_i W'_i(a_i)$ and so the irrelevance property holds regardless of the returns to scale.

(b) *Covered market*. Suppose a symmetric, covered market. Then equilibrium scale is invariant to competition ($q_i = 1/n$) and so is the moral action. Irrelevance holds again.

3.2.4 Non-consequentialist preferences

Last, focusing on supplier ethics, we compare the implications of consequentialism with those of the two main alternatives to consequentialism in moral philosophy: deontologism and categorical imperative. In the former, the supplier cares about her selected action rather than about its consequences; the impact of competition depends on the way in which it affects the profit stake of moral actions. In the latter, each supplier assumes everybody will mimic her action choice and so the suppliers' optimum always occurs. We thus obtain testable differences in the predictions of consequentialist and alternative moral criteria.

(a) *Deontologism*. Deontologism postulates that the morality of an action is based on whether the action is in itself right or wrong, irrespective of its scale and its consequences. Suppose therefore that supplier i values the act per se rather than its consequences. For instance, supplier i 's payoff could be $V_i = [p_i - c_i(a_i)]D_i(\hat{\mathbf{p}}) + \alpha_i \mathcal{W}_i(a_i)$, where $\mathcal{W}_i(a_i)$ (satisfying $\mathcal{W}'_i > 0 > \mathcal{W}''_i$) is an increasing and concave function of a_i . Such preferences are only partly deontological, as they reflect a material component (unless α_i is large).²⁶ The first-order condition for the moral choice under flexible prices writes: $\frac{\phi'_i(a_i) + c'_i(a_i)}{\alpha_i \mathcal{W}'_i(a_i)} = \frac{1}{D_i(\hat{\mathbf{p}})}$,

²⁵We have not undertaken a general study for non-constant returns to scale. Let us just add that, besides these three irrelevance situations, competition makes the market more immoral in the symmetric, multiplicative form: $C(q, a) = c(a)d(q)$, assuming that average cost (and so $d(q)/q$) is increasing in q .

²⁶Such preferences exhibit the Kahneman and Knetsch (1992)'s "embedding effect". Contingent valuations surveys tend to deliver stated willingnesses to pay that neglect scale.

and so the irrelevance property associated with consequentialist preferences in general does not hold.²⁷ If competition results in an expansion of the per-firm production (D_i increases), the profit motive is magnified relative to the ethical one and morality is *eroded*. A stricter enforcement of antitrust laws is an example in which increased competition is associated with an expansion of per-firm output. In contrast, if increased competition results from an increase in the number of licenses (an increase in n) and the market is covered, more competition is associated with a decline in per-firm output D_i ; in this case, competition *boosts* the ethical behavior of firms with deontological preferences.

(b) *Categorical imperative.* Suppose that suppliers follow Kant’s categorical imperative. If the market is covered,²⁸ then each supplier, behaving as if her choice was to be mimicked by other suppliers, selects the socially optimal action (assuming symmetry, otherwise it is not clear what the categorical imperative means). Suppliers behave fully ethically (as if $\alpha_i = +\infty$) regardless of the intensity of competition, and so raise no moral concerns.

In conclusion, not only does the model accommodate a range of variations on the moral criterion, but the irrelevance result is also valid under the categorical imperative criterium, although with a highly moral outcome. As for deontology, the impact of the intensity of competition is nonzero, but it is ambiguous and depends on the precise driver of the increase in competition.

4 Limits to price flexibility

Is the widespread opinion that competition erodes morality groundless? Proposition 1 suggests answers to this question. An impact of competition on (consequentialist) moral behavior may be related to prices not being flexible. If so, should we expect market morality to increase or decrease with the intensity of competition?

4.1 Determinants of moral choices for given prices

Assuming that suppliers wage differentiated-product price competition (as we will do in the rest of Section 4) and that internalized welfare is ethical welfare, this subsection shows that ethical choices are strategic complements for two reasons: an elasticity effect (which is inherited from the strategic complementarity in the price space) and, in the case of UPI consumers, a social responsibility effect (an increase in rivals’ ethical behavior makes it less

²⁷The irrelevance property still holds when the market is symmetric and covered ($D_i = 1/n$) and the increase in competition comes from an increase in substitutability, keeping the number of firms constant; in contrast, if the market is covered, but the increase in competition comes from entry of new firms (n increases), the increase in competition fosters moral behavior: Competition limits financial stakes and makes it more appealing to “do the right thing”.

²⁸We are agnostic as to the meaning of the categorical imperative in the presence of outside options, as the latter have no reason to obey the imperative and align the moral content with the suppliers’ moral choice.

desirable to steal market share away from them by cutting ethical corners). Furthermore, while the equilibrium ethical behaviors are uniquely determined when prices are flexible, they may not be when prices are not flexible, which requires either making assumptions guaranteeing equilibrium uniqueness (which we will do) or pursuing monotone comparative statics. Online Appendix D shows that similar, but differentiated results hold for other forms of consequentialist internalization (for example, for narrow internalization, strategic complementarity is driven by the sole elasticity effect).

Proposition 3 (*sufficient conditions for strategic complementarity*).

- (i) *For given prices, ethical choices are strategic complements if (a) consumers are UPI, or (b) suppliers do not internalize the social impact of their ethical choices too much.*²⁹
- (ii) *While the equilibrium moral actions are unique under price flexibility, there may be multiple equilibrium moral actions for fixed prices.*

We only provide the intuition for the proof of Proposition 3 here (see Online Appendix D for a formal proof). The maximization of $V_i = [p_i - c_i(a_i)]D_i(\hat{\mathbf{p}}) + \alpha_i \mathcal{W}_i(\hat{\mathbf{p}}, \mathbf{a})$ w.r.t. a_i (taking the total derivative, as explained in Section 3.1) yields

$$\frac{\alpha_i W'_i(a_i) - c'_i(a_i)}{\phi'_i(a_i)} = \eta_i \frac{p_i - (c_i - \alpha_i S_i)}{p_i} \quad (13)$$

The LHS of (13) is the per-unit-of-output benefit for the supplier of behaving more morally (expressed in monetary terms through the division by $\phi'_i(a_i)$). Our assumptions imply that it is locally decreasing in a_i (independently of the sign of ϕ'_i) and is independent of the competitive pressure. The RHS of (13) is the familiar product of the elasticity of demand by the firm's Lerner index, except that the marginal cost c_i is corrected for supplier i 's social responsibility index. Condition (13) points at two factors of strategic complementarity (how a_j affects the choice of a_i):

- *Elasticity effect:* Because price and moral choices jointly determine the net price ($\hat{p}_i = p_i + \phi_i(a_i)$), a strategic complementarity of moral choices is inherited from the strategic complementarity in the price domain.
- *Social responsibility effect:* Because $\partial S_i / \partial a_j < 0$ (stealing market share away from a moral supplier j is less morally attractive) under UPI consumers, a higher a_j increases the incentive to raise a_i , creating a second source of strategic complementarity.

Given strategic complementarity, it is straightforward to construct examples with multiple equilibria in the choice of actions (even symmetric ones): See Online Appendix D. For this reason, the following analysis requires conditions ensuring the existence of a unique equilibrium (alternatively, we could obtain monotone comparative statics).

²⁹This can be captured by scaling the internalization parameters by some λ : $\lambda \alpha_i$. Then for $\lambda \leq \bar{\lambda}$ for some $\bar{\lambda} > 0$, ethical choices are strategic complements.

4.2 Regulated prices in symmetric oligopoly

Definition (*symmetric oligopoly*). The oligopolistic market is symmetric if

- (i) the functions ϕ , c and W are the same for all firms;
- (ii) suppliers have symmetric demand functions ($D_i(\hat{p}_i, \hat{\mathbf{p}}_{-i})$ is invariant to permutations of $\hat{\mathbf{p}}_{-i}$ and $D_j(\hat{p}_i, \hat{\mathbf{p}}_{-i}) = D_i(\hat{p}_i, \hat{\mathbf{p}}_{-i})$ for all $(\hat{p}_i, \hat{\mathbf{p}}_{-i})$), and the same social preferences ($\alpha_i = \alpha$ for all i); and
- (iii) the market is covered.³⁰

Suppose that prices are regulated at the same level p . By “symmetric equilibrium”, we will mean an equilibrium in which all suppliers pick $a_i = a^*$ for some a^* , and the market is covered. The strategic complementarity between moral choices (Proposition 3) that always obtains under UPI consumers and may obtain under ethical consumers makes multiple equilibrium moral norms common. Online Appendix E establishes assumptions that guarantee equilibrium uniqueness and allow us to prove the comparative statics stated in the following proposition. For the purpose of Proposition 4, we will assume that on $[0, \bar{a}]$ the function $\alpha W'(a) - c'(a) - \phi'(a)[\eta(p, a)L(p)]$ is decreasing in a ; this assumption results from our previous assumptions under flexible prices (which guarantee that $\eta L \equiv 1$), but must be added to obtain a unique solution when prices are regulated.

Proposition 4 (*impact of competition on ethics under regulated prices*). Consider a symmetric, covered-market oligopoly equilibrium. The symmetric-equilibrium level of ethics is given by

$$\frac{\alpha W'(a^*) - c'(a^*)}{\phi'(a^*)} = \eta(p, a^*)L(p), \quad (14)$$

where $L(p) \equiv \frac{p-c}{p}$.

- (i) (*Elasticity of demand*). Suppose that an exogenous parameter $\sigma \in (0, \infty)$ (e.g. a substitutability parameter³¹ or the number of firms) moves the elasticity $\eta(p, a^*; \sigma)$ with $\partial\eta/\partial\sigma > 0$ and $\lim_{\sigma \rightarrow \infty} \eta = +\infty$. Then, with UPI (resp. ethical) consumers, the equilibrium level of ethics a^* is decreasing (resp. increasing) in the intensity of competition (σ).

³⁰A covered market combined with the symmetry among the n suppliers will imply that the firm’s social responsibility index is equal to 0 in equilibrium. In contrast, if the outside option has positive market share, there is no reason why the associated welfare, w_0 , be equal to the welfare generated by the suppliers, w^* , and that $S_i = 0$ in equilibrium. In general, there cannot be symmetry between the options chosen by the consumers if the outside option has positive market share.

³¹Suppose for instance a linear demand system: $D_i = \frac{1}{n} - \sigma[\hat{p}_i - \frac{\sum_{j \neq i} \hat{p}_j}{n-1}]$. Then, at a symmetric equilibrium, $\eta(p, a^*)L(p) = [\sigma n p][\frac{p-c}{p}] = \sigma n(p-c)$. So the substitutability parameter σ and the number of firms n are alternative measures of how competitive the industry is.

(ii) (*Regulated price level*). Under the (weak) condition that $\partial(\eta(p, a)L(p))/\partial p > 0$, the equilibrium level of ethics a^* with UPI (resp. ethical) consumers is decreasing (resp. increasing) in the fixed price (p).³²

Proof of Proposition 4. Differentiating equation (14) yields

$$\begin{aligned} \frac{\partial}{\partial a} [\alpha W'(a) - c'(a) - \phi'(a)\eta(p, a)L(p)] da - \frac{\partial}{\partial p} [\eta(p, a)L(p)] dp \\ - \frac{\partial}{\partial \sigma} [\eta(p, a)L(p)] d\sigma + W'(a) d\alpha = 0. \end{aligned}$$

The comparative statics then follow. ■

Proposition 4 demonstrates the sharp contrast between the case of UPI consumers (for which an increase in competition or in price makes the market less moral) and that of ethical consumers (where they make the market more moral). Intuitively, a supplier who cannot lower price is left with a single margin: the only margin the supplier can use to compete for customers is to reduce the morality of the offering when facing UPI consumers. When competition becomes more intense, the morality of offerings therefore declines. Market outcomes are also likely to be immoral when profit opportunities are sizeable, i.e., when the regulated price p is high. When the market is almost perfectly competitive, the only possibility for a supplier to keep market share is to select the most immoral action to attract some UPI consumers.

4.3 Asymmetries and financial viability

Regulation is only one reason why a firm's price may not be flexible. A break-even constraint may prevent the supplier from cutting price below cost. Such a downward price rigidity may in turn originate from an asymmetry in preferences (say, a firm is more virtuous than its rivals and may lose money) or in corporate charter (the firm may have the nonprofit status, say). Regarding the latter possibility, it is often argued that industries that are highly exposed to ethical choices, such as health and education, are particularly suited to the non-profit paradigm.³³ Is this so? Should we expect not-for-profit hospitals or schools to behave differently when in competition with for-profit entities? Furthermore, one would want to understand how competition among for-profits with different ethical objectives plays out. To contrast it with Section 4.2, we assume in the entire section that prices are unregulated.

³²With ethical consumers, when $p = c$, we also have the socially-efficient level of ethics \bar{a} ; with higher prices, a^* can thus only go down; but for very high prices, raising a_i is very attractive, since it is the only way to gain market share and so there is an incentive to go all the way to \bar{a} . Of course, for very large p , the assumption that the market is covered becomes much less plausible.

³³Indeed, many health and school providers around the world are not-for-profit entities, when not state-owned.

To encompass both forms of asymmetry within a single framework, we allow suppliers to differ in their corporate forms and/or their ethical values. To avoid compounding multiple sources of heterogeneity, we assume that the suppliers face symmetric demand and cost functions. Suppliers $i \in \{1, \dots, n_1\}$ are for-profit suppliers ranked by the intensity of their social preferences:³⁴

$$0 < \alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_{n_1}.$$

We keep assuming ethical welfare internalization (even though the results are much more general). As earlier, we denote by a_i^\dagger the supplier-specific level of morality given by condition (3) ($\alpha_i W'(a_i^\dagger) \equiv c'(a_i^\dagger) + \phi'(a_i^\dagger)$), with $a_1^\dagger \leq \dots \leq a_{n_1}^\dagger$. Suppliers $i \in \{n_1 + 1, \dots, n\}$ are not-for-profits; note that the absence of profit motive implies that their objective function is $\alpha_i \mathcal{W}_i$, and so their social preferences do not matter whenever $\alpha_i > 0$, which we will assume.³⁵

To handle such asymmetric environments, we further strengthen our assumptions:

Assumption 4 (*linear demand system, covered market*) *In the relevant-prices range, the demand system is $D_i = \frac{1}{n} - \sigma \left[\hat{p}_i - \frac{\sum_{j \neq i} \hat{p}_j}{n-1} \right]$, and so the market is covered.*

An important property of this linear demand system is that firm i 's change of behavior impacts other suppliers symmetrically.³⁶ Indeed, letting $\bar{w}_{-i} \equiv \frac{\sum_{j \neq i} w_j}{n-1}$ denote the average welfare footprint of i 's rivals, the social responsibility index is

$$S_i = \sigma(w_i - \bar{w}_{-i}).$$

Assumption 5 (*financial viability*) *Suppliers must be financially viable: $p_i \geq c(a_i)$ for all i .*

To illustrate the rationale for Assumption 5 in the case of for-profits, consider an otherwise symmetrical duopoly situation in which one supplier is more ethical than its

³⁴ $\alpha_1 = 0$ is allowed as well (taking the limit as $\alpha_1 \rightarrow 0$). We assume $\alpha_1 > 0$ for expositional simplicity.

³⁵We could assume that firms with different corporate status attract employees with different social preferences (see e.g. Besley-Gathak (2005), Prendergast (2007), Brekke-Nyborg (2008), Kosfeld-von Siemens (2011), Lazear et al (2012), Barigozzi-Buranib (2019); for field experiments on sorting and prosociality, see Ashraf et al (2020) and the references therein). A motivation for this assumption on the empirical side is assortative matching (not-for-profits attract more ethical employees), although it is not clear that working for a not-for-profit is necessarily the moral thing to do for someone who wants to have a strong ethical impact (Singer 2015). The same holds for entry decisions into an industry. It may well be that entering an immoral industry in which one can make a difference is more moral than entering an ethical one (Moisson (2020) shows that the moral pecking order is highly context specific; a known example of this general point concerns socially responsible investment, for which best-in-class strategy may have a bigger impact than the exclusion of sin stocks. See also Green-Roth (2023)).

Of course, there may be no such thing as a pure not-for-profit. Insiders may manage to convert profits into private benefits; private benefits are an inefficient currency, but more to the point, such conversion of profits would reinstate a role for the not-for-profit suppliers' exact level of altruism.

³⁶For example, one can construct strongly asymmetric linear demand systems for which Proposition 6 does not hold.

rival, prices are flexible, cutting ethical corners boosts demand (UPI consumers) and not cost ($c(a_i) = c$ for all a_i) and finally there exists an interior welfare-maximizing action \bar{a} (such that $W'(\bar{a}) = 0$).³⁷ Supplier 1 is selfish ($\alpha_1 = 0$), and therefore selects $a_1 = 0$; supplier 2 is a saint ($\alpha_2 = +\infty$), and therefore, in the absence of financial constraint, selects $a_2 = \bar{a}$ and is willing to set any price that will take market share away from firm 1: A deep-pocket, very ethical supplier would lose money when facing a much less ethical rival.³⁸

Returning to the general n -firm model, Assumption 5 deserves a couple of further comments. First, ignoring the issue of access to capital, Assumption 5 is irrelevant when differentials in social preferences are “not too large”; what this exactly means depends on the intensity of competition.³⁹ Second, Assumption 5 is innocuous in the absence of investors who have strong social preferences and are willing to foot the bill for virtuous actions. To be certain, one can think of undertakings that are financed by such investors (like some big NGOs or foundations), but the thrust of the debate on market morality is on firms that must at the very least break even (whether for-profits or not-for-profits). Third, we distinguish three groups when describing equilibrium behavior. Unconstrained for-profits select a positive mark-up ($p_i > c(a_i)$). Proposition 1 then implies that $a_i = a_i^\dagger$. For these suppliers, $w_i \equiv W(a_i^\dagger)$. Constrained for-profits have no mark-up ($p_i = c(a_i)$) and therefore behave like suppliers in the third group, the not-for-profits. We gather the latter two groups under the heading “constrained suppliers”.

We will say that there is a *race to the supplier ethical bottom* if

$$\lim_{\sigma \rightarrow +\infty} a_i = a_1^\dagger \quad \text{for all } i.$$

In particular, in the limit in which one of the suppliers is a pure profit maximizer, a race to the supplier ethical bottom implies that competition prevents any pro-social behavior originating from *supplier* social preferences; a_1^\dagger still reflects the stakeholders’ social preferences, though, and therefore ethical behavior need not converge to 0.

Proposition 5 (*behavioral convergence and race to the supplier ethical bottom*). *Assume n_1 for-profits with social preferences $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_{n_1}$ and $n_2 = n - n_1$ not-for-profits, and flexible prices. Under Assumptions 4-5,*

- (i) *Not-for-profits behave more ethically than for-profits (there exists a^* such that $a_i = a^* \geq a_j$ if $i > n_1 \geq j$). Furthermore, there exists $1 < m \leq n_1$ such that $a_i = a_i^\dagger$ for $i \leq m$ and $p_i = c(a_i)$ and $a_i = a^*$ for $i > m$. That is, constrained suppliers (all not-for-profits and those for-profits who are financially constrained) adopt the same moral behavior and are more virtuous than the financially unconstrained for-profits.*

³⁷ \bar{a}_i is finite in all examples provided as microfoundations in Online Appendix A.

³⁸For perfect substitutes, supplier 2 loses $\phi(\bar{a}) - \phi(0) > 0$.

³⁹For example, for a duopoly (with $\alpha_2 > \alpha_1$) and demand-based benefits from unethical behavior (UPI consumers), a sufficient condition for the financial constraint not to bind is $1 \geq 2\sigma\alpha_2(w_2 - w_1)$, where $\alpha_i W'(a_i) = \phi'(a_i)$ and $w_i \equiv W(a_i)$.

Assume UPI consumers. Then:

- (ii) The behaviors of all suppliers converge when competition (as indexed by σ) is intense: The for-profits mimic the not-for-profits' low price ($p_i \rightarrow c(a_i)$ for all i as $\sigma \rightarrow +\infty$), while the latter behave no more ethically than for-profits: There is a race to the supplier ethical bottom: $a_i \rightarrow a_1^\dagger$ for all i as $\sigma \rightarrow +\infty$.
- (iii) Suppose that initially there are only not-for-profits. Under intense competition, the entry of a single for-profit changes the not-for-profits' moral behavior from the socially optimal level \bar{a} to the low level a_1^\dagger (and maintains the price close to marginal cost).

Part (i) of the Proposition (proved in Online Appendix F) says that the more virtuous among the for-profits are financially constrained and therefore behave like not-for-profits. Their scruples makes them less attractive (in the case of demand-based benefits of unethical behavior) or face a cost disadvantage (for cost-based benefits), making it more difficult to compete for market share and even to break even. This holds for any intensity of competition σ . Part (ii) looks at intense competition. Under intense competition for consumers, suppliers end up charging similar net prices. The for-profits must lower their markup toward 0 to not lose all demand, while with UPI consumers the not-for-profits must pander at (approximately) level a_1^\dagger for the same reason. Competition homogenizes behavior across corporate forms and ethical preferences. Convergence happens toward the low-price, low-ethics “anchor” ($p = c(a_1^\dagger)$, $a = a_1^\dagger$). Thus not-for-profits have no influence on the market when competition is intense.

Does intense competition crowd out moral behavior? Proposition 5 indicates that intense competition for UPI consumers crowds out supplier ethics. If firms are all for-profit, formula (10) derived in the absence of financial constraint in Section 3.1 for each supplier i :

$$(\alpha_i + \alpha_W + \alpha_I)W_i'(a_i) - \phi_i'(a_i) = \gamma_i'(a_i)$$

under financial viability becomes in the limit as competition becomes very intense:

$$(\alpha_1 + \alpha_W + \alpha_I)W_i'(a_i) - \phi_i'(a_i) = \gamma_i'(a_i).$$

While intense competition for UPI consumers crowds out supplier ethics, it does not do so for the ethical impact of stakeholders.

Should we expect α_1 to be close to zero when competition in the market is intense? In many countries “shareholder value” has become the key force determining firm behavior (with top management being largely paid in stocks and stock options). But, as argued by Broccardo et al. (2022), this need not imply a pure for-profit behavior without any ethical consideration, since shareholders may have social preferences too. Moreover, the emergence of very large asset managers voting “on behalf of” diversified shareholders and owning stakes in many big players of key markets in the economy is a trend that could reduce the variance of supplier altruism.

Finally, Online Appendix G analyzes when competition for UPI consumers should be expected to weed out suppliers with high ethical standards or a nonprofit status (Gresham’s law).

5 Relevance

5.1 Welfare

Whether markets increase or decrease ethical behavior is a positive question. For a normative analysis, we must remember that the drivers of the intensity of competition (industry structure and public policy) have effects of their own, even in an homo-economicus, no-social-preferences world. Moreover, taking a welfare stance in an environment with social preferences requires making some further assumptions as to how these preferences are accounted for in the social welfare function.

Online Appendix H discusses these choices in detail, and derives several insights, assuming as usual that social welfare accounts for consumer welfare as well as (with a smaller weight) supplier profit. In Propositions 4 and 5, the intensity of competition changes equilibrium ethics and therefore welfare. But what about Proposition 1? The following intuitions and associated results are robust.

First, when suppliers are identical⁴⁰ and the market is covered, Proposition 1 implies that the optimal public policies (merger reviews, trade arrangements, transparency requirements, antitrust oversight...) remain unchanged under flexible prices, as the intensity of competition is irrelevant to moral behavior. Put differently, there is no need for revisiting our economics corpus of knowledge to account for social preferences. Second, when the market is not covered, a new effect is at play: whether the outside option is more or less moral than the competing offers. As we noted for energy consumption, the alternative may be energy sobriety (more moral regarding CO2 emissions) or the use of coal or deforestation (less moral). Online Appendix H obtains the following intuitive result: *Ceteris paribus*, a pro-competitive policy that increases welfare in the absence of social preferences ($\alpha_i = 0$ for all $i \geq 1$) a fortiori increases welfare whenever the outside option is no more moral than the market ones (for example, $a_0 \leq a^\dagger$ if the equilibrium is symmetric). In contrast, when the outside option is more moral than the market ones (for example, $a_0 > a^\dagger$ if the equilibrium is symmetric), then the analysis is ambiguous: Unless social preferences are weak, a pro-competitive policy that increases welfare in the absence of social preferences may decrease welfare in their presence.

⁴⁰When the set of suppliers expand in reaction to the policy (we here have in mind a trade opening) and the new suppliers may face different policy environments – think about GHG emissions or child labor –, then a policymaker with their own social preferences may revisit this general principle.

5.2 Connection to real world markets

While the broad question of the morality of markets is ancient, it has been prominent in some recent policy discussions.

First in the matter of antitrust policy. The Biden administration’s heads of the DOJ and the FTC, and the White House advisor on competition policy have pledged to crack down on buyout groups and their “buy, strip and flip model”. Targeting private equity firms as deal sponsors is new territory, as it departs from the antitrust focus on conducts and transactions. This debate on private equity and antitrust would be meaningless, were all private entities pure profit maximizers as is assumed in much industrial organization. Instead, the underlying view here is that some entities are more assertive profit maximizers, which may create more collateral damages for some stakeholders. Proposals for the revision of the 2011 antitrust guidelines in Europe have also put moral issues at the center stage.⁴¹

Our theory sheds light on the link between the intensity of competition and equilibrium ethics in a world where “intrinsic ethical urges” (the α ’s, which among other things reflect the (inverse) power of the individual supplier’s incentives) can vary. However our irrelevance result shows that reducing competition per se is unlikely to solve the problem. In fact, Proposition 1 suggests that competition authorities can safely push for more competition without having to fear negative ethical consequences, at least as long as its actors do not significantly differ in terms of greed.

Ethical debates linked to the intensity of the pursuit of profit are, unsurprisingly, ubiquitous in the healthcare sector, an area where ethical stakes are very high as patients are often ill-informed.

Scholars have studied the competition between not-for-profit and for-profit hospitals. Proposition 5 is consistent with evidence on the hospital sector. As argued in classic work by Weisbrod (1988) and Hansmann (1996), not-for-profits have historically been an important commitment device against excesses associated with the profit motive (see also Besley-Ghatak 2005). In recent decades, though, for-profit hospitals have made inroads in the sector, and, unsurprisingly, have been shown to put more emphasis on profit-related managerial compensation (Ballou-Weisbrod 2003) in comparison with their not-for-profit peers, consistently with part (i) of Proposition 5. In support of part (ii) of that proposition, Arnould et al. (2005) show that more competition from for-profit hospitals leads to a higher importance of the “profit motive” (i.e. net financial income) among not-for-profit ones, both in terms of the structure of managerial compensation

⁴¹They “aim to make it easier for undertakings to cooperate in ways which are economically desirable and thereby, for example, contribute to the green and digital transitions and to fostering the resilience of the internal market”. This statement is directly connected with what the social responsibility of business should be. See also the sustainability chapter (chapter 9) of the recent EC Guidelines on the applicability of Article 101 of the Treaty on the Functioning of the European Union to horizontal co-operation agreements.

and of managerial turnover decisions (and this is understood by donors, who reduce their contribution as a result of this weakening of the not-for-profit mission).

Observers have also contrasted the ethical performances of for-profit entities in the pharmaceutical industry differing in their governance. For example, US generics drug maker Impax sale in 2015 of its US rights to the Daraprim brand to privately-held Turing Pharmaceuticals was blamed for the 56-fold increase in the price of this antiparasitic drug, hurting patients. Another spectacular example is Purdue Pharma, a privately-held family firm that became hugely profitable through aggressive and deceptive marketing.⁴² In conformity with Propositions 3 and 5, the strategy of Purdue Pharma, the undisputed leader in the opioid crisis, had a strong contamination effect on other players in the industry.

Observers have more broadly expressed concerns about private equity (PE) groups' impressive indent into the US hospital sector.⁴³ In this respect, the nursing-home-sector study by Gupta et al. (2021) concludes: *“PE ownership increases the short-term mortality of Medicare patients by 10%, implying 20,150 lives lost due to PE ownership over our twelve-year sample period. This is accompanied by declines in other measures of patient well-being, such as lower mobility, while taxpayer spending per patient episode increases by 11%. We observe operational changes that help to explain these effects, including declines in nursing staff and compliance with standards. Finally, we document a systematic shift in operating costs post-acquisition toward non-patient care items such as monitoring fees, interest, and lease payments”*.

Our model indicates that the governance of suppliers matters. While it is difficult to observe the moral preferences of managers, our theory indicates that high-powered incentive schemes tend to reduce market morality, consistently with Gupta et al. (2021). And that the existence of for-profit suppliers may be ethically “dominant” in that not-for-profit suppliers may have to mimic for-profits' low-ethics choices if consumers are UPI and competition is intense, consistently with Arnould et al. (2005). On a more positive note, the ethical urges of other stakeholders (responsible consumers, workers and investors) will not be hampered under such circumstances, and can be “encouraged” for example by the transparency of supplier ethical choices.

⁴²Downplaying the addiction risk of its blockbuster Oxycontin, tying half of the compensation of its representatives to the prescription behavior of “their” medical doctors, and offering Oxycontin samples that would be free only for a limited time periods. These various practices were already discussed in a 2003 official report (see US General Accounting Office, 2003).

⁴³Ethical concerns were for example relayed by Robeznieks (2022), who summarizes the conclusions of a recent American Medical Association roundtable as follows: *“PE funds can help spur innovations or provide stable funding for workers' pensions, but investor expectations for a quick return on investment may clash with a medical practice's long-term sustainability and physicians' ethical demands”*. And he adds that, in this roundtable, the immediate past chair of the AMA-PPPS (Private Practice Physician Section) Governing Council, noted that *“this group, more than others, is not anti-profit given that section members often view themselves as—among other things—small business owners. But the problem is if the profit is a beginning and an end to itself, added an AMA member. It has no empathy, sympathy or engagement with the consumer... which is the patient”*

Overall, our model stresses that, for markets where ethical worries are significant (due to externalities, internalities or incomplete information) and hard to regulate away, there is value in mitigating the pursuit of profit, a concern of the by-now large CSR/ESG literature. Our value added here is to say that policy instruments on this dimension, from transparency to board composition to the choice of legal form (like that of a Benefit Corporation, which protects managers in case they diverge from pure profit maximization to pursue predefined societal objectives) should be the focus of attention, rather than trying to weaken product market competition.

5.3 Experiments

Although not initially designed for this purpose, the evidence on the replacement effect can be related to our theoretical framework. Bartling et al (2015) run experiments in settings similar to our “ethical consumers/cost benefits from cutting ethical corners/flexible prices” case. “Sellers” have two margins: They both set prices and a production technology: They choose between a costly & clean good and a cheap & polluting one. Bartling et al ask, will the cheapest, polluting good be delivered in a competitive market, as the replacement effect would suggest? They find, to the contrary, that “increased competition does not diminish the degree of concern exhibited toward externality-bearing parties outside of the market”, consistently with Proposition 1.

Falk et al (2020), in a one margin environment, show that (the perception of) pivotality is key to sustaining moral behavior. Their baseline experiment has full pivotality, with a single subject deciding between “killing a mouse” (not saving a surplus mouse) and forgoing 10 €. In the treatment, each subject can choose between unconditionally forgoing 10 € and giving the mouse a chance to survive, which will happen only if all 7 other subjects also abandon 10 € (which is unlikely). Many more choose to keep the 10 €. This experiment points at consequentialist preferences rather than deontological ones and at the potency of the replacement effect.⁴⁴ The Falk et al experiment is not set as a market but has the same features that (a) an unethical behavior boosts profit, (b) subjects have a single margin, their ethical choice, as in the fixed-price environment, (c) ethical behavior is less appealing to the supplier if others also behave unethically (strategic complementarity/Proposition 4), and (d) the probability that one’s morally

⁴⁴Bartling and Ozdemir (2021) demonstrate that the replacement effect is less prevalent when there is a strong social norm.

correct action delivers the morally right outcome decreases with the number of decision makers.^{45, 46} Their result is thus reminiscent of the forces at play in our Proposition 4.

6 Related literature

Dufwenberg et al (2011) and Sobel (2015), like us, assume non-standard preferences, and then derive conclusions regarding the nature of market outcomes. Both papers derive conditions under which other-regarding preferences (ORPs) make no behavioral difference relative to selfish preferences. In both papers, the absence of market power is key to the result. In Dufwenberg et al. (which only allow for consumer, and not producer, ORPs), this is in the context of a Walrasian setting, in Sobel perfect competition emerges as the limit of a standard double auction (with one good and money) with anonymous trading in large economies. Sobel also extends the analysis to the existence of market power and identifies conditions under which ORPs do not make a difference either.

At first glance our results may seem to be drastically at odds with those in Dufwenberg et al and Sobel. With flexible prices, our firms adopt the same behavior regardless of the intensity of competition: In many circumstances, a monopolist behaves as morally as firms under intense competition; what is irrelevant for moral behavior in our model is the intensity of competition, not social preferences like in these two papers. The difference in conclusions naturally can be traced to the different assumptions.

Dufwenberg et al and Sobel assume that one can only affect others' utilities through one's impact on others' quantities traded or the market price. Dufwenberg et al study a standard multi-good Walrasian setting. Indeed, next to a separability assumption (consumers' ordering over feasible consumptions is independent of other's choices, an assumption we also make), they assume that consumer i 's preferences can be represented by a utility function $V_i(m_i(x_i), x_{-i}, B)$ where $m_i(x_i)$ is the material utility from consumption vector x_i , x_{-i} is the vector of consumptions by others, and B are the agents' budget sets. This framework allows for externalities as well as inequality aversion (Velez 2017),

⁴⁵Let $a_i \in \{0, 1\}$ denote subject i 's moral action, and $a_{-i} \equiv X_{j \neq i} a_j$. Letting v denote the value of a mouse's life, subject i 's payoff is (proportional to) the sum of a material payoff and social preferences: $-10 \frac{a_i}{a_{-i}} + \alpha v a_i$. Note that a_i and a_{-i} are complements in the material payoff term. The non-material payoff in this writing of the preference function is of the narrow bracketing type. Note also that this expression shows that an increase in the number of subjects (which induces a reduction in a_{-i}) is similar to an increase in the subject's power of incentive scheme.

⁴⁶This impact of "pivotality", which decreases with a higher number of competitors, on ethics is also discussed in the context of voting models. Feddersen et al (2009) find that "ethically expressive" motives become more important relative to material self-interest in larger populations. In contrast, Kamenica and Egan Bard (2014) make a distinction between two non-material-self-interest forms of preferences: (a) the benefits of outcome-based ideology/social preferences (I care about my ideology being implemented, i.e., being voted for by a majority of voters) and (b) the expressive utility brought about by the match between my vote and my ideology (a form of warm-glow). While the former, as well as material preferences, matter only if the voter is pivotal, the latter does not. Kamenica and Egan Bard find that it is the former rather than the latter that matters, so that pivotality does not really affect the tradeoff between self-interest and ethics in their experiment.

but they exclude some key consequentialist internalizations: in particular, they do not consider as ORP the fact that an individual may want to change her consumption basket just because it is objectionable to others, even if this does not affect their ability to trade. Another difference with their framework is that decisions are interdependent in ours: A supplier’s moral action conditions the support of its stakeholders and therefore affects the supplier indirectly as well as directly. Finally the additional assumption on preferences which guarantees irrelevance of ORPs is that individuals prefer to make a desirable trade themselves rather than let another individual make exactly the same trade, an assumption which we also make but is not consequential in our framework.

Next to our general irrelevance result, we provide a precise identification of environments in which the intensity of competition makes markets more or less moral. While in the limit supplier ethics may be crowded out fully (i.e. only the lowest α_i matters), stakeholders’ ethical urges remain relevant even under these circumstances.

The paper also has a strong connection with the corporate social responsibility (CSR) literature.⁴⁷ A prominent view of CSR equates it with “delegated philanthropy”. The firm is a channel for the expression of citizen values; as in our model, consumers may be willing to pay a bit more for fair coffee,⁴⁸ investors may accept getting a smaller return from green funds, and workers may take a wage cut when employed by an NGO. A profit-maximizing company then maximizes profit as they pass through the higher cost or the lower return to the stakeholders. This view is embraced in Aghion et al. (2023), Bagnoli-Watts (2003), Besley-Ghatak (2007), Besley-Persson (2020), Green-Roth (2023), Kotchen (2006), Landier-Lovo (2020), Moisson (2020), Oehmke-Opp (2023), Barigozzi and Tedeschi (2015, 2019) and Weber and Zhang (2021). Weber and Zhang find experimental support for our result that when consumers are willing to pay more when the supplier stands for their values, competition fosters ethical behavior; they show that the suppliers react by incurring costs to express support to the causes that are favored by the buyers. Aghion et al show, theoretically and empirically, that competition pushes profit-maximizing suppliers toward greener innovation. The latter result does not contradict Proposition 1, as it is based on an IO mechanism (escaping competition effect) and not on the crowding out of supplier morality (in their model only consumers have social preferences).

⁴⁷See, e.g., the taxonomy in Bénabou-Tirole (2010). The suppliers’ role in shaping the morality of markets is in line with Henderson (2020)’s view of managers as key engines for “reimagining capitalism”. That economic agents in general may behave altruistically has received support in experimental economics and is a common assumption in the theoretical literature on social responsibility (see e.g. Besley-Ghatak 2018, Broccardo et al 2022, Green-Roth 2023, Hart-Zingales 2017, Landier-Lovo 2020, Oehmke-Opp 2023).

⁴⁸Conversely, ethical consumers can boycott firms that behave unethically, in the tradition of Baron (2001) and subsequent papers of his and Egorov and Harstad (2017) in a dynamic context. Feddersen and Gilligan (2001) show that “activist intermediaries”, who are better informed than consumers about supplier behavior, can help coordinate such boycott strategies and thereby push supplier actions towards more ethical behavior.

An alternative view of CSR is “insider-initiated corporate philanthropy”, namely philanthropy that clashes with profit-maximization.⁴⁹ This is the approach taken in Hart-Zingales (2017) and Broccardo et al (2022), where shareholders compare their monetary gains with the ethical impact of their actions. This tradeoff has “bite” when they vote at the general assembly or board of directors, since both impacts are non-zero only if their vote is pivotal. By contrast, this leads them to focus solely on monetary gains when they buy shares (there is no socially responsible investment), since they rationally expect not to be pivotal and therefore affect the company’s future actions only with a tiny probability, a “leakage” that is also present in Green-Roth (2023) and Moisson (2020). Broccardo et al (2022) extend the analysis in a model where they endogenize investor divestments and consumer boycotts (which they call “exit” mechanisms) where individual investors and consumers internalize their (nonzero) impact on firm behavior on aggregate social surplus. In their model, under social preference parameters consistent with experimental evidence, divestments and boycotts are insufficient and shareholder engagement through voting (“voice”) is socially preferable.⁵⁰ This “insider-initiated corporate philanthropy” literature can be seen as an ‘input’ to our model in that it focuses on how institutions shape suppliers’ ethical urges, i.e. their α_i ’s where we then look at how equilibrium ethics results from the α_i ’s and the intensity of competition.

To sum up, our paper belongs to these two literatures, as we allow both the supplier and the stakeholders to have social preferences and allow ethical choices to maximize corporate profits or to reduce them. Its unique focus is on the impact of the intensity of competition on market morality and on the predictions of heterogeneity in preferences and corporate form for moral behavior.

Finally, we have a model with two strategic variables, p (or q) and a , and we look at the interplay between the two as a function of the intensity of competition. Some models in the literature similarly have effort or quality instead of a . The multi-task incentive literature (e.g., Holmström and Milgrom 1991) stresses that high-powered incentives by a principal may compromise the agent’s provision of non-contractable quality.⁵¹ Relatedly,

⁴⁹Even leaving aside the agency literature, there is of course a long tradition of analyses of non-profit-maximization goals: Beckerian discrimination theory, labor-managed firms, etc. Becker (1957) made the point that (perfect) market competition weeds out those suppliers who have a preference for discrimination. There is complementary with our results, since he considers situations where suppliers “enjoy” an immoral behavior that raises the cost of business, namely the wage bill. He also argues that purely profit-maximizing (and thus unprejudiced) suppliers will “cater” to the prejudices of consumers. This is consistent with the results derived from limit results of our model when $\alpha_i = 0$.

⁵⁰Oehmke-Opp (2023) also emphasizes the benefits of voice exerted by socially responsible investors; in their paper, the latter relax the firm’s borrowing constraint conditional on choosing a clean production process. A recent paper by Herweg and Schmidt (2022) makes the point that managers’ ability to express their social responsibility depends on the institutions designed by the state. They compare cap-and-trade mechanisms and carbon taxes as alternatives for putting a price on carbon. Consequentialist managers behave solely in function of their material interests under a cap-and-trade as total pollution is fixed.

⁵¹Where quality here is viewed from the principal’s standpoint. In Lazear (1989), two workers are engaged in a tournament. The relative performance determines individual pay raises, which is conducive to “sabotaging”. Itoh (1991) studies optimal incentives for team workers who have individual performance measures but help each other.

the paper connects to the literature on not-for-profit firms. This literature emphasizes that the absence of profit motive reduces the incentive to cut on unobservable quality (Hansmann 1980, Glaeser-Shleifer 2001, Bubb-Kaufman 2013, Besley-Malcomson 2018).⁵² Our paper is complementary: It mostly assumes by contrast that the “quality” assessed by consumers is observable (directly, or through word of mouth or reputation) but not necessarily socially desirable; and it looks at a different set of issues (e.g., the convergence of behavior of for- and not-for profit firms as a function of the degree of competition).⁵³

7 Summing up

Critics of market economies have long emphasized that the institutional context may frame our ethical choices. Does that mean that competition- understood as an increase in the number of competitors or in their substitutability or a decrease in search costs- may strengthen incentives to cut ethical corners in order to please the consumer or to cut costs? The paper develops theoretical foundations for this concern, providing its rationale, the reasons why moral choices are often strategic complements, and an exact identification of the environments in which intense competition affects ethical choices.

The paper embodies two main contributions. First, and importantly for the public debate and public policies, it offers a strong warning against a sweeping condemnation of the market based on the ground that it promotes immoral behavior. Indeed, our central irrelevance result robustly shows that the intensity of competition does not affect behavior as long as (1) suppliers and stakeholders are consequentialists, (2) prices are flexible, an assumption that describes well many markets, and (3) technology is characterized by constant returns to scale (understood as the marginal cost of ethical choices being proportional to output). What determines equilibrium ethics in a market is then the set of ethical urges of the players, not the intensity of competition. Overall, the presumption should be that competition, unlike the values of the players, cannot be the overriding source of moral problems in trade; at the very least, it is ill-advised to blame the market

⁵²For instance, Besley and Malcomson posit that not-for-profits internalize the benefits of various dimensions of quality, although maybe in a paternalistic fashion. Their focus is on the ease of entry by a non-profit facing a for-profit incumbent, and variations thereof, to match the observations on entry in the school and hospitals sectors. Bubb and Kaufman show how ownership of the firm by its customers, as well as nonprofit status, can prevent firms from using contractual terms that take advantage of consumer biases in consumer financial services.

⁵³The IO literature on competition and incentives does stress the role of product market competition on firms’ non-price behavior. In that literature, a firm’s manager picks an effort under profit-based compensation, in the same way our suppliers pick a moral action and not solely a price. The “principal” of the IO literature corresponds to the stakeholders in our model, who demand a higher moral action; but there is no counterpart in the IO literature to our UPI consumers, who play a key role in the replacement effect literature. Besides the rather distinct motivations, the mechanisms described in the literature whereby competition may enhance effort (or not) are different from those in our paper : for example, the information or benchmarking route in Hart (1983) and the desire to avert bankruptcy in Schmidt (1997) which both positively link competition and effort, and the “scale effect”, the idea that effort is a fixed cost which is less valuable under lower market share, which does the opposite in Raith (2003).

for immoral behavior and to question the appropriateness of competition policy, anti-bottleneck regulation, competitive procurement, and competition through trade, without specifying in detail the nature of competition.

The second contribution is to analyze environments where price flexibility does not apply, either because of regulation, or because of “corner solutions” due to a zero-profit constraint linked with asymmetric competition or a not-for-profit status. When prices are fixed by regulation and consumers are UPI, critics of the market are vindicated: more competition among symmetric suppliers fosters immoral behavior. In contrast, an increase in competition fosters moral behavior under fixed prices and ethical consumers.

When suppliers differ in their ethics, either intrinsically or because of their corporate mission (some actors’ prices being *de facto*, although not *de jure*, constrained, as they must equate revenue with cost), competition can also erode equilibrium ethics. Not-for-profits behave more ethically than for-profit suppliers; and among the latter, more ethical suppliers tend to behave more ethically than less ethical ones. But the key lesson is that intense competition in a UPI-consumers market leads to a race to the for-profit-supplier ethical bottom (without however changing the impact of stakeholders’ ethical urges on equilibrium ethics). This suggests in particular not mixing corporate forms within the same competitive markets if the goal is to encourage moral behavior.

We saw that the competitive pressure may leave morality unaffected, reduce it or increase it. Does anything go or is the theory testable? The answer is that it is testable, because it makes clear predictions within each situational context. Under consequentialism and flexible prices, we should expect little relation between ethics and the intensity of competition. Under regulated prices and ethics-contingent demand, consumer attitudes will instead be crucial. Take fair trade, say: rich-world consumers enjoy no direct gain from poor farmers’ getting a higher income. Their demand is entirely driven by social responsibility and so the prediction is that competition will improve moral behavior by empowering morally conscious consumers. In contrast, in the bribing, performance-enhancing drug, unneeded prescription of opioids or sick days, or product misrepresentation examples, immoral behavior boosts demand. The context offers a clue as to the relevant prediction; this also provides guidance for experimental work on ethical behavior.

References

- Aghion, P., Bénabou, R., Martin, R., and A. Roulet (2023), “Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty?,” *American Economic Review: Insights*, 5(1): 1–20.
- Anderson, E. (1993), *Value in Ethics and Economics*, Harvard University Press.
- Arnould, R., Bertrand, M., and K.F. Hallock (2005), “Does Managed Care Change the Mission of Nonprofit Hospitals? Evidence from the Managerial Labor Market,” *Industrial and Labor Relations Review*, 58: 494–514.
- Ashraf, N., Bandiera, O., Davenport, E. and S. Lee (2020), “Losing Prosociality in the Quest For Talent? Sorting, Selection, And Productivity in The Delivery of Public Services,” *American Economic Review*, 110(5): 1355–1394.
- Bagnoli, M., and S. Watts (2003), “Selling to Socially Responsible Consumers: Competition and the Private Provision of Public Goods,” *Journal of Economic Management and Strategy*, 12(3): 419–445.
- Ballou, J.P., and B.A. Weisbrod (2003), “Managerial Rewards and the Behavior of For-Profit, Governmental, and Nonprofit Organizations: Evidence from the Hospital Industry,” *Journal of Public Economics*, 87: 1895–920.
- Barigozzi, F., and N. Buranib (2019), “Competition for Talent when Firms’ Mission Matters,” *Games and Economic Behavior*, 116: 128–151.
- Barigozzi F., and P. Tedeschi (2015), “Credit Markets with Ethical Banks and Motivated Borrowers,” *Review of Finance* 19(3): 1223–1279.
- Barigozzi F., and P. Tedeschi (2019), “On the Credibility of Ethical Banking,” *Journal of Economic Behavior & Organization*, 166(C): 381–402.
- Baron, D. (2001), “Private Politics, Corporate Social Responsibility, and Integrated Strategy,” *Journal of Economics and Management Strategy*, 10(1): 7–45.
- Bartling, B., Weber, R. and L. Yao (2015), “Do Markets Erode Social Responsibility?,” *Quarterly Journal of Economics*, 130(1): 219–66.
- Bartling, B., and Y. Özdemir (2021), “The Limits to Moral Erosion in Markets: Social Norms and the Replacement Excuse,” mimeo.
- Becker, G. (1957), *The Economics of Discrimination*, University of Chicago Press.
- Bénabou, R., and J. Tirole (2006), “Incentives and Prosocial Behavior,” *American Economic Review*, 96(5): 1652–1678.
- Bénabou, R., and J. Tirole (2010), “Individual and Corporate Social Responsibility,” Coase lecture, *Economica*, 77: 1–19.
- Besley, T., and J. Malcomson (2018), “Competition in Public Service Provision: The Role of Not-for-Profit Providers,” *Journal of Public Economics*, 162: 158–172.

- Besley, T., and M. Ghatak (2005), “Competition and Incentives with Motivated Agents,” *American Economic Review*, 95: 616–636.
- Besley, T., and M. Ghatak (2007), “Retailing Public Goods: The Economics of Corporate Social Responsibility,” *Journal of Public Economics*, 91(9): 1645–1663.
- Besley, T., and M. Ghatak (2018), “Prosocial Motivation and Incentives,” *Annual Review of Economics*, 10: 411–438.
- Besley, T., and T. Persson (2020), “Escaping the Climate Trap: Values, Technology and Politics,” mimeo.
- Brekke, K.A., and K. Nyborg (2008), “Attracting Responsible Employees: Green Production as Labor Market Screening,” *Resource and Energy Economics*, 30(4): 509–526.
- Broccardo, E., Hart, O., and L. Zingales (2022), “Exit vs. Voice,” *Journal of Political Economy*, 130(12): 3101–3145.
- Bubb, R., and A. Kaufman (2013), “Consumer Biases and Mutual Ownership,” *Journal of Public Economics*, 105(C): 39–57.
- Dufwenberg, M., Heidhues, P., Kirchsteiger, G., Riedel, F. and J. Sobel (2011), “Other-Regarding Preferences in General Equilibrium,” *Review of Economic Studies*, 78: 613–639.
- Dufwenberg, M., Johansson-Stenman, O., Kirchler, M., Lindner, F. and R. Schwaiger (2022), “Mean Markets or Kind Commerce?,” *Journal of Public Economics*, 209, 104648.
- Egorov, G., and B. Harstad (2017), “Private Politics and Public Regulation,” *Review of Economic Studies*, 84(4): 1652–1682.
- European Commission (2011), *Guidelines on the Applicability of Article 101 of the Treaty on the Functioning of the European Union to Horizontal Co-operation Agreements*.
- Falk, A., Neuber, T., and N. Szech (2020), “Diffusion of Being Pivotal and Immoral Outcomes,” *Review of Economic Studies*, 87(5): 2205–2229.
- Feddersen, T., Gailmard, S., and A. Sandroni (2009), “Moral Bias in Large Elections: Theory and Experimental Evidence,” *American Political Science Review*, 103(2): 175–192.
- Feddersen, T., and T. Gilligan (2001), “Saints and Markets: Activists and the Supply of Credence Goods,” *Journal of Economics and Management Strategy*, 10:149-171.
- Gabaix, X., and D. Laibson (2006), “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 121(2): 505–540.
- Glaeser, E.L., and A. Shleifer (2001), “Not-for-Profit Entrepreneurs,” *Journal of Public Economics*, 81(1): 99–115.
- Green, D., and B. Roth (2023), “The Allocation of Socially Responsible Capital,” mimeo.

- Gupta, A., Howell, S.T., Yannelis, C. and A. Gupta (2021), “Does Private Equity Investment in Healthcare Benefit Patients? Evidence from Nursing Homes,” conditionally accepted, *Review of Financial Studies*.
- Hansmann, H., (1980), “The Role of Nonprofit Enterprise,” *Yale Law Journal*, 89: 835–901.
- Hansmann, H., (1996), *The Ownership of Enterprise*, Cambridge MA: Belknap Press.
- Hart, O., (1983), “The Market Mechanism as an Incentive Scheme,” *Bell Journal of Economics*, 14(2): 366–382.
- Hart, O., and L. Zingales (2017), “Companies Should Maximize Shareholder Welfare Not Market Value,” *Journal of Law, Finance, and Accounting*, 2: 247–274.
- Henderson, R. (2020), *Reimagining Capitalism in a World on Fire*, New York: Public Affairs.
- Herweg, F., and K. Schmidt (2022), “How to Regulate Carbon Emissions with Climate-conscious Consumers,” CEPR Discussion Paper No DP16985.
- Hirschman, A.O. (1977), *The Passion and the Interest: Political Arguments for Capitalism Before Its Triumph*, Princeton University Press, Princeton.
- Holmström, B., and P. Milgrom (1991), “Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design,” *Journal of Law, Economics, and Organization*, 7: 24–52.
- Itoh H. (1991), “Incentives to Help in Multi-Agent Situations,” *Econometrica* 59(3): 611–636.
- Kahneman, D., and J. Knetsch (1992), “Valuing Public Goods: The Purchase of Moral Satisfaction,” *Journal of Environmental Economics and Management*, 22: 57–70.
- Kamenica, E., and L. Egan Brad (2014), “Voters, Dictators, and Peons: Expressive Voting and Pivotality,” *Public Choice*, 159(1/2): 159–176.
- Kosfeld, M., and F. Von Siemens (2011), “Competition, Cooperation and Corporate Culture,” *Rand Journal of Economics*, 42: 23–43.
- Kotchen, M. (2006), “Green Markets and Private Provision of Public Goods,” *Journal of Political Economy*, 114(4): 816–834.
- Landier, A., and S. Lovo (2020), “ESG Investing: How to Optimize Impact?,” mimeo.
- Lazear, E.P. (1989), “Pay Equality and Industrial Politics,” *Journal of Political Economy*, 97(3): 561–580.
- Lazear, E.P., Malmendier, U., and R.A. Weber (2012), “Sorting in Experiments with Application to Social Preferences,” *American Economic Journal: Applied Economics*, 4(1): 136–163.
- McCloskey, D.N. (2006), *The Bourgeois Virtues: Ethics for an Age of Commerce*, Chicago: University of Chicago Press.

- Moisson, P.-H. (2020), “Ethics and Impact Investment,” mimeo.
- Mokyr, J. (2016), *A Culture of Growth: The Origins of the Modern Economy*, Princeton: Princeton University Press.
- Oehmke, M., and M. Opp (2023), “A Theory of Socially Responsible Investment,” *R&R Review of Economic Studies*.
- Prendergast, C. (2007), “The Motivation and Bias of Bureaucrats,” *American Economic Review*, 97: 180–196.
- Raith, M. (2003), “Competition, Risk, and Managerial Incentives,” *American Economic Review*, 93(4): 1425–1436.
- Robeznieks, A. (2022), “Physicians Warned of the Pitfalls behind Private Equity Promises,” *American Medical Association*, <https://www.ama-assn.org/practice-management/private-practices/physicians-warned-pitfalls-behind-private-equity-promises>.
- Sandel, M. J. (2012), *What Money Can't Buy: The Moral Limits of Markets*, Macmillan.
- Satz, D. (2010), *Why Some Things Should Not Be For Sale: The Moral Limits of Markets*, Oxford University Press.
- Schmidt, K. (1997), “Managerial Incentives and Product Market Competition,” *Review of Economic Studies*, 64(2): 191–213.
- Singer, P. (2015), *The Most Good You Can Do*, Yale University Press.
- Sobel, J. (2015), “Do Markets Make People Selfish?,” Mimeo.
- US General Accounting Office (2003), “OxyContin Abuse and Diversion - Efforts to Address the Problem,” Report to Congressional Requesters, December.
- Velez, R. (2017), “Inequity-Averse Preferences in General Equilibrium,” *Journal of Mathematical Economics*, 70: 166–175.
- Walzer, M. (2008), *The Free Market and Morality* on YouTube.
- Weber, R. and S. Zhang (2021), “What Money Can Buy: How Market Exchange Promotes values,” mimeo.
- Weisbrod, B. (1988), *The Nonprofit Economy*, Cambridge MA: Harvard University Press.

Online Appendix: The Morality of Markets

by Mathias Dewatripont and Jean Tirole

A Unethical/Present-biased/Influenceable consumers

We here provide three reasons why demand-side considerations may drive immoral choices. In all cases, there is a wedge between the quality perceived by the customer and that assessed by a social planner; and being more ethical means reducing one's demand.

Example 1: Internalities (painkiller prescriptions). The supplier (a doctor) decides whether to prescribe an opioid to the client (the patient).¹ The fee p_i paid by the patient is either regulated or competitive, and is paid for the visit, regardless of what the doctor will prescribe. The painkiller brings known benefit b , but has side effects with cost γ . This cost is observed only by the doctor (who learns who is at risk) and is distributed according to distribution $G(\gamma)$ and density $g(\gamma)$. The doctor chooses a threshold γ^* under which she prescribes the painkiller. Assume that the patient knows γ^* ; one may have in mind that patients know the doctor's reputation for being easy ("pill mill doctor") or tough on prescriptions (alternatively, the patient might be searching until she gives up or a doctor supplies the opioid). Welfare is $b - \gamma$, but clients have hyperbolic preferences with coefficient $\beta < 1$: They long for quick relief and value the prescription at $b - \beta\gamma$. And so the surplus of the short-term self, u_i , and welfare, w_i , are:

$$u_i = \int_0^{\gamma_i^*} (\beta\gamma - b)g(\gamma)d\gamma \quad \text{while} \quad w_i = \int_0^{\gamma_i^*} (b - \gamma)g(\gamma)d\gamma.$$

The maximum gross surplus of the short-term self corresponds to $\gamma_i^* = b/\beta$ and is equal to $u^* \equiv \int_0^{b/\beta} (\beta\gamma - b)g(\gamma)d\gamma$, yielding $\underline{w} \equiv \int_0^{b/\beta} (b - \gamma)g(\gamma)d\gamma$. The welfare optimum corresponds to $\gamma_i^* = b$, $\bar{u} \equiv \int_0^b (\beta\gamma - b)g(\gamma)d\gamma$, and $\bar{w} \equiv \int_0^b (b - \gamma)g(\gamma)d\gamma$. Letting $a_i \equiv -u_i$, one has² $W' > 0 > W''$ over the relevant range $a_i \in [-u^*, -\bar{u}]$.

Instances of overconsumption due to imperfect self-control or biases in predicting one's future behavior are many outside the health domain (excessive indebtedness, gambling, videogaming, impulsive clicking on privacy consent forms...).

¹Opioid overconsumption illustrates internalities, given the addictive nature of such painkillers. Opioids represent both a useful treatment for acute pain (e.g. in case of terminal cancer) but also run the risk of addiction without proven medical benefits in the case of chronic pain (e.g. back pain). Opioid overdoses have been called the worst drug epidemic in the history of the United States (McGreal 2018). The crisis has multiple dimensions, including the role of companies like Purdue Pharma in inducing doctors to prescribe their opioid OxyContin. Our paper focuses on doctors' decision when facing patient demands for opioid prescriptions (see Schnell 2019 for an assessment of policies aimed at keeping opioid prescriptions in check).

²One has

$$W''(a) = \frac{d}{d\gamma^*} \left(\frac{b - \gamma^*}{\beta\gamma^* - b} \right) / \frac{da}{d\gamma^*} < 0.$$

Example 2: Externalities (vaccines, overprescription of branded drugs). This example replaces the internality of Example 1 by an externality. Patients have heterogeneous probabilities x of getting sick in the absence of vaccination, in which case they suffer damage d and contaminate an expected number e of other people. Patients are selfish and value being vaccinated at $E[b - \gamma]$, where $b = xd$ is the benefit and γ is a cost of vaccination. The social planner attaches value $E[(1 + e)b - \gamma]$. It is easy to check that this externality example is mathematically akin to the internality example, Example 1. After all, an internality is just an externality of the short-term self on the long-term one.³

Underconsumption, unlike overconsumption, raises the question of how the supplier can provide a quantity that exceeds the client's desired consumption: A doctor cannot physically vaccinate a patient who refuses to be inoculated. A first interpretation of the underconsumption model goes as follows: When the state mandates children to be vaccinated in order to be able to go to school or public sport facilities, parents may try to obtain a complacent (fake) vaccination certificate from the doctor (underconsumption of vaccines). Similarly, in some countries, occupational physicians may routinely deliver fake medical certificates allowing employees to take paid sick leave (underprovision of work). In both examples, unethical supplier behavior is associated with a fraudulent report to a third party. A second interpretation applies when no law or rule mandates a level of consumption in excess to that desired by the client. Ethical/unethical behavior then relates to the intensity with which the doctor puts pressure on the patient, say to be vaccinated; it may range from attempts at persuasion to outright refusal to keep seeing a patient who refuses the vaccination.

Overconsumption occurs in the case of antibiotics. Another case in point is the overconsumption of branded drugs when generics are available, imposing an externality on the social security system.⁴

Example 3: Influencing purchases through product misrepresentation. Product misrepresentation implies that the choice of a_i is unobservable. Yet consumers are not "indifferent" in our terminology: A lower a_i increases demand. Tobacco companies' advertising failed to warn against the harmful effects of smoking. More generally, firms typically emphasize positive attributes of their goods and services and rarely their flaws. To be certain, consumer protection agencies' and courts' mission is to combat inappropriate statements or frauds. But there is a thin line between outright misrepresentation and fraud on the one

³Underconsumption of vaccines may also be driven by a misperception of their side-effects. For example, a triple jab of the measles vaccine was falsely accused in *The Lancet* of causing autism, which led to a substantial drop in MMR vaccination. Such misperceptions may be captured as an underestimation of the net value of the vaccine, independently of contagion considerations.

⁴A fraction of French patients has always viewed generics as inferior products. Until the mid-90s French doctors faced no cost of prescribing branded drugs instead of generics (and pharmacists' compensation was proportional to the price of the drug!). Lo and behold, doctors pandered to their patients and generics' market share was about 2%. A reform introduced incentives for doctors to prescribe generics, and also gave pharmacists the ability to replace a branded drug by an equivalent generic. The share of generics' prescriptions improved, especially with general practitioners (whose patients are more loyal than for specialists, in conformity with the theory developed below).

hand, and omission, opaque language and the exploitation of consumer inattention on the other hand.⁵ Unexpected obsolescence, vague recommended usage or the downplaying of side effects may not be illegal or else hard to regulate given their ubiquity and the limited means of the agencies. Like internalities and externalities, product misrepresentations leave scope for moral judgment.

One way to formalize this within our model goes as follows: Suppose that the good actually delivers gross surplus \hat{u} to the consumers. Supplier i can inflate this surplus and claim it is $u_i \in [\hat{u}, \bar{u}]$. Gullible consumers take the announcement at face value (see below for a more rational version) and plan around the announced value, leading to later inconvenience (complementary investments miscalibration, misleading claims made to downstream users...) cost $\Gamma(u_i - \hat{u})$, with $\Gamma(0) = 0$, $\Gamma'(0) = 0$, $\Gamma'(u_i - \hat{u}) > 0$ and $\Gamma''(u_i - \hat{u}) > 0$ for $u_i > \hat{u}$, and $\Gamma'(\bar{u} - \hat{u}) = +\infty$ for some maximum exaggeration level $\bar{u} - \hat{u}$. Again let $a_i = -u_i$. Then $W(a) \equiv -\hat{u} - \Gamma(\hat{u} - a)$ satisfies the general assumptions. Finally, note that a more rational, asymmetric-information, version of the model would have consumers not know about the misreporting function. For example, with some probability they believe that misreporting is infeasible; the “irrational version” is just the limit of the “rational version” as this probability goes to 1.

In much of the recent “shrouded attributes” literature building on Gabaix-Laibson (2006),⁶ a_i refers to a disclosure decision ($a_i = 1$ if the supplier discloses some bad news for the consumer and $a_i = 0$ if he does not). The lack of disclosure of bad news (say, the unexpected need for the consumer to purchase an add-on later, then sold at an inflated price) can be viewed as decreasing the unit cost (by exactly the ex-post profit on the add-on if consumers are naive): $c'_i(a_i) > 0$, and furthermore the moral high-ground ($a_i = 1$) reduces demand (ϕ_i is decreasing in a_i). So the moral action in this particular product misrepresentation application impacts both cost (positively) and demand (negatively), and our framework accounts for the type of environments considered in the Gabaix-Laibson literature.

⁵More generally, and like all other works building on social preferences and social responsibility, all interpretations assume that tort law does not provide a perfect Pigovian correction of the wedge between market and society’s demands (which is the case for the applications envisioned throughout the paper). This may hold for multiple familiar reasons: (a) Limited liability, risk aversion or managerial turnover may prevent the collection of the Pigovian tax; (b) there may be no political will to levy or enforce such a tax (for example due to lobbying); (c) laws may embody loopholes; (d) the behavior may not be verifiable (e.g. corruption, doctors’ “judgement”, interpretation of advertising); (e) law enforcement is too costly for minor misdemeanors (Kaplow-Shavell 2007); (f) a last reason is the inability to tax externalities, such as the use of child labor, corruption or pollution, exerted by suppliers operating in different jurisdictions.

⁶See Heidhues-Koszegi (2018) for a detailed overview of this literature. Note also that we here build on the bare-bones Gabaix-Laibson model. Their model is actually richer, as supplier i ’s disclosing the existence of an add-on not only affects demands by changing supplier i ’s attractiveness, but also by revealing the possibility of an add-on for the rivals as well, affecting their own attractiveness.

B Verifying global second-order conditions under flexible prices in a covered-market symmetric oligopoly

Let us check that the tentative flexible-price equilibrium is a global optimum for each supplier in a covered-market symmetric oligopoly. Let $D_i(\hat{p}_i, \hat{p})$ denote the demand faced by supplier i when it charges net price \hat{p}_i and all others offers the same net price \hat{p} . Suppose that suppliers internalize \mathcal{E} (which implies that a symmetric equilibrium is still an equilibrium when they internalize $\mathcal{E} - \mathcal{M}$ where \mathcal{M} is the misallocation cost: See Section 3.2.1). Equilibrium behavior requires that there be no (p_i, a_i) such that

$$[p^* - c(a^*)]D_i(p^* + \phi(a^*), p^* + \phi(a^*)) < [p_i - c(a_i) + \alpha[W(a_i) - W(a^*)]] D_i(p_i + \phi(a_i), p^* + \phi(a^*)) \\ \equiv \mathcal{V}(p_i, a_i).$$

The concavity of $[\alpha W - c - \phi]$ and condition (3) imply that

$$\alpha[W(a_i) - W(a^*)] \leq [\phi(a_i) + c(a_i) - \phi(a^*) - c(a^*)].$$

So

$$\mathcal{V}(p_i, a_i) \leq [p_i - c(a^*) + \phi(a_i) - \phi(a^*)]D_i(p_i + \phi(a_i), p^* + \phi(a^*)).$$

The maximization w.r.t. the net price for a given moral behavior a^* by supplier i implies that for all \tilde{p}_i , $[p^* - c(a^*)]D_i(p^* + \phi(a^*), p^* + \phi(a^*)) \geq [\tilde{p}_i - c(a^*)]D_i(\tilde{p}_i + \phi(a^*), p^* + \phi(a^*))$. Applying this to $\tilde{p}_i = p_i + \phi(a_i) - \phi(a^*)$ yields

$$[p^* - c(a^*)]D_i(p^* + \phi(a^*), p^* + \phi(a^*)) \geq [p_i - c(a^*) + \phi(a_i) - \phi(a^*)]D_i(p_i + \phi(a_i), p^* + \phi(a^*)),$$

a contradiction. ■

C Further robustness checks on Proposition 1

Incomplete consumer information

We here study two classic environments exhibiting strategic complementarity, in which consumers are uninformed about the products, and augment them with moral choices. In the first environment, consumers search sequentially for offers $\{p_i, a_i\}$. In the second, some consumers are equipped to learn the moral content of offers and others do not observe this content (but all observe prices). We show that the irrelevance property holds in both environments.

(a) *Search.* The irrelevance result is also robust to a different model of strategic complementarity. An alternative to differentiated Bertrand is the search model, a classical version of which we here extend to incorporate moral choices. Suppose that there is a large- actually infinite- number of suppliers; so, in this model competition is indexed not

by the number of suppliers, but by the level of the search cost. We will say that competition becomes more intense when the search cost s decreases. Suppliers are identical (same cost and demand function, same moral preferences). At each search with a new supplier, the unit-demand consumer draws a valuation v from distribution $F(v)$ (independently of previous draws). She can take the offer and stop the search or continue searching. An extra search involves known search cost s . We assume that s is not too large so that search occurs in equilibrium (in this case all consumers purchase and the market is therefore covered). In a symmetric equilibrium, all offers are the same and so the cutoff v^* for the acceptance of an equilibrium offer is given by the sequential search condition:

$$s = \int_{v^*}^{+\infty} (v - v^*) dF(v)$$

Each firm chooses (\hat{p}, a) given other firms' strategy (\hat{p}^*, a^*) , so as to solve:

$$\max_{(\hat{p}, a)} [\hat{p} - \phi(a) - c(a) + \alpha[W(a) - W(a^*)]][1 - F(v^* - (p^* - p) - (\phi(a^*) - \phi(a)))]$$

The FOC w.r.t. a is again independent of the demand function (while in contrast p^* decreases with s): $a^* = a^\dagger$, where

$$c'(a^\dagger) + \phi'(a^\dagger) = \alpha W'(a^\dagger),$$

so the irrelevance property holds (a^\dagger is independent of s). Once again, consequentialism and constant returns to scale imply that everything is linear in demand, and thanks to flexible prices we can adjust the ethical choice while leaving demand unchanged.

(b) *Unobserved attributions.* We have assumed that consumers are fully informed about the suppliers' moral choice (the latter is a search good). Alternatively, they could be uninformed (the moral choice is a credence good and so $\phi'_i = 0$) or imperfectly informed (the supplier's choice is revealed with some probability). Assume that, with probability $1-x$ consumers do not observe supplier i 's ethical choice and rationally expect equilibrium choice a_i^* , and with probability x the actual choice a_i is publicly revealed (say, there is a scandal revealed in the media). Assume linear and symmetric demand. The average net price charged by firm i , which determines the demand it faces, is

$$\hat{p}_i \equiv p_i - x\phi(a_i) - (1-x)\phi(a_i^*),$$

where a_i^* is the equilibrium behavior (anticipated by customers). So the elasticity of demand with respect to a_i is now smaller, and supplier i 's first-order condition under flexible prices is:

$$\alpha W'(a^*) = c'(a^*) + x\phi'(a^*).$$

With ethical consumers (we argued that $x = 1$ is the natural assumption for unethical consumers), a reduction in x implies an increase in the RHS of this condition, and thus a decrease in the level of ethics (as is the case for a decrease in α), which is also intuitive

since the supplier gains customers by behaving more ethically only when this is observed by them.⁷ The irrelevance result still holds.

Non-linear pricing

Proposition 1 does not require unit demands. The consumers might be consuming multiple units. As long as pricing is linear and economic agents are consequentialists, the irrelevance result carries over. But what about non-linear pricing? One can consider two types of second-degree price discrimination, one based on volume and the other on moral intensity.

(a) *Volume-based price discrimination.* Suppose that consumers all have the same moral preferences, $-\phi(a)q$ for consumption q , and differ in their marginal utility of consumption as in Mussa-Rosen (1978) and Maskin-Riley (1984): $u(q, a, \theta) = v(q, \theta) - T(q) - \phi(a)q$ where $T(q)$ is the tariff to be paid for q units. Consumer h , when buying from supplier i , has utility $U_i(\theta) + \varepsilon_{hi}$ (where, as earlier, the vector $\{\varepsilon_{hi}\}_{i \in \{0, \dots, n\}}$ is distributed according to some smooth joint distribution), where

$$U_i(\theta) = \max_q \{v(q, \theta) - T_i(q) - \phi_i(a_i)q\}.$$

As in Armstrong-Vickers (2001), Rochet-Stole (2002) or Bénabou-Tirole (2016), one can view competition as a competition in utilities rather than transfers: Supplier i chooses $\{U_i(\theta), q_i(\theta), a_i\}$ subject to incentive compatibility and individual rationality.⁸ Letting $\mathbf{U}(\theta) = \{U_i(\theta)\}_{i \in \{1, \dots, n\}}$ let $X_i(\mathbf{U}(\theta))$ denote the market share of supplier i among types θ . Letting \mathbb{E} denote the expectation with respect to θ , supplier i 's objective function is:

$$\mathbb{E}[X_i(\mathbf{U}(\theta)) [v(q_i(\theta), \theta) - [\phi_i(a_i) + c_i(a_i)]q_i(\theta) - U_i(\theta)] + \alpha_i \sum_j W_j(a_j) X_j(\mathbf{U}(\theta)) q_j(\theta)].$$

A subprogram consists in minimizing cost

$$\min \mathbb{E} \left[[\phi_i(a_i) + c_i(a_i) - \alpha_i W_i(a_i)] q_i(\theta) \right]$$

with respect to a_i , yielding again the irrelevance result.⁹

(b) *Price discrimination based on moral preferences.* Suppose now that consumers have unit demands, but are heterogeneous in their moral preferences: They differ in their internalization $\phi_\theta(a)$, where a higher θ indicates a more moral consumer ($\phi'_\theta(a)$ grows with θ). The results go as follows: (i) Competition affects moral choices. (ii) Whether competition makes the allocation more or less moral depends on the nature of the outside option, i.e. on whether $\phi_\theta(a)$ grows or decreases with θ .¹⁰

⁷Bonneton (2020) studies the provision of information about supplier moral behavior in the form of a binary standard. In his paper, suppliers are intrinsically motivated and have image concerns as well.

⁸Namely $dU_i/d\theta = v_\theta(q_i(\theta), \theta)$ and $U_i(\theta) \geq U_i^0(\theta)$ (the utility from the outside option). Assuming that the ethical choice can be tailored to type $(a_i(\theta))$ would not affect the irrelevance result obtained shortly.

⁹The counterpart of price flexibility is that the choice of $U_i(\cdot)$ among mechanisms that are incentive compatible and individually rational is unconstrained.

¹⁰Readers familiar with mechanism design will here recognize the distinction between classic and countervailing incentives.

Suppose for instance that there are two types, a low-morality and a high-morality types. Note first that in this private-value model, competition delivers each type the symmetric-information moral outcome, given by $c'(a_\theta^\dagger) + \phi'_\theta(a_\theta^\dagger) = \alpha W'(a_\theta^\dagger)$, so $a_H^\dagger > a_L^\dagger$. To illustrate how monopoly power affects moral choices, suppose first that the outside option involves no pollution (case of ethical consumers) or no corruption (unethical consumers); then $a_0 = \bar{a}$. More moral types then get a lower net utility from trade with the monopolist than less moral ones. The participation constraint under monopoly is then binding on the most moral type. There is no distortion for the less moral type ($a_L = a_L^\dagger$). The moral type's allocation is distorted upwards ($a_H > a_H^\dagger$) so as to prevent mimicking by the less moral type. So the monopoly allocation is more moral than the competitive one. Conversely, suppose that the outside option involves maximal pollution (say, coal-based electricity generation) or maximal corruption: $a_0 = 0$. Then the participation constraint under monopoly is binding on the less moral type, and we are in the Maskin-Riley/Mussa-Rosen configuration, with no distortion for the high type and a downward distortion for the low type, making the competitive market more moral than monopoly.

More generally suppose that there are two types of consumers: type H (moral), in proportion ρ , and type L (less moral), in proportion $1 - \rho$. The H type cares more about morality than the L type:

$$\phi'_H(a) < \phi'_L(a). \quad (\text{C.1})$$

Let us look at the polar cases of pure monopoly and perfect competition (more generally we could consider all degrees of imperfect competition as in Bénabou-Tirole 2016 and Garrett et al 2019). The following holds for both the ethical and UPI consumers.

Perfect competition

Equilibrium conditions are:¹¹

$$p_L + \phi_L(a_L) = p_H + \phi_L(a_H) \quad (\text{C.2})$$

$$a_H = a_H^\dagger \quad (\text{C.3})$$

$$a_L = a_L^\dagger. \quad (\text{C.4})$$

Monopoly

The incentive compatibility condition does not determine who has a rent, which is crucial in the case of a monopoly (but not for perfect competition, as the participation constraints are then not binding).

Countervailing incentives. Assume, first, that the L type has a rent, i.e. moral preferences are individually a nuisance (reduce utility): for $a < \bar{a}$

$$\phi_H(a) > \phi_L(a). \quad (\text{C.5})$$

¹¹The concavity of the objective functions guarantees incentive compatibility.

Consumer utility is type-independent at $a = \bar{a}$:

$$\phi_H(\bar{a}) = \phi_L(\bar{a}) \quad (\text{C.6})$$

Thus, we posit that if there is no moral issue, consumers are not differentiated according to their moral preferences.¹² (a) Ethical consumers: all consumers have the same preferences if there is no pollution, i.e. $a = \bar{a}$, but the H type loses more utility from pollution. (b) Unethical consumers: all consumers have the same preferences if there is no corruption, i.e. $a = \bar{a}$, but the H type gains less utility from corruption.

In either case (ethical or unethical consumers), the L type has a rent and the incentive constraint implies that $a_H \geq a_L$, implying in turn that $p_H \geq p_L$ (ethical consumers) and $p_H \leq p_L$ (unethical consumers)

As usual, we will denote by a_θ^\dagger the symmetric-information moral action (assuming it is interior):

$$c'(a_\theta^\dagger) + \phi'_\theta(a_\theta^\dagger) = \alpha W'(a_\theta^\dagger)$$

with

$$a_H^\dagger > a_L^\dagger.$$

The IR and IC conditions are

$$p_H + \phi_H(a_H) = v \quad (\text{C.7})$$

and

$$p_L + \phi_L(a_L) \leq p_H + \phi_L(a_H). \quad (\text{C.8})$$

The monopolist solves:

$$\max\{\rho[p_H - c(a_H) + \alpha W(a_H)] + (1 - \rho)[p_L - c(a_L) + \alpha W(a_L)]\}$$

There is no distortion at the top (here at type L) and distortion at the bottom:

$$a_H > a_H^\dagger \quad (\text{C.9})$$

and

$$a_L = a_L^\dagger \quad (\text{C.10})$$

where a_H depends on ρ according to a slight modification of the standard quality discrimination condition:

$$c'(a_H) + \phi'_H(a_H) - \alpha W'(a_H) = \frac{1 - \rho}{\rho} [\phi'_L(a_H) - \phi'_H(a_H)] > 0. \quad (\text{C.11})$$

There is more moral behavior under monopoly (that is linked with the fact that the immoral type has a rent). There is an interesting analogy with the existing literature

¹²For example, suppose that $a \in [0, \bar{a}]$ and $\phi_\theta(a) = \beta_\theta(\bar{a} - a)$ with $\beta_H > \beta_L$ and $\beta_\theta > 0$ for ethical consumers and $\beta_\theta < 0$ for UPI consumers.

here. Starting with Lewis-Sappington (1989), the latter has studied “countervailing incentives”, namely the situation that occurs when a monopolist price-discriminates, “higher types” value quantity or quality more, but have better outside opportunities, so the individual-rationality constraint binds for the high types. There is then an enlargement of the “quantity/quality spectrum”, rather than the traditional Mussa-Rosen (1978) and Maskin-Riley (1984) compression of that spectrum. That literature does not consider competitive price discrimination, but speaks to why under monopoly highly moral types consume with especially high moral intensity.

Classical incentives. Under countervailing incentives, monopoly delivers a higher morality than competition. This suggests the following result: suppose, say, that the outside option pollutes maximally or involves maximal corruption (for instance): $a_0 = 0$. Then the H type is the type who enjoys a rent when dealing with the firm; we are then in the Mussa-Rosen/Maskin-Riley conventional case and we have a lower morality under monopoly than under competition (classical incentives case). More formally, keep the sorting condition (C.1) and replace conditions (C.5) by: for $a > 0$,

$$\phi_H(a) < \phi_L(a) \tag{C.12}$$

and (C.6) by:

$$\phi_H(0) = \phi_L(0). \tag{C.13}$$

The standard proof then shows that under monopoly and if both types are served:¹³

$$a_H = a_H^\dagger > a_L^\dagger > a_L \tag{C.14}$$

while the outcome under competition is still $(a_H^\dagger, a_L^\dagger)$. Competition then yields a more moral outcome.

Full welfare internalization

(i) Internalization of consumer surplus

Let us allow the supplier to internalize (some of) the increase in consumer surplus when lowering net price.

- When the market is covered, market power generates no consumption distortion and aggregate welfare does not depend on (net) prices, except to the extent that the latter determine the allocation of consumers among products with different social footprints.
- More generally, the irrelevance results still holds as the consumers’ net surplus depends only on the net prices \hat{p} and so we can add an internalization of consumer surplus by the firms (with some coefficient $\beta_i < 1$ for supplier i), that does not affect the cost-minimization argument.

¹³If only type H is served, then type L consumes the outside option and $a_L = 0$; so again monopoly reduces morality.

(ii) *Misallocation of products to consumers*

Consumers may not be matched with their preferred supplier. This will be the case even in symmetric oligopoly when net prices $\{\hat{p}_i\}$ differ. To illustrate this misallocation, consider for notational simplicity constant, identical costs (the reasoning extends to heterogeneous costs). The misallocation cost for consumer h is the difference between the surplus she gets from her preferred supplier and that offered by the supplier $i(h)$ she ends up picking. Using the discrete-choice notation introduced above and aggregating over all consumers yields the misallocation cost:

$$\mathcal{M}(\hat{\mathbf{p}}) \equiv \int [(\max_i \varepsilon_{hi}) - \varepsilon_{hi(h)}] dh.$$

More generally, total welfare is¹⁴

$$\mathcal{W}(\hat{\mathbf{p}}, \mathbf{a}) = \mathcal{E}(\hat{\mathbf{p}}, \mathbf{a}) - \mathcal{M}(\hat{\mathbf{p}}).$$

Condition (11) in the text is satisfied since $\partial\mathcal{M}/\partial a_i = 0$, $\partial\mathcal{E}/\partial a_i = W'_i(a_i)D_i(\hat{\mathbf{p}})$ and so

$$\Gamma_i(a_i) \equiv W'_i(a_i)$$

satisfies $\Gamma'_i < 0$.

Note that in symmetric oligopoly $\mathcal{M}(\hat{\mathbf{p}}) = 0$ (and is therefore minimized) when all net prices are equal. This implies that misallocation losses are locally second-order: $\partial\mathcal{M}/\partial\hat{p}_i = 0$ at a symmetric equilibrium (for example, in the Hotelling-Lerner-Salop style models of differentiation, the misallocation cost is quadratic in price differences: $\mathcal{M} \propto \sum_{i < j} (\hat{p}_i - \hat{p}_j)^2$).

Lemma 1 *Consider a symmetric equilibrium of a symmetric-oligopoly game in which the firms internalize the ethical welfare \mathcal{E} . Then the resulting allocation is also an equilibrium of the game in which the firms internalize the full welfare $\mathcal{W} = \mathcal{E} - \mathcal{M}$.*

Proof. Even though a supplier's deviating from a symmetric-equilibrium behavior generates a misallocation of consumers to firms, the proof of Lemma 1 is straightforward: The resulting misallocation of consumers to firms makes such a deviation away from symmetric behavior even less attractive under full-welfare internalization than when only the ethical welfare is considered. ■

Cartel and broad bracketing

Cartel agreements, whether formal or enforced through tacit collusion, raise difficult questions under broad bracketing, that are not specific to this paper. To see why, suppose

¹⁴We could add the consumer net surplus from consumption (weighted by α), but this would not change the derivations below; for, a small change in i 's policy implies a shift, from or toward alternative suppliers, of marginal consumers who by definition are indifferent between supplier i and their best alternative supplier. So the marginal impact on consumer surplus is 0.

that the market is covered, and that the n firms can sign a formal cartel agreement specifying a symmetric per-firm strategy $\{\hat{p}, a\}$. Assuming broad bracketing (internalization of ethical welfare, $\mathcal{E} = \sum_{j=1, \dots, n} W_j(a_j) D_j(\hat{\mathbf{p}})$), the per-firm payoff in a symmetric collusive outcome is, if we follow the formalism of the paper:

$$V \equiv \max_{\{\hat{p}, a\}} \{[\hat{p} - \phi(a) - c(a)] \frac{D(\hat{p})}{n} + \alpha W(a) D(\hat{p})\}$$

In particular, the industry-optimal morality, maximizing nV , is given by

$$n\alpha W'(a) = \phi'(a) + c'(a)$$

So, unlike in Nash equilibria, the optimal cartel morality is increasing in n . [As we have seen, a in contrast does not depend on n under tacit collusion and narrow bracketing.] Broad bracketing induces double counting in the following sense: Each supplier takes credit (or discredit) for what others do, and not only for its own impact. Double counting does not matter as long as the analysis is positive; but when it becomes normative (or cooperative among suppliers), it faces this problem. We refer to Diamond (2006) and the references therein for related puzzles when agents have warm glow or social preferences.

D Proof of Proposition 3

Consider the first-order condition with respect to the ethical choice ($\partial V_i / \partial a_i = 0$), for given prices $\mathbf{p} = (p_1, \dots, p_n)$. Behaving more ethically (increasing a_i) has three effects on supplier i 's payoff function $V_i = [p_i - c_i(a_i)] D_i(\hat{\mathbf{p}}) + \alpha_i \mathcal{W}_i(\hat{\mathbf{p}}, \mathbf{a})$:

$$\frac{\partial V_i}{\partial a_i} = \underbrace{(p_i - c_i) \phi'_i \frac{\partial D_i}{\partial \hat{p}_i} - c'_i D_i}_{\text{impact on profit}} + \underbrace{\alpha_i W'_i D_i}_{\substack{\text{ethical} \\ \text{impact on} \\ \text{supplier } i\text{'s} \\ \text{inframarginal} \\ \text{consumers}}} + \underbrace{\alpha_i \phi'_i \frac{\partial \mathcal{W}_i}{\partial \hat{p}_i}}_{\substack{\text{ethical} \\ \text{impact} \\ \text{of gain/} \\ \text{loss in} \\ \text{market} \\ \text{share}}} = 0. \quad (\text{D.15})$$

using

$$\frac{\partial D_i}{\partial a_i} = \phi'_i \frac{\partial D_i}{\partial \hat{p}_i} = \phi'_i \frac{\partial D_i}{\partial p_i}. \quad (\text{D.16})$$

In (D.15), the impact of the choice of a_i on profit captures the demand and cost effects that are familiar from models of quality choice (e.g. Spence 1975). The other two effects are proportional to supplier i 's social preferences parameter α_i . The increase in a_i has a positive ethical impact on supplier i 's inframarginal consumers. Finally, supplier i also gains (resp. loses) market share at the detriment of the other suppliers. Letting $w_j \equiv W_j(a_j)$, note that $\partial \mathcal{W}_i / \partial \hat{p}_i = \sum_{j=0}^n w_j (\partial D_j / \partial p_i)$: The last term in (D.15) captures the ethical impact of the reallocation of market shares (including with the outside good).

Assuming, say, $\phi'_i > 0$ (UPI consumers), it is positive if supplier is “on average” less ethical than alternatives (where “average” refers to weights based on the intensities of market share displacements).

Strategic complementarity. One can rewrite the first-order condition with respect to a_i

$$H(\mathbf{a} \mid \mathbf{p}) \equiv \phi'_i(a_i) \left[\frac{\alpha_i W'_i(a_i) - c'_i(a_i)}{\phi'_i(a_i)} - \eta_i L_i \right] = 0, \quad (\text{D.17})$$

as:

$$\frac{\alpha_i W'_i(a_i) - c'_i(a_i)}{\phi'_i(a_i)} = \eta_i L_i, \quad (\text{D.18})$$

where η_i is the elasticity of demand facing supplier i , L_i is the “generalized Lerner index”¹⁵

$$L_i = \frac{p_i - [c_i - (\partial \mathcal{W}_i / \partial \hat{p}_i) / (-\partial D_i / \partial \hat{p}_i)]}{p_i} \equiv \frac{p_i - (c_i - \alpha_i S_i)}{p_i},$$

and $c_i \equiv c_i(a_i)$ for short. Under UPI consumers, $\phi'_i(a_i) > 0$ and the LHS of (D.18) is a decreasing function of a_i . Under ethical consumers, $\phi'_i(a_i) < 0$, and the LHS of (D.18) is an increasing function of a_i .¹⁶

The social responsibility index, $S_i \equiv \sum_{\substack{j \neq i \\ j \geq 0}} \sigma_{ij} (w_i - w_j)$, captures supplier i 's competitive impact on overall welfare and is equal to her ethical differentials with other suppliers weighted by her substitutability with these suppliers. S_i is positive if supplier i is on average more ethical than her rivals (including the outside option) and negative otherwise.

For example, in the symmetric equilibrium ($w_i = w_j$ for $i, j \geq 1$) of an Hotelling-Lerner-Salop model of product differentiation in which the outside option is not binding ($\sigma_{i0} = 0$), the social responsibility index is nil and the generalized Lerner index L_i is equal to the ordinary Lerner index:

$$S_i = 0 \quad \text{and} \quad L_i = L \equiv \frac{p - c}{p}.$$

How does an increase in supplier i 's rivals' morality affect her own moral choices, assuming strict quasi-concavity of the supplier's objective function in a_i ? For this, we must compute $\partial H / \partial a_j$: $\text{sgn}(\partial H / \partial a_j) = \text{sgn}(\phi'_i) \text{sgn}(\partial(\eta_i L_i) / \partial a_j)$. Equation (D.17) unveils two possible channels of strategic interaction when, say, a_j changes ($j \neq i$): through η_i (elasticity effect) and through L_i (social responsibility effect).

Elasticity effect. Because price and moral choices jointly determine the net price ($\hat{p}_i = p_i + \phi_i(a_i)$), a strategic complementarity of moral choices is inherited from the strategic complementarity in the price domain, regardless of whether consumers are UPI, ethical, or indifferent.¹⁷

¹⁵Using $\sum_{j=0}^n D_j(\hat{\mathbf{p}}) = 1$ and so $\sum_{\substack{j=0 \\ j \neq i}}^n \frac{\partial D_j}{\partial p_i} + \frac{\partial D_i}{\partial p_i} = 0$, $\frac{\partial \mathcal{W}_i}{\partial p_i} = (-\frac{\partial D_i}{\partial p_i}) S_i$.

¹⁶Indeed, in this case the partial derivative of the LHS wrt a_i is $[\phi'_i(\alpha_i W''_i - c''_i) - \phi''_i(\alpha_i W'_i - c'_i)] / (\phi'_i)^2$, which is positive since $\text{sgn}(\phi'_i) = \text{sgn}(\alpha_i W'_i - c'_i)$ from (D.18).

¹⁷Formally, $\partial \eta_i / \partial a_j = [\partial \eta_i / \partial \hat{p}_j] \phi'_j(a_j)$. Our assumption that prices are strategic complements ($\partial \eta_i / \partial \hat{p}_j < 0$), then implies that $\text{sgn}(\partial \eta_i / \partial a_j) = \text{sgn}(-\phi'_j(a_j))$. And so, because $\text{sgn}(\frac{d}{da_i} (\frac{\alpha_i W'_i(a_i) - c'_i(a_i)}{\phi'_i(a_i)})) = \text{sgn}(-\phi'_i(a_i))$, we have $\text{sgn}(\partial a_i / \partial a_j) = \text{sgn}(\phi'_i(a_i) \phi'_j(a_j)) \geq 0$.

Social responsibility effect. How is the social responsibility index S_i altered when a rival's action, a_j , changes? An increase in a_j increases w_j and thereby decreases supplier i 's social responsibility index: $\partial S_i / \partial a_j < 0$. Put simply, stealing market share away from supplier j becomes morally less attractive. Under UPI consumers, a higher a_j increases the incentive to raise a_i , creating a second source of strategic complementarity. Together with the previous remark on the elasticity effect, this proves part (i)(a) of Proposition 3.

Under ethical consumers, a higher a_j reduces the incentive to raise a_i : The social responsibility effect per se induces strategic substitutability. Finally, for α_i small enough, the elasticity effect dominates the social responsibility effect under ethical consumers, proving part (i)(b) of the proposition.

Because of strategic complementarity in moral choices, one should not be surprised by the possibility of multiple equilibria in the absence of price flexibility (while, strikingly, moral outcomes will be uniquely determined under flexible prices). Such multiplicity fits well with some informal discussions of multiple social norms regarding morality.

Example of multiple equilibria. Consider a symmetric duopoly with perfect substitutes and equal market shares in case of equally attractive offerings. Suppose that the price is regulated at level p , and that unit cost $c < p$ does not depend on moral choices $a_i \in \{a_L, a_H\}$, which in contrast affect demand: Consumers are UPI with $\phi(a_L) < \phi(a_H)$ while $a_L < a_H$. A high-morality equilibrium ($a_i = a_H$) exists if the material gain from immoral behavior is smaller than internalized social consequences:

$$\frac{p - c}{2} \leq \alpha[W(a_H) - W(a_L)]$$

The low-morality equilibrium always exists as raising one's moral offer to a_H implies profit loss $\frac{p-c}{2}$ and no ethical gain because of the replacement effect: all consumers keep consuming at the low-morality level a_L . This proves part (ii) of the proposition. ■

Other forms of consequentialism. Proposition 3 assumes that internalized welfare is ethical welfare. What about alternative forms of internalization? Consider first narrow ethical internalization ($\mathcal{W}_i = \mathcal{E}^n$). The modified Lerner index is now $L_i = [p_i - (c_i - \alpha_i w_i)] / p_i$. So, other firms' ethical choices no longer impact the social responsibility index $S_i (= w_i)$, and so strategic complementarity is driven solely by the elasticity effect.

Last, suppose that suppliers internalize the welfare including misallocation costs ($\mathcal{E} - \mathcal{M}$). Relative to the expression for the ethical welfare, there is an extra incentive to decrease a_i when a_j goes down, namely $-\alpha_i \phi'_i(a_i) \frac{\partial \mathcal{M}(\hat{\mathbf{p}})}{\partial \hat{p}_i}$. The derivative with respect to a_j of this incentive is proportional to $-\partial^2 \mathcal{M}(\hat{\mathbf{p}}) / \partial \hat{p}_i \partial \hat{p}_j$ since $\phi'_i(a_i) \phi'_j(a_j) \geq 0$. The consideration of the misallocation cost therefore adds another factor of strategic complementarity if $\partial^2 \mathcal{M} / \partial \hat{p}_i \partial \hat{p}_j < 0$. The latter property holds for instance for the Hotelling-style models, where the cross-term in \mathcal{M} is (proportional to) $(\hat{p}_i - \hat{p}_j)^2$. Intuitively, optimizing the matching of consumers to firms often requires aligning net prices and therefore, for a given common price, aligning ethical behaviors. This effect can be labelled the "misallocation minimization effect", or alternatively the "net price alignment effect".

E Sufficient conditions for equilibrium uniqueness under regulated prices

Sufficient conditions. Let us make an assumption which will ensure the uniqueness of equilibrium and the monotonicity of equilibrium behavior a^* with respect to the symmetric fixed price p :

Assumption 6 Suppose that the price is fixed at some level, $p_i = p \geq c$ for all i . Let $\eta(p, a^*) \equiv -\frac{\partial D_i(\hat{\mathbf{p}})}{\partial p_i} p / D_i(\hat{\mathbf{p}})$ denote the elasticity of demand in a symmetric equilibrium ($a_i = a^*$ for all i) and $L(p) \equiv (p - c)/p$ denote the ordinary Lerner index. Then, for all a^* ,

$$(i) \quad \frac{\partial \eta(p, a^*) L(p)}{\partial p} > 0,$$

$$(ii) \quad \frac{\partial \eta(p, a^*) / \partial a^*}{\eta(p, a^*)} \geq \frac{\alpha W''(a^*) - c''(a^*)}{\alpha W'(a^*) - c'(a^*)} - \frac{\phi''(a^*)}{\phi'(a^*)}.$$

Note first that $\frac{dL}{dp} > 0$. So for Assumption 6(i) to be satisfied, it suffices that the elasticity of demand does not decrease too fast with p . For example, the elasticity of demand is proportional to p (and so is increasing in p) when consumers' demand is linear and the market is covered. As for part (ii) (which is the condition stated in the text just prior to Proposition 4), it is satisfied for example in the discrete choice model when the market is covered (for which $\partial \eta / \partial a^* = 0$). Assumption 6 (ii) guarantees that there is a unique symmetric equilibrium.

We also assume that the second-order condition is satisfied; a sufficient condition for this is that the demand be linear or concave.

$W'(0) = +\infty$ implies that the continuous function

$$\alpha W'(a^*) - c'(a^*) - \eta(p, a^*) L(p) \phi'(a^*)$$

is positive at $a^* = 0$. It is strictly negative at $a^* = \bar{a}$ (as $W'(\bar{a}) = 0$). Its derivative when it takes value 0 is negative from Assumption 6 (ii). So, the function equals 0 at exactly one value of a^* . ■

F Proof of Proposition 5.

We allow for the internalization of the misallocation cost, which in general is positive due to the asymmetries among suppliers. Let $\mathbb{1}_{\mathcal{M}} = 1$ if suppliers internalize full welfare ($\mathcal{E} - \mathcal{M}$) and $\mathbb{1}_{\mathcal{M}} = 0$ if they internalize only the ethical welfare (\mathcal{E}). Let $\mathcal{M}(\hat{p}_i, \hat{\mathbf{p}}_{-i})$ denote the misallocation cost when firm i charges \hat{p}_i while its rivals charge $\hat{\mathbf{p}}_{-i}$.

The objective function of supplier i in any of these three groups is

$$V_i = [\hat{p}_i - \phi(a_i) - c(a_i)]D_i(\hat{p}_i, \hat{\mathbf{p}}_{-i}) + \alpha_i \left[w_i D_i(\hat{p}_i, \hat{\mathbf{p}}_{-i}) + \sum_{j \neq i} w_j D_j(\hat{p}_i, \hat{\mathbf{p}}_{-i}) - \mathcal{M}(\hat{p}_i, \hat{\mathbf{p}}_{-i}) \mathbb{1}_{\mathcal{M}} \right],$$

where we can without loss of generality adopt the convention that $\alpha_i = +\infty$ for not-for-profits and constrained for-profits.¹⁸

Let us show that at least one firm must be unconstrained. Suppose, a contrario, that all suppliers sell at cost. Then in a symmetric equilibrium with moral action a^* , $p_i = c(a^*)$ and $w_i = W(a^*)$ for all i . A small increase in the price p_i of a for-profit then has only a second-order effect on the misallocation (all net prices are equal) and no impact on ethical welfare (as all suppliers select the same a_i). This price increase raises profit by $D_i = 1/n$. And so not all suppliers can charge their marginal cost.

Using the identity $\hat{p}_i = \phi(a_i) + c(a_i)$ for all constrained suppliers $i > m$, their FOC with respect to a_i yields $a_i \equiv a^*$ (with associated welfare $w^* \equiv W(a^*)$ and individual demand D^*), where

$$\frac{W'(a^*)}{\phi'(a^*) + c'(a^*)} D^* = \sigma \frac{\sum_{j \neq i} (w^* - w_j)}{n - 1} + \frac{\partial \mathcal{M}}{\partial \hat{p}_i} \mathbb{1}_{\mathcal{M}}. \quad (\text{F.19})$$

Turning to unconstrained suppliers ($k \leq m$), their FOC with respect to \hat{p}_k is

$$D_k - [\hat{p}_k - \phi(a_k^\dagger) - c(a_k^\dagger)] \sigma - \alpha_k \left[\sigma \frac{\sum_{j \neq k} (w_k^\dagger - w_j)}{n - 1} + \frac{\partial \mathcal{M}}{\partial \hat{p}_k} \mathbb{1}_{\mathcal{M}} \right] = 0. \quad (\text{F.20})$$

The opportunity cost function,¹⁹ $c(a) + \phi(a)$, is always increasing in the case of UPI consumers; for ethical consumers, $c(a) + \phi(a)$ is first decreasing (for $a < \overset{\circ}{a}$) and then increasing (for $a > \overset{\circ}{a}$), where $\phi'(\overset{\circ}{a}) \equiv -c'(\overset{\circ}{a})$. Intuitively, choices $a_i < \overset{\circ}{a}$ are dominated for supplier i : They represent immoral actions that have a high cost. So, in the following we will focus on choices in $[\overset{\circ}{a}, \bar{a}]$, where $\overset{\circ}{a} = 0$ in the case of UPI consumers.

Suppose, first, that $\mathbb{1}_{\mathcal{M}} = 0$. To show that constrained firms behave more ethically than unconstrained for-profits, suppose a contrario that $a^* < a_m^\dagger \Leftrightarrow w^* < w_m^\dagger \Leftrightarrow \phi(a^*) + c(a^*) < \phi(a_m^\dagger) + c(a_m^\dagger)$ for choices a^* and a_m^\dagger in the relevant range (above $\overset{\circ}{a}$). Because $p^* = c(a^*)$ and $p_m \geq c(a_m^\dagger)$, $\hat{p}_m > \hat{p}^*$, and so $D_m < D^*$. Now, let

$$E^* \equiv \sum_{j \neq i} \frac{(w^* - w_j)}{n - 1} \quad \text{for } i > m$$

and

$$E_m \equiv \sum_{j \neq m} \frac{(w_m^\dagger - w_j)}{n - 1}.$$

¹⁸More precisely, a not-for-profit maximizes the term in brackets in the expression of V_i .

¹⁹A unit increase in $\phi(a)$ can be compensated by a unit decrease in price, implying de facto an increase in cost.

Conditions (F.19) and (F.20) can be rewritten (when $\mathbb{1}_{\mathcal{M}} = 0$) as

$$\frac{W'(a^*)}{\phi'(a^*) + c'(a^*)} D^* = \sigma E^* \quad \text{and} \quad \frac{W'(a_m^\dagger)}{\phi'(a_m^\dagger) + c'(a_m^\dagger)} D_m \geq \sigma E_m,$$

using $\alpha_m W'(a_m^\dagger) = \phi'(a_m^\dagger) + c'(a_m^\dagger)$. Because $W'/[\phi' + c']$ is decreasing, $\frac{W'(a^*)}{\phi'(a^*) + c'(a^*)} D^* > \frac{W'(a_m^\dagger)}{\phi'(a_m^\dagger) + c'(a_m^\dagger)} D_m$. On the other hand, $a^* < a_m^\dagger$ implies that $E^* < E_m$, a contradiction.

Condition (F.19) requires that as $\sigma \rightarrow \infty$, $[mw^* - \sum_{j \leq m} w_j^\dagger]/n$ tend to 0. And so $w_j \rightarrow w(a_1^\dagger)$ for all j : The equilibrium exhibits a race to the supplier ethical bottom.

Using condition (F.19), condition (F.20) requires that $\hat{p}_k \rightarrow \phi(a_k^\dagger) + c(a_k^\dagger)$ as $\sigma \rightarrow +\infty$, and so there is convergence to marginal cost pricing. Note also that only suppliers with $\alpha_k = \alpha_1$ can be unconstrained as σ goes to ∞ .

The equilibrium when $\mathbb{1}_{\mathcal{M}} = 0$ satisfies $\hat{p}_i = \phi(a_1^\dagger) + c(a_1^\dagger)$ in the limit as $\sigma \rightarrow +\infty$. And so, any small departure from the net price structure has only second-order effects: $\frac{\partial \mathcal{M}}{\partial \hat{p}_i}(\hat{p}_i, \hat{\mathbf{P}} - \mathbf{i}) \Big|_{\hat{p}_i = \hat{p}_1 = \phi(a_1^\dagger) + c(a_1^\dagger)} = 0$. So the equilibrium when only ethical welfare is internalized is still an equilibrium when full welfare is internalized, that is when $\mathbb{1}_{\mathcal{M}} = 1$.

Next, suppose that there are only not-for-profits in the industry ($n_1 = 0$). Condition (F.19) yields a uniform moral behavior with $a_i = \bar{a}$ for all i (such that $W'(\bar{a}) = 0$), $p = c(a_i)$ and no misallocation ($\mathcal{M} = \frac{\partial \mathcal{M}}{\partial \hat{p}_i} = 0$). This shows that when competition is intense ($\sigma \rightarrow +\infty$) and when suppliers internalize either the ethical welfare or the full welfare, a single “bad apple” (a for-profit) drastically changes the behavior of not-for-profits and morality (from \bar{a} to a_1^\dagger).

■

G Gresham’s law of ethical behavior with UPI consumers

Do unethical suppliers drive out ethical ones? The proof of Proposition 5 (i) only shows that constrained/ethical suppliers and not-for-profits have higher opportunity costs ($c(a^*) + \phi(a^*) \geq c(a_k^\dagger) + \phi(a_k^\dagger)$ for an unconstrained firm k). But they also have lower markups. To investigate the possibility of a “Gresham law of ethical behavior”, we consider the following simple environment:

Proposition 6 (*Gresham’s law*) *Consider a symmetric oligopoly with for-profit firms satisfying Assumption 4, $\phi' > 0$ (UPI consumers), and $c(a) = c$ for all a . There are n_A suppliers with social preferences α_A and n_B suppliers with social preferences $\alpha_B > \alpha_A$ (so $n_A + n_B = n$). Consider equilibria in which suppliers’ strategies are $\{p_A, a_A\}$ and $\{p_B, a_B\}$ (uniform within a group), and per-firm realized demands are D_A and D_B . Suppliers internalize ethical welfare.*

(i) Under a fixed price ($p > c$), less ethical suppliers command a larger market share:

$$D_A > D_B.$$

(ii) Under flexible prices and assuming that the financial-viability constraint is not binding (which requires that σ not be too large), there exists $n_A^* \geq 1$ such that

- $D_A > D_B$ for $n_A \leq n_A^*$ (this is therefore the case under duopoly)
- $D_A \leq D_B$ for $n_A > n_A^*$.

Proof of Proposition 6 We study sequentially fixed and flexible prices. Suppose that suppliers differ solely in their ethical concerns. Assuming that $\mathcal{W}_i = \mathcal{E}$ (ethical welfare), supplier i solves:

$$\max \{(p_i - c)D_i + \alpha_i [w_i D_i + \sum_{j \neq i} w_j D_j]\}.$$

Regulated prices. Let $w_i = W(a_i)$ and $\bar{w}_{-i} = \frac{\sum_{j \neq i} w_j}{n-1}$. In the fixed-price context, $p_i \equiv p$ for all i , and the only decision variable is a_i . The maximization of supplier i 's objective function with respect to a_i yields first-order condition

$$\sigma[(p - c) + \alpha_i(w_i - \bar{w}_{-i})]\phi'(a_i) = \alpha_i W'(a_i)D_i. \quad (\text{G.21})$$

The intuition behind condition (G.21) goes as follows: A unit decrease in a_i (and so in \hat{p}_i) attracts σ new clients, bringing markup $(p - c)$ on each of them. A unit increase in market share further improves welfare by $w_i - \bar{w}_{-i}$ (decreases it if $w_i < \bar{w}_{-i}$). Finally, the decrease in the welfare corresponding to the D_i clients of supplier i has an ethical cost and a monetary benefit for supplier i .

Consider two firms i and j with types B for firm i and A for firm j . Suppose that $a_i \leq a_j$, implying for UPI consumers $\phi'(a_i) \leq \phi'(a_j)$. Then $w_i \leq w_j$ and so $(w_i - w_{-i}) - (w_j - w_{-j}) = \frac{n}{n-1}(w_i - w_j) \leq 0$, so the LHS of (G.21) is weakly smaller for firm i than for firm j . Furthermore $\hat{p}_i \leq \hat{p}_j$ and the symmetry of the demand functions imply that $D_i \geq D_j$. Finally, $W'(a_i) \geq W'(a_j) > 0$. This implies that $\alpha_i W'(a_i)D_i > \alpha_j W'(a_j)D_j$. And so (G.21) cannot be satisfied for both i and j , a contradiction. This proves (i) for this case: *High-ethics firms have a lower market share*, and so $\hat{p}_i > \hat{p}_j$ or $a_i > a_j$.

Next, let us check whether the solution to the FOCs satisfies financial viability and positive demand. Because p must exceed c for firms to be financially viable, the financial constraint is not binding. Remembering that more ethical firms have a lower market share, if $D_B = 0$, type- α_B suppliers could set ethical choice $a_A + \varepsilon$ for ε small, command a positive market share, make a financial profit and improve overall morality. Hence all suppliers have a strictly positive equilibrium market share.

Flexible prices. The first-order condition with respect to p_i is

$$-\sigma[(p_i - c) + \alpha_i(w_i - \bar{w}_{-i})] + D_i = 0. \quad (\text{G.22})$$

The intuition for (G.22) is similar to that underlying condition (G.21). Furthermore, combining (G.21) and (G.22) yields the irrelevance property in our context:

$$\alpha_i W'(a_i) = \phi'(a_i). \quad (\text{G.23})$$

Do firms with higher ethical concerns still command a lower market share? Condition (G.23) implies that $a_B > a_A$ and so their ethical choices will make them unattractive to clients. But this is not the end of the story. Their lack of attractiveness calls for lower prices. And their ethical concerns also make them eager to capture market shares away from less scrupulous suppliers, who conversely do not want to gain market share for that specific reason. Subtracting the first-order conditions (G.22) for the two groups and using Assumption 4 yields

$$\left(\frac{2n-1}{n-1}\right)(\hat{p}_B - \hat{p}_A) = \phi(a_B) - \phi(a_A) + (w_A - w_B) \left[\alpha_A \frac{n_B}{n-1} + \alpha_B \frac{n_A}{n-1} \right].$$

The concavity of W , the convexity of ϕ and condition (G.23) yields

$$\phi(a_B) - \phi(a_A) + \alpha_A [W(a_A) - W(a_B)] > 0 > \phi(a_B) - \phi(a_A) + \alpha_B [W(a_A) - W(a_B)].$$

Thus there exists $n_A^* < n$ such that $\hat{p}_B < \hat{p}_A$ if and only if $n_A \geq n_A^*$. In the duopoly case ($n_A = n_B = n - 1$), $\hat{p}_B < \hat{p}_A$ and so the ethical firm commands a higher market share. Finally, the ethical group's financial viability constraint is not binding for $\sigma \leq \sigma^*$ for some σ^* . ■

Part (i) of Proposition 6 is intuitive. Under regulated prices, a more ethical firm is less attractive to UPI consumers. This handicap in the market place translates into a smaller market share. Part (ii) stresses a force in the opposite direction; namely, under flexible prices, a more ethical supplier can lower price to offset her “quality” disadvantage, and gain market share in particular at the expense of less ethical suppliers; the ethical impact of such undercutting hinges on the “market’s morality”. With few unethical suppliers, the ethical gain is low and an ethical firm still commands a lower market share than an unethical one. By contrast, in a low-morality market, the ethical firm has a big impact when undercutting and ends up commanding a higher market share.

The bottom line of Proposition 6 is that, although ethical suppliers are at a competitive disadvantage due to their scruples, they need not command a smaller per-firm market share: Their moral obligation to make the market ethical makes them fierce competitors in the market place. Indeed, they command a higher market share when there a few of them in proportion to unethical ones.

H Welfare

Preliminaries

Our analysis so far has been positive: how is the morality of the market affected by the intensity of competition? Any normative analysis of market morality requires taking a stance on the following four issues (the first two unrelated to ethical behavior):

1. *Is more competition good for given moral choices?* While an important industrial-organization body of theory and empirical work extols the virtues of competition, we also know that there are exceptions to the rule. A possible stance, which we will adopt here, is the presumption that, for antitrust cases handled in the real world, the answer is “competition is good for given moral choices”.
2. *What drives the increase in competition?* Relatedly, we will posit that there is no direct cost for firms or benefit for consumers from increased competition. An increase in competition is indeed relatively costless when, in a search environment, a comparison website increases competition, in the sense of increasing the within-elasticity η . Another example of costless increase in competitive intensity is a strengthening of antitrust enforcement that has the effect of reducing tacit collusion. A third example is an increase in the number of licenses (e.g. of taxi medallions). Of course, it may also cost to increase competition. For example, an entry subsidy reducing the fixed cost of entry for firms would have to be accounted for in the welfare analysis. But these are standard considerations in antitrust and industrial organization, and therefore likely to be accounted for by analysts.
3. *Should the drivers of UPI behavior be included in the social welfare function?* In the case of an externality or an internality, the decision-maker does not internalize the damage done to someone else or to the future self. Should the decision-maker’s benefit from selfish/impatient behavior be included in social welfare? We will take the answer to be “no”. In the case of an internality, we thereby follow the standard approach in public finance of looking at welfare from the point of view of the long-term self.
4. *Should the drivers of ethical behavior be included in the social welfare function?* This is a more difficult question. There is some disagreement among economists as to whether warm glows (here the suppliers’, consumers’, workers’ and investors’ internalization of welfare) should be counted as part of welfare. The case in favor of doing so is that warm glows drive individual behavior and that social welfare should reflect individual preferences. On the other hand, including warm glow gives rise to some paradoxes.²⁰ See Atkinson (2009), Bergstrom (2006), and Diamond (2006) for discussions of the pros and cons. In the following, we will not account for warm glow, but our conclusions would not be affected if we did: The irrelevance result implies that we can take ethical behavior as a given when altering the intensity of competition: See footnote 21 below.

²⁰For example, depriving someone of income and giving another person the opportunity to help the former may create social value.

For conciseness we will focus on a symmetric equilibrium (p^*, a^\dagger) . We take the antitrust authority's objective function to be $U + \chi\Pi$ where U is the consumers' material surplus and Π is industry monetary profit. The parameter $\chi \in [0, 1]$ is the weight put on profit; it is equal to 1 in the usual definition of welfare and to 0 under the consumer-standard mission given to many antitrust authorities.

The covered market case

In the covered-market case, total output is $Q = 1$. Letting v denote the consumers' (average) valuation for the good or service, $U = v - p^* + W(a^\dagger)$ and $\Pi = p^* - c(a^\dagger)$. So total welfare is

$$v - (1 - \chi)p^* - \chi c(a^\dagger) + W(a^\dagger).$$

Using the equilibrium condition, a^\dagger is constant (Proposition 1, applied to the symmetric case, yields $c'(a^\dagger) + \phi'(a^\dagger) = \alpha W'(a^\dagger)$), and welfare increases with (is constant in) the intensity of competition if $\chi \in [0, 1)$ (resp. $\chi = 1$). The antitrust presumption that competition is good for welfare carries over under ethical concerns.²¹

The non-covered market case

More generally, it is important to distinguish between the within-elasticity (between individual products in the industry), and the across-elasticity (giving the overall elasticity relative to the outside good). The latter plays a role when the market is not covered. In that case, there is a horserace between the standard welfare gain from more competition and for example a higher or lower externality from market expansion.

Let's take the Cournot example. Let $F(v)$ denote the distribution of valuations (so the across-semi-elasticity is $f/(1 - F)$). Letting $v^* \equiv p^* + \phi(a^\dagger) - \phi(a_0)$ denote the cutoff, welfare is equal to

$$\int_{v^*}^{+\infty} [v - (1 - \chi)p^* - \chi c(a^\dagger)] dF(v) + [1 - F(v^*)]W(a^\dagger) + F(v^*)W(a_0).$$

Suppose that competition becomes more intense, leading to a decrease in p^* . The derivative of welfare w.r.t. p^* is:

$$-(1 - \chi)[1 - F(v^*)] - f(v^*) \left[\chi[p^* - c(a^\dagger)] - [[W(a^\dagger) - W(a_0)] - [\phi(a^\dagger) - \phi(a_0)]] \right].$$

The first term of this derivative was already present in the covered-market case and captures the inframarginal consumers' monetary gain from increased competition. This term is necessarily dominant if the across-semi-elasticity $f/(1 - F)$ is not too large.

²¹ Suppose we accounted for warm glows in the social welfare function. The latter would then write:

$$[v - [p^* + \phi(a^\dagger)] + (\alpha_I + \alpha_W)W(a^\dagger)] + \chi[p^* - (\gamma(a^\dagger) - (\alpha_I + \alpha_W + \alpha)W(a^\dagger))]$$

where $\alpha = \alpha_i$ for all i is the suppliers' common internalization parameter. So the derivative of social welfare with respect to p^* is again $-(1 - \chi)$.

The first element in the “market expansion” term (proportional to $f(v^*)$) is the weight χ on profit times the markup, $p^* - c(a^\dagger)$, which is always positive. It reinforces the direct price decrease impact, $(1 - \chi)[1 - F(v^*)]$. To make the last term more concrete, assume that $W(a) = -\psi(a)e$, where e is a pollution externality, and $\phi(a) = \alpha_C W(a)$, with $\alpha_C < 1$ (less than full internalization). Then

$$[W(a^\dagger) - W(a_0)] - [\phi(a^\dagger) - \phi(a_0)] = [1 - \alpha_C][W(a^\dagger) - W(a_0)].$$

This term’s sign depends on whether the market is more or less moral than the outside option: For example, using natural gas is bad for greenhouse emissions and certainly worse than not using energy or using a clean one (then $a^\dagger < a_0$), but is good if the outside option is to reopen coal mines (then $a^\dagger > a_0$). So a more intense competition increases welfare unless the demand is very elastic and production is much less moral than non-production.

References for Online Appendix

- Armstrong, M., and J. Vickers (2001), “Competitive Price Discrimination,” *Rand Journal of Economics*, 32: 579–605.
- Atkinson, A. (2009), “Giving Overseas and Public Policy,” *Journal of Public Economics*, 93: 647–653.
- Bénabou, R., and J. Tirole (2016), “Bonus Culture: Competitive Pay, Screening and Multitasking,” *Journal of Political Economy*, 124(2): 305–370.
- Bonneton, N. (2020), “The Market for Ethical Goods,” mimeo.
- Bergstrom, T. (2006) “Benefit-Cost in a Benevolent Society,” *American Economic Review*, 96 (1): 339–351.
- Diamond, P. (2006), “Optimal Tax Treatment of Private Contributions for Public Goods with and without Warm Glow Preferences,” *Journal of Public Economics*, 90: 897–919.
- Gabaix, X., and D. Laibson (2006), “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *Quarterly Journal of Economics*, 121(2): 505–540.
- Garrett, D., Gomes, R. and L. Maestri (2019), “Competitive Screening under Heterogeneous Information,” *Review of Economic Studies*, 86: 1590–1630.
- Heidhues, P. and B. Köszegi (2018), “Behavioral Industrial Organization,” *Handbook of Behavioral Economics: Applications and Foundations* 11: 517–612.
- Kaplow, L., and S. Shavell (2007), “Moral Rules, the Moral Sentiments, and Behavior: Toward a Theory of an Optimal Moral System,” *Journal of Political Economy*, 115: 494–514.

- Lewis, T., and D. Sappington (1989), “Countervailing Incentives in Agency Problems,” *Journal of Economic Theory*, 49: 294–313.
- Maskin, E., and J. Riley (1984), “Monopoly with Incomplete Information,” *Rand Journal of Economics*, 15(2): 171–196
- McGreal, C. (2018), *American Overdose: The Opioid Tragedy in Three Acts*, Public Affairs.
- Mussa, M., and S. Rosen (1978), “Monopoly and Product Quality,” *Journal of Economic Theory*, 18(2): 301–317.
- Rochet, J.C., and L. Stole (2002), “Nonlinear Pricing with Random Participation,” *Review of Economic Studies*, 69(1): 277–311.
- Schnell, M. (2019), “The Opioid Crisis: Tragedy, Treatments and Trade-offs,” Stanford Institute for Economic Policy Research Policy Brief, February.
- Spence, A.M. (1975), “Monopoly, Quality, and Regulation,” *The Bell Journal of Economics*, 6(2): 417–429.