A new version of Edgeworth's taxation paradox

Robert A. Ritz*
Faculty of Economics
University of Cambridge
rar36@cam.ac.uk

Revised version: October 2012

Abstract

Edgeworth's taxation paradox states that a unit tax can *decrease* the market price of a good. This paper presents a new version of the paradox in which a tax reduces price — and increases industry output — because it attracts additional entry into the market. It is particularly striking that the demand conditions under which cost pass-through exceeds 100% for a fixed number of firms are also those for which pass-through can turn negative with endogenous entry. A novel application to the environment shows that a Pigouvian emissions tax can lead to an increase in industry emissions. A basic principle of environmental policy therefore fails under the conditions of the paradox.

Keywords: Cost pass-through, Edgeworth's paradox, endogeneous entry, environmental regulation, excise taxation, market structure.

JEL classifications: D43 (oligopoly & imperfect markets), H22 (tax incidence), L51 (economics of regulation), Q58 (environmental policy).

*My thanks are due to the Editor and two anonymous referees, Toke Aidt, Simon Cowan, Robert Hahn, Stephen Hamilton, John Quah, and John Vickers for helpful comments and advice, to Glen Weyl for an insightful discussion on the economics of pass-through, and to seminar participants at EARIE 2012 (Rome), especially Gianni de Fraja and Gerhard Clemenz, for their feedback. All views expressed and any errors are my own.

1 Introduction

Can a unit tax on a product decrease its price? Standard economic analysis suggests that the answer is "no". However, Edgeworth (1925) showed that, under certain conditions, a multi-product monopolist may indeed respond to a tax on one of its goods by reducing price: "When the supply of two or more correlated commodities — such as the carriage of passengers by rail first class or third class — is in the hands of a single monopolist, a tax on one of the articles — e.g., a percentage of first class fares — may prove advantageous to the consumers as a whole... The fares for *all* classes might be reduced" (p. 139). This result has become known as Edgeworth's taxation paradox.[1]

The intuition for the result is that — in contrast to a single-product setting — the price of the other (untaxed) good need not remain unchanged. In particular, the price of the other good may decrease, which, under conditions of complementarity, can exert downward pressure on the price of the taxed good to the extent that both prices end up falling in response to the tax.

Edgeworth's surprising result prompted contributions by several distinguished economists. Hotelling (1932) showed that price-reducing unit taxes can also obtain under perfect competition with substitute goods and diseconomies of scope in production; see also Bailey (1954). Vickrey (1960) presents conditions for the case of perfect competition with two products, and gives a variety of practical examples. Coase (1946) provides a useful graphical analysis of how the paradox can occur; see also the discussions in Bowley (1924), Wicksell (1934), and Creedy (1988).

More recently, Salinger (1991) uses the logic underlying Edgeworth's paradox to show that vertical integration may *decrease* welfare in multi-product settings (whereas it is always beneficial with a single product). The standard argument that vertical integration is welfare-enhancing — because it eliminates a double-marginalization problem — turns out to depend on a (strictly) positive rate of cost pass-through by the downstream firm. With multiple products, pass-through turns negative under the conditions of the paradox, and so vertical integration can end up reducing welfare.

In this paper, I present a new version of Edgeworth's paradox of taxation. The basic setup involves a market with a low-cost incumbent and a higher-cost

[1] The original version of Edgeworth's paper was in Italian (and first published in 1897). See Moss (2003) for further historical background on the taxation paradox.

potential entrant. Initially, that is, before a unit tax is introduced, the entrant decides not to enter the market because of an entry cost. The incumbent firm initially acts as a monopolist, charging the monopoly price. However, the introduction of the tax "levels the playing field" between the incumbent and the entrant in the post-entry game. In particular, the tax increases the operating profits of a sufficiently small entrant (Lemma 1), and can therefore attract additional entry.

The basic trade-off is clear: Although the unit tax as such leads to an increase in price (as in a standard model), it can also induce the higher-cost firm to enter the market, which, all else equal, reduces price. My main result gives conditions under which *any* sufficiently small tax induces entry and leads to a decrease in price for a range of values for the entry cost (Proposition 1).

Of course, the conditions required for the paradox to arise are special — just as they are for Edgeworth's original formulation. The two key elements of my approach are that (i) an input price increase (e.g., due to an excise tax) may increase profits, and (ii) entry reduces price. The first element applies more widely than is often realized — see also Seade (1985) and Anderson, de Palma and Kreider (2001) — while the second element is, of course, common to many models of competition.[2]

The main result in the paper is derived from a single-product setting with a generalized form of Cournot competition (following entry). The industry demand curve is log-convex (i.e., $\log D(p)$ is convex in price), a condition which is satisfied, for example, by any constant-elasticity demand curve. In such settings, a unit tax always increases a sufficiently small firm's market share as well as its operating profits, thus making the paradox possible.

My version of the paradox is stronger than Edgeworth's formulation in two respects. First, it shows that a unit tax can reduce price even in a single-product setting. Second, it shows that an input tax can increase the output of a taxed good. By contrast, other formulations of the paradox rely on a multi-product setting, and, although the price of the taxed good falls, its output does not increase; see also Salinger (1991). (However, my version of the paradox is weaker than some existing versions in that it relies on imperfect competition in the product market.)

I also present a new application of the paradox — based on the novel

[2] Some exceptions to this central message of economics have recently been highlighted by Bulow and Klemperer (2002), Chen and Riordan (2008) and Cowan and Yin (2008).

output-increasing effect — to environmental economics. In particular, with imperfect competition in the product market, a per-unit tax on industrial emissions may lead to an *increase* in emissions because it induces additional entry of firms. A basic principle of environmental policy, originally due to Pigou (1920), therefore fails under the conditions of my version of Edgeworth's paradox. For example, the standard Pigouvian tax set at the social marginal damage of emissions may fail to control emissions. Amongst other things, this finding suggests a novel reason for why cap-and-trade schemes might be preferable to emissions taxes: They can be less vulnerable to unforeseen changes in market structure induced by environmental policy.

A similar yet distinct mechanism to the one presented here operates in the multi-product framework of Hamilton (2009), in which firms choose prices as well as the breadth of their product variety. If demand conditions are such that cost pass-through would be high for a fixed product breadth, this induces firms to widen their product portfolios and increases competition — so pass-through, in equilibrium, is low. However, pass-through always remains positive in this framework and Edgeworth's taxation paradox does not arise.

Also related is Kind, Koethenbuerger and Schjelderup (2008) who show that *ad valorem* (but not per-unit) taxes can lead to *output* increases in models of "two-sided" markets such as newspapers and credit cards. In particular, a higher value-added tax on one side of the market can induce a firm to shift revenue to the other side — which in turn raises the marginal profitability of taxed side of the market under conditions of complementarity — and thus can lead to increases in equilibrium output (on one or both sides of the market).

My version of the tax paradox is less likely to hold with ad valorem taxes. It is well-known that such taxes can be re-interpreted as the combination of a unit tax and a profit tax.[3] Moreover, with asymmetric firms, the corresponding unit tax is higher for a firm with higher marginal cost. In my setting, the potential entrant would thus both face a higher unit tax than the incumbent, and also incur a direct tax on operating profits. These features work against the tax's profit-increasing effect, making the entry-induced paradox less likely.

The present analysis considers per-unit taxes because these (i) are the focus

[3] Write firm k's operating profits as $\Pi_k(s) = (1-s)p_k X_k - c_k X_k$, where p_k is price, X_k is output, c_k is unit cost, and $s \in (0,1)$ is an ad valorem tax. Equivalently, write $\Pi_k(s) = (1-s)(p_k - c_k)X_k - sc_k X_k$, showing that the ad valorem tax corresponds to a profit tax s on operating profits $(p_k - c_k)X_k$ together with a firm-specific per-unit tax sc_k (which is higher for a firm with higher unit cost).

of existing discussions of the Edgeworth paradox, and (ii) correspond directly to Pigouvian taxes in the application to the environment.

The plan for the remainder of the paper is as follows. Section 2 sets up the basic model and derives the initial equilibrium before introduction of the tax. Section 3 shows how a unit tax can increase profits, and Section 4 shows when a tax together with additional entry reduce price. Section 5 contains the main result in form of a new version of the taxation paradox. Section 6 presents the application to environmental policy. Section 7 gives concluding remarks.[4]

2 Model

□ **Model setup**. I use a simple model to show how the taxation paradox can arise. Consider an industry with two firms, the incumbent, firm A, with unit cost c^A, and a potential entrant, firm B, with unit cost c^B ($c^B > c^A$). The industry faces an inverse demand curve $p(X)$, where X is industry output. The price elasticity of demand $\eta \equiv -p(X)/Xp'(X) > 0$, and suppose that demand is price-elastic, $\eta > 1$, and that its elasticity is non-decreasing in price, $d\eta/dp \geq 0$.[5] Define an index of demand curvature $\xi \equiv -Xp''(X)/p'(X)$, where $\xi < 2$ so the monopoly problem is well-behaved.[6] The key assumption is that the demand curve is sufficiently convex with $\xi > 1$, which holds if and only if the corresponding direct demand $D(p)$ is log-convex (i.e., $\log D(p)$ is convex in price). These conditions are satisfied, for example, by any constant-elasticity demand curve (for which $\xi = (1 + \eta^{-1})$, with η a constant).

There is a fixed cost of K that firm B incurs by entering the market. If firm B decides to enter, there is a generalized version of Cournot competition in which firms choose their outputs X^A and X^B respectively. Let σ^A and σ^B denote the associated market shares (with $\sigma^A + \sigma^B = 1$). In particular, post-entry industry outcomes are determined by a conjectural-variations equilibrium with conduct parameter θ: When firm j changes its output by, say, dX^j, it believes that industry output will change by $dX = \theta(dX^j)$ as a result. The parameter θ serves as a useful summary statistic of the intensity of competition in the industry; it may be regarded as a reduced-form of an underlying (unmodelled) dynamic game (e.g., Cabral, 1995). Lower values of θ correspond to more com-

[4] All proofs are collected in the Appendix.

[5] This latter condition is sometimes referred to as "Marshall's second law of demand".

[6] In other words, industry revenue $p(X)X$ is strictly concave in industry output.

petitive outcomes. I assume that $\theta > 0$ so that competition is imperfect, and that $\theta < 2$ so that industry conduct is more competitive than the (symmetric) perfectly collusive outcome. The standard case of Cournot-Nash competition, where each firm takes its rival's output as given, is nested where $\theta = 1$.

If firm B does not enter, the incumbent firm A acts as a (profit-maximizing) monopolist against the demand curve $p(X)$.

The timing of the model is as follows. First, the regulator sets the level of an excise tax of t per unit of each firm's output. Second, firm B decides whether or not to enter the market. Third, firm A and firm B (if it has entered) choose their outputs and the market price is determined.

Initially, the tax is set to zero ($t = 0$), firm B decides whether to enter (with firm A already in the market as the incumbent), and the equilibrium market price is determined. Then, an excise tax of t per unit of output is introduced, followed by firm B's entry decision, and the market outcome. How does the market price under the tax compare to the initial market price in its absence?[7]

☐ **Initial equilibrium**. I am interested in a situation where firm B is a "serious" potential entrant. In other words, its unit cost is sufficiently low to allow it to gain positive market share in competition against the incumbent.[8] However, firm B initially does not enter because its operating profits are not sufficient to cover the entry cost. Specifically, consider the (hypothetical) outcome of generalized Cournot competition following entry in the absence of the tax. Let $\Pi^A(0)$ and $\Pi^B(0)$ denote, respectively, the incumbent's and the potential entrant's operating profits at a zero tax rate, $t = 0$. Assume that the entry cost is such that $K > \Pi^B(0)$, thus making entry unprofitable for firm B. It follows that the initial equilibrium has firm A as a monopolist, with market price

$$p_0 = \left(\frac{\eta_0}{\eta_0 - 1}\right) c^A,$$

where $\eta_0 > 1$ is the initial equilibrium value of the price elasticity of demand.

[7] The two equilibria of the competition-with-entry game that I compare (with and without the tax) are both subgame-perfect. To bring out the paradox as clearly as possible, I do not consider any entry-deterrence strategies that firm A may be able to pursue as a first-mover. Such strategies probably make the tax paradox more difficult to obtain. If deterrence is successful, then the price-reducing effect of entry disappears. Conversely, if entry deterrence is attempted but unsuccessful, the initial, pre-entry market price is (weakly) lower than it otherwise would have been, so the paradox becomes less likely.

[8] A precise statement for firm B to indeed be a serious entrant is provided as condition A2 in conjunction with the main result in Section 5.

3 Profit-increasing taxes

Now consider the situation after the introduction of the excise tax. The key issue is how the tax affects the competitive balance in the industry, especially firm B's operating profits following entry.

Lemma 1 *The post-entry impact of a change in a unit tax t is such that:*
(i) The rate of pass-through exceeds 100%, $dp(t)/dt > 1$;
(ii) Firm B's market share increases, $d\sigma^B(t)/dt > 0$;
(iii) Firm B's operating profits increase if only if its market share is sufficiently small,

$$\frac{d\Pi^B(t)}{dt} > 0 \text{ if and only if } \sigma^B(t) < \left(1 - \frac{1}{\xi}\right).$$

Lemma 1(iii) shows that firm B actually benefits from an increase in the tax as long as its market share is sufficiently small, $\sigma^B(t) < (1-\xi^{-1})$. Note that this condition requires that the industry demand curve is log-convex $\xi > 1$, and is more easily satisfied for more convex demand (i.e., higher values of ξ).

Furthermore, the results from Lemma 1 do not depend on the intensity of competition in the post-entry market, that is, on the precise value of the conduct parameter θ. The qualitative nature of the tax's impact is thus determined solely by market structure and demand conditions (while its magnitude is also influenced by industry conduct).

The intuition is that the tax "levels the playing field" between the incumbent and the entrant by making their market shares more symmetric. For a sufficiently small firm, this effect outweighs the negative impact on profits implied by the reduction in industry revenue.

In particular, the fact that firm B gains market share implies, all else equal, that it also makes higher profits. Of course, the tax also increases price, and thus reduces industry revenue (since demand is price-elastic), *in the post-entry game*. All else equal, lower industry revenue implies lower profits. For a sufficiently small firm, however, reduced industry revenue due to a small increase in the tax has a second-order effect on profits — while the gain in market share is first-order. Lemma 1(iii) gives a simple condition for the tax's overall impact to increase profits. For example, with constant-elasticity demand, the coefficient of demand curvature satisfies $\xi = (1 + \eta^{-1})$, so the condition becomes $\sigma^B(t) < 1/(\eta + 1)$. This latter condition is virtually always satisfied as the price elasticity $\eta \to 1$ (so industry revenue is approximately constant).

Another perspective on the result is that, *for a fixed number of firms*, the rate at which an input tax is passed through to consumers exceeds 100% under generalized Cournot competition with a log-convex demand curve. So operating profit margins increase — by an equal amount (in dollars) for all firms — in response to the tax. Put differently, since the consumer price p rises by more than the tax, the producer price $(p - t)$ increases due to the introduction of the tax. In relative terms, this helps a smaller, lower-margin firm more, and its overall profits may rise as a result.

Since the tax can increase firm B's profits, it can also induce it to enter the market by making its operating profits sufficient to cover the entry cost, $\Pi^B(t) \geq K > \Pi^B(0)$. (Using the same arguments as those underlying Lemma 1, a unit tax always reduces the incumbent firm A's operating profits as it loses market share.)

4 Price-reducing taxes

The basic trade-off is clear: Although the tax as such leads to an increase in price, it can also induce additional entry, which, all else equal, leads to a decrease in the market price.

Suppose for a moment that the tax does induce firm B to enter the market. The following result shows when the overall effect of an input tax plus additional entry is to reduce price.

A1. (Price reduction) Firms' unit costs satisfy $\dfrac{(c^B - c^A)}{c^A}(\eta_0 - 1) < (2-\theta)$.

Lemma 2 *Suppose that condition A1 holds. The post-entry impact of any tax $t < t_{\max}$ is such that the final price is lower than the initial price $p(t) < p_0$, where $t_{\max} \equiv \dfrac{1}{2}\left[\dfrac{[\eta_0 + (1-\theta)]}{(\eta_0 - 1)}c^A - c^B\right] > 0$.*

So the overall impact of a unit tax is price-reducing as long as the entrant's cost disadvantage is not too large, and the tax itself is also not too large. (Condition A1 ensures that the critical value for the unit tax is strictly positive, $t_{\max} > 0$.)

5 A new version of Edgeworth's paradox

Obtaining a new version of Edgeworth's paradox now only requires ensuring that the tax indeed increases firm B's operating profits — thus inducing additional entry and a reduction in the market price. Clearly, such an increase in firm B's profits occurs for some unit tax, say $\widetilde{t}$, if and only if $\left[\Pi^B(\widetilde{t}) - \Pi^B(0)\right] = \int_{z=0}^{\widetilde{t}} \left[\frac{d\Pi^B}{dt}(z)\right] dz > 0$, which in turn holds whenever $d\Pi^B/dt > 0$ for a sufficiently large portion of the interval $[0, \widetilde{t}]$. While this is, as such, straightforward, it is not possible to obtain the necessary-and-sufficient conditions for $\Pi^B(\widetilde{t}) > \Pi^B(0)$ in a simple, explicit form.

Instead, I present a sharper result showing that the tax paradox can hold for *any* unit tax $t \in (0, t_{\max})$. Two further ingredients are needed to deliver this main result. First, it requires that firm B is indeed a "serious" potential entrant with positive market share. Second, building on this, it requires that firm B's operating profits increase for any $t \in (0, t_{\max})$. Put differently, its market share is such that $0 < \sigma^B(t) < \left(1 - \xi^{-1}\right)$ from Lemma 1(iii) holds for all $0 \leq t < t_{\max}$.

A2. (Serious entrant) Firms' unit costs satisfy $\dfrac{(c^B - c^A)}{c^A}(\eta_0 - \theta) < \theta$.

A3. (Profit increase) Firms' unit costs satisfy $\dfrac{(c^B - c^A)}{c^A}(\eta_0 - 1) > \theta\dfrac{(2 - \underline{\xi})}{\underline{\xi}}$,
where $\underline{\xi} \equiv \inf_{t\in[0,t_{\max})} \xi(t)$ defines a lower bound on demand curvature.[9]

Condition A2 guarantees that firm B's market share $\sigma^B(t) > 0$ for any $t \geq 0$; its proof uses the assumption that the price elasticity of demand is non-decreasing in price ($d\eta/dp \geq 0$) to obtain a statement, as in condition A1 above, in terms of the *initial* price elasticity η_0. (The elasticity, of course, varies with the level of the tax whenever demand has non-constant elasticity.)

Condition A3 ensures that $\sigma^B(t) < \left(1 - \xi^{-1}\right)$ from Lemma 1(iii) holds for any $t \in [0, t_{\max})$; its proof uses the result from Lemma 1(ii) that $d\sigma^B(t)/dt > 0$ as well as the fact that demand curvature satisfies $\xi(t) \geq \underline{\xi}$ by the definition of its lower bound.

[9]This notation is short for $\xi(t) \equiv \xi(X(t))$ where $X(t)$ is equilibrium post-entry industry output as a function of the tax rate, t. There is no necessary relationship between demand curvature and industry output (and hence the tax rate); put differently, the sign of $\frac{d}{dX}\xi(X)$ is, in general, ambiguous. (Some familiar demand curves (including constant-elasticity) satisfy $\frac{d}{dX}\xi(X) = 0$, so demand curvature is constant — and thus invariant to the tax rate.)

Proposition 1 *Suppose that conditions A1–A3 hold. For any unit tax* $t \in (0, t_{\max})$*, there is a range of values for the entry cost* K *such that:*
(i) Firm B *enters the market,* $\Pi^B(t) \geq K > \Pi^B(0)$*;*
(ii) The final price is lower than the initial price, $p(t) < p_0$*.*

Proposition 1 offers a new version of Edgeworth's paradox of taxation in form of an per-unit tax that reduces price because it attracts additional entry into the market. The bounds on firm B's unit cost from conditions A2 and A3 imply that its market share is such that the *any* unit tax that is not too large leads to an increase in its profits (see Lemma 1). Under these conditions, there is always a range of values for K such that the tax induces entry and leads to a decrease in price.

Put differently, choose the entry cost in a way that firm B is initially not too far away from entering the market. Then, for a sufficiently high-cost (but still "serious") potential entrant, a small unit tax will increase (post-entry) profits, induce entry, and decrease the market price under these conditions.

It is also clear that the tax here increases consumer surplus, and increases total welfare insofar as sufficiently high weight is placed on consumers relative to producers.[10] The new paradox also implies that equilibrium tax revenue exceeds the "naïve" revenue forecast that uses firm A's initial monopoly output as the tax base (because industry output rises).[11]

I now discuss in more detail when conditions A1–A3 which underlie Proposition 1 are satisfied, as well as several important special cases of the model.

While the lower bound on firm B's unit cost from A3 is fairly straightforward, the upper bound comes from the more stringent condition among A1 ("price reduction") and A2 ("serious entrant"). It is easily checked that condition A1 is more stringent if the conduct parameter $\theta \geq 1$, while A2 is more stringent whenever $\theta < 1$. Intuitively, if competition is relatively soft, then it is more difficult to get the market price to decrease because entry has less of an impact. (Condition A2 is automatically satisfied if the conduct parameter

[10] To see why industry profits must be lower with the tax, let Π^M denote monopoly profits (by firm A), and observe that, by revealed preference, $\Pi^M(t) \leq \Pi^M(0) \equiv \Pi_0^M$. Moreover, duopoly profits under generalized Cournot competition with the tax are certainly lower than monopoly profits, $\Pi^A(t) + \left[\Pi^B(t) - K\right] < \Pi^M(t)$. So industry profits fall from Π_0^M down to $\left[\Pi^A(t) + \left[\Pi^B(t) - K\right]\right]$.

[11] In this sense, too, the paradox "will surely be a grateful boon to the perplexed and weary secretaries of the Treasury and ministers of finance around the world," Seligman (1921, p. 214).

$\theta > \eta_0$.) By contrast, if post-entry competition is relatively tough, then it is more difficult for the potential entrant to be "serious" and gain positive market share. With Cournot-Nash competition ($\theta = 1$), conditions A1 and A2 coincide as $\left[(c^B - c^A)/c^A\right] < 1/(\eta_0 - 1)$.[12]

The conditions for the paradox can therefore be thought of as two cases:

□ **"Soft" competition** $\theta \in [1, \underline{\xi})$. Conditions A1–A3 together become

$$\frac{(c^B - c^A)}{c^A}(\eta_0 - 1) \in \left(\theta \frac{(2 - \underline{\xi})}{\underline{\xi}}, (2 - \theta)\right),$$

where $\theta < \underline{\xi} \in (1, 2)$ ensures that the upper bound on the entrant's relative cost exceeds the lower bound. In the limiting case of Cournot-Nash competition, the conditions simplify further to $\left[(c^B - c^A)/c^A\right](\eta_0 - 1) \in ((2 - \underline{\xi})/\underline{\xi}, 1)$. Moreover, if demand curvature $\underline{\xi}$ lies well above unity, this case admits competitive conditions significantly "more collusive" than in a Cournot-Nash equilibrium.

□ **"Tough" competition** $\theta \in (\underline{\theta}, 1)$. Conditions A1–A3 together become

$$\frac{(c^B - c^A)}{c^A} \in \left(\frac{\theta}{(\eta_0 - 1)} \frac{(2 - \underline{\xi})}{\underline{\xi}}, \frac{\theta}{(\eta_0 - \theta)}\right),$$

where $\underline{\theta} \equiv \max\left\{0, \left[\underline{\xi} - 2\eta_0(\underline{\xi} - 1)\right]/(2 - \underline{\xi})\right\}$ and $\theta > \underline{\theta}$, similar to above, ensures that an interval exists — post-entry competition between firms cannot be too intense. The precise value of $\underline{\theta}$ depends on the shape of the demand curve. However, note that $\underline{\theta} = 0$ is entirely possible; indeed, this *always* holds with constant-elasticity demand (for which $\underline{\xi} = \left(1 - \eta_0^{-1}\right)$). So the "tough" case can, in principle, cover everything from almost-perfect to almost-Nash competition.

Taken together, these two cases show that the new version of the tax paradox can occur for a very wide range of competitive conditions in an industry (facing a sufficiently convex demand curve).

□ **Numerical example**. The simplest conditions for the paradox obtain with Cournot-Nash competition ($\theta = 1$) and constant-elasticity demand, for which conditions A1–A3 together become $(c^B - c^A)/c^A \in ((\eta + 1)^{-1}, (\eta - 1)^{-1})$. For

[12] Loosely put, with a Cournot-Nash conjecture, the incumbent firm A's output choice varies smoothly between monopoly and duopoly equilibrium.

example, if firm A's unit cost $c^A = 1$ and the elasticity $\eta = 2$, then the initial price $p_0 = 2$. As long as firm B's unit cost $c^B \in (\frac{4}{3}, 2)$, the paradox can occur for any tax $t \in (0, t_{\max})$, where $t_{\max} = \left(1 - \frac{1}{2}c^B\right) \in (0, \frac{1}{3})$, for a range of values of the entry cost K. The final price $p(t) \in (1\frac{5}{9}, 2)$ can thereby fall (almost) to $p(t)|_{t\to 0} = 1\frac{5}{9}$, or up to around 22% lower than the initial price.

6 Application to environmental policy

Following Pigou (1920), a basic principle of environmental economics is that industrial pollution can be controlled with a per-unit tax on the emissions from a polluting activity. Such an emissions tax makes pollution costly to firms (so they internalize the external damages from pollution), and, under standard conditions, leads to a decrease in emissions. However, an application of the taxation paradox shows that this conclusion is not necessary.

Consider the model from Section 2, but now write a firm's operating profits as $\Pi^j(\tau) = (p - c^j)\, X^j - \tau E^j$ for $j \in \{A, B\}$, where $\tau \geq 0$ is a tax on its emissions E^j. For simplicity, suppose that emissions are a fixed proportion of output, $E^j = \lambda X^j$, where $\lambda > 0$ is the emissions intensity of production. (This may be a reasonable assumption for industries in which cleaner production is technologically infeasible or unprofitable at the prevailing tax rate on emissions.)

This setup maps into the above analysis by letting the emissions tax $\tau = t/\lambda$. When emissions are unpriced, firm B does not enter the market because of the entry cost, so the initial market price $p_0 = [\eta_0/(\eta_0 - 1)]c^A$, with associated output X_0. Under the conditions of Proposition 1, *any* emissions tax $\tau < \dfrac{1}{2\lambda}\left[\{[\eta_0 + (1-\theta)]/(\eta_0 - 1)\}\, c^A - c^B\right] \equiv \tau_{\max}$ induces firm B to enter and decreases the market price for a range of values for K. Since total industry emissions $E(\tau) = \lambda\left[X^A(\tau) + X^B(\tau)\right]$ are proportional to industry output, *the emissions tax leads to an increase in emissions*, $E(\tau) > E_0$, where $E_0 = \lambda X_0$ is initial emissions (at $\tau = 0$).

The paradox demonstrates that regulation may have unintended consequences when it not only affects firms' decision-making at the margin but also induces changes in market structure. Levin (1985) has shown that, with Cournot-Nash competition, an emissions tax may increase industry emissions when firms have sufficiently asymmetric emissions intensities (also assumed to be fixed). Similarly, environmental regulation that applies only to a subset of firms in an industry can increase emissions due to "emissions leakage" to dirtier, unregulated firms; see, e.g., Fowlie (2009). By contrast, the above argument requires neither that firms have asymmetric emissions intensities nor that only a subset of firms in the industry is subject to regulation.

With imperfect competition in product markets, the optimal second-best emissions tax typically deviates from a Pigouvian tax that is set at the social marginal damage of emissions, as first pointed out by Buchanan (1969). In particular, it is generally lower than social marginal damages in order to compensate for monopolistic (see, e.g., Barnett, 1980) or oligopolistic (see, e.g., Katsoulacos and Xepapadeas, 1996) underproduction. Nonetheless, Oates and Strassman (1984) argue that ignoring market structure imperfections leads only to small inefficiencies, and thus suggest that policymakers employ a simple (third-best) Pigouvian tax on emissions. My version of Edgeworth's paradox shows that this recommendation *may* be environmentally counterproductive. The reason is that the social marginal damage from emissions could easily lie within the interval between zero and $\tau_{\max}$. If so, Proposition 1 shows that the Pigouvian tax can inadvertently attract additional entry into the market (for some values of K), and thus lead to an increase in overall emissions.

The underlying point is that standard analyses of environmental policy are based on the premise that a higher emissions tax always reduces emissions, that is $E(\tau'') \leq E(\tau')$ for any $\tau'' > \tau'$. Although by no means universal, this relationship holds under *fairly* weak conditions for a given market structure, whether perfect or imperfect (with a fixed number of firms). However, my analysis here shows that this relationship can break down where the market structure is not only imperfect, but also *endogenous* (in the sense that the number of active firms is endogenously determined).[13]

This finding also has potential implications for the choice of environmental policy instruments. In particular, it suggests a novel reason for why an emissions trading scheme may be preferable to an emissions tax when product markets are imperfectly competitive.[14] The reason is simply that a cap-and-trade scheme for this industry would guarantee that the environmental objective, say

[13] The standard relationship would hold in the model of Section 2 if there is no entry, that is, if firm A remains a monopolist after the introduction of the emissions tax (because single-product monopoly pass-through is positive).

[14] Emissions taxes and trading are, of course, equivalent with perfectly competitive markets and no uncertainty. The classic analysis of instrument choice for perfect competition under uncertainty is due to Weitzman (1974).

$\overline{E} < E_0$, is met.[15] In the model of Section 2, this would imply — regardless of whether firm B enters or not — a decrease in industry output due to the policy, so the paradox could not arise. With imperfect competition in product markets, cap-and-trade schemes may thus be less vulnerable to unforeseen changes in market structure than an emissions tax.

Finally, the result is also of interest in that it runs counter to the standard intuition that tighter environmental policy tends to reduce industry output and decrease consumer welfare, see, e.g., Maloney and McCormick (1982). Under the conditions of my version of the paradox, consumers and the entrant benefit from environmental policy, while the incumbent and the environment both lose out. In this case, the incumbent firm's likely resistance to environmental regulation might be beneficial insofar as it prevents the introduction of an emissions-increasing emissions tax. This is a distinct variation on the standard theory of the political economy of instrument choice due to Buchanan and Tullock (1975), in which the incumbent wants to block an emissions tax that is socially desirable.

Another related strand of the literature, albeit with a somewhat different focus, examines the design of a second-best emissions tax in oligopoly settings with free entry. Here, also with an endogenous number of firms, the second-best tax may actually exceed the Pigouvian tax in order to compensate for excess entry due to business-stealing effects; see, e.g., Katsoulacos and Xepapadeas (1995). The key difference is that this literature looks at free-entry equilibria with a large pool of symmetric potential entrants in which number of active firms is pinned down by a zero-profit condition, while my model has a small number of asymmetric firms making positive profits.[16] Moreover, Edgeworth's paradox does not arise in such free-entry models since the emissions tax makes entry less attractive to symmetric firms, which in turn further increases price — so the standard relationship $E(\tau'') \leq E(\tau')$ for any $\tau'' > \tau'$ still holds.

7 Concluding remarks

This paper has presented a new version of Edgeworth's taxation paradox: A per-unit excise tax can reduce price — and increase output — because it

[15] As is standard in the literature, this assumes compliance by regulated firms.

[16] I also effectively restrict attention to examining third-best emissions taxes of the kind suggested by Oates and Strassman (1984).

"levels the playing field" between firms and attracts additional entry into the market. In contrast to existing formulations of the paradox, my version applies in single-product settings.

The analysis shows that price-reducing taxes (Section 4) can result from profit-increasing taxes (Section 3) when the market structure is endogenously determined. It is particularly striking that the demand conditions (i.e., log-convexity of demand) under which cost pass-through exceeds 100% for a fixed number of firms are also those for which cost pass-through can turn negative due to entry of additional firms in equilibrium.[17]

Weyl and Fabinger (2009) suggest that empirical estimates of pass-through can be used to infer the curvature of demand (which is generally unobservable and also difficult to estimate empirically), and thus sign — or even quantify — many comparative statics in models of imperfect competition.[18] My analysis makes clear that this technique, while promising, may be quite sensitive to the underlying assumption that the number of firms in the market is fixed.

The application of the new version of the paradox to environmental policy (Section 6) shows that regulation in form of a per-unit emissions tax may have unintended consequences when it not only affects firms' decision-making at the margin but also induces changes in market structure. In particular, a Pigouvian emissions tax may inadvertently attract additional entry into the market, and thus cause industrial emissions to rise. This has a number of potential implications for the choice and design of environmental policy instruments, as well as for their political economy.

Of course, the conditions required for the paradox to occur are special, and ultimately the question of how a tax affects industry prices, outputs and emissions is an empirical one. The analysis presented in this paper shows that induced changes in market structure can play an important role, and may overturn standard results that are usually taken for granted by economists.

[17] I conjecture that the paradox can also exist under price competition with differentiated products. Anderson, de Palma and Kreider (2001) show that, with a fixed number of symmetric firms, an input tax is passed through by more than 100% and raises operating profits under very similar conditions to those above (i.e., sufficiently convex demand curves). I am, however, not aware of any existing results on the profit impact of excise taxes under price competition with *asymmetric* firms — which would be needed to be able to extend the present results.

[18] See also Hepburn, Quah and Ritz (2007) for an application to environmental economics that uses a similar basic idea.

Appendix

Proof of Lemma 1. The proof begins by deriving post-entry equilibrium conditions and then addresses the three parts of the result. In a generalized Cournot equilibrium with a unit tax t, firm $j = \{A, B\}$ chooses its output according to the first-order condition

$$\left[(p-t)-c^j\right] = \theta[-p'(X)X^j]. \tag{1}$$

Since $c^B > c^A$, it follows that $X^B < X^A$ and so $\sigma^B < \frac{1}{2} < \sigma^A$ for any unit tax $t \geq 0$. Summing the two first-order conditions for firms A and B gives

$$2(p-t) - (c^A + c^B) = \theta[-p'(X)X]. \tag{2}$$

Part (i). Differentiating the expression from (2), while noting $d[-p'(X)X]/dt = (\xi - 1)(dp/dt)$, yields $[2\,(dp/dt - 1)] = \theta(\xi - 1)(dp/dt)$. This can be rearranged to obtain the rate of pass-through

$$\frac{dp(t)}{dt} = \frac{1}{[1-(\theta/2)(\xi-1)]}. \tag{3}$$

Note that the denominator $[1-(\theta/2)(\xi-1)] > 0$ given the maintained assumptions $\xi < 2$ and $\theta < 2$. (These assumptions also ensure that the post-entry equilibrium is unique and stable.) It follows immediately that $dp(t)/dt > 1$ if and only if demand is log-convex, $\xi > 1$, as assumed.

Part (ii). Rewrite the first-order condition for firm B from (1) in terms of its market share:

$$\sigma^B(t) = \frac{\left[(p-t)-c^B\right]}{\theta[-p'(X)X]}. \tag{4}$$

Differentiating this expression and some rearranging shows that

$$\frac{d\sigma^B(t)}{dt} = \frac{1}{\theta[-p'(X)X]}\left[\left(\frac{dp}{dt}-1\right) - \theta(\xi-1)\sigma^B\frac{dp}{dt}\right]. \tag{5}$$

Using the formula for the rate of pass-through from (3) and simplifying yields

$$\frac{d\sigma^B(t)}{dt} = \frac{(\xi-1)}{[-p'(X)X]}\frac{\left(\frac{1}{2}-\sigma^B\right)}{[1-(\theta/2)(\xi-1)]}. \tag{6}$$

Since $\sigma^B < \frac{1}{2}$ and $\xi > 1$, it follows that $d\sigma^B/dt > 0$.

Part (iii). Write firm B's operating profits after the introduction of the unit tax as $\Pi^B(t) = \left[(p-t) - c^B\right] X^B$. Differentiating and some rearranging yields

$$\begin{aligned}
\frac{d\Pi^B(t)}{dt} &= X^B \left[\left(\frac{dp}{dt} - 1 \right) + \frac{\left[(p-t) - c^B\right]}{X^B} \frac{dX^B}{dt} \right] \\
&= X^B \left[\left(\frac{dp}{dt} - 1 \right) + \theta[-p'(X)] \frac{dX^B}{dt} \right],
\end{aligned} \tag{7}$$

where the second line uses the fact that the first-order condition for firm B from (1) implies that

$$X^B(t) = \frac{\left[(p-t) - c^B\right]}{\theta[-p'(X)]}. \tag{8}$$

Differentiating this expression for X^B also shows that the change in output

$$\begin{aligned}
\frac{dX^B(t)}{dt} &= \frac{1}{\theta[-p'(X)]} \left[\left(\frac{dp}{dt} - 1 \right) - \frac{\left[(p-t) - c^B\right]}{\theta[-p'(X)X]} \theta \frac{[-Xp''(X)]}{p'(X)} \frac{dp}{dt} \right] \\
&= \frac{1}{\theta[-p'(X)]} \left[\left(\frac{dp}{dt} - 1 \right) - \theta\xi\sigma^B \frac{dp}{dt} \right],
\end{aligned} \tag{9}$$

where the second line uses the expression for firm B's market share from (4) and the definition of demand curvature $\xi \equiv -Xp''(X)/p'(X)$. Using this result in the expression for the change in firm B's operating profits from (7) gives

$$\frac{d\Pi^B(t)}{dt} = X^B \left[2 \left(\frac{dp}{dt} - 1 \right) - \theta\xi\sigma^B \frac{dp}{dt} \right]. \tag{10}$$

Finally, using the formula for pass-through from (3) and simplifying yields

$$\frac{d\Pi^B(t)}{dt} = \frac{\theta X^B}{[1 - (\theta/2)(\xi - 1)]} \left[(\xi - 1) - \xi\sigma^B \right]. \tag{11}$$

So $d\Pi^B(t)/dt > 0$ if and only if $\sigma^B(t) < \left(1 - \xi^{-1}\right)$, thus completing the proof.

Proof of Lemma 2. Rearranging the combined first-order condition from (2) yields an expression for the equilibrium market price

$$p(t) = \frac{\eta(t)}{\left(\eta(t) - \dfrac{\theta}{2} \right)} \left[t + \frac{1}{2}(c^A + c^B) \right], \tag{12}$$

where the notation $\eta(t)$ makes explicit that the price elasticity of demand may, in general, vary with the level of the unit tax. Define $t_{\max}$ as the unit tax which satisfies $p(t_{\max}) \equiv p_0$, and note that then also $\eta(t_{\max}) = \eta_0$ (on a given demand curve, equal prices imply equal elasticities) such that

$$p(t_{\max}) = \frac{\eta_0}{\left(\eta_0 - \dfrac{\theta}{2}\right)} \left[t_{\max} + \frac{1}{2}(c^A + c^B)\right]. \tag{13}$$

Recalling from Section 2 that the initial equilibrium price $p_0 = [\eta_0/(\eta_0 - 1)]c^A$, it is easily checked that

$$t_{\max} = \frac{1}{2}\left[\frac{[\eta_0 + (1-\theta)]}{(\eta_0 - 1)}c^A - c^B\right]. \tag{14}$$

Observe that $t_{\max} > 0$ if and only if unit costs satisfy $[(c^B - c^A)/c^A](\eta_0 - 1) < (2 - \theta)$, which is condition A1. Since duopoly pass-through is positive, $dp(t)/dt > 0$, from (3), it follows that $p(t) < p_0$ for any $t < t_{\max}$ as long as c^B/c^A is not too large in the sense of A1.

Proof of Proposition 1. The proof of the main result proceeds in three steps: First, it uses first-order conditions to obtain an expression for firm B's market share in terms of the price elasticity. Second, it uses the assumption that the price elasticity of demand is non-decreasing in price, $d\eta/dp \geq 0$, to find a condition (A2) on firms' costs — in terms of the initial price elasticity η_0 — which guarantees that firm B is a serious potential entrant. Third, it uses the result of Lemma 1(iii) to find a sufficient condition (A3) for firm B's profits to increase for any tax $t \in (0, t_{\max})$.

Step 1. From (1), firm B's market share $\sigma^B(t) > 0$ if and only if the market price $p(t) > \left(t + c^B\right)$. Using the post-entry price $p(t)$ from (12) in the expression for $\sigma^B(t)$ from (4) gives the following formula for firm B's market share in terms of the price elasticity $\eta(t)$:

$$\sigma^B(t) = \frac{\eta(t)c^A - [\eta(t) - \theta]c^B + \theta t}{\theta\left[(c^A + c^B) + 2t\right]}. \tag{15}$$

Therefore, a sufficient condition for $\sigma^B(t) > 0$ is that $[(c^B - c^A)/c^A][\eta(t) - \theta] < \theta$. Note that this condition is more difficult to satisfy with a higher price elasticity, $\eta(t)$.

Step 2. Given condition A1 and the result from Lemma 2, the post-entry market price satisfies $p(0) < p(t) < p(t_{\max}) \equiv p_0$ for any $t \in (0, t_{\max})$. With the maintained assumption that the price elasticity of demand is non-decreasing in price, $d\eta/dp \geq 0$, this translates into $\eta(0) \leq \eta(t) \leq \eta(t_{\max}) \equiv \eta_0$ for the associated elasticities. Combining this with the result from Step 1 shows that $[(c^B - c^A)/c^A](\eta_0 - \theta) < \theta$ is sufficient for $\sigma^B(t) > 0$ for any tax $t \in [0, t_{\max}]$, which is condition A2. Then firm B is indeed a serious potential entrant, and so results from Lemma 1 apply for any tax $t \geq 0$.

Step 3. From Lemma 1(iii), $d\Pi^B(t)/dt > 0$ if and only if $\sigma^B(t) < 1 - [\xi(t)]^{-1}$, where the notation $\xi(t)$ here makes explicit that, in general, demand curvature may vary with the unit tax. Since $\underline{\xi} = \inf_{t \in [0, t_{\max}]} \xi(t)$ defines a lower bound on demand curvature, it follows that, if $\sigma^B(t) < (1 - \underline{\xi}^{-1})$ for all $t \in [0, t_{\max}]$, then $\Pi^B(t) > \Pi^B(0)$ for all $t \in (0, t_{\max}]$. Moreover, since log-convexity $\xi > 1$ implies that $d\sigma^B(t)/dt > 0$ from Lemma 1(ii), it follows that, if $\sigma^B(t_{\max}) < \left(1 - \underline{\xi}^{-1}\right)$, then certainly $\Pi^B(t) > \Pi^B(0)$ for all $t \in (0, t_{\max}]$. Now rewrite the expression for firm B's market share from (4) as

$$\sigma^B(t_{\max}) = \frac{\eta_0}{\theta} \left(\frac{p_0 - t_{\max} - c^B}{p_0} \right), \tag{16}$$

which again uses that the fact that $\eta(t_{\max}) \equiv \eta_0$. Inserting the expressions for $p_0 = [\eta_0/(\eta_0 - 1)]c^A$ and for $t_{\max}$ from (14) and simplifying yields

$$\sigma^B(t_{\max}) = \frac{1}{2} \left[1 - \frac{(\eta_0 - 1)}{\theta} \frac{(c^B - c^A)}{c^A} \right]. \tag{17}$$

Straightforward manipulations show that $\sigma^B(t_{\max}) < (1 - \underline{\xi}^{-1})$ if and only if $[(c^B - c^A)/c^A](\eta_0 - 1) > \theta[(2 - \underline{\xi})/\underline{\xi}]$, which is condition A3. So $\Pi^B(t) > \Pi^B(0)$ for any $t \leq t_{\max}$ as long as c^B/c^A is not too small in the sense of A3.

To summarize, for any unit tax $t \in (0, t_{\max})$, and under conditions A1–A3, (i) firm B enters the market for any entry cost K such that $\Pi^B(t) \geq K > \Pi^B(0)$, and (ii) the final price is lower than the initial price $p(t) < p_0$.

References

Anderson, Simon P., André de Palma, and Brent Kreider (2001). Tax incidence in differentiated product oligopoly. *Journal of Public Economics* **81**, pp. 173–192.

Bailey, Martin J. (1954). Edgeworth's taxation paradox and the nature of demand functions. *Econometrica* **22**, pp. 72–76.

Barnett, A. H. (1980). The Pigouvian tax rule under monopoly. *American Economic Review* **70**, pp. 1037–1041.

Bowley, A.L. (1924). *The Mathematical Groundwork of Economics.* Oxford: Clarendon Press.

Buchanan, James M. (1969). External diseconomies, corrective taxes, and market structure. *American Economic Review* **59**, pp. 174–177.

Buchanan, James M. and Gordon Tullock (1975). Polluters' profits and political response: Direct controls versus taxes. *American Economic Review* **65**, pp. 139–147.

Bulow, Jeremy and Paul Klemperer (2002). Prices and the winner's curse. *RAND Journal of Economics* **33**, pp. 1–21.

Cabral, Luís M. B. (1995). Conjectural variations as a reduced form. *Economics Letters* **49**, pp. 397–402.

Chen, Yongmin and Michael H. Riordan (2008). Price-increasing competition. *RAND Journal of Economics* **38**, pp. 1042–1058.

Coase, Ronald H. (1946). Monopoly pricing with interrelated costs and demands. *Economica* **13**, pp. 278–294.

Cowan, Simon and Xiangkang Yin (2008). Competition can harm consumers. *Australian Economic Papers* **47**, pp. 264–271.

Creedy, John (1988). Wicksell on Edgeworth's tax paradox. *Scandinavian Journal of Economics* **90**, pp. 101–112.

Edgeworth, Francis Y. (1925). The pure theory of monopoly. In: *Papers Relating to Political Economy* — Volume I, pp. 111–142. London: Macmillan & Co.

Fowlie, Meredith L. (2009). Incomplete environmental regulation, imperfect competition and emissions leakage. *American Economic Journal: Economic Policy* **1**, pp. 72–112.

Hamilton, Stephen F. (2009). Excise taxes with multiproduct transactions. *American Economic Review* **99**, pp. 458–471.

Hepburn, Cameron J., John K.-H. Quah, and Robert A. Ritz (2007). Emissions trading and profit-neutral grandfathering. Department of Economics, Oxford University (February 2007).

Hotelling, Harold H. (1932). Edgeworth's paradox of taxation and the nature of demand and supply functions. *Journal of Political Economy* **40**, pp. 577–615.

Katsoulacos, Yannis and Anastasios Xepapadeas (1995). Environmental policy under oligopoly with endogenous market structure. *Scandinavian Journal of Economics* **97**, pp. 411–420.

Katsoulacos, Yannis and Anastasios Xepapadeas (1996). Emission taxes and market structure. In: Carlo Carraro, Yannis Katsoulacos, and Anastasios Xepapadeas (editors), *Environmental Policy and Market Structure*, pp. 3–22. Dordrecht/Boston/London: Kluwer Academic Publishers.

Kind, Hans Jarle, Marko Koethenbuerger and Gutthorm Schjelderup (2008). Efficiency-enhancing taxation in two-sided markets. *Journal of Public Economics* **92**, pp. 1531–1539.

Levin, Dan (1985). Taxation within Cournot oligopoly. *Journal of Public Economics* **27**, pp. 81–90.

Maloney, Michael T. and Robert E. McCormick (1982). A positive theory of environmental quality regulation. *Journal of Law & Economics* **25**, pp. 99–123.

Moss, Laurence S. (2003). The Seligman-Edgeworth debate about the analysis of tax incidence: The advent of mathematical economics, 1892–1910. *History of Political Economy* **35**, pp. 205–240.

Oates, Wallace E. and Diana L. Strassman (1984). Effluent fees and market structure. *Journal of Public Economics* **24**, pp. 29–46.

Pigou, Arthur C. (1920). *The Economics of Welfare.* London: Macmillan & Co.

Salinger, Michael A. (1991). Vertical mergers in multi-product industries and Edgeworth's paradox of taxation. *Journal of Industrial Economics* **39**, pp. 545–556.

Seade, Jesús (1985). Profitable cost increases and the shifting of taxation: Equilibrium responses of markets in oligopoly. Department of Economics, Warwick University (April 1985).

Seligman, Edwin R. A. (1921). *Shifting and the Incidence of Taxation.* New York: Columbia University Press.

Vickrey, William S. (1960). Can excises lower prices? In: Ralph W. Pfouts (editor), *Essays in Economics and Econometrics*, pp. 165–177. Chapel Hill: University of North Carolina Press.

Weitzman, Martin L. (1974). Prices vs. quantities. *Review of Economic Studies* **41**, pp. 477–491.

Weyl, E. Glen and Michal Fabinger (2009). Pass-through as an economic tool. Harvard University (October 2009).

Wicksell, Knut (1934). *Lectures on Political Economy.* London: Routledge.