The role of information and market concentration in regulation [a]

Larry Karp[b]

December 20, 2025

Abstract

Each of a finite number of strategic agents benefits from their emissions and incurs damages from aggregate emissions. Each receives a private signal about their abatement costs that is correlated, via a market fundamental, with a damage shock and other agents' private signals. I consider demand function equilibria under a tax and a quota. Mapping the n-agent economy into its representative-agent analog produces a welfare decomposition that facilitates comparative statics of welfare under each policy and of the ranking criterion between policies. These functions depend on: the number of agents; the correlation between cost and damage shocks; the precision of the private signal; and a familiar ratio of damage and abatement cost slopes. I then discuss a "smart tax" that supports the full-information optimum in a setting with strategic agents.

Keywords: regulation, asymmetric information, market power, pollution control, policy ranking

JEL, classification numbers: Q000, Q500, H200, D800

[a]Prepared for the 2026 Winter ASSA conference.

[b]Department of Agricultural and Resource Economics, University of California, Berkeley, email: karp@berkeley.edu.

1 Introduction

Market concentration and the distribution of information jointly shape both the performance of standard environmental instruments and the scope for more efficient regulatory design. This paper studies pollution control when each of n strategic agents benefits from their own emissions but incurs damages from aggregate emissions. Each agent observes a private signal about its abatement costs, the sum of a stochastic market fundamental and an idiosyncratic shock. The market fundamental also affects marginal damages, implying that private cost information is correlated across agents and with the economy-wide damage shock.

Taxes and quotas alter behavior through different strategic channels. Under a binding quota implemented via a permit market, agents submit demand schedules for permits and an auctioneer clears the market. Because aggregate emissions are fixed at the cap, agents treat pollution damages as given, but they shade their bids to reduce the equilibrium permit price. With heterogeneous and imperfectly observed costs, this shading creates allocative inefficiency: dispersion of agents' realized marginal abatement costs. Under a tax, by contrast, agents treat the tax rate as given but understand that their own emissions affect aggregate emissions and therefore their own damages. This partial internalization can raise welfare, but with heterogeneous information it also generates inefficiency via dispersion of equilibrium marginal abatement costs. Imperfect competition and correlated information change the nature and magnitude of the distortions, and therefore the relative performance of instruments.

A direct welfare comparison of instruments in this setting can be complicated. I therefore map the n-agent fragmented economy (FE) into a representative-agent economy (RAE) analog that yields a simple decomposition: welfare under either instrument equals (i) a component that depends only on expected aggregate emissions and (ii) a "β-function" that collects welfare effects of second moments (volatility and dispersion) generated by strategic behavior and imperfect information. For certainty-equivalent policy pairs—policies that induce the same expected aggregate emissions—ranking reduces to comparing the associated β-functions. This decomposition makes comparative statics and instrument ranking tractable.

Four parameters are key. The number of agents n indexes market concentration; higher n means more competition, not a larger economy.[1] The parameter τ captures the relative precision of public versus private information; holding public-signal variance fixed, higher τ corresponds to a less precise private signal. The parameter λ is the regression coefficient of the damage shock on the average of agents' private cost signals, capturing the sign and strength

[1]Suppose that there are 1000 factories, each with its production (or abatement) technology. In once scenario there are 10 agents, each owning 100 factories, and in a second scenario there are 100 agents, each with 10 factories. A change in n from 10 to 100 does not alter the technology, but possibly changes agents' incentives, altering the outcome.

of cost–damage correlation induced by the market fundamental. Finally, W is the Weitzman ratio: the slope of marginal damages relative to the slope of marginal abatement costs.

The paper contains three sets of results. First, the welfare decomposition provides transparent comparative statics under each instrument. Under a quota, increased competition raises welfare by weakening incentives to shade bids; this gain is especially important when private signals are imprecise, since bid shading then produces larger allocative inefficiencies through dispersion in marginal abatement costs. Under a tax, increased competition weakens incentives to internalize damages, tending to raise expected emissions and expected damages; the effect on emissions volatility is generally ambiguous and depends on W, λ, and τ.

Second, I show how market power and correlated information reshape the tax–quota ranking. In the competitive limit ($n \to \infty$), the ranking of certainty-equivalent policies returns Weitzman's (1974) and Stavins's (1996) and ranking criteria: taxes welfare-dominate quotas if and only if $1 - 2\lambda > W$. Positive correlation between private costs and social damages tilts the ranking toward quantity instruments. The extension shows, for finite n, how the criterion depends on market concentration, the precision of the public and private signals and the correlation of abatement cost and damages. Changes in the relative precision of public versus private information can have opposite effects on equilibrium outcomes under the two instruments, and hence on the ranking.

Third, I study a more efficient instrument. As in Angeletos & Pavan's (2007) "team problem," suppose agents choose actions to maximize aggregate welfare instead of their own payoff, but without sharing information. The team equilibrium coincides with the outcome of a planner who observes the average private signal. I show how a tax schedule conditioned on aggregate emissions—a "smart tax"—can support this team outcome as a noncooperative Nash equilibrium when n is finite. This smart tax is impractical when $n = \infty$, where it makes unrealistic demands on agents' information (Karp & Traeger 2025). However, with moderate or small n strategic agents, the informational demands under the smart tax are the same as in the standard tax or quota games; in this case, the smart tax is a practical policy.

Roadmap. Section 2 derives the demand-function equilibria under a quota and a tax and constructs the FE→RAE mapping and welfare decomposition. Section 3 presents comparative statics under each instrument. Section 4 derives the tax–quota ranking and its comparative statics in market concentration and information precision. Section 5 analyzes smart instruments and shows how the smart tax can implement the team outcome for finite n. Section 6, in discussing the relation between static models of flow pollutants and dynamic models of stock pollutants (e.g., greenhouse gasses), argues that positively correlated cost and damage shocks should be considered the leading case, not a special case. Section 7 concludes.

1.1 Related literature

This paper intersects four literatures. A first strand studies how dispersed information shapes market outcomes and welfare, often focusing on the possibility that better information lowers welfare. Vives, a major contributor, surveys this literature in Vives (2008). Morris & Shin (2002) analyze environments with higher-order beliefs in a game that resembles a Keynesian "beauty contest". Angeletos & Pavan (2007) provide a unifying framework, emphasizing the distinction between information dispersion and aggregate volatility. This distinction is important in ranking the tax and quota (Section 4). They also study the team problem, used in Section 5. Applications of these ideas to environmental policy are less common, but Lemoine (2022) is an important exception: he shows how a "carbon share" (deposit–refund) policy can aggregate dispersed information about both climate damages and abatement costs. Other natural resource applications include Cantillon & Slechten (2018) on information aggregation in dynamic permit markets, Marcoul (2020) on information sharing in fisheries, and Englander et al. (forthcoming) on estimating imperfect-information fishery models.

A second strand studies imperfect competition in a cap-and-trade market. Malueg & Yates (2009) explain why demand-function competition for a fixed permit supply generates inefficiency due to unequal marginal abatement costs. Montero (2009) surveys market power in pollution permit markets, and Wirl (2009) studies demand-function equilibrium in deterministic imperfectly competitive permit markets. My paper contributes by embedding correlated information about costs and damages into an n-agent strategic environment and by providing a welfare decomposition that delivers transparent comparative statics in both market concentration and information precision.

A third strand compares welfare under taxes and cap and trade with uncertainty. Weitzman (1974) provides the ranking criterion when cost and damage shocks are uncorrelated, and Stavins (1996) extends the ranking to include correlated shocks. Newell & Pizer (2003), in comparing the tax and quota when the regulator uses an open-loop decision rule to control a stock pollutant, note that the same criterion applies to certainty-equivalent (equal-mean) policy pairs, even when the instruments are not set optimally. I extend this literature by providing the ranking criterion and its comparative statics whem there are a finite number of strategic agents. My discussion of the relation between the ranking criterion in the static (flow pollutant) versus the dynamic (stock pollutant) problems is based on Karp & Zhang (2005), Karp & Traeger (2024) and Karp (2025).

Taxes and cap-and-trade dominate applied policy debate because they are widely used, but they are typically second best. A fourth literature proposes more efficient instruments that enable a regulator to acquire dispersed information. Montero (2008) extends an earlier literature using mechanism design, constructing an information-revealing mechanism that supports the full-information outcome. Lemoine (2022) suggests a "carbon share" policy. Karp & Traeger (2025) propose a "smart cap" that aggregates dispersed information about abatement costs and is implementable as a competitive equilibrium. The smart cap is related to an extensive earlier literature that examines hybrid policies (Roberts & Spence 1976, Weitzman 1978, Pizer 2002, Hepburn 2006, Fell & Morgenstern 2010, Grüll & Taschini 2011, Fell et al. 2012). This paper shows that when there are a finite number of strategic agents, the simpler "smart tax" implements the solution to agents' team problem; that solution is the same as under a planner who knows agents' average cost shock when choosing policy

2 Model and the Welfare Decomposition

Agents might be firms within a sector or countries within a global agreement. Each agent obtains a benefit from its own emissions (reduced abatement costs) but incurs a cost from aggregate emissions. The economy is "fragmented" because it consists of multiple agents who receive diferent cost shocks. When agents face an aggregate quota, each announces a demand function for permits, resulting in a demand function equilibrium. If agents face an emissions tax, they choose their quantity of emissions, resulting in a different demand function equilibrium. A model with multiple agents captures the distinct externalities agents face under the two policies – a pecuniary externality under a quota and a pollution externality under a tax. After determining the equilibria in the two settings, I aggregate the fragmented economy to obtain a Representative Agent Economy that simplifies the analysis.

2.1 The Fragmented Economy

Agent i's abatement cost shock is $\theta_i \equiv \alpha + \eta_i$, with the market fundamental $\alpha \sim (0, \sigma_\alpha^2)$ and the firm-specific shocks $\eta_i \sim \left(0, \sigma_\eta^2\right)$, with $\mathbb{E}\left(\alpha\eta_i\right) = \mathbb{E}\left(\eta_i\eta_j\right) = 0$ for $i \neq j$. The shock to pollution-related damages is $\gamma \equiv \varphi\alpha + \varepsilon$ with $\varepsilon \sim (0, \sigma_\varepsilon^2)$ and $\mathbb{E}(\alpha\varepsilon) = \mathbb{E}(\eta_i, \varepsilon) = 0$. The covariance between i's cost shock and the pollution damage shocks is $\mathbb{E}\left(\alpha + \eta_i\right)\left(\varphi\alpha + \varepsilon\right) = \varphi\sigma_\alpha^2$. The market fundamental α affects both abatement costs in the polluting sector and the damages in the general economy; it generates correlation between the abatement cost and the damage shocks. Section 6 argues that positive correlation should be considered the leading case, not a special case.

Agents' information about α arises from a public signal (not explicitly modeled), so $1/\sigma_\alpha^2$ is the precision of that signal. The precision of the agent's private signal increases with $1/\sigma_\eta^2$, so the ratio $\tau \equiv \sigma_\eta^2/\sigma_\alpha^2$ is a measure of the relative precision of the two sources of information. A higher σ_α^2 increases the volatility of cost shocks; that change would appear in a time series of costs. A higher σ_η^2 (i.e., larger τ) increases the dispersion of idiosyncratic shocks; that change would appear in a cross section (over agents) of shocks. Changes in these two parameters sometimes have different comparative static effects. For most of the analysis I hold σ_α^2 constant, so an increase in τ represents both an increase in the relative precision of public information and an absolute decrease in the precision of private information.

It is common to assign (usually normal) distributions to the cost and damage shocks. With normal (or some other) distribution, we can speak of a Bayesian Nash equilibrium. However, in the quadratic setting, payoffs depend on only the first two moments of the random variables. For that reason, I describe the distributions using only these moments. This parsimony means that I can speak of a Nash equilibrium, but not a Bayesian Nash equilibrium. Agents use their signals to compute the Best Linear Unbiased Estimator (BLUE) of the underlying state of nature. Under normality, the BLUE is the Bayesian posterior.[2] The unbounded support of the normal distribution is another reason to avoid that choice.

Assumption 1 *Outcomes are interior: Under the tax, each agent emits a positive amount in every state of nature; in every state of nature the quota is binding and at the equilibrium quota price all agents emit a postive amount.*

I use this assumptions in evaluating agents' first order conditions and to determine the equilibrium quota price. With unbounded support of the random variables, there may be a high probability that Assumption 1 is violated, in which case the calculated equilibrium would be far from the actual equilibrium. Even with a finite support for the random variables, for some policies (possibly including the optimal ones) Assumption 1 fails. For example, a quota close to the unregulated average level of emissions would be slack in many states of nature; a very high tax, or the quota price associated with a very small quota, would cause an agent's emissions to fall to zero in many states. Thus, although I consider a range of policy levels, possibly including the optimal level, the analysis does not accommodate very tight or very lax policies. Policies must (at least approximately) satisfy Assumption 1.[3]

If Agent i's emissions are e_i and aggregate emissions are $\sum_j e_j$, i's payoff is[4]

$$\mathbb{E}\left[\underline{(b_0+\theta_i-p)\,e_i-\frac{b}{2}e_i^2}-\left((d_0+\varphi\alpha+\varepsilon)\left(\sum_j e_j\right)+\frac{d}{2}\left(\sum_j e_j\right)^2\right)\right]. \tag{1}$$

[2] The reader can convert "Nash" into "Bayesian Nash" by assigning normal distributions to random variables. It does not matter whether we interpret agent i's signal, $\alpha+\eta_i$ as i's actual cost parameter, or as a signal that improves i's estimate of that parameter.

[3] Karp (2025) explains how the violation of Assumption 1 changes the analysis.

[4] The following generalizations are straightforward: heterogeneous intercepts (b_{0i} instead of b_0), heterogeneous damage shocks (ε_i instead of ε), an unequal distribution of total damages across agents (d_{0i} instead of d_0), and the possibility that a non-polluting agents suffers some of the pollution-related damages (Footnote 7).

The emissions price, p is endogenous under the quota and it equals the exogenous tax under a tax. The expectation is conditioned on the agent's private signal, θ_i and it is taken over all random variables, $e_j, j \neq i, \alpha, \varepsilon$, and in the case of a quota, p. The underlined term in expression 1 is the agent's benefit from emissions minus the cost of permits. The remaining term is the agent's damages caused by aggregate pollution (the same for all agents).

The shocks $\theta_i = \alpha + \eta_i$ and $\varphi\alpha + \varepsilon$ affect the intercepts of marginal abatement cost and marginal damage. The function $b_0 + \theta_i - be_i$ is agent i's marginal benefit of emissions, equivalent to its marginal abatement cost, expressed as a function of emissions instead of abatement. For that reason, I refer to θ_i as the shock to marginal abatement costs.

2.1.1 Nash equilibria under a quota or a tax

Suppose that the regulator announces a quota Q and agents then submit a demand function for permits – a function of the endogenous permit price, p. The regulator sets the price to clear the market and then allocates each agent the number of permits on their demand function, at the equilibrium price. With a finite number, n, of agents, each has an incentive to decrease their demand for emissions permits to reduce the equilibrium price. Hereafter I assume $n \geq 3$ to avoid the need to consider special cases.

The symmetric linear demand function equilibrium, conditional on the aggregate quota, Q, is unique. With a binding quota, $\sum_j e_j = Q$, the agent takes damages as given, so I drop the damage component from the agent's payoff. With a quota price p, the agent's payoff is

$$P_i \equiv (b_0 + \theta_i)\, e_i - \frac{b}{2} e_i^2 - pe_i. \tag{2}$$

In the demand function equilibrium, agent j announces the demand function for permits $e_j = g + h\theta_j - fp$.[5] Equilibrium constants g, h, f are given by

Lemma 1 *In the unique symmetric linear demand function equilibrium,*

$$f = \frac{n-2}{b\,(n-1)} = h \quad \textit{and} \quad g = \frac{(n-2)}{b\,(n-1)} b_0. \tag{3}$$

To avoid additional notation (and to emphasize the relation between the equilibrium quota price and a tax), I denote the tax as p in the tax game. Agents treat the tax as given, but they understand that their action affects aggregate emissions and thus affects their own pollution-related damages. Agent i's payoff under the tax p is given by expression 1.

[5] The regulator needs to know the sum $g + h\theta_j$, not the separate components of the sum.

Lemma 2 *The coefficients of the unique linear-in-information Nash equilibrium in the tax game, $e_i = q + r\theta_j - sp$, are*

$$\text{(i) } r = \frac{1-\varphi}{b+dn}\text{; (ii) } s = \frac{1}{b+dn}\text{, and (iii) } q = \frac{b_0 - d_0}{b+dn}. \tag{4}$$

Aggregate emissions E, equal

$$E \equiv \sum e_j = n\left(q - sp\right) + r\sum \theta_j = n\left(\frac{b_0-d_0}{b+dn} - \frac{1}{b+dn}p\right) + \frac{1-\varphi}{b+dn}\sum \theta_j. \tag{5}$$

From the perspective of the regulator (who does not observe the cost shocks, θ_i), the expectation and the variance of aggregate emissions are

$$\mathbb{E}(E) = \frac{n}{b+dn}\left(b_0 - d_0 - p\right) \text{ and } var(E) = \left(\frac{1-\varphi}{b+dn}\right)^2 n\left(n\sigma_\alpha^2 + \sigma_\eta^2\right). \tag{6}$$

2.2 Evaluating welfare using the representative agent analog

The fragmented economy (FE) involves heterogeneous agents. The parameter n here potentially has two distinct types of effects that can confuse the analysis. First, an increase in n always means that each agent is a smaller fraction of the economy and therefore has less incentive to internalize the pecuniary externality in the case of quota or the pollution externality in the case of the tax. A second possible interpretation is that a change in n alters the size of the economy, possibly changing technology. For example, an increase in n might increase aggregate pollution simply because it increases the number of polluters.

With both interpretations, changing n alters strategic incentives, changing equilibrium payoffs. The difference between the two interpretations is that with the first, the change in n alters payoffs only via the change in agents' behavior. With the second interpretation, a change in n alters aggregate payoffs even if there is no change in individual behavior. I want a change in n to pick up the first but not the second type of effect: the change in payoffs arises only via a change in equilibrium behavior due to the change in strategic incentives, as the example in Footnote 1 illustrates. The RAE characterization achieves this goal.

I need the FE to model non-cooperative behavior and to obtain agents' equilibrium decision rules under a tax and a quota. Substituting those decision rules into agents' payoffs and aggregating over the n agents produces the aggregate payoff. The next lemma provides a simple way to evaluate aggregate payoffs, perform comparative statics, and to rank policies.

To nest the models with a tax or a quota, I write an agent's FE decision rule as $e_i = x - yp + z\theta_i$, where x, y, z are constants. By specializing the constants x, y, z using Lemma 1 or 2, I obtain the equilibrium decision rules under a quota or a tax. (In other contexts, the coefficients take other values.) Under a tax, p is a constant, and under a quota, p is the solution to a market clearing condition. Aggregate emissions equal $E \equiv \sum_i e_i = X - Yp + Z\theta$, with

$$X \equiv nx, Y \equiv ny, Z \equiv nz, \text{ and } \theta \equiv \frac{\sum_i \theta_i}{n}. \tag{7}$$

Z plays a crucial role, so I use Lemma 1 and 2 to write its value under a quota or a tax

$$Z^{quota} \equiv nf = \frac{n(n-2)}{b(n-1)}; \text{ and } Z^{tax} \equiv nr = \frac{n(1-\varphi)}{b+dn}. \tag{8}$$

The aggregate FE payoff is the expected sum of agents' payoffs in expression 1, exclusive of quota rents or tax payments; these are a pure transfer.

$$\text{FE: } \mathbb{E}\left[\sum_i^n \left[(b_0 + \theta_i)\, e_i - \frac{b}{2} e_i^2\right] - n\left((d_0 + \varphi\alpha + \varepsilon)\left(\sum_{j=1}^n e_j\right) + \frac{d}{2}\left(\sum_{j=1}^n e_j\right)^2\right)\right]. \tag{9}$$

In the RAE, the representative agent faces the damage shock $\Gamma \equiv \lambda\alpha + w$ and the abatement cost shock $\theta \equiv \alpha + \mu \sim (0, \sigma^2)$ (no subscript), with $w \sim (0, \sigma_w^2)$, $\mathbb{E}\alpha^2 = \sigma_\alpha^2$, $\mathbb{E}\alpha\mu = \mathbb{E}(\mu w) = 0$ and $\mathbb{E}\mu^2 = \sigma_\mu^2$; α is a market fundamental that affects both abatement costs and pollution-related damages; μ affects only abatement costs, and w affects only damages. The parameter φ in the FE is the coefficient of the regression of the damage shock on a single firm's cost shock; the RAE analog, λ, is the coefficient of the regression of the damage shock on the aggregate cost shock.[6] These parameters have the same sign as the correlation between the damage and the cost shocks, but they can be of any magnitude.

I define the RAE payoff when the agent uses the decision rule $E = X - Yp + Z\theta$ as

$$\text{RAE: } \mathbb{E}\left[(B_0 + \theta)\, E - \frac{B}{2} E^2 - \left((D_0 + \lambda\theta + w)\, E + \frac{D}{2} E^2\right)\right]. \tag{10}$$

Expressions 9 and 10 are both quadratic in emissions and linear in shocks.[7]

[6] I can write $\Gamma = \lambda\alpha + w = \lambda(\alpha + \frac{\sum_i \eta_i}{n}) + w - \frac{\sum_i \eta_i}{n} = \lambda\theta + w^*$, with $w^* = w - \lambda\frac{\sum_i \eta_o}{n} = w - \lambda\mu$, and $\theta = \frac{\sum_i \theta_i}{n}$. This observation justifies interpreting λ as the coefficient of the regression of the damage shock on the "aggregate cost shock" (defined as the average of agents' shocks).

[7] Perhaps the most important of the several generalizations mentioned in Footnote 4 recognizes that some of the pollution-related damages fall on non-polluting agents. If the polluting agents bear the share $\frac{1}{1+\Lambda}$ of pollution-related damages, the RAE damage parameters should be inflated by $1 + \Lambda$; for example, equation 12.(i) becomes $nd_0(1 + \Lambda) = D_0$.

Lemma 3 *If agents in the FE use the decision rule $e_i = x - yp + z\theta_i$, the aggregate payoff is*

$$\beta^j + \left[B_0(\mathbb{E}E) + \frac{B}{2}(\mathbb{E}E)^2 - \left((D_0(\mathbb{E}E) + \frac{D}{2}(\mathbb{E}E)^2 \right) \right] \ , j \in \{quota, tax\}, \tag{11}$$

with $\mathbb{E}(E) = X - Yp$ (using the definitions in equation 7) and the following restrictions:

$$(i) \ \ nd_0 = D_0; \ (ii) \ \ nd = D; \ (iii) \ b_0 = B_0; \ (iv) \ \tfrac{b}{n} = B; \ (v) \ \ n\varphi = \lambda$$

$$(vi) \ \beta^{quota} = Z\left(\tfrac{n-1}{n}\right)\left(1 - \tfrac{B}{2}Z\right)\tau\sigma_\alpha^2, \text{ and} \tag{12}$$

$$(vii) \ \beta^{tax} = \left((1+\tau)\left(1 - \tfrac{B}{2}Z\right) - \left(\lambda + \tfrac{D}{2}Z\left(1 + \tfrac{\tau}{n}\right)\right)\right) Z\sigma_\alpha^2.$$

In the case of the tax, in writing the payoff we can replace the stochastic E with its expectation, $\mathbb{E}E = X - Yp$. All of the components of the payoff that are related to the second moments are incorporated into the function β^{tax}. In the case of a quota, where $E = Q$ is non-stochastic, we simply replace $\mathbb{E}E$ with Q and incorporate the payoff components related to second moments into β^{quota}. These two functions depend on Z^{quota} and Z^{tax}, which differ when n is finite. Using equations 8 and 12 (iv) & (v):

$$Z^{quota} = \frac{n-2}{B(n-1)} \text{ and } Z^{tax} = \frac{1 - \frac{\lambda}{n}}{B + \frac{D}{n}}, \tag{13}$$

which both converge to B^{-1} as $n \to \infty$.

3 Comparative statics in the games

Some of the results below use the "Weitzman ratio", $W \equiv \frac{D}{B}$. In Weitzman"s (1974) setting, taxes welfare-dominate quotas if and only if $W < 1$. I first compare outcomes in the quota and tax games with the outcome in a familiar benchmark where agents facing a quota take the quota price as given, and where agents facing a tax take aggregate emissions as given. Next, I show how the parameters n, λ, τ and W affect welfare in the two games.

I use Lemmas 1 and 3 to write the equilibrium quota price in the demand function equilibrium ("Nash") in terms of the RAE parameters, the quota, and the cost shock:

$$\text{Nash quota price: } p = B_0 + \theta - \frac{B(n-1)}{n-2}Q. \tag{14}$$

I use Lemmas 2 and 3 to write equilibrium aggregate emissions in the tax game ("Nash") , in terms of the RAE parameters and the tax.

$$\text{Nash emissions rule: } E = \left(\frac{1}{B\left(1+\frac{W}{n}\right)}\right)\left[B_0 - \frac{D_0}{n} - p + \left(1-\frac{\lambda}{n}\right)\theta\right]. \tag{15}$$

In the familiar benchmark, agents are price takers under the quota and they ignore the effect of their emissions on damages. In this benchmark, the emissions price under a quota and aggregate emissions under a tax, are:

$$p^{\text{quota price}} = B_0 + \theta - BQ \ ; \ E^{\text{emission w/ tax}} = \frac{B_0}{B} - \frac{1}{B}p + \frac{1}{B}\theta. \tag{16}$$

Equations 14 – 16 imply:

Remark 1 *(i) The Nash quota price and emissions rule approach their benchmark analogs as $n \to \infty$. (ii) For a given quota, Q, the Nash equilibrium price is lower than the benchmark price. The price volatility falls with n, so the benchmark price is less volatile than the Nash quota price. (iii) In the benchmark, emissions under the tax do not depend on λ or τ. With finite n (i.e., in the game) and a tax, the volatility of emissions: (a) falls in W, (b) increases with λ and (c) increases with n if and only if $W > \lambda$; when this inequality holds, emissions are more volatile in the benchmark than in the game with $n < \infty$.*

Greater W increases the convexity of pollution damages relative to abatement costs, making volatile emissions more costly (by Jensen's inequality). Because agents in the game internalize some of these costs, a larger W reduces equilibrium volatility in the game. Higher correlation of abatement cost and pollution damage shocks (higher λ) increases the importance of adjusting emissions in response to a cost shock; this greater incentive to adjust increases emissions volatility in the game. The explanation for the effect (on emissions volatility) of a larger n turns on the fact that W and λ have opposite equilibrium effects on volatility. A larger n moves outcomes in the game toward the familiar benchmark, where the representative agent ignores both W and λ in responding to the tax (because this agent ignores the consequence of their emissions on damages). Therefore, the equilibrium effect of both W and λ diminish with the increase in n. If $W > \lambda$, the convexity of damages is "more important" than the cost correlation; therefore, as n increases and outcomes in the game move toward those in the benchmark, the net effect is to increase emissions volatility.

In general, when agents face a tax they adjust emissions in response to a cost shock. This adjustment lowers expected abatement costs and increases expected damages (again, due to Jensen's inequality). This tension is the basis for Weitzman's (1974) ranking. Emissions under a tax, in the benchmark, are always more volatile than is socially optimal. Therefore, if a more concentrated industry (smaller n) lowers volatility (i.e., for $W > \lambda$), the decrease in industry concentration tends to raise aggregate welfare. For a given tax, a lower n also lowers expected emissions, another source of welfare gains when emissions are excessive.

The next two propositions formally examine the welfare effect of parameter changes.

Proposition 1 *Equilibrium welfare in the quota game with quota Q is*

$$V^{quota}(Q) \equiv \frac{(n-2)}{2B(n-1)}\tau\sigma_\alpha^2 + \left[B_0 Q - \frac{B}{2}Q^2 - \left(D_0 Q + \frac{D}{2}Q^2\right)\right], \tag{17}$$

an increasing concave function of n, and an increasing function of both τ and σ_α^2 (i.e. of the variances of both the market fundamental and the idiosyncratic shock). The welfare gain from greater competition (larger n) increases as these variances increase.

Welfare under a quota is higher when competition is greater or the market fundamental or the idiosyncratic shocks are more variable.. In the quota game, agents lower the equilibrium price by announcing a demand function lower than their true marginal valuation. This self-interested behavior benefits other agents in the polluting sector, shifting quota rents from the general treasury to polluting firms. However, because agents have different cost shocks, their price shading causes them to have different levels of marginal abatement cost in equilibrium. This inefficiency is the source of the welfare loss in the game (Malueg & Yates 2009). Greater competition is especially important when private signals are imprecise (τ is large), because then agents' demand shading leads to large inefficiencies.

I now turn to the tax. Inspection of equation 15 shows that both a larger n and a smaller W increase expected emissions under a tax. When emissions exceed the socially optimal level, this change lowers the term in square brackets in equation 11, tending to lower welfare. Because this effect is obvious, the next proposition considers only β^{tax}, the component of welfare that depends on second moments. This function is a ratio of cubics in n. Given numerical values it is easy to evaluate the derivatives of this ratio, but in general they provide little insight. Therefore, the proposition reports the sign of the derivative for "for large n"; that is, I consider the leading term (in n) of the derivatives.[8]

Proposition 2 *Under the tax p, β^{tax} is*

$$\beta^{tax} = \frac{f(n;\lambda,\tau,W)}{2Bg(n;W)}\tau\sigma_\alpha^2,$$

where f and g are cubics in n; the proof reports the formulae. (i) For large but finite n, β^{tax} increases with τ if and only if

$$2(\tau-\lambda)+1 > W. \tag{18}$$

[8]The functions $f(n)$ and $g(n)$ in this proposition are unrelated to the constants f, g in equation 3.

(ii) For large n, β^{tax} decreases with λ. (iii) For large but finite n, welfare decreases with n if and only if

$$\tau < 4\lambda + 2W + 2\frac{\lambda^2}{W}. \tag{19}$$

Because the Nash emissions rule, equation 15, does not depend on τ, inequality 18 is sufficient for an increase in τ to raise equiibrium welfare. With fixed σ_α^2, an increase in τ is equivalent to an increase in $\sigma_\eta^2 = \tau\sigma_\alpha^2$. Thus, inequality 18 is sufficient for less precise private or public signals to raise welfare in the tax game when n is large. Less informative signals are more likely to raise welfare, the lower is the correlation between abatement cost and damage shocks, and the lower is the Weitzman ratio (the slope of marginal damages to marginal costs.)

Remark 1(iiib) notes that a larger λ increases the volatility of emissions, lowering the expected payoff (by Jensen's inequality). Because higher λ also lowers β^{tax} (at least for large n), an increase in the covariance of the cost and damage shocks lowers welfare.

An increase in n raises expected emissions, lowering welfare when emissions are excessive. If inequality 19 holds (so larger n lowers β^{tax}) and in addition $W > \lambda$ (so that larger n increases the volatility of emissions, then all components of the welfare change operate in the same direction: larger n lowers welfare in the tax game. Reversal of this conclusion would require a special configuration of parameters.

4 Policy ranking

Certainty equivalent taxes and quotas give the same expected level of aggregate emissions. Setting emissions equal to the quota, Q, and using equation 15 to solve for p, gives the certainty equivalent tax for the quota Q: $p = B_0 - \frac{D_0}{n} - \left(B + \frac{D}{n}\right) Q$. With a certainty equivalent tax and quota pair, the terms in square brackets in equation 11 are equal under the two policies ($\mathbb{E}E = Q$), so the payoff difference is $\beta^{tax} - \beta^{quota}$. Evaluating this difference produces:[9]

Proposition 3 *(i) For any certainty equivalent tax and quota pair satisfying Assumption 1, the payoff under taxes minus the payoff under quotas is (asymptotically, as $n \to \infty$)*

$$h(n) = \frac{(1 - W - 2\lambda)}{2B}\sigma_\alpha^2 + O\left(\frac{1}{n}\right). \tag{20}$$

For large finite n, taxes welfare-dominate quotas if and only if $1 - 2\lambda > W$.

(ii) A first-order finite-n approximation of the payoff difference is

$$h(n) = \frac{(1 - W - 2\lambda) n + ((3 - \tau) W + (2\lambda^2 + 2\lambda + \tau - 1))}{2B(n + (2W - 1))}\sigma_\alpha^2 + o(1). \tag{21}$$

[9]The function $h(n)$ used in this proposition is unrelated to the constant h in equation 3.

Corollary 1 *Using the finite n approximation of the welfare difference:*

(i) An increase in τ favors taxes if and only if $1 > W$.

(ii) An increase in n favors taxes if and only if $2(\lambda + W)^2 + \tau(1 - W) < 0$. For $\lambda = 0$ this inequality is satisfied if and only if $\frac{1}{4}\tau - \frac{1}{4}\sqrt{\tau(\tau - 8)} < W < \frac{1}{4}\tau + \frac{1}{4}\sqrt{\tau(\tau - 8)}$. This interval has positive measure if and only if $\tau > 8$.

(iii) An increase in λ favors taxes if and only if $\lambda > 0.5(n - 1)$.

Stavins (1996) obtains the ranking criterion beginning with the representative agent model (implicitly, with $n = \infty$). Part (i) of the proposition confirms this conclusion, showing (not surprisingly) that it holds if we begin with a multiple agent model and take the limit as the number of agents approaches infinity (while maintaining a constant "size" of the aggregate economy).[10] Equation 20 shows that if $\lambda > 0.5$, $h(n) < 0$. In this case, quotas dominate taxes regardless of the value of W.

Corollary 1 uses the finite n approximation in Proposition 3 ii to obtain comparative statics. The familiar result states that for $\lambda = 0$ taxes dominate quotas if and only if $1 > W$. The familiar explanation is that the disadvantage of the tax (relative to the quota) is that it makes emissions volatile, increasing expected damages; greater convexity of damages (higher D) increases this disadvantage. The advantage of the tax is that it enables firms to reduce expected abatement costs, by having emissions respond to the cost shock; a steeper marginal abatement cost (larger B) increases this advantage. A larger D increases the disadvantage of taxes and a larger B increases the advantage of taxes, so a smaller $W = \frac{D}{B}$ favors taxes.[11]

Part (i) of the corollary states that $1 > W$ implies that a less precise private signal (larger τ) favors taxes. The explanation echoes that of the previous paragraph. When marginal abatement costs are relatively steep, and marginal damages relatively flat, a higher variance in the idiosyncratic shock makes it more valuable to obtain the advantage provided by the tax (reduction in expected abatement costs).

This result also illustrates the possibility that increased precision of public versus private information can have opposite effects; the papers cited in the opening paragraph of Section 1.1 offer many such examples. Suppose that $1 > W$ but $f(n) < 0$ (i.e., quotas welfare-dominate taxes because of the correlation between abatement costs and damages). The first inequality implies that a less precise private signal favors taxes (i.e., it erodes the advantage of quotas), and the second inequality implies that a less precise public signal increases the advantage of quotas. A less precise public signal increases the volatility of emissions, raising

[10] Appendix A.1 shows how to obtain this ranking for the limiting case, $n = \infty$, in a few lines.

[11] When I state that a parameter change "favors taxes" I mean that it increases the advantage of taxes, or erodes the advantage of quotas. The statement does not mean that the parameter change causes taxes to welfare-dominate quotas.

expected damages and favoring quotas. A less precise private signal increases the dispersion of costs, increasing firms' benefit from being able to respond to their private information.

Part (ii) of the corolllary shows that if $W < 1$ (so that taxes dominate quotas in Weitzman's setting), larger n always favors quotas. For small λ and sufficiently noisy private signals, larger n favors taxes when W takes intermediate values. Thus, although it is possible to find constellations of parameters for which greater competition favors taxes, at least for $\lambda \geq 0$ such a constellation seems unlikely: greater competition tends to favor quotas.

Because $n \geq 3$, the inequality in part (iii) requires $\lambda > 1$. When this inequality holds, quotas welfare-dominate taxes; a larger λ (at some point) merely reduces the dominance of quotas. Especially for moderate or large n, the inequality $\lambda > 0.5\,(n-1)$ is "unlikely" to hold. Thus, in cases of interest, it is reasonable to conclude that a larger λ favors quotas.

5 A different perspective, and first best policies

This section uses results from Karp & Traeger (2025) both to provide a different perspective on the ranking criterion and also to suggest first best policies. I begin with the limit of the FE as $n \to \infty$, and then consider the finite agent FE.

In Figure 1, the solid curve labeled MB is expected marginal benefit of emissions, and the parallel dotted curve represents marginal benefits under a negative cost shock, θ. The line labeled MD is expected marginal damages, and the parallel dashed line with vertical displacement $\lambda\theta$ is expected marginal damages, conditional on the realization of θ. The solid line connecting the intersections of these pairs of lines, labeled S, is the "smart tax": $S(E) \equiv g + mE$ with $g = \frac{D_0 - \lambda B_0}{1-\lambda}$ and $m = B\frac{W+\lambda}{1-\lambda}$.[12]

This tax is "smart" because if the price-taking representative agent were to face it, then equilibrium emissions would occur at the optimal level, conditional on the realization of θ.[13] The slope of marginal benefits, B, is greater than the slope of the smart tax if and only if $B > B\frac{W+\lambda}{1-\lambda}$. Rearranging this inequality produces the ranking criterion in Proposition 3 (i). Weitzman's (1974) famous result is that, with uncorrelated cost and damage shocks, the optimal tax welfare-dominates the optimal quota if and only if the marginal benefit of emissions is steeper than marginal damages. The generalized result is[14]

[12]This figure is based on Figure 1 in Karp & Traeger (2025), where it is used both to illustrate the smart tax and to show graphically (by comparing welfare triangles) that positive correlation of shocks favors quotas.

[13]The tax does not achieve the full information outcome; it is not conditioned on the realization of the damage shock.

[14]Karp & Traeger (2025) show that, for a stock pollutant, the optimal tax welfare-dominates the optimal quota if and only if the slope of marginal benefits is greater than the slope of the dynamic analog of the smart tax. The proof in the dynamic setting is more complicated than in the static setting here. Remark 2 compares welfare under the optimal tax and quota, whereas Proposition 3 compares welfare under certainty

Figure 1: Solid lines show expected marginal benefits (MB) and marginal damages (MD). Dashed lines show MB under the shock θ and expected MD, conditional on this shock. The locus of intersections of these curves, as θ varies, is S, the smart tax.

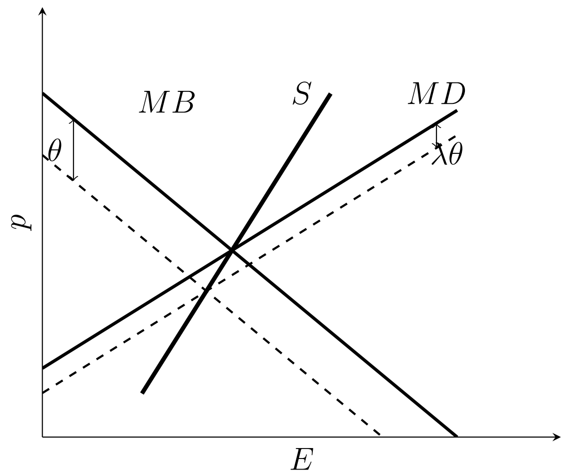

Remark 2 *With (possibly) correlated cost and demand shocks, the optimal tax welfare-dominates its certainty quota (given that these satisfy Assumption 1) if and only if the slope of marginal benefits of emissions is steeper than the slope of the smart tax.*

The smart tax is pedagogically useful (because of Remark 2), but it is not a practical policy when n is large. Neither firms nor the regulator have real-time information of aggregate emissions, so agents do not know the equilibrium value of the smart tax when they make their emissions decision. However, the smart tax is a stepping stone in designing the policy-relevant "smart cap". Under this policy, the regulator distributes $C > 0$ tradable certificates, and announces a "conversion function" $q(p^c)$, where p^c is the endogenous price of a certificate. One certificate entitles the firm to emit $q(p^c)$ units, so the endogenous price of a unit of emissions is $p = \frac{p^c}{q(p^c)}$, and total emissions equal $E = Cq(p^c)$. Under the smart tax, $p = S(E)$. In order that the smart cap return the same outcome as the (likely impractical) smart tax, use $p = \frac{p^c}{q(p^c)}$ and $E = Cq(p^c)$ to rewrite $p = S(E)$ as

$$\frac{p^c}{q(p^c)} = S(Cq(p^c)). \tag{22}$$

Equation 22 translates the arbitrary constant C and the known smart tax $S(E)$ into the conversion function $q(p^c)$. Normalizing by setting $C = 1$, and (to avoid a taxonomy) assuming

equivalent policies. The smart tax is defined only for optimal policies.

that $\lambda < 1$, so $S'(E) > 0$, the unique positive solution to equation 22 is[15]

$$q(p^c) = -\frac{1}{2m}\left(g - \sqrt{4p^c m + g^2}\right). \tag{23}$$

Under the smart cap, agents need only know the certificate price to choose their emissions. Here, the informational requirements for equilibrium are (essentially) the same as in a standard market, where agents have to know the permit price to choose their action.

When $n < \infty$, and especially when n is small, the informational requirements of the smart tax are exactly the same as those of the noncooperative Nash equilibrium under either a constant tax or a quota: in all three cases, agents have to form beliefs about the other agents' actions. Here, the smart tax is a practical policy.

However, when n is finite, the interpretation of "first best" is potentially ambiguous, because it depends on the amount of information that the planner has (as discussed below). I therefore use the (unambiguous) "team problem" (Angeletos & Pavan 2007). In this setting, each agent chooses its emissions to maximize global welfare, rather than its individual welfare. As in the non-cooperative Nash equilibrium, agents base their beliefs about the market fundamental, α, and on other agents' actions, on their private signal, θ_i.

Proposition 4 *(i) In the FE with n agents, the smart tax*

$$S = \left(\frac{n-1}{n}\right)\left(\frac{D_0 - \lambda B_0}{1-\lambda}\right) + B\left(\frac{(W+\lambda)(n-1)}{(1-\lambda)(n+1)}\right)E$$

supports the the equilibrium in the team problem.

(ii) The team problem equilibrium is identical to the solution of the planner who knows the realization of θ (the average of agents' signals) when setting policy.

(iii) The comparative statics of the smart tax are

$$(i)\;\; \frac{dS}{d\lambda} > 0, \;\; (ii)\;\; sign\left(\frac{dS}{dW}\right) = sign\,(1-\lambda) \;\; and$$
$$(iii)\;\; \frac{dS}{dn} = \frac{-1}{n^2(1-\lambda)}\left((\lambda B_0 - D_0) - 2B\frac{n^2}{(n+1)^2}(\lambda + W)\,E\right)(> 0 \text{ for small } \lambda). \tag{24}$$

Proposition 4 (ii) shows the relevance of the team problem: it's equilibrium is identical to that of the planner who can condition policy on the realization of the average signal, θ. That is, the smart tax enables a planner who knows only the distribution of the signals to

[15] The slope of the smart tax is positive if and only if $\lambda < 1$, In this case, a lower cost shock corresponds to a lower level of emissions and a lower unit tax. With $\lambda > 1$, in contrast, the slope of the smart tax is negative. In this case, a lower cost shock (i.e., cheaper abatement) implies that expected marginal damages have fallen by so much that it is optimal to increase emissions.

obtain the same outcome as the planner who knows the average signal.

The potential ambiguity mentioned above is that knowing the average signal is not equivalent to knowing all of the signals. For example, suppose that the planner knows the average signal and agents know their own private signal. The planner could announce the average signal. In the resulting "enhanced" team problem, agents condition their beliefs about other agents' actions on the average signal, and they continue to use their private signal as information about their own costs. It is easy to compare the original and the enhanced team problem equilibria.[16] However, the equilibrium to the enhanced team problem cannot be supported by a *linear* smart tax. The linear equilibrium in the enhanced team problem has three parameters (a constant, and the coefficients of the own signal and the average signal); the linear smart tax, with only two free parameters, cannot support the three-parameter equilibrium.

The results in equation 24 are intuitive. More highly correlated shocks (larger λ) make it more important to offset a likely higher cost of damages resulting from high emissions; the policy response is a higher tax. A higher relative damage slope (larger W) also makes it important to reduce emissions (when $\lambda < 1$), increasing the tax. As n increases, agents have less incentive to internalize the damage externality. Therefore (at least for small λ), an increase in n raises the smart tax. For example, with $-W < \lambda < \frac{D_0}{B_0}$ (so that $\lambda < 1$) a larger n increases both the vertical intercept and the slope of $S(E)$.

6 Insights for stock pollutants

Even when the underlying economic problem is dynamic, a static model often provides useful intuition. A dynamic model can modify and in some cases correct this intuition. In the dynamic setting, positively serially correlated abatement cost shocks cause the social cost of carbon to be perfectly positively correlated with marginal abatement costs.[17] (The perfect positive correlation means that $\lambda > 0$; it has no implications concerning the magnitude of λ.) This correlation favors quotas under a stock pollutant, for (essentially) the same reason as positive correlation favors quotas in the static model. However, with a stock pollution, the positive correlation is mechanical, depending only on the assumption of positively serial

[16] For example, given any pair of signals, θ_i, θ_j, the difference between agent i and j's emissions in the enhanced team problem is greater than in the original team problem if and only if $\frac{(n-1)W}{n+W} > -\lambda$. Thus, positive correlation between the damage shock and the abatement cost shocks is a sufficient condition for the additional information in the enhanced team problem to lead to greater dispersion of agents' emissions.

[17] For example, a negative abatement cost shock in the current period makes current abatement cheaper. With positively serially correlated cost shocks, the current negative shock makes it more likely that future cost shocks will also be negative, and therefore future abatement costs will be lower. Especially when policy is set optimally, but *even in the absence of any policy intervention*, the lower future abatement costs will lower future emissions, thus lowering the current social cost of carbon.

correlation of cost shocks. Moreover, the effect can be large. Therefore, the consensus belief that the social cost of carbon has a very small slope does not imply that taxes are likely to welfare-dominate quotas: the "smart tax" can be much steeper than the social cost of carbon (as Figure 1 illustrates in a simpler setting) (Karp & Traeger 2024).

Correlation between abatement costs and damage is important in the analysis here, so it is reasonable to ask how significant it is in practice. Stavins (1996) considers this question in the context of flow pollutants. He provides several examples of (for the most part) positive correlation. For example, it is more costly to prevent the formation of ozone on sunny days, and on these days people are more likely to be outside, where they are more vulnerable to the pollution. However, as already noted, the Weitzman-type analysis has been so important in environmental economics because people have used the intuition it provides to think about stock pollution problems, principally, climate change. If we are interested in the static problem primarily as a guide for thinking about climate change, a model with positively correlated cost and damage shocks is a natural starting place: the dynamic problem is likely to exhibit this correlation, so a static simplification should also have it, even if the source of the correlation is different. Positive correlation is the leading case, not a special case.

For a given linear approximation of marginal damages and marginal abatement costs, the welfare-ranking of taxes and quotas in the static (flow pollutant) setting does not depend on the policy level. A change in policy stringency shifts the relevant welfare triangles up and down, but without changing their magnitudes and thus without changing the policy ranking for a flow pollutant.[18] In a dynamic setting (i.e., for a stock pollutant) the meaning of "policy stringency" is possibly ambiguous. Does a "laxer policy" mean "a laxer policy today" or "a laxer policy today and also in the future"? I take it to mean the latter: climate activists are concerned primarily about the (likely) laxity of policies over the next several decades, not principally because today's policies are weak.

With this interpretation, a change in policy stringency alters the ranking criterion, even holding fixed the linear approximation. A laxer policy environment (both today and in the future) increases the slope of the social cost of carbon. This increase favors quotas, for (essentially) the same reason that a steeper marginal damage function favors quotas under a flow pollutant. Thus, the invariance of the policy ranking, to the stringency of the tax and quota, that obtains under a flow pollutant does not hold with a stock pollutant (Karp 2025).

In a dynamic setting, the evaluation of a current policy and the comparison of two policies often depend on beliefs about future behavior. Those beliefs depend on whether we

[18]The policy level under consideration likely does determine the point on general cost and damage functions at which the linear approximation is taken. Thus, indirectly, the policy level affects the ranking via its effect on slopes of the linear approximation.

assume policies are open loop or feedback. With open loop policies, we act today as if future policies are chosen today. With feedback policies, we act today as if future policies will be conditioned on information that becomes available only in the future. The (more rational) feedback policies produce higher welfare in a non-strategic setting. In a static setting the distinction between open loop and feedback is meaningless. In moving to a dynamic setting, it is important to be clear about the distinction, and to consider how a choice between equilibrium concepts might influence policy ranking (Karp & Zhang 2005).

7 Conclusion

This paper studies instrument choice under uncertainty when agents possess market power, and a market fundamental affects both abatement costs and damages from pollution. An agent's private signal is therefore correlated with the damage parameter and also with other agents' signals. Starting from a fragmented n-agent economy, I derive welfare expressions that separate the role of expected emissions from the contribution of second moments. This structure allows transparent comparative statics with respect to market concentration, the precision and correlation of private information, and the Weitzman slope ratio (the ratio of marginal damages to marginal abatement costs).

The comparative statics show that the same change in market structure or information can have qualitatively different welfare effects under the two instruments. Under a quota, increased competition raises welfare because it lowers agents' incentive to shade bids. That shading creates inefficiency by raising dispersion of agents' marginal abatement cost levels. Greater competition is especially valuable when agents' private signals are imprecise. Under a tax, greater competition reduces incentives to internalize damages, resulting in higher emissions and damages, tending to lower welfare. Greater competition has an ambiguous effect on the volatility of emissions. Volatility falls with increased competition if and only if the Weitzman ratio is small, i.e., in circumstances where taxes tend to welfare-dominate quotas. A less informative private signal is more likely to raise welfare, the lower are Weitzman ratio and the correlation between abatement cost and damage shocks. These differences rationalize why empirical or institutional features that shift concentration or information quality can change not only policy performance but also the direction of the policy ranking.

Confirming earlier results, I show that as the market becomes perfectly competitive ($n \to \infty$) taxes welfare-dominate quotas if and only if $1 - 2\lambda > W$, where λ is the regression coefficient of the damage shock on the average abatement cost shocks. Positive correlation between cost and damage shocks therefore systematically tilts the ranking towards quantity instruments. The new result is that for finite n, lower precision of private signals favors taxes if and only if $W < 1$ (so that taxes would welfare-dominate quotas, absent cost correlation). The analysis also provides another example of the possibility that the decreased precision of public versus private information produce different comparative statics. The explanation is that (under a tax) a less precise public signal creates greater volatility of (the time series of) emissions, whereas the primary affect of a less precise private signal is to increase the cross section of dispersion (over agents).

With a finite (and moderate) number of strategic agents, a "smart" tax requires agents to have the same amount of information, and ability to use it, as assumed in constructing the equilibria (in the games) under the standard tax and cap. A smart tax, i.e., a schedule mapping aggregate emissions into the tax, implements the soluton to the "team problem", where agents do not share information, but each seeks to maximize collective welfare. The team equilibrium produces the same aggregate expected emissions as under a planner who knows the average of agents' signals. A higher correlation of cost and damage shocks and (for $\lambda < 1$) a higher Weitzman ratio (e.g., a steeper marginal damage) raise the tax schedule.

A Appendix

Proof. (Lemma 1) Denote $\tilde{\theta}_i = \sum\limits_{j \neq i} \theta_j$, the sum of the $j \neq i$ firms' abatement cost shocks. In the demand function equilibrium, the price satisfies

$$Q = (n-1)\, g + h\tilde{\theta}_i - (n-1)\, fp + e_i,$$

so the equilibrium price is

$$p = \frac{e_i + h\tilde{\theta}_i + (n-1)\, g - Q}{(n-1)\, f}. \tag{25}$$

Agent i expects that $\frac{dp}{de_i} = \frac{1}{(n-1)f}$. The firm views the equilibrium price as stochastic, because it depends on other firms' cost shocks, $\tilde{\theta}_i$. However, because the demand functions are (assumed to be) linear, the firm treats the change in price due to a change in its quantity as the constant, $\frac{dp}{de_i} = \frac{1}{(n-1)f}$. The expected change in i's profits due to a change in its demand, conditional on its observed realization of θ_i, is

$$\mathbb{E}\left(\frac{dP_i}{de_i}|\theta_i\right) = \mathbb{E}\left(b_0 + \theta_i - \left(p + \frac{1}{(n-1)\, f} e_i\right) - be_i\right) = \mathbb{E}\left(b_0 + \theta_i - p - \left(b + \frac{1}{(n-1)\, f}\right) e_i\right).$$

At the time the firm announces its supply function, the price is unknown. Given the firm's announced supply function, its value of e_i is also stochastic. Substituing the firm's supply function into $\mathbb{E}\frac{dP_i}{de_i}$ produces

$$\begin{aligned}
& \mathbb{E}\left(b_0+\theta_i-p-\left(b+\frac{1}{(n-1)\,f}\right)(g+h\theta_i-fp)\right) \\
= \quad & \mathbb{E}\left(\left(f\left(b+\frac{1}{f\,(n-1)}\right)-1\right)p+\left(1-h\left(b+\frac{1}{f\,(n-1)}\right)\right)\theta_i+\left(b_0-g\left(b+\frac{1}{f\,(n-1)}\right)\right)\right).
\end{aligned}$$

Conditional on the behavior of the other firms, i's supply function maximizes its profits for all realizations of p, θ_i (and thus maximizes the firm's expected profits, conditional on θ_i), if and only if the above expression is identically zero. This condition requires

$$\begin{aligned}
\left(f\left(b+\frac{1}{f\,(n-1)}\right)-1\right) &= 0 \\
\left(1-h\left(b+\frac{1}{f\,(n-1)}\right)\right) &= 0 \\
\left(b_0-g\left(b+\frac{1}{f\,(n-1)}\right)\right) &= 0.
\end{aligned}$$

Solving these equations recursively yields

$$f=\frac{n-2}{b\,(n-1)}=h \quad \text{and} \quad g=\frac{(n-2)}{b\,(n-1)}b_0.$$

■

Proof. (Lemma 2) Denote $\tilde{e}_i \equiv \sum_{j\neq i} e_j$ as cumulative emissions of the other $n-1$ agents. The agent's first order condition is[19]

$$\mathbb{E}\left(b_0+\theta_i-p-be_i-(d_0+\varphi\alpha+v)-d\left(\tilde{e}_i+e_i\right)\right)=0.$$

Its equilibrium decision rule is (using $\mathbb{E}v=0$)

$$e_i=\frac{1}{b+d}\left(b_0-d_0+\theta_i-p-\mathbb{E}\left(\varphi\alpha+d\tilde{e}_i\right)\right). \tag{26}$$

(In choosing current emissions, the agent knows its realization of θ_i but still must form expectations over the damage shock and other agents' emissions levels.) I substitute the linear decision rules $e_j=q+r\theta_j-sp$ into this first order condition and use $\mathbb{E}\alpha=\mathbb{E}\theta_j=\theta_i$

[19] When there is no danger of ambiguity, I simplify notation by not specifying what variables the expectations operate on. For example, in this equation the random variables, from i's perspective, are $\alpha, \nu, \tilde{e}_i$, and the firm's information is θ_i

(because i knows θ_i) to write equation 26 as

$$e_i = q + r\theta_i - sp = \frac{1}{b+d}\left(b_0 - d_0 + \theta_i - p - \varphi\theta_i - (n-1)\,d\,(q + r\theta_i - sp)\right). \tag{27}$$

Equating coefficients implies

$$\begin{aligned}
r &= \left(-\frac{1}{b+d}\left(\varphi + dr\,(n-1) - 1\right)\right) \Longrightarrow r = \frac{1-\varphi}{b+dn} \\
s &= \frac{ds\,(n-1)-1}{b+d} \Longrightarrow s = \frac{1}{b+dn} \\
q &= -\left(\frac{1}{b+d}\left(d_0 - b_0 + dq\,(n-1)\right)\right) \Longrightarrow q = \frac{b_0 - d_0}{b+dn},
\end{aligned}$$

establishing equation 4. Equation 5 follows from replacing e_i with its equilibrium decision rule and summing over i. This equation leads directly to the mean and variance of aggregate emissions (from the regulator's perspective). ■

Proof. (Lemma 3)

When the policy is a quota, Q, the quota price p is the solution to $X - Yp + Z\theta = Q$. Here, $E = Q$ is nonstochastic as are aggregate damages, equal to $n\left(d_0 Q + \frac{d}{2}Q^2\right)$. However, here p is stochastic. The market clearing condition that defines p requires $Q - X - Z\theta = -Yp$, or $-yp = \frac{Q}{n} - x - z\theta$. (Recall that $\theta = \frac{1}{n}\sum_i \theta_i$.) Substituting this relation into i's decision rule implies $e_i = \frac{Q}{n} + z\,(\theta_i - \theta)$. Agent i's expected benefit of emissions is

$$\mathbb{E}\left[(b_0 + \theta_i)\left(\tfrac{Q}{n} + z\,(\theta_i - \theta)\right) - \tfrac{b}{2}\left(\tfrac{Q}{n} + z\,(\theta_i - \theta)\right)^2\right] =$$

$$b_0\tfrac{Q}{n} - \tfrac{b}{2}\left(\tfrac{Q}{n}\right)^2 + \mathbb{E}\left[\left[\theta_i\left(z\,(\theta_i - \theta)\right) - \tfrac{b}{2}z^2\left((\theta_i - \theta)\right)^2\right]\right]$$

$$\tfrac{b_0}{n}Q - \tfrac{b}{2n^2}Q^2 + z\left(\tfrac{n-1}{n}\right)\sigma_\eta^2 - \tfrac{b}{2}z^2\left(\tfrac{n-1}{n}\right)\sigma_\eta^2 =$$

$$\tfrac{b_0}{n}Q - \tfrac{b}{2n^2}Q^2 + \tfrac{Z}{n}\left(\tfrac{n-1}{n}\right)\sigma_\eta^2\left(1 - \tfrac{b}{2n}Z\right).$$

Aggregate expected benefits equals n times this quantity,

$$b_0 Q - \frac{b}{2n}Q^2 + Z\left(\frac{n-1}{n}\right)\sigma_\eta^2\left(1 - \frac{b}{2n}Z\right).$$

Thus, aggregate expected benefits minus damages in the FE under the quota Q equals

$$b_0Q - \frac{b}{2n}Q^2 - \left(n\left(d_0Q + \frac{d}{2}Q^2\right)\right) + Z\left(\frac{n-1}{n}\right)\sigma_\eta^2\left(1 - \frac{b}{2n}Z\right).$$

Using equations 12 (i) – (iv), we can write this expression as

$$B_0Q - \frac{B}{2}Q^2 - \left(\left(D_0Q + \frac{D}{2}Q^2\right)\right) + Z\left(\frac{n-1}{n}\right)\sigma_\eta^2\left(1 - \frac{B}{2}Z\right).$$

Using $\sigma_\eta^2 = \tau\sigma_\alpha^2$ we can write the last term as

$$Z\left(\frac{n-1}{n}\right)\sigma_\eta^2\left(1 - \frac{B}{2}Z\right) = Z\left(\frac{n-1}{n}\right)\left(1 - \frac{B}{2}Z\right)\tau\sigma_\alpha^2 = \beta^{quota}$$

The establishes the lemma for the case where the policy is a quota.

Now consider a tax, where p is fixed. Using $e_i = x - yp + z\theta_i$, i's benefit from emissions (excluding the damages) in the FE is

$$\mathbb{E}\left[(b_0 + \theta_i)(x - yp + z\theta_i) - \tfrac{b}{2}(x - yp + z\theta_i)^2\right] =$$

$$b_0(x - yp) + \mathbb{E}z\theta_i^2 - \tfrac{b}{2}(x - yp)^2 - \tfrac{b}{2}\mathbb{E}(z\theta_i)^2 =$$

$$b_0(x - yp) + \mathbb{E}z\theta_i^2 - \tfrac{b}{2}(x - yp)^2 - \tfrac{b}{2}\mathbb{E}(z\theta_i)^2 =$$

$$b_0(x - yp) + \mathbb{E}z\theta_i^2 - \tfrac{b}{2}(x - yp)^2 - \tfrac{b}{2}\mathbb{E}(z\theta_i)^2 =$$

$$b_0(x - yp) + z\left(\sigma_\alpha^2 + \sigma_\eta^2\right) - \tfrac{b}{2}(x - yp)^2 - \tfrac{b}{2}z^2\left(\sigma_\alpha^2 + \sigma_\eta^2\right).$$

The aggregate benefits from agents' emission in the FE is n times this quantity:

$$n\left(b_0(x - yp) + z\left(\sigma_\alpha^2 + \sigma_\eta^2\right) - \tfrac{b}{2}(x - yp)^2 - \tfrac{b}{2}z^2\left(\sigma_\alpha^2 + \sigma_\eta^2\right)\right)$$

$$\left(b_0(X - Yp) + Z\left(\sigma_\alpha^2 + \sigma_\eta^2\right) - \tfrac{nb}{2n^2}(X - Yp)^2 - \tfrac{nb}{2}\tfrac{Z^2}{n^2}\left(\sigma_\alpha^2 + \sigma_\eta^2\right)\right) =$$

$$\left(b_0(X - Yp) + Z\left(\sigma_\alpha^2 + \sigma_\eta^2\right) - \tfrac{b}{2n}(X - Yp)^2 - \tfrac{b}{2n}Z^2\left(\sigma_\alpha^2 + \sigma_\eta^2\right)\right) =$$

$$\left(b_0(\mathbb{E}(E)) + Z\left(\sigma_\alpha^2 + \sigma_\eta^2\right) - \tfrac{b}{2n}(\mathbb{E}(E))^2 - \tfrac{b}{2n}Z^2\left(\sigma_\alpha^2 + \sigma_\eta^2\right)\right).$$

The last equality uses $X - Yp = \mathbb{E}(E)$

The aggregate expected damages in the FE equals

$$n\mathbb{E}\left(\left(d_0+\varphi\alpha+\varepsilon\right)\left(X-Yp+Z\theta\right)+\tfrac{d}{2}\left(X-Yp+Z\theta\right)^2\right)=$$

$$n\left[d_0\left(X-Yp\right)+\tfrac{d}{2}\left(X-Yp\right)^2\right]+n\varphi Z\sigma_\alpha^2+n\tfrac{d}{2}Z^2\left(\sigma_\alpha^2+\tfrac{\sigma_\eta^2}{n}\right)=$$

$$n\left[d_0\left(\mathbb{E}\left(E\right)\right)+\tfrac{d}{2}\left(\mathbb{E}\left(E\right)\right)^2\right]+n\varphi Z\sigma_\alpha^2+n\tfrac{d}{2}Z^2\left(\sigma_\alpha^2+\tfrac{\sigma_\eta^2}{n}\right).$$

Thus, the aggregate payoff (the expectation of benefits minus damages) in the FE is

$$\begin{aligned}&\left(b_0\left(\mathbb{E}E\right)-\tfrac{b}{2n}\left(\mathbb{E}E\right)^2\right)-n\left[d_0\left(\mathbb{E}E\right)+\tfrac{d}{2}\left(\mathbb{E}E\right)^2\right]\\&+Z\left(\sigma_\alpha^2+\sigma_\eta^2\right)\left(1-\tfrac{b}{2n}Z\right)-\left[n\varphi Z\sigma_\alpha^2+n\tfrac{d}{2}Z^2\left(\sigma_\alpha^2+\tfrac{\sigma_\eta^2}{n}\right)\right]\end{aligned} \tag{28}$$

Equations 12 (i) - (iv) imply that the first line of expression 28 can be written as $\left(B_0\left(\mathbb{E}E\right)-\frac{B}{2}\left(\mathbb{E}E\right)^2\right)-\left[D_0\left(\mathbb{E}E\right)+\frac{D}{2}\left(\mathbb{E}E\right)^2\right]$. I use $\sigma_\eta^2=\tau\sigma_\alpha^2$ to write the second line as

$$\beta^{tax}=\left(\left(1+\tau\right)\left(1-\frac{B}{2}Z\right)-\left(\lambda+\frac{D}{2}Z\left(1+\frac{\tau}{n}\right)\right)\right)Z\sigma_\alpha^2.$$

This argument establishes the lemma for the tax policy (where p is nonstochastic). ■

Proof. (Remark 1) I first derive equations 14 and 15. Using Lemma 1, aggregate demand in the FE is

$$\sum_i e_i=n\left(g+h\frac{\sum_i\theta_i}{n}-fp\right). \tag{29}$$

Setting aggregate demand equal to the quota, Q, produces the equilibrium price in terms of the FE parameters. Using Using Lemmas 1 and 3 to replace the FE parameters with the RAE parameters produces equation 14. This step uses the fact that the average across firms (in the FE) of the cost shock ($\frac{\sum_i\theta_1}{n}$) equals the cost shock in the RAE, θ.

From Lemma 2 aggregate emissions in the FE equal

$$\sum_i e_i=n\left(\frac{b_0-d_0}{b+dn}-\frac{1}{b+dn}p\right)+\frac{(1-\varphi)n}{b+dn}\bar{\theta} \tag{30}$$

Replacing the FE parameters by the RAE parameters produces equation 15.

Statement (i) in the Remark follows from inspection. Statement (ii) follows from comparing the expected quota prices in equations 14 and 16 and taking the derivative (with respect to n) of the coefficient of θ in equation 14. Statement (iii) follows from taking the derivative of $\left(\frac{1+\frac{\lambda}{n}}{B\left(1+\frac{W}{n}\right)}\right)^2$ (the square of the coefficient of the shock) with respect to λ, W, n. ■

Proof. (Proposition 1) Substituting Z^{quota} from equation 13 into the formula for β^{quota} in Lemma 3, and then simplifying by using the equalities in equation 12, produces equation 17. The other statements in the proposition follow from taking derivatives or by inspection. ■

Proof. (Proposition 2) Inserting Z^{tax} from equation 13 into equation 12 and simplifying produces

$$\beta^{tax} = \frac{h(n)}{2Bg(n;W)}\tau\sigma_\alpha^2.$$

This expression uses the two functions

$$h(n:\lambda,\tau,W) \equiv (\tau - 2\lambda - W + 1)n^3 +$$

$$\left(2\lambda^2 + W(\tau+2)\right)n^2 + \left(\lambda W(\lambda-2) - \lambda^2(\tau+1)\right)n - \lambda^2\tau W,$$

$$g(n;W) \equiv n^3 + 2n^2W + nW^2.$$

The derivatives, with respect to n, τ, λ, W, of the ratio of cubics (in n) are too complicated to admit general analysis. However, because these derivatives are also ratios of polynomials, it is easy to determine the sign (for large n) of each derivative by considering only the leading terms in the numerator and denominator.

(i) I have

$$\tfrac{d\beta^{tax}}{d\tau} = \tfrac{k}{2Bn(n+W)^2}\sigma_\alpha^2 \quad \text{with}$$

$$k \equiv (2(\tau-\lambda) - W + 1)n^3 + (\lambda + 2W + \lambda(2\lambda - 2\tau + W - 1) + 2\lambda\tau - \lambda W + 2\tau W)n^2$$

$$+ (2\lambda\tau W - \lambda(\lambda + 2W + 2\lambda\tau - \lambda W + 2\tau W))n - 2\lambda^2\tau W.$$

The denominator of the derivative is positive and for large n the numerator is positive iff the coefficient of the leading term is positive, i.e., iff inequality 18 holds.

(ii) The derivative of β^{tax} with respect to λ is

$$\frac{d\beta^{tax}}{d\lambda} = -\frac{n^3 + (-2\lambda)n^2 + (\lambda + W + \lambda\tau - \lambda W)n + \lambda}{Bn(n+W)^2}\tau\sigma_\alpha^2.$$

This derivative is negative for large n.

(iii) The derivative of β^{tax} with respect to n is

$$\frac{d\beta}{dn} = \frac{m}{2Bn^2(n+W)^3}\tau\sigma_\alpha^2 \text{ with}$$

$$m \equiv \left(\tau W - 4\lambda W - 2W^2 - 2\lambda^2\right) n^3 +$$

$$\left(2\lambda^2 + 2W^2 + 4\lambda W + 2\lambda^2\tau + \tau W^2\right) n^2 + 3\lambda^2\tau W n + \lambda^2\tau W^2$$

The denominator of this expression is positive, and for large n the numerator has the same sign as the coefficient of the leading term. This term, $\tau W - 4\lambda W - 2W^2 - 2\lambda^2$, is positive if and only if inequality 19 holds. ■

Proof. (Proposition 3) With certainty equivalent policies, the payoff under taxes minus the payoff under quotas is $\beta^{tax} - \beta^{quota}$. Using the formulae for these two functions in equation 12, and keeping in mind that the two functions are evaluated using different coefficients Z (shown in equation 13), the difference in payoffs can be written as the ratio of two quartics in n. The ratio of leading terms of these quartics is

$$h(n) = \frac{(1 - W - 2\lambda) n^4 + ((3-\tau) W + (2\lambda^2 + 2\lambda + \tau - 1)) n^3 + o(n^3)}{2Bn^4 + 2B(2W-1) n^3 + o(n^3)}\sigma_\alpha^2. \tag{31}$$

Equation 20 follows from from inspection of equation 31, retaining only the first leading term in the numerator and denominator to obtain the asymptote of f. Equation 21 retains both of the two leading terms to obtain a first-order finite-n approximation of the payoff difference. We need this approximation to identify the effect of τ on the payoff difference. ■

Proof. (Corollary 1). Denote $\widetilde{h} \equiv \frac{\frac{1}{n^3}}{\frac{1}{n^3}} h(n)$. The three parts of the corollary follow from taking derivatives of $\widetilde{h}$.

(i) The derivative of $\widetilde{h}$ with respect to τ is $\frac{1}{2B}\frac{1-W}{n+2W-1}$. Because $n + 2W - 1 \geq n - 1 > 0$, the denominator is positive. Therefore, this derivative is positive if and only if $1 > W$. When this inequality holds, an increase in τ favors taxes.

For part (ii) use

$$\frac{d\widetilde{h}}{dn} = -\frac{2(\lambda + W)^2 + \tau(1-W)}{2B(n+2W-1)^2}.$$

The denominator is positive so the derivative is positive if and only if $2(\lambda+W)^2 + \tau(1-W) < 0$. For $\lambda = 0$, the quadratic on the left is negative only for W between the two roots $\frac{1}{4}\tau + \frac{1}{4}\sqrt{\tau(\tau-8)}, \frac{1}{4}\tau - \frac{1}{4}\sqrt{\tau(\tau-8)}$, which are real (and positive) if and only if $\tau > 8$.

For part (iii) use

$$\frac{d\widetilde{h}}{d\lambda} = \frac{1}{2B}\left(\frac{4\lambda - 2n + 2}{n + 2W - 1}\right).$$

The denominator is positive so the derivative is positive if and only if the numerator is positive, i.e., when $\lambda > 0.5\,(n-1)$. ■

Proof. (Proposition 4) Part i The first step generalizes Lemma 2 to obtain the Nash equilibrium decision rule when agents in the FE face a linear tax $G + M\sum_j e_j$. The second step finds the equilibrium emissions rule in the team problem, where agents choose their emissions to maximize aggregate welfare (not their individual welfare). Agents base their expectations of the market fundamental and of other agents' emissions on their private signal. The third step chooses the constants G, M in the tax rule to induce the Nash equilibrium decision rule to be the same as the equilibrium to the team problem. This step uses Lemma 3 to convert the FE parameters into their RAE equivalents.

Step 1. Agent i's maximand is the expectation of

$$\Pi = (b_0 + \theta_i - G)\,e_i - \tfrac{b}{2}\,(e_i)^2 - M\left(\textstyle\sum_j e_j\right)e_i -$$

$$\left((d_0 + \varphi\alpha + v)\textstyle\sum_j e_j + \tfrac{d}{2}\left(\sum_j e_j\right)^2\right).$$

If agent i anticipates that agents $j \neq i$ use the decision rule $e_j = \Psi - \Delta G + \Phi\theta_j$, the first order condition for i's problem is

$$0 = \mathbb{E}\left[(b_0 + \theta_i - G) - be_i - M\left(\textstyle\sum_{j\neq i}(\Psi - \Delta G + \Phi\theta_j) + 2e_i\right)\right] -$$

$$\mathbb{E}\left[\left((d_0 + \varphi\alpha + v) + d\left(\textstyle\sum_{j\neq i}(\Psi - \Delta G + \Phi\theta_j) + e_i\right)\right)\right].$$

Using $\mathbb{E}\theta_j = \mathbb{E}\alpha = \theta_i$, I write the first order condition as

$$(-2M - b - d)\,e_i + \theta_i - G + b_0 - d_0 - \varphi\theta_i - M\,(n-1)\,(\Psi - G\Delta + \Phi\theta_i)$$

$$-d\,(n-1)\,(\Psi - G\Delta + \Phi\theta_i) = 0.$$

The hypothesis that the equilibrium decision rule is $e_j = \Psi - \Delta G + \Phi\theta_j$ for all j implies $e_i = \Psi - \Delta G + \Phi\theta_i$. Substituting this expression into the equality above, and collecting terms in θ_i produces

$$\begin{gathered}(1-\Phi(2M+b+d)-M\Phi(n-1)-d\Phi(n-1)-\varphi)\,\theta_i\\ +b_0-G-d_0-(\Psi-G\Delta)(2M+b+d)\\ -M(\Psi-G\Delta)(n-1)-d(\Psi-G\Delta)(n-1)=0.\end{gathered} \tag{32}$$

Because this equality must hold for all θ_i, the coefficient of θ_i must equal zero, implying

$$\begin{gathered}(1-\Phi(2M+b+d)-M\Phi(n-1)-d\Phi(n-1)-\varphi)=0\\ \Rightarrow\\ \Phi=\tfrac{1-\varphi}{M+b+Mn+dn}.\end{gathered}$$

Collecting terms in G in the remaining two lines of the first order condition, equation 32, produces

$$\begin{gathered}(\Delta(2M+b+d)+M\Delta(n-1)+d\Delta(n-1)-1)\,G+\\ (b_0-d_0-\Psi(2M+b+d)-M\Psi(n-1)-d\Psi(n-1))=0.\end{gathered} \tag{33}$$

This equality must hold for all G, implying that the coefficient of G is zero:

$$\begin{gathered}(\Delta(2M+b+d)+M\Delta(n-1)+d\Delta(n-1)-1)=0\\ \Rightarrow\\ \Delta=\tfrac{1}{M+b+Mn+dn}.\end{gathered}$$

The term in the second line of equation 33 must also vanish, implying

$$\Psi=\frac{b_0-d_0}{M+b+Mn+dn}.$$

Step 2. This step obtains the equilibrium decision rule in the team problem, where agent i chooses its emissions, e_i, to maximize

$$\mathbb{E}\left[\sum_i\left((b_0+\theta_i)\,e_i-\frac{b}{2}e_i^2\right)-\left((D_0+\lambda\alpha+w)\,E+\frac{D}{2}E^2\right)\right].$$

This agent uses its signal θ_i to form expectations of other agent' signals and their actions. Suppose agents use decision rule $e_j=U+V\theta_j$, where the constants U,V are to be determined.

The agent's first order condition is

$$\mathbb{E}\left[b_0 + \theta_i - be_i - \left(D_0 + \lambda\alpha + w + D\left(\sum_{j\neq i} e_j + e_i\right)\right)\right] =$$

$$\mathbb{E}\left[b_0 + \theta_i - be_i - \left(D_0 + \lambda\alpha + w + D\left(\sum_{j\neq i} (U + V\theta_j) + e_i\right)\right)\right] =$$

$$\left[b_0 + \theta_i - be_i - \left(D_0 + \lambda\theta_i + D\left(\sum_{j\neq i} (U + V\theta_i) + e_i\right)\right)\right] =$$

$$\left[b_0 + \theta_i - be_i - \left(D_0 + \lambda\theta_i + D\left((n-1)(U + V\theta_i) + e_i\right)\right)\right] =$$

$$\left[b_0 + \theta_i - b(U + V\theta_i) - \left(D_0 + \lambda\theta_i + D\left((n-1)(U + V\theta_i) + (U + V\theta_i)\right)\right)\right] =$$

$$(1 - D(V + V(n-1)) - Vb - \lambda)\theta_i + (b_0 - D_0 - D(U + U(n-1)) - Ub) = 0.$$

Setting the coefficient of θ_i equal to zero and solving for V implies

$$V = \frac{1-\lambda}{b + nD}.$$

Setting the intercept equal to zero and solving for U implies

$$U = \frac{1}{b + nD}(b_0 - D_0).$$

Step 3. This step chooses the coefficients of the tax rule, $p = G + ME$, to induce the non-cooperative Nash equilibrium to support the equilibrium to the team problem. The two decision rules are $e_j = \Psi - \Delta G + \Phi\theta_j$ and $e_j = U + V\theta_j$, respectively. Setting $\Phi = V$, using Lemma 3 to transform the FE parameters into their RAE equivalents, and solving for M, gives

$$M = \frac{n-1}{(n+1)(1-\lambda)} B(\lambda + W).$$

Using this definition of M, setting $\Psi - \Delta G = U$, again using Lemma 3 to convert FE parameters to their RAE equivalents, and solving for G gives

$$G = \frac{n-1}{n(1-\lambda)}(D_0 - \lambda B_0).$$

This argument establishes the formula for the smart tax when $n < \infty$.

Part ii A straightforward calculation establishes that the optimal policy for the planner who knows the realization of the average signal θ (but not the individual signals) results in the same aggregate level of emissions as in the equilibrium to the team problem. Part iii The comparative statics in equation 24 follow from taking derivatives and simplifying. ■

A.1 A quick derivation of the ranking criterion

We can interpret the standard formulation of social welfare in this setting,

$$W = \mathbb{E}\left[\left(B_0 + \theta\right)E - \frac{B}{2}E^2 - \left(\left(d + \Gamma\right)E + \frac{D}{2}E^2\right)\right], \tag{34}$$

as the RAE with $n = \infty$. As in the text $\theta \sim (0, \sigma^2)$ and $\Gamma = \lambda\theta + u$, where $u \sim iid(0, \sigma_u^2)$, so $\mathbb{E}\theta\Gamma = \lambda\sigma^2$. The agent who chooses emissions ignores the damage externality. Under a binding quota, Q, so $E = Q$. Taking expectations of W gives welfare under this quota:

$$W^Q = \overline{W} \equiv B_0 Q - \frac{B}{2}Q^2 - \left(dQ + \frac{D}{2}Q^2\right).$$

The agent facing a tax, p, maximizes expected benefits minus tax payments by setting emissions equal to $E = \chi\left(p\right) + \frac{\theta}{B}$, with $\chi = \chi\left(p\right) \equiv \frac{b-p}{B}$ equal to expected emissions under the tax. The tax and quota pair, (Q, p) are certainty equivalents, i.e., $\chi(p) = Q$. Substituting $E = \chi\left(p\right) + \frac{\theta}{B}$ into the expression for welfare, using $\chi(p) = Q$, and taking expectations, gives welfare under the tax:

$$W^{tax} = \overline{W} + \frac{\sigma^2}{2B}\left(1 - W - 2\lambda\right)$$

with $W = \frac{D}{B}$. Welfare under the tax minus welfare under the quota equals $\frac{\sigma^2}{2B}\left(1 - W - 2\lambda\right)$: taxes dominate quotas if and only if $1 - 2\lambda > W$.

References

Angeletos, G.-M. & Pavan, A. (2007), 'Efficient use of information and social value of information', *Econometrica* **75**, 1103–1142.

Cantillon, E. & Slechten, A. (2018), 'Information aggregation in emissions markets with abatement', *Annals of Economics and Statistics* (132), 53–79.

Englander, G., Karp, L. & Simon, L. (forthcoming), 'The value of information in a congested fishery', *Review of Economic Studies* . Forthcoming, Review of Economic Studies.

Fell, H., Burtraw, D., Morgenstern, R. D. & Palmer, K. L. (2012), 'Soft and hard price collars in a cap-and-trade system: A comparative analysis', *Journal of Environmental Economics and Management* **64**, 183–198.

Fell, H. & Morgenstern, R. D. (2010), 'Alternative approaches to cost containment in a cap-and-trade system', *Environmental Resource Economics* **47**, 275–297.

Grüll, G. & Taschini, L. (2011), 'Cap-and-trade properties under different hybrid scheme designs', *Journal of Environmental Economics and Management* **61**(1), 107–118.

Hepburn, C. (2006), 'Regulation by Prices, Quantities, or Both: A Review of Instrument Choice', *Oxford Review of Economic Policy* **22**(2), 226–247.

Karp, L. (2025), 'The gains from international carbon trade and the ranking of suboptimal climate policies', Working paper.

Karp, L. S. & Traeger, C. P. (2025), 'Smart cap', *Journal of the European Economic Association* **23**(2), 554–593.

Karp, L. & Traeger, C. (2024), 'Prices versus quantities reassessed', *J. Environmental Economics and Management* **125**, 1–19.

Karp, L. & Zhang, J. (2005), 'Regulation of stock externalities with correlated abatement costs', *Environmental and Resource Economics* **32**, 273–299.

Lemoine, D. (2022), Informationally efficient climate policy: Designing markets to measure and price externalities, NBER Working Paper w30535, National Bureau of Economic Research. Revised June 2025.

Malueg, D. A. & Yates, A. J. (2009), 'Bilateral oligopoly, private information, and pollution permit markets', *Environmental and Resource Economics* **43**(4), 553–572.

Marcoul, P. (2020), 'Fish abundance information and fishery's aggregate behavior', *Journal of the Association of Environmental and Resource Economists* **7**(5), 933–964.

Montero, J. (2009), 'Market power in pollution permit markets', *The Energy Journal* **30**(2_suppl), 115–142.
URL: *https://doi.org/10.5547/ISSN0195-6574-EJ-Vol30-NoSI2-6*

Montero, J.-P. (2008), 'A simple auction mechanism for the optimal allocation of the commons', *American Economic Review* **98**(1), 496–518.

Morris, S. & Shin, H. S. (2002), 'Social value of public information', *American Economic Review* (5), 1521–1534.

Newell, R. & Pizer, W. (2003), 'Regulating stock externalitiues under uncertainty', *Journal of Environmental Economics and Management* **45**, 416–32.

Pizer, W. A. (2002), 'Combining price and quantity controls to mitigate global climate change', *Journal of Public Economics* **85**, 409–434.

Roberts, M. J. & Spence, M. (1976), 'Effluent charges and licences uncer uncertainty', *Journal of Public Economics* **5**, 193–208.

Stavins, R. (1996), 'Correlated uncertainty and policy instrument choice', *Journal of Environmental Economics and Management* **30**(2), 218–232.

Vives, X. (2008), *Information and Learning in Markets: The Impact of Market Microstructure*, Princeton University Press, Princeton and Oxford. Course Book Edition.

Weitzman, M. L. (1974), 'Prices versus quantities', *Review of Economic Studies* **41**, 477–491.

Weitzman, M. L. (1978), 'Optimal rewards for economic regulation', *American Economic Review* **68**(4), 683–691.

Wirl, F. (2009), 'Oligopoly meets oligopsony: The case of permits', *Journal of Environmental Economics and Management* **58**, 329–337.
URL: *https://doi.org/10.1016/j.jeem.2009.04.006*