Regulating Global Externalities

Roweno J.R.K. Heijmans and Reyer Gerlagh*

December 18, 2019

Abstract

The question in which we are interested is how a market, inhabited by multiple agents about whom we are differentially uncertain and who exchange goods the use of which imposes a negative externality on society, is to be ideally regulated. We show that (posterior) observed trades, conditional on prior *asymmetric uncertainty* about agents' demand, is a rich source of information usable to reduce aggregate uncertainty. The observation implies that whereas asymmetric *information* usually entails a cost on welfare, asymmetric *uncertainty* can help achieve greater efficiency in regulation. We conclude with a welfare comparison between quotas that optimally exploit asymmetric uncertainty and prices.

JEL codes: D82, D83, H23

Keywords: asymmetric information, regulatory instruments, policy updating, asymmetric uncertainty, decision making under uncertainty

1 Introduction

One evening in a small bar, an individual going by the name of A sits by a table together with another whose name is B. The two smoke one cigarette after another and at the end of a long evening, the bar is black with smoke. As the smokers head home, several remaining customers voice their frustrations to the bartender. "We did not come here," they say, "to drink our beers in the smoke of others!" These complainants being regulars, the bartender takes their frustration seriously. Counting butts in the ashtray, s/he concludes the couple

*Heijmans (r.j.r.k.heijmans@uvt.nl) and Gerlagh (r.gerlagh@uvt.nl) are at the Economics Department of Tilburg University. Gerlagh is also affiliated to the Oslo Centre for Research on Environmental friendly Energy (CREE), Norway.

smoked 20 cigarettes in total. As it happens, smoker A has been frequenting this bar for years, and it is known that A always smokes 10 cigarettes on a night out. This leads to a simple arithmetic conclusion: smoker B, who is new in town, has also smoked 10 cigarettes. After elaborate calculations – into which we shall not dwell any further here – on the back of a coaster, the tapster deduces that half this number of cigarettes smoked would, in expectations, be economically efficient for the full clientèle.

Next week, A and B are back, and the barkeeper politely asks them to cut down on the smoking. "You can have five cigarettes each," s/he says, "though you can freely trade these between the two of you, as long as the total does not exceed ten." Then, something peculiar happens. As the couple spends a cozy evening, A smokes nine cigarettes whereas B vapes only one. What should the bartender do? Has s/he made a mistake? S/He ponders the matter long and hard, eventually concluding that B must, somehow, have had less of an appetite for smoking than initially thought. Using simple economics and in pursuit of the best overall interest, next time the couple is asked to decrease their common budget further. They can smoke six cigarettes only.

At a certain level of conceptual abstraction, the barkeeper's situation is similar to a signal extraction problem, which as a branch of statistical theory goes by the name of filtering. Economists have applied these methods – and most frequently the Kalman filter (which we know to be the optimal filter when signals are normally distributed) – in many corners of their discipline; from investment decisions (Townsend, 1983) to macroeconomic theory (Lucas, 1978), from finance (Makarov and Rytchkov, 2012) and time series econometrics (Hamilton, 1989; Baxter and King, 1999; Harvey and Trimbur, 2003; Talmon and Coifman, 2013) to behavioral economics (Mullainathan, 2002; Moore and Healy, 2008). Our mission will be to apply the insights so fruitfully used in said applications and further develop these for the regulation of traded 'goods' which impose an externality. In particular, we study how to maximally refine quota instruments to yield the highest possible level of social welfare under incomplete information, given the constraint that the instrument must be implementable and incentive-compatible.

As said, in refining policy instruments we focus on quantities only. It is well known that a combination of instruments, such as a hybrid between prices and quantities, can in principle do better than one of these instruments alone (Roberts and Spence, 1976; Weitzman, 1978; Pizer, 2002). We abstract away from such combinations of different instruments and instead focus on pure instruments and take this type of regulation to its utmost extreme. What we will see is that, when used smartly, endogenous cap and trade can mitigate welfare losses substantially compared to pure trading policies, motivating the idea that complicated hybrid instruments may not be needed if one is willing to carefully contemplate on and improve the pure instruments. However, in Section 4 we compare our refined quantity instruments with a straightforward tax (price) on smoking, reproducing the seminal contribution of Weitzman (1974). In this interpretation, our model is intimately related to Mideksa and Weitzman (2019) and Doda et al. (2019).

The classical argument to favor trade and markets is of remarkably individualistic a nature. Any distribution of production or consumption that comes about through free exchanges in well-functioning market environments, so it goes, will always be at least as efficient as had allocations been set to individual agents directly. This is why economists tend to favor market-wide constraints over individual-specific ones, as evidenced by many large-scale real-world examples such as fishing quotas, tradeable pollution permits for sulfur and carbon-dioxide in the U.S. and Europe, milk quotas in the EU, and the Agricultural Supply Management Scheme in Canada. We do not doubt the rightness of this conviction, yet we do point out it is in a sense incomplete. Though markets foster efficiency of the *individual* distribution of some undertaken activity, they do not resolve the problem of joint efficiency, i.e. in the aggregate.

Our central thesis builds on the logic presented above: when left to free exchange, agents in a market will trade and barter until no more mutually beneficial deals can be made, ensuring an efficient allocation of, say, cigarettes across individual agents. However, and this is our central insight, there is more to be inferred, in particular about aggregate efficiency. The logic is simple. Suppose we observe some final allocation of cigarettes comes about and suppose also this one is different from the outcome we had initially assumed would occur. We then realize our initial estimation of smokers' preferences is in need of some correction. There's nothing new under the sun, some would argue, for this possibility is the very reason a market was established in the first place. However, after we applaud the market for having achieved greater welfare, we should get back to business and realize that the mistaken initial individual estimates may fairly well imply our aggregate estimate was, in fact, also wrong; a better estimate can be constructed on the basis of observed market behavior. Ex post we might wish to adapt the aggregate allocation of cigarettes to the market.

Consider the example with which we started. Suppose we know the preferences of smoker A while those of smoker B are uncertain, though we initially expect them to be the same as for A. We allocate a total of ten cigarettes to 'the market', which we know is optimal given our initial estimate of preferences. This estimate also tells us that in fact each smoker will consume five cigarettes. Upon noting that in reality A smokes nine out of ten cigarettes, leaving B with only one, we realize our estimate of B's preferences must have been wrong all along, for B has revealed to be in fact a less devout follower of the tobacco religion than A. Given we know the preferences of A, and given we learned B to enjoy smoking less, we conclude 'aggregate' preferences for smoking are lower than expected. Standard economics then dictates a lower aggregate cigarette consumption would have been socially optimal. Instead of ten, the bar tender should have given them only six cigarettes.

In practical terms we are saying that the market cap on some activity with an associated externality should be endogenous to market behavior as the latter represents information regarding relevant characteristics of the market we are regulating and about which we are uncertain. If our exercise is to not only make the cap endogenous but in fact make it optimally so, we introduce a technical challenge. The standard problem, meaning with a fixed cap, imposes total consumption of cigarettes as exogenously given and lets the market sort out the individual allocations, subject to a simple constraint: total smoking is not to exceed the cap. An optimal exogenously fixed cap is therefore efficient in expectations, meaning before any observation on individual outcomes is available.

With an endogenous cap, on the other hand, individual smokers still trade freely, but these trades have an effect on how much can be freely traded in the first place, meaning the cap is now a function of trades. Our question is simple but fundamental: what should be the properties of an efficient endogenous cap, i.e. what should this function look like? The answer is surprisingly straightforward. We propose to make use of the most refined source of information available from trading behavior, namely individual consumption. We also identify the precise conditions under which trade improves information about *aggregate* preferences and thus creates some scope for improving the efficiency of regulation. The intuition was presented earlier in this introduction. What we want to emphasize in particular is that we are exploiting asymmetries in *uncertainty*. If the bar tender had exactly the same prior information about A as about B, he could update the estimates for relative individual preferences but *not* for absolute aggregate preferences. Aggregate revisions require *asymmetric uncertainty*, the second moment of asymmetric information.

Our work bears some superficial resemblance with the literature on optimal taxation (for instance Mirrlees, 1971) which generally speaking studies situations where agents hold private information about, for instance, their preferences and in such a framework derives an optimal (from the social welfare point of view) tax or policy. Though this be similar in spirit to our current undertaking, the main distinction is rather fundamental: under the optimal taxation paradigm *individual* types or preferences are indeed mysterious to the regulator, yet one still assumes the distribution of individual types to somehow be known – i.e. individual preferences behave in a predictable way. Our framework does not incorporate the latter assumption. Hence, whereas optimal taxationists may still uphold the idea that *aggregate* preferences are known, we, additionally, have to deal with uncertainty even in the aggregate.

It is clear that our adventure shall exist in constructing regulatory policies dealing with both individual and aggregate uncertainty. We believe this to be a novelty, and as a stepping stone we shall focus attention to the case of regulating a divisible good or activity (e.g. cigarettes or smoking) with two smokers causing a global externality (smoke). For optimal regulation this practically implies we need only address aggregate uncertainty in our model; free trade will, in the Coasean fashion, ensure an individually optimal allocation of what is in the aggregate available.

Of course, to think about regulation under asymmetric information is not altogether new. We have already mentioned the optimal taxationists, such as Mirrlees (1971), but also mechanism design comes to mind, i.e. the branch of game theory where menus of individual contracts are sought for which it is in best interest of an individual to select the contract which reveals its private information. Now, clearly, the construction of such menus occurs under certain assumptions, most importantly for the current case that individual types (i.e. that quantity about which knowledge is privately held by the regulated agent) are drawn from a *commonly known* distribution. Again, therefore, aggregate uncertainty is carefully abstracted away from. However, we nonetheless learn something important from this literature even for the current project: somehow we can seduce agents to reveal their private information, namely by offering a set of contracts that is fine-tuned to this exact purpose. Our regulation structure is set up to invoke the same mechanism, mutatis mutandis, for aggregate private information.

Jointly these observations and insights set out the agenda for our project: we will seek a policy that regulates a market in the aggregate (as in the literature on optimal taxation), but this policy must also force the market to reveal its private information *and* act upon this information in an ideal manner (as in the literature on mechanism design). Remarkably, we shall prove this policy exists. That is, we show the above considerations imply certain 'desirable properties' of a policy, and then proceed to showing a policy that satisfies these properties can be found. What is more, the policy rule that comes about this way is extremely simple.

2 Model

Our universe is a bar, inhabited by two atomic or price-taking smokers $i \in \{1, 2\}$ and a handful of other clients. A price-taking duo may read as unrealistic. One possible interpretation is that our environment is not a bar, but instead a large-scale disco in New York City where smoking is allowed but controlled, and where there are two large groups of individuals which have identical smoking habits and preferences among individuals within a group, but not between groups. In this case, every individual smoker only has a negligible effect on the eventual price of cigarettes. As we will illustrate later on, there is no relevant information available through trade within a group of prior-identical agents, but there is through trade between groups of different prior characteristics. We may therefore, out of convenience and without loss of generality, refer to two atomic smokers indexed i, possibly representative agents. To every smoker i, smoking $\widetilde{s}_i$ cigarettes yields private benefits $B_i(\widetilde{s}_i; \theta_i)$. The parameter θ_i can be thought of as a preference shifter, affecting how much pleasure is derived from smoking a given number of cigarettes and is known to smoker i but not to any other agent. Although the actual realization of θ_i is private information, it is common knowledge that $\mathbb{E}[\theta_i] = 0$, $\mathbb{E}[\theta_i^2] = \sigma_i^2$, and $\mathbb{E}[\theta_1\theta_2] = \rho\sigma_1\sigma_2$. Because this universe is uncertain to the extent that θ_i cannot be properly predicted, σ_i^2 is a logical measure for the uncertainty about a agent i's preferences. To say that uncertainty is asymmetric in our terminology is equivalent to saying that $\sigma_1 \neq \sigma_2$.

As far as smoking emits smoke, it is disliked by other customers. We will assume the severity of this externality to only depend on the total amount of smoke in the bar, independent of which smoker exhaled it, so we are dealing with a global externality. Because there is (in this simplified universe) a one-to-one relationship between the amount of cigarettes consumed and the amount of smoke produced, we can treat the externality as a cost, broadly interpreted, $C(\widetilde{s}_1 + \widetilde{s}_2)$ depending solely on the aggregate number of cigarettes smoked.

As the other clients are bothered by the amount of smoke emitted, the barkeeper faces the task of finding quantities $\widetilde{s}_1$ and $\widetilde{s}_2$ that maximize social welfare:

$$W = B_1(\widetilde{s}_1; \theta_1) + B_2(\widetilde{s}_2; \theta_2) - C(\widetilde{s}_1 + \widetilde{s}_2). \tag{1}$$

In the absence of asymmetric information, the fully knowledgeable barkeeper can set these quantities directly or else put a price on cigarettes that will make the individual smokers consume the same quantities, and these two instruments are perfectly equivalent, see Montgomery (1972). As was first shown by Weitzman (1974), this formal equivalence between instruments breaks down once we introduce an informational disparity, captured here by θ_i. Note that we will not be dealing with what price to put on cigarettes in this manuscript, our focus being confined to quantity-based regulations only.

Note that program (1) takes into account the utility of both smokers as well as the global disutility due to their smoking behavior. In this sense, the barkeeper implicitly equates the (expected) shadow value of cigarettes for each smoker individually to the shadow value of 'fresh air' to all other visitors. Relating our general model to the regulation of greenhouse gas emissions, the resulting market price of cigarettes in this bar therfore mirrors what Kotchen (2018) calls the 'Global Social Cost of Carbon'. Economic inefficiencies related to the more egocentric 'Domestic Social Cost of Carbon' do not arise.

It will serve the analysis to make some restrictive assumptions regarding the forms benefits and costs take. In particular, ket benefits to smoker i are given by:

$$B_i(\widetilde{s}_i;\theta_i) = (p_i^* + \theta_i)(\widetilde{s}_i - s_i^*) - \frac{\beta_i}{2}(\widetilde{s}_i - s_i^*)^2, \tag{2}$$

where the vector (p_i^*, s_i^*) is common knowledge and will be elaborated upon soon. Marginal benefits are therefore linear in cigarettes with the intercept determined by θ_i:

$$MB_i(\widetilde{s}_i) = p_i^* - \beta_i(\widetilde{s}_i - s_i^*) + \theta_i. \tag{3}$$

Costs as a result of disutlity from smoke (the externality) are described by the functional form:

$$C(\widetilde{s}_1 + \widetilde{s}_2) = p^*(\widetilde{s}_1 + \widetilde{s}_2 - s_1^* - s_2^*) + \frac{\gamma}{2}(\widetilde{s}_1 + \widetilde{s}_2 - s_1^* - s_2^*)^2. \tag{4}$$

Marginal costs due to smoke are then seen to be:

$$MC(\widetilde{s}_1 + \widetilde{s}_2) = p^* + \gamma(\widetilde{s}_1 + \widetilde{s}_2 - s_1^* - s_2^*). \tag{5}$$

For brevity of notation, where convenient we may write $\widetilde{S} = \widetilde{s}_1 + \widetilde{s}_2$ and $S^* = s_1^* + s_2^*$. Our model is now characterized by eight parameters $(\beta_i, \gamma, p_i^*, s_i^*, p^*)$ describing slopes and intercepts of three linear curves. We need only two per curve (slope and level) for three curves (2 marginal benefits 1 marginal costs), six in total. Consequently, we may take the freedom to reduce the number of parameters through defining $p^* = p_1^* = p_2^*$, with the convenient implication that (p^*, s_1^*, s_2^*) is the vector of welfare-maximizing prices and cigarettes for smoker i, given preferences turn out as expected ($\theta_1 = \theta_2 = 0$). We label this the ex-ante optimum. It is easily seen that global marginal costs and individual marginal benefits equal p^*. This is clearly not an assumption, nor even a normalization; it is a definition. Entailing no more than a simplification of otherwise cumbersome notation, we introduce it here for our own convenience without any implicated loss of generality.

Before proceeding to the analysis, we introduce some further notation. Superscripts will be scenario (instrument) labels for equilibrium outcomes. Moreover, let $\widetilde{x}^k$ denote the value of a variable x under policy k, then let $x^k := \widetilde{x}^k - x^*$ be the deviation of x under policy k from the ex-ante expected optimal value x^*, and let $\Delta^k x := \widetilde{x}^k - x^{SO}$ denote the difference between the value of x under scenario k and its ex post socially optimal value (to be derived shortly).

Our game has the following stages:

1. The barkeeper chooses an instrument to regulate the market for cigarettes.
2. Smokers observe their individual preference shock θ_1 and θ_2.
3. Trade clears the market. Prices and quantities are chosen, jointly for both smokers, consistent with utility maximization by each smoker,

$$-\beta_i s_i^k + \theta_i = p_i^k, \tag{6}$$

   while the policy rules determine the relation between quantities and prices within and across smokers.

Finally a note on the uncertainty of smoker preferences. It has been noted that when $\sigma_1 \neq \sigma_2$, uncertainty is asymmetric. An extreme case thereof occurs when the preferences of one smoker are uncertain while those of the other are not, e.g. $\sigma_1 = 0, \sigma_2 > 0$. In this case, it is clear that behavior of the second, uncertain smoker is informative about his true preferences (θ_2). Preferences drive behavior, and so if smoker 2 behaves differently from smoker 1, this indicates the direction and magnitude in which their preferences are different. Since those for smoker 1 are known, θ_2 is determined perfectly. In this sense, the smoker with certain preferences is a perfect *anchor* for the Bayesian updating of smoker preferences. Similarly, we can say that for any two smokers with asymmetrically uncertain preferences, the more certain smoker is a relative *anchor* for Bayesian updating of priors.

2.1 Social Optimum: Common Knowledge

By standard arguments, it is immediately clear marginal benefits of smoking should equal the marginal costs of smoke in an efficient outcome, which implies $MB_1 = MB_2$. Since marginal benefits also equal prices, these are the same, so $p_1^{SO} = p_2^{SO} = p^{SO}$. Labeling the symmetric information equilibrium as Social Optimum, we have the profit-maximization condition (6) for $k = SO$ and

$$\gamma(s_1^{SO} + s_2^{SO}) = p^{SO}$$

so the Social Optimum is fully characterized:

$$p^{SO} = \frac{\gamma(\beta_2\theta_1 + \beta_1\theta_2)}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2}, \tag{7}$$

$$s_i^{SO} = \frac{\beta_{-i}\theta_i + \gamma(\theta_i - \theta_{-i})}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2}, \tag{8}$$

$$S^{SO} = \frac{\beta_2\theta_1 + \beta_1\theta_2}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2}, \tag{9}$$

where $i \in \{1, 2\}$ and $-i$ is the complement of i. Thus, a positive preference shift induces increased consumption of cigarettes by the smoker to whom it occurred, and decreases it for the other, though aggregate consumption and the common price always increase for a positive shift to either smoker.

The variance of prices (eq. 7) is:

$$\begin{aligned}\mathbb{E}\left[\left[p^{SO}\right]^2\right] &= \left(\frac{\gamma}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2}\right)^2 \mathbb{E}\left[\beta_2^2\theta_1^2 + \beta_1^2\theta_2^2 + 2\beta_1\beta_2\theta_1\theta_2\right] \\ &= \left(\frac{\gamma}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2}\right)^2 \left[\beta_2^2\sigma_1^2 + \beta_1^2\sigma_2^2 + 2\beta_1\beta_2\rho\sigma_1\sigma_2\right].\end{aligned}$$

We note that increasing uncertainty translates in a more volatile price.

2.2 Social Optimum: Private Information

A straightforward mechanism that implements the Social Optimum is a simple ascending clock auction. In its most basic form, the barkeeper offers a supply curve, in price-cigarette space, that coincides with the marginal cost curve. This way, utility-maximizing smokers necessarily incorporate the externality caused by their smoke into smoking behavior, guaranteeing implementation of the Social Optimum as a first-best equilibrium.

That such an easy way of implementing the Social Optimum exists is reason for optimism. After all, it implies that the Social Optimum can be reached without great effort or complication. Nonetheless, practice suggests instruments such as the ascending clock auction are not feasible for reasons outside the scope of our model. The barkeeper is in such a case constrained to other instruments, accepting a second-best regulated equilibrium as at least superior to no regulation at all. It is that constrained context we are about to define and analyze. Before doing so, however, we will first derive in general terms the aggregate welfare loss under a given policy relative to the first-best, Social Optimum welfare level.

3 Policies

3.1 Welfare Costs of Policies

By definition of the difference under policy k with the social optimum and considering a smoker's equilibrium behavior (6), it is immediate that smoking deviations from the Social Optimum scale with price deviations:

$$\Delta^k p_i = -\beta_i \Delta^k s_i. \tag{10}$$

Welfare losses are then given by:

$$\begin{aligned} \Delta^k W &= \mathbb{E}\left[\Delta^k B_1 + \Delta^k B_2 - \Delta^k C_1 - \Delta^k C_2\right] \\ &= \frac{\gamma}{2}\mathbb{E}\left[\left(\Delta^k S\right)^2\right] + \sum_i \frac{\beta_i}{2}\mathbb{E}\left[\left(\Delta^k s_i\right)^2\right]. \end{aligned} \tag{11}$$

Given realized preferences, individual prices and cigarette consumption map injectively. Policies featuring equal prices across smokers thus admit the property that individual and aggregate consumption scale with the common price. Consequently, for such policies welfare losses can be written as a function of the price gap:

$$\Delta^k W = \frac{1}{2}\frac{(\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2)(\beta_1 + \beta_2)}{\beta_1^2\beta_2^2}\mathbb{E}\left[\left(\Delta^k p\right)^2\right]. \tag{12}$$

3.2 Quotas

The simplest possible policy simply sets quotas for each smoker individually:

Definition 1 (Quotas)**.** To both smokers individually, the barkeeper allocates the ex-ante optimal amount of cigarettes s_i^*:

$$s_1 = s_2 = 0, \tag{13}$$

while prices adjust to reach equilibrium on the market for cigarettes (6).

We readily obtain expected welfare losses:

$$\Delta^Q W = \frac{1}{2}\frac{(\gamma+\beta_2)\sigma_1^2 + (\gamma+\beta_1)\sigma_2^2 - 2\gamma\rho\sigma_1\sigma_2}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2}. \tag{14}$$

For future reference, it is important to note that Quantities under autarky as regulation can be considered the execution of the welfare program

$$\max_{s_1,s_2} \mathbb{E} \quad W(s_1, s_2; \theta_1, \theta_2) \tag{15}$$

That is, Quantities is the optimal choice under the information constraint that both quantities must be set before any information is revealed, and without the use of any information extracted from markets. It admits the desirable property that expected marginal benefits to each smoker equals marginal costs:

$$\mathbb{E}[MB_i|s_i] = MC \tag{16}$$

where we note that the RHS marginal costs are perfectly known, when quantities are set whereas the prices at the LHS are stochastic variables due to unknown preferences θ_i.

3.3 Trade

The barkeeper may well understand that through establishing a market for cigarettes even better outcomes, in terms of social welfare, can be obtained. The basic idea is very simple and familiar to every economist in the modern tradition: when a market for cigarettes exists, smokers can freely exchange their smokes, and since such exchanges will only take place as long as they are mutually beneficial, aggregate welfare will be larger as compared to a scenario where no market exists. The policy where a market with the essential feature of free exchange is created will be called Trading.

Definition 2 (Trading)**.** The barkeeper allocates the ex-ante optimal number of cigarettes s_i^* to both smokers, who can freely exchange cigarettes, subject to:

$$s_1 + s_2 = 0. \tag{17}$$

Equilibrium on the market for cigarettes implies (6). Optimization and free trading ensures cigarettes are exchanged until marginal benefits are equal for both smokers:

$$p_1 = p_2. \tag{18}$$

Trading allows smokers the flexibility to efficiently redistribute cigarette allocations in response to preference shifts, subject to the constraint that total smoking is fixed. It follows:

$$\Delta^T W = \frac{1}{2}\frac{1}{\beta_1+\beta_2}\frac{\beta_2^2\sigma_1^2+\beta_1^2\sigma_2^2+2\beta_1\beta_2\rho\sigma_1\sigma_2}{\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2}. \tag{19}$$

The following proposition is now immediate:

Proposition 1. *Trading always outperforms Quotas in terms of welfare.*

Proof. In Appendix. Q.E.D.

It is quite trivial that Trading improves welfare compared to autarkic Quotas. Under the former, aggregate smoking, and therefore costs, are fixed. Yet through a process of mutually beneficial exchanges, benefits (including revenues from cigarette sales) increase for both smokers, raising welfare overall.

The conceptual quality of trade as regulatory principle can be seen more elegantly when taking a more principal look at the policy and its rules. Quotas under autarky as policy can be considered the implementation of welfare program (15). Trading, on the other hand, shifts the expectations operator outside the maximization operator; it is the execution of:

$$\max_{G} \quad \mathbb{E}\Big[\max_{s_1,s_2} W(s_1,s_2;\theta_1,\theta_2) \text{ s.t. } s_1+s_2=G\Big] \tag{20}$$

where the regulator chooses G, the cap on aggregate smoking. That is, Trading effectively delays the choice for optimal smoking per smoker until the point information is revealed, which makes it preferable to static Quotas. It not only admits the desirable property that marginal costs equal expected marginal benefits (which is true also for Quotas, see (16)) – realized marginal benefits are also equal for both smokers:

$$MB = MB_1 = MB_2 \tag{21}$$

$$\mathbb{E}[MB|s_1+s_2] = MC \tag{22}$$

One might wonder about price volatility in the market for cigarettes. After all, if welfare under Trading is higher than under Quotas, as Proposition 1 establishes, and if we know from condition (18) prices will be equal for both smokers, given also it is known in markets with constant prices welfare losses scale with price deviations (equation (12)), should it not follow price volatility is lower under Trading than under Quotas? The proposition maintained a purposeful silence about aggregate price volatility in a Trading market. It did so because results are ambiguous. One might at first suspect trade to reduce global price volatility compared to autarkic Quotas, but a simple thought experiment is illustrative for the wrongness of this intuition. To that end, we note first that prices in equilibrium equate marginal benefits, so that the volatility of prices is also equal, in equilibrium, to the volatility of marginal benefits. Consider then two smokers, the second with more uncertain preferences than the first, $\sigma_2 \gg \sigma_1$, as well as a larger 'absorptive capacity', $\beta_2 \gg \beta_1$. Consequently, if under Quotas, say, five cigarettes are allocated to both of them, we expect price volatility to be almost zero for smoker 1 but much larger for smoker 2, for only the latter experiences large 'preference shifts'. When we introduce Trading, the globally efficient allocation of cigarettes may be such that many flow to smoker 2, inducing a much stronger price fluctuation for smoker 1 than under Quotas. This means the first smoker has practically *imported* part of the price volatility from the second. As the first is characterized by little absorption, prices in equilibrium under Trading will mainly be driven by the second smoker. Aggregate price volatility has in fact increased.

This point is neither just a footnote nor a mere theoretical curiosity. For the success of a Trading regime, it is of fundamental importance. After all, would a stable smoker be willing to expose itself directly to the risk of its unstable comrades? In real-world application this problem is even more severe. Can we expect a stable country to engage in trade with unpredictable nations, thereby possibly damaging itself? It appears unlikely. There is no trivial solution to this fundamental problem. Yet we are about to show how it can be substantially mitigated.

3.4 Rates of substitution

True as it may be that Trading is good for welfare, we argue that it is not best, and therefore the barkeeper can do better. In fact, he is able to do so rather easily, although some delicate insights must be developed first, for else one may not fully understand the subtle mechanisms at work.

The primary notion we need to establish is there exists no such thing as a single marginal rate of substitution for permits; rather, there are two. One operates at the individual smoker's level, the other on the aggregate or global level, the market at large.

The individual rate, labeled MRS_i, appears to be the rate one most frequently has in mind when speaking loosely of marginal substitution, it being the rate at which cigarettes can change hands between individual smokers. Although it might seem intuitive that manipulating this 'trading ratio' could improve welfare, it can in fact be shown that letting it deviate from unity is never optimal. This follows from first principles. An efficient allocation of cigarettes equates marginal benefits across smokers. Smokers, however, only have an incentive to perfectly equate marginal benefits (or the value of smoking) when cigarettes can be traded one-to-one, that is, when one cigarette lost by a smoker translates into exactly one cigarette gained by the other. Therefore, only if the MRS_i is unity will smokers have no incentive to trade cigarettes other than the full equalization of marginal benefits.

These observations do not imply, as might come across at first, that the barkeeper's arsenal of instruments is left depleted. Individual trades must be left untouched, it is true, but *aggregate* trades can still be manipulated. Indeed, one cannot seriously think of any proper economic reasoning contra such operations. It combines the best of two worlds: individual smokers consider their impact on aggregate trade flows to be negligible and will therefore trade freely on the basis of one-to-one exchanges, steering cigarette consumption toward a perfect equalization of marginal benefits, while at the global level smoking is nonetheless adapted to correct for the possibly flawed and suboptimal initial endowment of cigarettes revealed by the market through its self-chosen trade flows. In its simplest form, this manipulation of aggregate permit trades operates through a fixed ratio, called the aggregate marginal rate of substitution, MRS_a for short. To find a handle on the aggregate rate of substitution, we will formulate regulation as the solution to a welfare-maximization problem. We then replicate Trading as regulation and find a natural generalization, labeled Stabilized Trading (because indeed, this policy will be seen so stabilize prices through trade). This approach will tell us how to model such aggregate manipulations and how these should occur in an 'optimal bar'.

3.5 Stabilized Trading

Definition 3 (Stabilized Trading)**.** The barkeeper adapts the aggregate allocation of cigarettes based on observed trade for fixed $MRS_A = \delta$:

$$\delta s_1 + s_2 = 0. \tag{23}$$

Profit maximization and free trading with $MRS_i = 1$ ensures that smokers allocate cigarettes so that marginal productivity is equal for both of them:

$$p_2 = p_1. \tag{24}$$

We observe that for $\delta = 1$, Stabilized Trading is one-to-one also at the aggregate level and thus equivalent to traditional Trading. This observation immediately suggests Stabilized Trading always outperforms traditional Trading, since the barkeeper is free to choose a stabilization rate δ equal to unity but not imposed to do so, which added freedom cannot deteriorate global welfare. We can derive the following result:

Proposition 2. *The optimal stabilization rate is given by:*

$$\delta^* = \frac{\beta_1[\sigma_2^2 - \rho\sigma_1\sigma_2] + \gamma[\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2]}{\beta_2[\sigma_1^2 - \rho\sigma_1\sigma_2] + \gamma[\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2]}. \tag{25}$$

Stabilized Trade equaling Trade is a measure-zero event. In particular,

$$\delta^* \lesseqgtr 1 \iff \frac{\beta_1}{\beta_2} \lesseqgtr \frac{\sigma_1^2 - \rho\sigma_1\sigma_2}{\sigma_2^2 - \rho\sigma_1\sigma_2}. \tag{26}$$

Before we turn our attention to the particular properties of the stabilization rate, we want to appreciate its principles more fundamentally at the informational level. If we consider instruments from the Quantity-based class in general, we understand these can be interpreted as allowing, in principle, full observation of each smoker's cigarette consumption and setting restrictions to these. Thus, what we aim for as the most efficient quantity-based instrument is an equalization of marginal benefits and marginal costs, given all observed quantities (cf (22))! Formally:

$$\mathbb{E}[MB|s_1, s_2] = MC \tag{27}$$

The information on which we condition is both available and finer than that used in (22), in the sense that a very large number of combinations (s_1, s_2) yields the same sum $s_1 + s_2$. Seen this way, it is intuitive that focusing on aggregate smoking only implies learning an element from a relatively coarse partition, which is of course less insightful than learning an element from one that is further refined.

Note that upon observing both quantities, the difference in preference shifts can be constructed:

$$\mu \equiv \theta_2 - \theta_1 = \beta_1 s_1 - \beta_2 s_2 \tag{28}$$

Using the demand equation and plugging in μ, we find

$$\mathbb{E}[MB|s_1, s_2] = \mathbb{E}[\theta_1|\mu] - \beta_1 s_1 \tag{29}$$

$$= \mu\frac{\mathbb{E}[\mu\theta_1]}{\mathbb{E}[\mu^2]} - \beta_1 s_1 \tag{30}$$

$$= \mu\frac{\rho\sigma_1\sigma_2 - \sigma_1^2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2} - \beta_1 s_1 \tag{31}$$

$$= -\frac{\sigma_2^2\rho - \sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}\beta_1 s_1 - \frac{\sigma_1^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}\beta_2 s_2 \tag{32}$$

The RHS is straightforwardly equated to marginal damages:

$$\gamma(s_1 + s_2) = -\frac{\sigma_2^2\rho - \sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}\beta_1 s_1 - \frac{\sigma_1^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}\beta_2 s_2, \tag{33}$$

which for convenience we rewrite as

$$\delta s_1 + s_2 = 0, \tag{23}$$

with δ given by

$$\delta = \frac{\gamma + \frac{\sigma_2^2\rho - \sigma_1\sigma_2}{\sigma_1^2+\sigma_2^2-2\rho\sigma_1\sigma_2}\beta_1}{\gamma + \frac{\sigma_1^2 - \rho\sigma_1\sigma_2}{\sigma_1^2+\sigma_2^2-2\rho\sigma_1\sigma_2}\beta_2}, \tag{34}$$

and (25) follows. That is, whereas the proof of Proposition 2 in the appendix mechanically derives the aggregate marginal rate of substitution from welfare optimization, here we derive it from a fundamental property: Stabilized Trading with an optimal stabilization rate is not 'just another' regulation rule that is optimized, it is *the* most efficient implementable quantity-based regulation. It equalizes marginal costs and expected marginal benefits given all information available in observed trades.

Figure 1, which plots the optimal stabilization rate as a function of relative uncertainty of smokers and the correlation in their preference shocks, visualizes several interesting properties of the optimal stabilization rate (25). First, as stated in the Proposition, hardly ever will this rate be unity. Thus, Stabilized Trading outperforms traditional Trading in terms of welfare almost always. Second, the optimal stabilization rate may well be negative, meaning higher-than-expected smoking of one smoker translates into higher-than-expected smoking by the other too. Closer inspection reveals such is more likely to occur for strongly positively correlated preferences and very asymmetric uncertainty. This is easily understood: if we do not know the second smoker too well but we do know he is very much alike the first, then smoking should be increased for both, or for none. The strong increase in smoking appetite by the latter is, by their incredible similarity, almost surely followed suit by an equally strong increase in appetite by the former. A negative stabilization rate bears some resemblance with putting negative weights on observations in making (econometric) predictions (see Bunn, 1985; Elliott and Timmermann, 2004; Timmermann, 2006).

Third, the share of preference shifts absorbed by a smokers decreases in the smoker's responsiveness of benefits to smoking, that is, in its slope parameter β. For any adaption of global smoking to shocks, marginal costs change accordingly. Since trade leads to the ex-post equality of individual marginal benefits, and since an optimal mechanism equates individual marginal benefits to global marginal costs, for any realized pair of shifts smoking changes relatively less for the smoker with steeper marginal benefits.

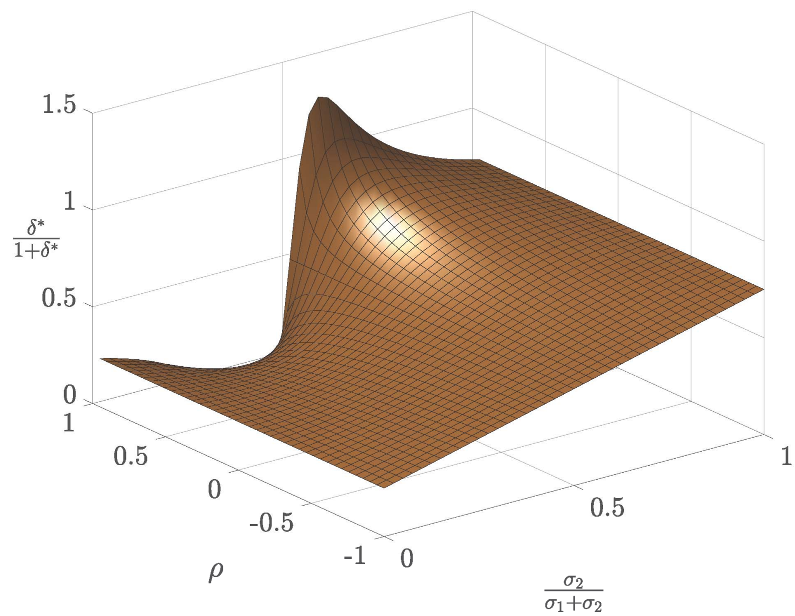

Figure 1: The (normalized) optimal stabilization rate δ^*, as a function of preference-shock correlation ρ and relative uncertainty $\sigma_2/(\sigma_1+\sigma_2)$.

3.6 Asymmetric Uncertainty

Finally, our potentially most interesting observation is that the stabilization rate tends to increase, all else equal, if the uncertainty of smoker 2 increases. This is clearly driven by our choice of definition; had we alternatively set $s_1 + \delta s_2 = 0$, the result would be reversed. The intuitive meaning remains unaffected, though, being that the regulator anchors supply to demand of the best-known smoker, and most variability or flexibility is warranted for the smoker about whom we know least and therefore can learn most.

This intuition is most clearly understood by considering a rather extreme example. Suppose smoker A is so frequent a visitor of this bar that the barkeeper in fact knows A's preferences perfectly and without error, i.e. $\sigma_A = 0$. Assume moreover that this is not true for smoker B, the new smoker in town, meaning $\sigma_B > 0$. The barkeeper thus clearly faces a situation with asymmetric uncertainty about the smokers' preferences.

Suppose now the barkeeper observes trade and so establishes demand. Since the barkeeper knows the preferences of smoker A perfectly, the barkeeper can thus infer the precise marginal benefits of smoker B, as trade will equalize marginal benefits across the smokers. But this, in turn, implies the barkeeper is able to identify the *exact* preferences of smoker B, including, that is, the initially unobserved shift θ_B.

We learn something important here. Without paying attention to an asymmetry of uncertainty, as in the classical argument in favor of trade and markets, upon seeing a flow of cigarettes from smoker B to smoker A all the barkeeper could conclude is preferences had shifted in such a way that the value of an extra cigarette was higher for smoker 1 than for smoker B. The barkeeper may feel happy, for has not trade resulted in an efficient exchange of cigarettes? True, these trades are efficient, but only at the *individual* smoker's level. The new allocation of cigarettes is, in the economist's parlance, constrained Pareto optimal: given a potentially suboptimal aggregate consumption of cigarettes, individual consumption levels are optimal. The barkeeper could have stopped here. However, we argue it could have been known the aggregate amount of smoking was in fact inefficient, a piece of knowledge that could, and maybe should, have been acted upon. Taking into account the fact that smoker A has far less uncertain preferences than smoker B, when seeing a flow of cigarettes from the latter to the former, it can be concluded that the 'aggregate' or 'average' preference for smoking is in fact lower than anticipated. After all, this sale of cigarettes tells us more than simply that our initial expectations about individual preferences were off – it tells us that smoker B has a weaker preference for cigarettes than smoker A. Since the latter's preferences are known with certainty, we can conclude what we may loosely write as $\theta_2 < \mathbb{E}[\theta_2]$. For the economist, this knowledge is enough to conclude that aggregate smoking shall be further restricted. The lesson drawn from this extreme example is of course more generally true and applicable: there is *more* scope for learning about agents whose preferences are *more* uncertain.

Given an optimal stabilization rate, we can solve for the associated level of expected global welfare. We do so in the our Theorem:

Theorem 1. *Stabilized Trading is strictly welfare-superior to Quotas and Trading, with welfare given by:*

$$\Delta^{ST}W = \frac{1}{2}\frac{\beta_1+\beta_2}{\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2}\frac{(1-\rho^2)\sigma_1^2\sigma_2^2}{\sigma_1^2+\sigma_2^2-2\rho\sigma_1\sigma_2}. \tag{35}$$

Proof. In Appendix. Q.E.D.

Since $\Delta^T W > \Delta^{ST} W$, we refer to (12) and conclude that Stabilized Trading moves prices closer to the Social Optimum than traditional Trading. We also see the illustration used above reflected in our theorem. If the barkeeper has perfect information on one smoker, $\sigma_1 = 0$, then all preferences are revealed through trade and no welfare losses occur, $\Delta^{ST}W = 0$. Moreover, it is straightforward to derive that price volatility relative to the *ex ante* price level for both smokers is also lower under Stabilized Trading:

Proposition 3. *Stabilized Trading admits lower price volatility than Trading:*

$$\forall i : \mathbb{E}\left[\left(p^{ST}\right)^2\right] \leq \mathbb{E}\left[\left(p^{T}\right)^2\right]. \tag{36}$$

Proof. In appendix. Q.E.D.

3.7 Symmetric Uncertainty: No Information

The discussion above for the time being focused on the information value of trade conditional on asymmetries in uncertainty. We can further shape our intuition by reversing the argument: There is no information in trade when uncertainties are perfectly symmetric. If we are uncertain about all smokers, but equally so about each, observed aggregate behavior does not allow to learn *anything* about aggregate preferences. For inferences about aggregate preferences to possibly be made, differential uncertainty about smokers is a prerequisite.

Suppose we have a group of N identical individuals, which we split in two groups of size n and $N-n$ respectively. Trade between these groups cannot contain information on aggregate preferences; after all, we do not have any reason to assume that prior knowledge of one (sub-)group was more precise than the prior knowledge of the other (sub-)group (apart from the obvious scaling of uncertainty). Uninformative trade between the groups means that no information on 'aggregate preferences' of the population or market as a whole can be filtered, and consequently that the allocation of cigarettes to the market should be independent of any such observed trades. This independence is achieved precisely when $\delta = 1$.

Formally, let a market be inhabited by n identical independent smokers for whom $\mathbb{E}[\theta_i^2] = \sigma_1^2$, with $i \in \{1, ..., n\}$. We want to describe the market as one representative agent that satisfies

$$-\beta_n s_n + \theta_n = p, \tag{37}$$

and to assess how β_n and σ_n scale with n. Note that s_n is the number of aggregate cigarettes consumed by the market, $s_n = \sum_{i=1}^n s_i$.

For all individuals i within the group (representative agent), we have:

$$-\beta_1 s_i + \theta_i = p, \tag{38}$$

which is simply the demand equilibrium condition (6). Summing over all individuals i and dividing by group size n, we get:

$$-\frac{\beta_1}{n} s_n + \frac{1}{n} \sum_{i=1}^{n} \theta_i = p, \tag{39}$$

which immediately gives

$$\beta_n = \beta_1 / n, \tag{40}$$

$$\mathbb{E}[\theta_n^2] = \sigma_1^2 / n \tag{41}$$

The same thought experiment also shows that for a group of N split into two groups of size n and $N-n$, we have $\beta_n = \beta_1/n, \sigma_n = \sigma_1/n$ and $\beta_{N-n} = \beta_1/(N-n), \sigma_{N-n} = \sigma_1/N - n$.

Plugging all this into equation (25), we see that δ^* continues to be 1 after aggregating individual identical smokers into any two sets of n and $N - n$ identical smokers. This confirms the intuition that no information can be obtained from trade between *symmetrically uncertain* smokers.

Corollary 1. *No aggregate information can be obtained from trade between smokers about whom we are symmetrically uncertain. That is, only trade between smokers with asymmetrically uncertain preferences allows for aggregate information filtering. Labeling* *smokers or groups as '1' and '2':*

$$(\rho = 0 \text{ and } \sigma_1/\beta_1 = \sigma_2/\beta_2) \Rightarrow \delta = 1. \tag{42}$$

4 Prices vs. Quantities Revisited

In this section, we reconsider Marty Weitzman's seminal contribution on the relative efficiency of price and quantity instruments (Weitzman, 1974). As an interpretation, smokers may be considered different countries or jurisdictions, whereas smoking represents greenhouse gas emissions. This relates our model to recent contributions by Mideksa and Weitzman (2019) and Doda et al. (2019).

Thus far, our discussion has focused on quantity instruments only. To make a juxtaposition of prices and quantities possible, we define a price instrument as the policy that solves:

$$\max_{p_1,p_2} \quad \mathbb{E}\left\{B_1(s_1(p_1);\theta_1) + B_2(s_2(p_2);\theta_2) - C(s_1(p_1) + s_1(p_2))\right\} \tag{43}$$

$$\text{s.t.} \quad B_i' = p_i, \tag{44}$$

where the constraint follows from utility-maximizing smokers choosing to smoke until the point where their marginal utility of a cigarette equals the marginal cost of smoking, as represented by the price of cigarettes. In other words, the barkeeper chooses shadow values of smoking (p_1, p_2) to implement (expected) optimal smoking behavior. In the greenhouse gas emission interpretation of the model, these prices or shadow values of emissions equal what Kotchen (2018) calls the Global Social Cost of Carbon. Inefficiencies due to regulation according to the Domestic Social Cost of Carbon do not arise.

Parametric solutions to this comparison are easily obtained but require tedious algebra, which is relegated to the Appendix. They are plotted in Figure 2. The canonical result due to Weitzman (1974) is reproduced in this setup by the indifference plane between prices and trade, since in his original contribution the one-firm-market can be re-interpreted as a market with many small firms that trade until the point at which marginal profits are equal for all, which, when translated back, is precisely our Trade instrument. Stabilized Trade always performs weakly better than Trade, meaning the indifference plane for Stabilized Trade is always below the plane for Trade. When either preference shocks are very strongly positively or negatively related, or uncertainty is heavily asymmetric, Stabilized Trade comes very close to implementing the Social Optimum. As a consequence, it is almost impossible for Prices to outperform Stabilized Trade. Of course, this argument is somewhat unfair, as no optimal responsive version of the classic tax instrument is being considered. One can easily see that for the price-equivalent of our Stabilized Trade instrument, Weitzman's classic result on the comparison between the two is again obtained.

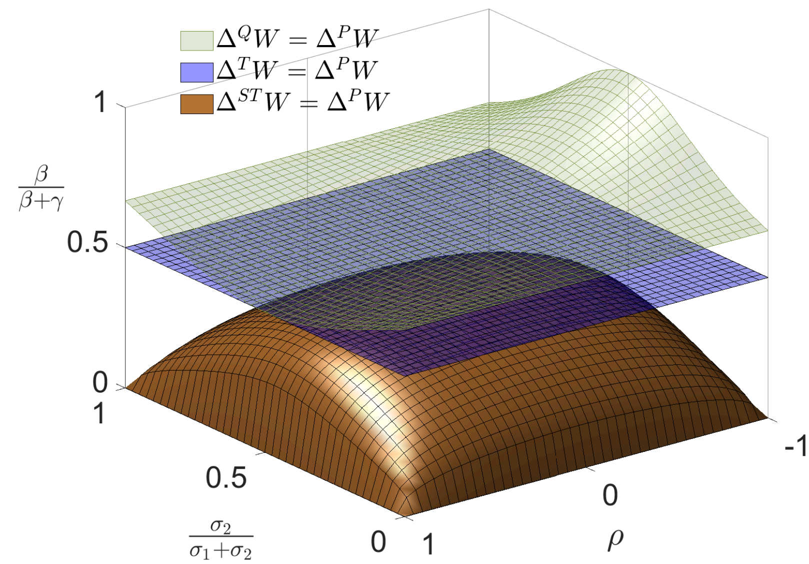

Figure 2: Indifference planes for Quantities, Trade, and Stabilized Trade versus Prices. For parameter values 'below' the surface of a given plane, the respective quantity instrument outperforms prices in terms of welfare.

5 Summary

It has herein been demonstrated that fairly simple and straightforward manipulations of traditional free trade in a market under regulation can yield substantial welfare gains. A key concept used thereto is that of asymmetric uncertainty and the implied differential learning potentials about different regulated agents. Our results can be directly applied to all activities where a global externality creates market imperfections, including but not restricted to international trade of pollutants, such as CO_2. Given the volatility and large sums involved in these newly developed markets, it seems important to develop regulation mechanisms that mitigate those price volatility and reduce welfare losses as a result of suboptimal policy.

References

Baxter, M. and King, R. G. (1999). Measuring business cycles: approximate band-pass filters for economic time series. *Review of economics and statistics*, 81(4):575–593.

Bunn, D. W. (1985). Statistical efficiency in the linear combination of forecasts. *International Journal of Forecasting*, 1(2):151–163.

Doda, B., Quemin, S., and Taschini, L. (2019). Linking permit markets multilaterally. *Journal of Environmental Economics and Management*, 98:102259.

Elliott, G. and Timmermann, A. (2004). Optimal forecast combinations under general loss functions and forecast error distributions. *Journal of Econometrics*, 122(1):47–79.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica: Journal of the Econometric Society*, pages 357–384.

Harvey, A. C. and Trimbur, T. M. (2003). General model-based filters for extracting cycles and trends in economic time series. *Review of Economics and Statistics*, 85(2):244–255.

Kotchen, M. J. (2018). Which social cost of carbon? a theoretical perspective. *Journal of the Association of Environmental and Resource Economists*, 5(3):673–694.

Lucas, R. E. (1978). Asset prices in an exchange economy. *Econometrica: Journal of the Econometric Society*, pages 1429–1445.

Makarov, I. and Rytchkov, O. (2012). Forecasting the forecasts of others: Implications for asset pricing. *Journal of Economic Theory*, 147(3):941–966.

Mideksa, T. K. and Weitzman, M. L. (2019). Prices versus quantities across jurisdictions. *Journal of the Association of Environmental and Resource Economists*, 6(5):883–891.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The review of economic studies*, 38(2):175–208.

Montgomery, W. D. (1972). Markets in licenses and efficient pollution control programs. *Journal of Economic Theory*, 5(3):395–418.

Moore, D. A. and Healy, P. J. (2008). The trouble with overconfidence. *Psychological review*, 115(2):502.

Mullainathan, S. (2002). A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, 117(3):735–774.

Pizer, W. A. (2002). Combining price and quantity controls to mitigate global climate change. *Journal of public economics*, 85(3):409–434.

Roberts, M. J. and Spence, M. (1976). Effluent charges and licenses under uncertainty. *Journal of Public Economics*, 5(3-4):193–208.

Talmon, R. and Coifman, R. R. (2013). Empirical intrinsic geometry for nonlinear modeling and time series filtering. *Proceedings of the National Academy of Sciences*, 110(31):12535–12540.

Timmermann, A. (2006). Forecast combinations. *Handbook of economic forecasting*, 1:135–196.

Townsend, R. M. (1983). Forecasting the forecasts of others. *Journal of Political Economy*, 91(4):546–588.

Weitzman, M. L. (1974). Prices vs. quantities. *Review of Economic Studies*, 41(4):477–491.

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

A Derivations and Proofs

DERIVATION OF (19):
Combining the definition with the firms' FOCs, (6), we find the change in permit use by region:

$$\Delta^T s_1 = \frac{\theta_1 - \theta_2}{\beta_1 + \beta_2} \tag{45}$$

$$\Delta^T s_2 = \frac{\theta_2 - \theta_1}{\beta_1 + \beta_2}. \tag{46}$$

PROOF OF PROPOSITION 1:

Proof. For the first part, note that Trading outperforms Quantities if and only if the following condition is satisfied:

$$\frac{\beta_2^2\sigma_1^2 + \beta_1^2\sigma_2^2 + 2\beta_1\beta_2\rho\sigma_1\sigma_2}{\beta_1+\beta_2} < (\gamma+\beta_2)\sigma_1^2 + (\gamma+\beta_1)\sigma_2^2 - 2\gamma\rho\sigma_1\sigma_2$$

$$\Longleftrightarrow$$

$$\frac{2\rho\sigma_1\sigma_2}{\sigma_1^2+\sigma_2^2} < \frac{\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2}{\gamma\beta_1+\gamma\beta_2},$$

which is always true. Q.E.D.

PROOF OF PROPOSITION 2; DERIVATION OF (25):
Regional and global deviations from Socially Optimal permit use are given by:

$$\Delta^{ST}s_1 = \frac{\beta_2}{\beta_1+\delta\beta_2}\frac{[\delta\beta_2-\gamma(1-\delta)]\theta_1+[\beta_1+\gamma(1-\delta)]\theta_2}{\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2} \tag{47}$$

$$\Delta^{ST}s_2 = \frac{\beta_1}{\beta_1+\delta\beta_2}\frac{[\delta\beta_2-\gamma(1-\delta)]\theta_1+[\beta_1+\gamma(1-\delta)]\theta_2}{\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2} \tag{48}$$

$$\Delta^{ST}Q = \frac{\beta_1+\beta_2}{\beta_1+\delta\beta_2}\frac{[\delta\beta_2-\gamma(1-\delta)]\theta_1+[\beta_1+\gamma(1-\delta)]\theta_2}{\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2}. \tag{49}$$

Define

$$\xi := \frac{\beta_1+\gamma(1-\delta)}{\beta_1+\delta\beta_2} \implies 1-\xi := \frac{\delta\beta_2-\gamma(1-\delta)}{\beta_1+\delta\beta_2}. \tag{50}$$

Welfare losses can now be written as:

$$\begin{aligned}\Delta^{ST}W &= \frac{1}{2}\frac{\gamma(\beta_1+\beta_2)^2+\beta_1^2\beta_2+\beta_1\beta_2^2}{(\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2)^2}\mathbb{E}\left[(1-\xi)\theta_1+\xi\theta_2\right]^2 \\ &= \frac{\beta_1+\beta_2}{2}\frac{(1-\xi)^2\sigma_1^2+\xi^2\sigma_2^2+2\xi(1-\xi)\rho\sigma_1\sigma_2}{\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2}.\end{aligned} \tag{51}$$

If for notational convenience, we define:

$$\psi := \frac{1}{2}\frac{\beta_1+\beta_2}{\gamma\beta_1+\gamma\beta_2+\beta_1\beta_2}, \tag{52}$$

it is straightforward to derive:

$$\frac{\partial}{\partial\xi}\frac{\Delta^{ST}W}{\psi} = 2\xi\sigma_2^2 - 2(1-\xi)\sigma_1^2 + 2(1-\xi)\rho\sigma_1\sigma_2 - 2\xi\rho\sigma_1\sigma_2. \tag{53}$$

The welfare-maximizing ξ^* therefore satisfies:

$$\xi^* = \frac{\sigma_1^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}. \tag{54}$$

From the definition of ξ, the optimal stabilization rate δ^* follows:

$$\delta^* = \frac{(\beta_1 + \gamma)[\sigma_2^2 - \rho\sigma_1\sigma_2] + \gamma[\sigma_1^2 - \rho\sigma_1\sigma_2]}{(\beta_2 + \gamma)[\sigma_1^2 - \rho\sigma_1\sigma_2] + \gamma[\sigma_2^2 - \rho\sigma_1\sigma_2]}, \tag{55}$$

as stated.

PROOF OF THEOREM 1:

Proof. Plugging (54) in (51), we find:

$$\begin{aligned}
\frac{\Delta^{ST}W}{\psi} &= \left[\frac{\sigma_2^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}\right]^2 \sigma_1^2 + \left[\frac{\sigma_1^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}\right]^2 \sigma_2^2 \\
&+ \left[\frac{\sigma_2^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}\right]\left[\frac{\sigma_1^2 - \rho\sigma_1\sigma_2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2}\right]\rho\sigma_1\sigma_2 \\
&= \frac{(1-\rho^2)\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2} \\
&\Longrightarrow \\
\Delta^{ST}W &= \frac{1}{2}\frac{\beta_1 + \beta_2}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2}\frac{(1-\rho^2)\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2},
\end{aligned}$$

as stated. This is strictly lower than the welfare loss under traditional Trading if and only if:

$$\begin{gathered}
2\Delta^T W - 2\Delta^{ST}W \geq 0 \\
\Longrightarrow \\
\frac{1}{\beta_1 + \beta_2}\frac{\beta_2^2\sigma_1^2 + \beta_1^2\sigma_2^2 + 2\beta_1\beta_2\rho\sigma_1\sigma_2}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2} - \frac{\beta_1 + \beta_2}{\gamma\beta_1 + \gamma\beta_2 + \beta_1\beta_2}\frac{(1-\rho^2)\sigma_1^2\sigma_2^2}{\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2} \geq 0 \\
\Longrightarrow \\
(\sigma_1^2 + \sigma_2^2 - 2\rho\sigma_1\sigma_2)(\beta_2^2\sigma_1^2 + \beta_1^2\sigma_2^2 + 2\beta_1\beta_2\rho\sigma_1\sigma_2) - (1-\rho^2)(\beta_1^2 + \beta_2^2 + 2\beta_1\beta_2)\sigma_1^2\sigma_2^2 \geq 0 \\
\Longrightarrow \\
[(\beta_2\sigma_1^2 - \beta_1\sigma_2^2) + (\beta_1 - \beta_2)\rho\sigma_1\sigma_2]^2 \geq 0,
\end{gathered}$$

which is always true. Q.E.D.

PROOF OF PROPOSITION 3:

Proof. We derived quantity derivations under both policies. Prices are equal in both regions, so without loss of generality we can solve for price deviations in region 1:

$$\Delta^T p_1 = \frac{\beta_2\theta_1 + \beta_1\theta_2}{\beta_1 + \beta_2}$$
$$\Delta^{ST} p_1 = \frac{\delta\beta_2\theta_1 + \beta_1\theta_2}{\beta_1 + \delta\beta_2}.$$

Thus:

$$\mathbb{E}\left[\left(\Delta^T p\right)^2\right] = \frac{\beta_2^2\sigma_1^2 + \beta_1^2\sigma_2^2 + 2\beta_1\beta_2\rho\sigma_1\sigma_2}{\beta_1^2 + \beta_2^2 + 2\beta_1\beta_2}$$
$$\mathbb{E}\left[\left(\Delta^{ST} p\right)^2\right] = \frac{\delta^2\beta_2^2\sigma_1^2 + \beta_1^2\sigma_2^2 + 2\delta\beta_1\beta_2\rho\sigma_1\sigma_2}{\beta_1^2 + \delta^2\beta_2^2 + 2\delta\beta_1\beta_2}.$$

Writing these out, we obtain:

$$\mathbb{E}\left[\left(\Delta^{ST} p\right)^2\right] < \mathbb{E}\left[\left(\Delta^T p\right)^2\right] \iff (\delta - 1)\left[\beta_2\left(\sigma_1^2 - \rho\sigma_1\sigma_2\right) - \beta_1\left(\sigma_2^2 - \rho\sigma_1\sigma_2\right)\right] < 0.$$

We now invoke Proposition 2 and establish that this condition is always satisfied. Q.E.D.