Regulating Stock Externalities

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

Tilburg University

October 2, 2020

Abstract

We develop a dynamic framework for the regulation of stock externalities under incomplete and asymmetric information. For the case of quadratic costs and benefits, optimal policies are derived and ordered explicitly, and we characterize convergence to the first best when there are many periods. In a more general framework, we define and prove existence of optimal regulation as the implementation of a welfare maximization program conditional only on informational constraints and instrument class.

JEL codes: D47, D62, D82, H23, L51

Keywords: asymmetric information, regulatory instruments, policy updating, dynamic regulation

1 Introduction

Stock externalities are the unintended byproduct of cumulative economic activity in a market over the course of time. To a planner burdened with the control of this market, the question arises whether traditional tax or quota instruments could – or indeed should – be adjusted to the dynamic properties intrinsic to the stock externality and, if so, which instrument performs best. This paper studies these questions for environments with asymmetric information about market fundamentals.

The issue of instrument choice touches upon an influential literature that originates with Weitzman (1974). Papers in this tradition usually take an exogenously given set of

*Gerlagh: r.gerlagh@uvt.nl. Heijmans: r.j.r.k.heijmans@uvt.nl

policy instruments and assess their relative merits (*c.f.* Hoel and Karp, 2001; Zhao, 2003; Newell and Pizer, 2003; Fell et al., 2012; Weitzman, 2019; Mideksa and Weitzman, 2019; Pizer and Prest, 2020; Heutel, 2020). On top of these comparisons, this literature also proposes policy refinements by formulating intuitive policies that are optimized over their exogenous set of parameters. In this spirit, Kling and Rubin (1997), Newell et al. (2005), and Pizer and Prest (2020) allow the planner to depreciate or top up banked – unused and saved for future use – allowances. Similarly, Yates and Cronshaw (2001) consider banking with a discount rate for allowances (a kind of intertemporal trading ratio à la Holland and Yates (2015)). Newell et al. (2005) and Lintunen and Kuusela (2018) discuss adjusting quota in response to the quantity of outstanding allowances. Finally, Karp and Traeger (2017, 2018) study a cap on emissions that changes in response to aggregate private information inferred from price signals.

While we too propose and compare new instruments, ours is not a typical prices vs. quantities exercise. Instead of formulating intuitive policies based on some combination of existing policies, we *derive* them from primitives of the problem like market fundamentals and the externality under investigation. Rather than maximize welfare over a policy's exogenously given parameters, we maximize over policies.

Our approach means that the planner considers key characteristics of the externality carefully. For example, as discussed in Hasegawa and Salant (2014), when regulating a stock externality with a cap-and-trade program the planner can afford to be less strict on compliance in any given period as long as overall compliance is guaranteed. The planner can be lenient in this way, but why stop there? At least in theory it is possible that further improvements of static policy instruments exist. The driving force behind any such improvement is the fact that observed outcomes in the market provide valuable information (i.e. signals, see Harstad and Eskeland, 2010) to the planner, who, in response, can adapt future policies. In this spirit, Pizer and Prest (2020) develop a dynamically adjusted quantity instrument while Heutel (2020) extends static taxes over time.[1] The basic idea of these novel instruments is the same: a well-devised instrument allows the planner to extract private information from the market and uses this information to make future regulation more efficient. If done well, such policy updating may be very efficient.

Yet while a dynamic framework creates opportunities for the planner to learn and update policies accordingly, it also breeds problems not encountered in static environments. The distinction between *stock* and *flow* externalities cuts to the core of these. At the risk

[1]Pizer and Prest also consider policy uncertainty from which, though relevant, both Heutel and this manuscript abstract away.

of oversimplification, a flow externality is the byproduct of economic activity at a given point in time. A stock externality, in comparison, is the consequence of economic activity *over the course of time*. The difference matters. For stock externalities that we consider, past activity does affect the marginal externality of today's business.

What we do can now be summarized in three simple steps: (1) we build a model with asymmetric information about market fundamentals where productive activity causes a pure stock externality, (2) we show that the planner can construct instruments that drain *all* private information from the market apart from information that is only revealed through the market at the end of the last period, and (3) we derive pure price and quantity instruments that use this private information and implement the best achievable welfare levels very generally. The second step suggests a regulatory framework can approach the first-best by decreasing its regulatory time-windows. Indeed, we describe and order instruments with respect to convergence to first-best. Note that our last step is fundamentally constructive. We start from optimal social welfare and copy its conditions so that the market optimally adjusts production to changing market fundamentals. That is, we require that profit-maximizing firms are *always* incentivized to produce optimal amounts when regulated. It turns out that such instruments can always be found and are intimately related.

We apply our framework to climate change, caused mainly by the cumulative stock of emitted CO_2 (Kolstad, 1996; Allen et al., 2009; Golosov et al., 2014). It is by now widely accepted that the only way to avoid severe climate change is a large-scale reduction in global CO_2 emissions. There are fundamental uncertainties inherent to the problem, though. One will be the focus of attention here: the aggregate costs of abatement. While some fear that any emission reduction threatens thousands of jobs, others foresee that a reduced use of fossil fuels has no significant bearing on neither economic growth nor employment. We develop policies that optimally respond when the market learns about these uncertainties. The emphasis on abatement costs is allegorical; other types of uncertainty resolved before the closing of markets, and relevant for welfare and thus regulation could be used in their stead.

Our approach may seem comparable to both Pizer and Prest (2020) and Heutel (2020). The similarities are mostly superficial. Like Pizer and Prest, we develop a dynamically updated quantity instrument. Like Heutel, we introduce a new dynamic price instrument. The crucial difference is our choice of externality. While Pizer and Prest (2020) and Heutel (2020) study the dynamic regulation of a flow externality, our focus is on stock externalities.

This distinction is important and may even reverse the ordering of instruments.[2] Details of the externality are crucial in a dynamic framework. The implicit assumption of a flow externality in the climate context is that historic emissions of greenhouse gases do not, at all, affect the marginal damages caused by climatic consequences of emissions today. We think this assumption does not square well with the natural science of climate change (see for example Allen et al., 2009).

In the application, our proposed policy instruments can be thought of as an advanced cap-and-trade system. The planner allows firms to freely bank and borrow allowances between periods. Under our quantity instrument, the planner adapts future injections of new allowances in response to the amount of periodic over- or under-compliance. Importantly, we do *not* propose that banked emissions allowances be apprenticed or depreciated between periods![3] This is a very subtle, but also very important difference. The argument boils down to the crucial distinction between flow and stock externalities. When climate change is modeled as a pure flow externality, an extra ton of emissions last year may, in principle, have a different effect on climate change than an extra ton of emissions this year. If there is reason to believe this is true, that is a strong argument to appreciate or depreciate banked emission allowances. Yet when climate change is modeled as a pure stock externality the marginal climate damage from an extra ton of emissions is exactly the same whether they are emitted this or any other year (because all that matters is total emissions over time). An efficient instrument therefore treats the marginal climate damage of emissions in any period as equal and should not touch banked allowances. To nonetheless make aggregate emissions responsive to market fundamentals, new injections of emission allowances can instead be adjusted. Our price instrument deviates in the last period, when it does not set a quota but fixes the emission *tax* based on previous demand for allowances. Either of our two instruments may constitute an efficient means of a country or group of countries to implement emissions reductions.

2 Model

2.1 Benefits, Costs, and Welfare

Consider a two-period world (relaxed in Section 2.8) and a representative profit-maximizing firm producing a homogeneous good. At every time $t \in \{1, 2\}$, producing an amount $\widetilde{q}_t$

[2]As a point in case, Weitzman's 2019 surprising ordering of instruments originates in his flow externality model and is reversed, as we show, for a stock externality.

[3]Compare Yates and Cronshaw (2001) and Rubin (1996).

of the good yields benefits $B_t(\widetilde{q}_t;\theta_t)$ to the firm (we abstract away from broader social benefits of production). The parameter θ_t captures market fundamentals observed by the firms at time t but not known to the planner. We conveniently write $B'_t = \partial B_t/\partial \widetilde{q}_t$ and normalize θ_t such that $\partial B_t/\partial\theta_t = 1$. We assume concave benefits $B''_t < 0$. It is common knowledge that $\mathbb{E}[\theta_t] = 0$, $\mathbb{E}[\theta_t^2] = \sigma_t^2$, and $\mathbb{E}[\theta_1\theta_2] = \rho\sigma_1\sigma_2$. In a market with free and competitive trade of production rights, an equilibrium price will emerge, denoted $\widetilde{p}_t = B'_t$. For the purposes of our study, the variance σ_t^2 provides a natural measure of uncertainty in the market.

Cumulative production imposes a cost on society in the form of a stock externality, given by $C(\widetilde{q}_1 + \widetilde{q}_2)$. We assume convex costs, $C' > 0, C'' > 0$.

The planner's problem is to find policies such that production levels $\widetilde{q}_1$ and $\widetilde{q}_2$ maximize welfare:

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

The timing of regulation and equilibrium follows these stages:

1. The planner chooses a policy instrument;
2. Firms observe first-period ($t = 1$) fundamentals θ_1;
3. First-period markets open and production $\tilde{q}_1$ is realized;
4. Firms observe second-period ($t = 2$) fundamentals θ_2;
5. Second-period markets open and production $\tilde{q}_2$ is realized;
6. Costs due to aggregate production $\widetilde{Q} = \sum_t \widetilde{q}_t$ are realized.

Note that market outcomes are public information; i.e. they are observed by the planner. With complete information on θ_t, the fully knowledgeable planner can set these quantities $\widetilde{q}_1, \widetilde{q}_2$ directly or else charge a price on production that will make the profit-maximizing firm produce the same quantities, and these two instruments are perfectly equivalent, see Montgomery (1972). However, this formal equivalence between instruments breaks down once we introduce an informational disparity, captured here by θ_t (Weitzman, 1974).

2.2 Optimal Response

We study an environment with asymmetric information *and* imperfect foresight. Because of the unpredictable element in future market conditions, the ex post first best is unattainable: it would require the planner to be aware of firms' private knowledge about market fundamentals even before firms themselves are. Instead, the best instrument a planner could aim for is one that reacts to any innovations in market fundamentals as soon as they are revealed to the firms. Since such a hypothetical instrument responds optimally to new information, we call it the Optimal Response.

In terms of our model, the Optimal Response determines the cap on emissions in any period t only after θ_t, the market fundamentals in period t, has been drawn. It sets $\tilde{q}_1$ and $\tilde{q}_2$ that implement:

$$\max_{\tilde{q}_1} \mathbb{E}_1 \max_{\tilde{q}_2} \mathbb{E}_2 W(\tilde{q}_1, \tilde{q}_2; \theta_1, \theta_2), \tag{2}$$

where $\mathbb{E}_t$ is shorthand for the expected value of W conditional on θ_s for all $s \leq t$. (Note that this instrument is equivalent to one where prices are chosen conditional on market fundamentals).

While the Optimal Response is a hypothetical instrument, it provides a useful benchmark for policy performance. As we shall show, a smart choice of pure price or quantity instrument allows the planner to implement the Optimal Response solution in all regulatory periods but the last. When there are many periods and each period is relatively short (see Section 2.8), this result is remarkably strong. Simple pure price or quantity instruments suffice to let the planner implement welfare levels almost as though there were no asymmetric information. For all but the last period, complicated and multi-dimensional hybrid instruments (Roberts and Spence, 1976; Weitzman, 1978; Pizer, 2002; Abrell and Rausch, 2017) cannot do better than our pure price and quantity instruments.

2.3 Linear-Quadratic Specification

The main body of our analysis will focus on a simplified model where benefits and costs are linear-quadratic in emissions. This simplest possible case will allow us to derive our new instruments – Responsive Quotas and Endogenous Taxes – constructively and in a precise parametric form. This formulation also permits an intuitive implementation of either instrument. Section 3 briefly returns to the general model and is devoted to proving the existence and implementability of Responsive Quotas and Endogenous Taxes generally.

Let marginal benefits and costs in period t be of the form

$$B'_t(q_t) = p^* - \beta(\widetilde{q}_t - q^*) + \theta_t, \tag{3}$$

$$C'(q_1 + q_2) = p^* + \gamma(\widetilde{q}_1 + \widetilde{q}_2 - Q^*). \tag{4}$$

Note that we take the *average* of production q_t for cumulative production Q. This adaption facilitates a common interpretation for marginal costs γ independent of the number of periods (see Section 2.8).

For convenience, we normalize our notation such that variables q_t and p_t denote deviations from the ex-ante expected optimum: $p_t \equiv \widetilde{p}_t - p^*$, and similarly for $q_t \equiv \widetilde{q}_t - q^*$. In a competitive market, production is so allocated that prices satisfy:

$$p_t = -\beta q_t + \theta_t, \tag{5}$$

which is a first-order condition for profit-maximization by firms.

For simplicity, we start with the 2-period case. We assume that fundamentals θ_t follow an AR(1) process according to:

$$\theta_2 = \alpha\theta_1 + \mu, \tag{6}$$

with commonly known $\alpha \in [-1, 1]$ and μ white noise, so that $\sigma_2^2 = \alpha^2\sigma_1^2 + \sigma_\mu^2$, and $\rho = \alpha\sigma_1/\sigma_2$.

2.4 Classic Banking and the Waste of Information

Before delving into our new instruments, we quickly revisit a standard policy for dynamic cap-and-trade systems: banking (and borrowing). Under such a policy, the planner allocates an amount of emissions allowances to the market in each period but is lenient with respect to periodic compliance, as long as aggregate compliance is safeguarded (Hasegawa and Salant, 2014). This is called banking or bankable quantities because unused allowances can be "banked" for future use. In our notation, the planner sets $Q = 0$, the ex ante expected optimal stock of emissions, while firms choose their periodic emissions levels q_t subject to the constraint that $Q = q_1 + q_2$. Since the market is still free to choose $q_1 = q_2 = 0$ but is not required to do so, a basic argument establishes right away that banking always outperforms fixed periodic quantities.

Banking contains information. The decision to use an extra emission allowance in one period at the cost of emissions in the other signals some of the market's private information to the planner. To see why, consider the following simple example. Let market fundamentals be imperfectly persistent over time, i.e. $\alpha < 1$, and assume that firms decide to bank a positive amount of allowances for use in the second period, i.e. $q_1 < 0$. The planner then learns that $\theta_1 < 0$: firms maximize expected profits, which means emissions in each period are chosen so that $p_1 = \mathbb{E}p_2$, or $\theta_1 - \beta q_1 = \alpha\theta_1 - \beta q_2$, which (for $q_2 = -q_1$) is consistent with $q_1 < 0$ if and only if $\theta_1 < 0$.

But if $\theta_1 < 0$, the initial (aggregate) allocation of $Q = 0$ allowances is too loose and the planner knows it. After the first period market has cleared, the planner who implements a pure banking policy is stuck with a known-to-be inefficient allocation, forced to disregard the valuable information that the market's banking decision has made readily available. Surely the planner can do better?

2.5 Responsive Quotas

We first develop our optimal pure quantity instrument. This instrument resembles the many cap-and-trade systems operative across the globe to reduce greenhouse gas emissions (including EU ETS, RGGI, UK ETS, Chine ETS, South-Korean ETS, California ETS), though with an important modification: new permit injections are a function of the outstanding amount of allowances banked. We call this policy Responsive Quotas.

Starting from classic cap-and-trade, we show how the planner can construct a policy that filters all private information about fundamentals from the market and, exploiting that firms maximize profits and anticipate the planner's response to any observed first-period behavior, implements the ex post efficient level of emissions in the first period. Formally, this instruments yields the following welfare maximization program:

$$\max_{q_1, q_2} \quad \mathbb{E}_1 W(q_1, q_2), \tag{7}$$

that is, emissions in *both* periods are determined only after market fundamentals in the first period are observed by the firms. Much different from classic banking, the total cap on emissions is endogenous to market fundamentals as revealed through banking under a Responsive Quotas regime. We may therefore define the planner's policy-response R such that the $Q = R(q_1)$, i.e. the function R translates first-period emissions q_1 into an endogenous aggregate cap on emissions Q. The problem of our planner is then to find an optimal response function R.

Since firms maximize expected profits, after observing θ_1 they will choose q_1 such that $p_1 = \mathbb{E}_1 p_2$, or $\theta_1 - \beta q_1 = \alpha\theta_1 - \beta(R(q_1) - q_1)$, which uses that firms anticipate the planners policy of setting $Q = R(q_1)$ and the AR(1) development of market fundamentals. The planner, in turn, wants to maximize welfare and therefore equates (expected) marginal benefits to marginal climate damages, given by $\gamma(q_1 + q_2) = \gamma R(q_1)$. An optimal response function R^* therefore solves:

$$\theta_1 - \beta q_1 = \alpha\theta_1 - \beta(R^*(q_1) - q_1) = \gamma R^*(q_1), \tag{8}$$

for all first-period fundamentals θ_1. See Section 3 for the existence result of an optimal response function R^* for general costs and benefits. In our linear model, the optimal response function R^* specifies a simple linear relationship:

$$Q = \delta^* q_1, \tag{9}$$

where δ^*, the optimal response rate, is given by:

$$\delta^* := \frac{Q}{q_1} = \frac{(1+\alpha)\beta}{(1-\alpha)\gamma + \beta}. \tag{10}$$

Note that the response of cumulative allowances equals the injection of allowances in the second period. Looking at (9), Repsonsive Quotas coincides with a standard banking and borrowing policy when $\delta^* = 0$. From (10), this will be the case only if marginal damages rise very sharply with emissions ($\gamma \to \infty$) or if marginal benefits are constant ($\beta = 0$). In all other cases, Responsive Quotas strictly outperforms standard banking and borrowing.

The optimal response rate is increasing in the persistence α of market fundamentals. The more fundamentals are expected to persist, the likelier it becomes that an increase in the marginal value of emissions in the first period is matched in the second. If $\alpha = -1$, fundamentals are perfectly negatively correlated, and any first-period decrease in demand offsets an equal increase in second-period demand; there is no reason to adjust the cap, $\delta^* = 0$. At the other extreme, if fundamentals are perfectly and positively correlated, a first-period decrease in demand is matched by an equal decrease in second-period demand; the adjustment of the cap doubles the observed adjustment of first period demand, $\delta^* = 2$.

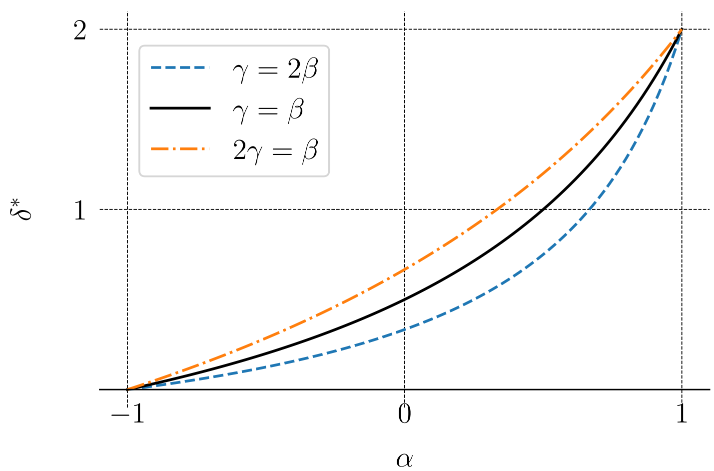

Figure 1: Optimal Response Rate δ^*, for different ratios γ/β, dependent on the correlation between fundamentals α.

If we plug the optimal R^* (or δ^*) back into the planner's welfare maximization problem, we see that a Responsive Quotas policy forces profit maximizing firms to choose first- and second-period emission levels that are ex post efficient with respect to first-period market fundamentals. By exploiting rational firms' anticipation of the planner's policy updating through R^*, it is almost as though the first period is retro-actively regulated, after it has cleared. This is also evidenced by the expected welfare losses under an Responsive Quotas policy relative to the ex post social optimum, which derive solely from unforeseen innovations in second-period market fundamentals:

$$\mathbb{E}W^{OR} - \mathbb{E}W^{RQ} = \frac{1}{4}\frac{1}{\beta+\gamma}\sigma_\mu^2. \tag{11}$$

We conclude our discussion of Responsive Quotas with two key observations. First, depreciating or topping up privately banked allowances à la Yates and Cronshaw (2001) and Kling and Rubin (1997) will not work for a pure stock externality. The stabilization rate δ^* *cannot* be interpreted as an intertemporal trading ratio for emission allowances. An efficient policy lets firms internalize the marginal damage caused by its emissions. But for a stock externality, the marginal damage is the same in each period. Firms' decisions to bank allowances should therefore be driven exclusively by (expected) market fundamentals. An intertemporal trading ratio different from 1 distorts this tradeoff.

Second, a direct comparison of our Responsive Quotas with Pizer and Prest's 2020 optimal dynamic quantities is not necessary. Our approach has been constructive: we let the structure of our problem dictate the ideal quantity instrument. The instrument implements first-best emission levels in both periods if there are no innovations (i.e. $\mu = 0$) in market fundamentals. If Pizer and Prest's flow externality instrument were optimal for stock externalities as well, our method would have reproduced it. Since it did not, Responsive Quotas outperform Pizer and Prest's quantity instrument for regulating stock externalities (and theirs outperforms Responsive Quotas for flow externalities).

2.6 Endogenous Taxes

Responsive Quotas is the optimal quantity instrument to regulate a pure stock externality. We can similarly devise an optimal price policy. This is the exercise undertaken here. We call our new instrument Endogenous Taxes.

If the planner taxes emissions at a rate of p_1 in the first period, then after observing θ_1 profit-maximizing firms choose q_1 so that $p_1 = \theta_1 - \beta q_1$. Since the chosen level of emissions reveals first-period market fundamentals θ_1, the planner can adjust the second-period tax according to some tax-response function $p_2 = T(q_1)$. As for Responsive Quotas, the planner's problem is to find an optimal tax-response function T^* that maximizes welfare conditional on first-period market fundamentals θ_1. In our model's language, the tax-response function T^* should implement the solution to:

$$\max_{p_1,p_2} \quad \mathbb{E}_1 W(q_1(p_1), q_2(p_2)), \tag{12}$$

In our linear model, the optimal tax-response function T^* is linear with slope τ^*:

$$p_2 = \tau^* q_1. \tag{13}$$

But, to implement the best price instrument, the first-period price must autonomously adjust to the fundamentals θ_1, so that the optimal first-order condition $p_1 = \mathbb{E}_1 p_2$ is satisfied. The regulator can achieve this feat by allocating allowances, fixing the second-period price to its optimal level based on its first-period information, and let free banking and borrowing link the markets. A recursive tax policy cannot reach this outcome, which is why we label it 'Endogenous Taxes'. Under this construction, the optimal price response becomes

$$\tau^* = \frac{(1+\alpha)\beta\gamma}{(1-\alpha)\gamma + \beta}. \tag{14}$$

Upon closer inspection of (10) and (14), it turns out the Responsive Quota and Endogenous Taxes instrument are fundamentally related. We observe that

$$\tau^* = \gamma \cdot \delta^*. \tag{15}$$

The normalized (i.e. unit-less) optimal updating rule τ^*/γ coincides with δ^*, the optimal stabilization rate for Responsive Quotas. While perhaps surprising at first, this is in fact intuitive. If first-period emissions are q_1, then the aggregate cap is adjusted to $\delta^* q_1$ under a RQ regime, and the adjustment of marginal damages is therefore $\gamma(\delta^* q_1)$. By construction, this response is efficient in expectations, so $\Delta MB_1 = \Delta \mathbb{E}_1 MB_2 = \Delta MC = \gamma\delta^* q_1$, i.e marginal benefits in the first period equal expected marginal benefits (no innovations) in the second period, which equal marginal costs. But if $\Delta MB_2 = \gamma\delta^* q_1$ is in expectations efficient, then a *price* instrument should tax second-period emissions at a rate $p_2 = \gamma\delta^* q_1$ also. Since Endogenous Taxes is defined as $p_2 = \tau^* q_1$, see (13), it follows that a welfare-maximizing planner chooses $\tau^* = \delta^* \cdot \gamma$.

The expected welfare losses under an Endogenous Taxes regime derive solely from unforeseen innovations in the second period:

$$\mathbb{E}W^{OR} - \mathbb{E}W^{ET} = \frac{1}{4}\left(\frac{\gamma}{\beta}\right)^2 \frac{1}{\beta+\gamma}\sigma_\mu^2 \tag{16}$$

Note that our approach has been constructive: we let the fundamentals of our problem dictate the ideal price instrument. If Heutel's (2020) Bankable Prices were the optimal price instrument for stock externalities as well, our method would have told us so. Since it did not, Endogenous Taxes outperform Heutel's Bankable Prices for regulating stock externalities. Another way to see this to observe that Endogenous Taxes implements the same price in both periods for any innovation in fundamentals, and first-best emission levels in both periods if there are no innovations in market fundamentals. Bankable Prices does not preserve those two properties, which are efficient for stock externalities since marginal damages are period-independent.

2.7 Prices vs. Quantities

If there are no innovations in second-period market fundamentals, both Responsive Quotas and Endogenous Taxes implement the first best level of emissions in either period and neither instrument is favored over the other (Montgomery, 1972). When second-period innovations are possible, the instruments only deviate with respect to the effect of these innovations in the second period. The relative performance of our instruments is therefore determined solely by how unforeseen second-period fundamentals affect emissions in the second period. This boils down to the classic choice problem studied by Weitzman (1974).

Proposition 1 (Weitzman for Stock Externalities)**.** *Endogenous Taxes outperform Responsive Quotas in terms of welfare if and only if $\beta \geq \gamma$:*

$$\mathbb{E}W^{ET} \geq \mathbb{E}W^{RQ} \iff \beta \geq \gamma. \tag{17}$$

For constant marginal damages, $\gamma = 0$, the first best is relatively straightforwardly implemented by setting the price to match those. When knowledge about the true marginal damages evolve over time, a generalized version of the Endogenous Taxes policy uses all information available to the market to set the market price at the expected value for marginal damages, at each point in time. Importantly, a stock externality still requires that current prices equal expected prices, which is the core property of Endogenous Taxes not upheld by other instruments in the literature.

The jury is still out whether climate damages are convex ($\gamma > 0$) or proximately linear ($\gamma = 0$) in greenhouse gas emissions.[4] In either case, however, Endogenous Taxes is favored over Responsive Quotas if the number of periods N becomes large, as we see below.

2.8 A Finer Grid

Increasing the number of periods to $N > 2$, we also increase the number of market operations that can be regulated, effectively using each trading opportunity as an instrument. Much like in the two-period model, this allows the planner in every period but the last to implement emission levels that are first best given that period's market fundamentals. With more and more periods, the relative effect of the final period (where asymmetric information continues to plague the planner) on welfare becomes smaller and smaller and the welfare performance of our instruments increases.

Consider a time window of unit length, $t \in [0, 1]$, divided in N periods of equal duration $\varepsilon = 1/N$, so that the n^{th} period ($n \in 1, ..., N$) covers the interval $[(n-1)\varepsilon, n\varepsilon]$. Benefits and costs are given by[5]

$$B'_n = \theta_n - \beta q_n, \tag{18}$$

$$C' = \gamma Q = \frac{\gamma}{N} \sum_{n=1}^{N} q_n, \tag{19}$$

while demand shocks follow the AR(1) process

$$\theta_n = \alpha^{1/(N-1)} \theta_{n-1} + \mu_n \tag{20}$$

with $\theta_1 \sim N(0, \sigma)$ and $\mu_n \sim N(0, (1-\alpha^{2/(N-1)})^{1/2}\sigma)$ iid so that ex-ante demand uncertainty is independent of the grid, $\forall n : \theta_n \sim N(0, \sigma/N)$, and α measures the last-period demand shock correlation to first period demand shock: $\mathbb{E}_1\theta_N = \alpha\theta_1$.

To see how well our instruments can do, consider again the Optimal Response discussed in Section 2.2, i.e. the hypothetical policy where the planner chooses q_t only after observing

[4] See for example Schlenker and Roberts (2009) and Dietz and Stern (2015), who argue there are strong non-linearities, versus Burke et al. (2015), who argue that damages are at most linear.

[5] To facilitate the interpretation, we change from damages depending on cumulative emissions $Q = q_1+q_2$ (as in the previous section) to dependence on average emissions. This implies that the γ in the formulas for the N=2-period model is twice as large as the γ in the 2-period model of the previous section.

θ_t.[6] The Optimal Response turns out to provide a useful benchmark for instrument performance since, as we show, the difference in welfare between either Responsive Quotas or Endogenous Taxes and the Optimal Response becomes vanishingly small if there are many periods. This reminisces the result in Roberts and Spence (1976) and Weitzman (1978), who show that one can approximate the environmental marginal damage curve arbitrarily closely by combining an increasing number of specific quantity and price instruments. A formal characterization of the Optimal Response is given in the Appendix.

Yet while it is somewhat intuitive that welfare losses become vanishingly small with increasingly fine grids, we establish a substantially stronger result: Endogenous Taxes approaches the Optimal Response welfare level for an increasingly fine grid of trades two orders of magnitude faster than Responsive Quota. Let $W_N^{OR} - W_N^i$ be the welfare losses under policy i compared to the Optimal Response with N regulatory periods. Our interest is primarily in the performance of our new instruments Endogenous Taxes (ET) and Responsive Quotas (RQ). Because of its widespread occurrence in emission trading systems, we also include classic banking (B, see subsection 2.4) in the comparison.

Theorem 1. *Let N denote the number of periods. For sufficiently large N, polices are strictly ordered $OR \succ ET \succ RQ \succ B$. The welfare gap between the best possible allocation OR and the policies decreases with N according to*

$$\mathbb{E}W^{OR} - \mathbb{E}W^{ET} = O(N^{-4}), \tag{21}$$

$$\mathbb{E}W^{OR} - \mathbb{E}W^{RQ} = O(N^{-2}). \tag{22}$$

That is, Endogenous Taxes approaches the Optimal Response welfare level for an increasingly fine grid of trades two orders of magnitude faster than Responsive Quota. The welfare loss associated with Standard Banking does not vanish for many periods.

All claims made in this section are proven in Appendix B.

2.9 Implementation

We discuss possible approaches to implement our instruments here.

Responsive Quotas can be implemented by means of a cap-and-trade system where new allowances are periodically injected. The planner is lenient with respect to periodic

[6] In this N-period model, the Optimal Response is defined as:

$$\max_{q_1} \mathbb{E}_1 \max_{q_2} \mathbb{E}_2 \cdots \max_{q_N} \mathbb{E}_N W(q_1, q_2, \cdots, q_N; \theta_1, \theta_2, \cdots, \theta_N).$$

compliance, but aggregated over all regulatory periods firms must comply with their allocations. Periodic lenience can be achieved by allowing firms to bank and borrow emission allowances between periods, much like many emissions trading systems operative today do. The difference between cap-and-trade with banking and Responsive Quotas is that the number of new allowances injected in any given period becomes a function of the amount of banked allowances under a Responsive Quotas regime. Qualitatively, this implementation of Responsive Quotas comes remarkably close to the European Union's Emissions Trading System after its 2018 revision (Perino, 2018; Gerlagh and Heijmans, 2019).

Endogenous Taxes, though a pure price instrument, can be implemented by the same cap-and-trade system as Responsive Quotas, with the exception of the last period when the hard cap on emissions is abandoned and allowances are auctioned for a fixed price. The final-period auction price depends on cumulative surrendered allowances.

It may be puzzling that a pure price instrument is implemented, to a substantial extent, by a cap-and-trade system in all but the last period. We exploit one of our key results here: both Responsive Quotas and Endogenous Taxes implement Optimal Response (no asymmetric information) emissions levels in all but the final period. Hence, for all periods up to the last these instruments are essentially equivalent. The only difference arises in the last period, where our proposed implementation of an Endogenous Taxes regime indeed deviates from a pure Responsive Quotas policy by taxing emissions.

Importantly, the fact that Endogenous Taxes can be implemented by a combination of Responsive Quotas and a final-period tax does *not* mean it is a hybrid instrument. From the very definition, Endogenous Taxes is a pure price instrument, as it does not set any quantity constraints.

3 The General Case: Two Existence Results

In the preceding analysis, we develop our new instruments for the case of linear-quadratic costs and benefits. We did so for the ease of exposition. In this section, we briefly define the most general versions on Responsive Quotas and Endogenous Taxes. We then show that these general instruments can be implemented for any concave benefits and convex costs. To that end, we first need to give a general characterization of an instrument.

We characterize an instrument as the choice of policy variables for both periods, x_1, x_2 with $x = (x_1, x_2)$, such that they maximize expected welfare *given* an optimization program, where "given" means "for fixed points in time at which x_1 and x_2 are determined".

Formally, an instrument implements the solution to:

$$\max_{x_1} \mathbb{E}_{t_1} \max_{x_2} \mathbb{E}_{t_2} W(\widetilde{q}_1(x), \widetilde{q}_2(x); \theta_1, \theta_2), \tag{23}$$

where t_1 (t_2) is the point in time at which x_1 (x_2) is decided upon.

In this characterization, the defining element of any price or quantity instrument is the timing at which its levels are set, indicated in (23) by the subscripts $0 \leq t_1 \leq t_2 \leq 2$ of the expectations operators. When $t_i = 2$, the choice of policy variable x_i is decided after all information (θ_1 and θ_2) is collected and we can omit the expectations symbol. When $t_i = 1$, x_i is determined after θ_1 is observed but before θ_2 is observed, while $t_i = 0$ implies choosing x_i before any information is revealed.

With the general characterization (23) of an instrument in mind, we recall that our optimal pure quantity instrument Responsive Quotas implements (7), i.e. it fixes q_2 after q_1 has been realized. Thus we defined a response function R such that:

$$Q = R(q_1). \tag{24}$$

The defining characteristic of such a response function is that, by making second period quantities a function of first period emissions, it lets firms choose q_1 while knowing how this will affect q_2, their allocation in the second period. A smart choice of R therefore tries to set q_2 in such a way that firms choose both q_1 and $q_2 = R(q_1) - q_1$ optimally in light of the first-period fundamentals (θ_1). If such a R can be found, it implements the solution to (7). For the case of linear marginal costs and benefits, we saw in Section 2.5 that an optimal R^* implementing (7) exists. Remarkably, one can show that this result generalizes: response-function R implementing the solution to (7) exists for any concave benefits and convex costs.

Theorem 2. *For any concave benefits B_t and convex costs C, there exists an optimal response function R^* that implements the solution* (7).

Similarly to Responsive Quotas, Endogenous Taxes is defined as the instrument that implements (12). It fixes the emissions tax in both periods after first-period market fundamentals are realized. We defined a tax-response function T which sets prices in the second period in response to emissions in the first:

$$p_2 = T(q_1). \tag{25}$$

Section 2.6 illustrated that a policy-response function T^* implementing Endogenous Taxes can be found in a model with linear marginal benefits and costs. Theorem 3 establishes that such a T^* can be found for all concave benefits and convex costs.

Theorem 3. *For any concave benefits B_t and concave costs C, there exists an optimal response function T^* that implements the solution* (12).

4 Discussion and Conclusions

We study the optimal regulation of pure stock externalities in environments where the planner is asymmetrically aware of market fundamentals and the future is partly unpredictable. In its most general form, we define regulation as the implementation of a welfare maximization problem. This constructive approach yields two policies, Responsive Quotas and Endogenous Taxes, each of which always exists and is strictly welfare superior among all pure quantity and price instruments, respectively. Both instruments implement welfare levels that converge (faster than standard policies including recursive quota and recursive taxes) to the Optimal Response, the hypothetical level of welfare attained when a planner is informed about market fundamentals at the same time firms are. Our instruments therefore allow the planner to regulate the market almost as though there were no asymmetric information. This means our new instruments could yield substantial gains in social welfare compared to currently existing policies, for example in the mitigation of CO_2 emissions. With regard to the latter application, we suggest fairly simple ways in which our instruments can be implemented as an adaptation of existing cap-and-trade policies. Given the abundance of cap-and-trade programs across the global – think of the European Unions Emissions Trading System (EU ETS), the Regional Greenhouse Gas Initiative, the Chinese national carbon trading scheme, the UK ETS, California's cap-and-trade program, and South-Korea ETS – our results may provide useful guidance in the design and adaption of policies to mitigate climate change.

References

Abrell, J. and Rausch, S. (2017). Combining price and quantity controls under partitioned environmental regulation. *Journal of Public Economics*, 145:226–242.

Allen, M. R., Frame, D. J., Huntingford, C., Jones, C. D., Lowe, J. A., Meinshausen, M., and Meinshausen, N. (2009). Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, 458(7242):1163–1166.

Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239.

Dietz, S. and Stern, N. (2015). Endogenous growth, convexity of damage and climate risk: how nordhaus' framework supports deep cuts in carbon emissions. *The Economic Journal*, 125(583):574–620.

Fell, H., MacKenzie, I. A., and Pizer, W. A. (2012). Prices versus quantities versus bankable quantities. *Resource and Energy Economics*, 34(4):607–623.

Gerlagh, R. and Heijmans, R. J. R. K. (2019). Climate-conscious consumers and the buy, bank, burn program. *Nature Climate Change*.

Golosov, M., Hassler, J., Krusell, P., and Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82(1):41–88.

Harstad, B. and Eskeland, G. S. (2010). Trading for the future: Signaling in permit markets. *Journal of public economics*, 94(9-10):749–760.

Hasegawa, M. and Salant, S. (2014). Cap-and-trade programs under delayed compliance: Consequences of interim injections of permits. *Journal of Public Economics*, 119:24–34.

Heutel, G. (2020). Bankability and information in pollution policy. *Journal of the Association of Environmental and Resource Economists*, 7(4):779–799.

Hoel, M. and Karp, L. (2001). Taxes and quotas for a stock pollutant with multiplicative uncertainty. *Journal of Public Economics*, 82(1):91–114.

Holland, S. P. and Yates, A. J. (2015). Optimal trading ratios for pollution permit markets. *Journal of Public Economics*, 125:16–27.

Karp, L. and Traeger, C. (2017). Smart cap. *Paper presented at the CESifo Energy and Climate Economics Conference, Munich, October 2017.*

Karp, L. and Traeger, C. (2018). Prices versus quantities reassessed. *Paper presented at the CESifo Energy and Climate Economics Conference, Munich, October 2018.*

Kling, C. and Rubin, J. (1997). Bankable permits for the control of environmental pollution. *Journal of Public Economics*, 64(1):101–115.

Kolstad, C. D. (1996). Fundamental irreversibilities in stock externalities. *Journal of Public Economics*, 60(2):221–233.

Lintunen, J. and Kuusela, O.-P. (2018). Business cycles and emission trading with banking. *European Economic Review*, 101:397–417.

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.

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

Newell, R., Pizer, W., and Zhang, J. (2005). Managing permit markets to stabilize prices. *Environmental and Resource Economics*, 31(2):133–157.

Newell, R. G. and Pizer, W. A. (2003). Regulating stock externalities under uncertainty. *Journal of Environmental Economics and Management*, 45(2):416–432.

Perino, G. (2018). New eu ets phase 4 rules temporarily puncture waterbed. *Nature Climate Change*, 8(4):262.

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

Pizer, W. A. and Prest, B. C. (2020). Prices versus quantities with policy updating. *Journal of the Association of Environmental and Resource Economists*, 7(3):483–518.

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

Rubin, J. D. (1996). A model of intertemporal emission trading, banking, and borrowing. *Journal of Environmental Economics and Management*, 31(3):269–286.

Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to us crop yields under climate change. *Proceedings of the National Academy of sciences*, 106(37):15594–15598.

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.

Weitzman, M. L. (2019). Prices or quantities can dominate banking and borrowing. *Scandinavian Journal of Economics*.

Yates, A. J. and Cronshaw, M. B. (2001). Pollution permit markets with intertemporal trading and asymmetric information. *Journal of Environmental Economics and Management*, 42(1):104–118.

Zhao, J. (2003). Irreversible abatement investment under cost uncertainties: tradable emission permits and emissions charges. *Journal of Public Economics*, 87(12):2765–2789.

A General Model Existence of Response Functions that support Theorems 2 and 3

Proof. We only need to establish that information in q_1 and θ_1 are identical, that is, that q_1 is monotonic in θ_1. We prove this for Responsive Quotas explicitly. The same algebra can be applied to Endogenous Taxes.

Totally differentiating the condition that prices are constant in expectations (i.e. realized first-period prices are equal to expected second-period prices), we obtain:

$$B_1'' dq_1 + d\theta_1 - \mathbb{E}_1 B_2'' dq_2 - \mathbb{E}_1 d\theta_2 = 0. \tag{26}$$

Similarly, when we totally differentiate the first-order condition that prices in expectations equal marginal costs, we find:

$$B_1'' dq_1 + d\theta_1 - C''(dq_1 + dq_2) = 0. \tag{27}$$

We can multiply (26) by C'' and (27) by $\mathbb{E}_1 B_2''$ and subtract one from the other, to obtain:

$$[B_1'' \cdot C'' - B_1'' \cdot \mathbb{E} B_2'' + C'' \cdot \mathbb{E}_1 B_2'']dq_1 + [C'' - \mathbb{E}_1 B_2'']d\theta_1 - C'' \mathbb{E}_1 d\theta_2 = 0 \tag{28}$$

This in turn can be rewritten to yield:

$$\frac{\mathbb{E}_1 d\theta_2}{d\theta_1} = \frac{C'' - \mathbb{E}_1 B_2''}{C''} + \frac{B_1'' C'' - B_1'' \mathbb{E}_1 B_2'' + C'' \mathbb{E}_1 B_2''}{C''} \frac{dq_1}{d\theta_1} \tag{29}$$

Since $B_t'' < 0$ and $C'' > 0$ by assumption, the first term on the RHS is larger than one:

$$\frac{C'' - \mathbb{E}_1 B_2''}{C''} > 1. \tag{30}$$

Moreover:

$$\frac{B_1''C'' - B_1''\mathbb{E}_1 B_2'' + C''\mathbb{E}_1 B_2''}{C''} < 0. \tag{31}$$

Clearly, then, if $\mathbb{E}_1 d\theta_2/d\theta_1 \leq 1$, it is immediate that $dq_1/d\theta_1 > 0$ and any response $q_1 + q_2$ dependent on θ_1 can be written implicitly as dependent on q_1. Q.E.D.

B Linear Demand, N periods

In this section we describe in detail the allocations brought about by the various policies, and the welfare gaps. The elements of Theorem 1 are proven as part of the policy characterization.

B.1 One-period model for reference

It will turn out convenient, for the N-period model, to have the one-period Weitzman (1974) model at hand. In the competitive market, prices satisfy:

$$p = -\beta q + \theta. \tag{32}$$

The Social Optimum is characterized by:

$$p^{SO} = \frac{\gamma}{\beta + \gamma}\theta, \tag{33}$$

$$q^{SO} = \frac{1}{\beta + \gamma}\theta \tag{34}$$

When the regulator sets quota at its ex-ante optimal level $q^Q = 0$, prices given by market equilibrium (32), $p^Q = \theta$, and welfare losses are given by

$$\begin{aligned} \mathbb{E}W^{SO} - \mathbb{E}W^Q &= \mathbb{E}\left[\frac{1}{2}(p^{SO} + p^Q)(q^{SO} - q^Q) - \frac{1}{2}\gamma(q^{SO} + q^Q)(q^{SO} - q^Q)\right] && (35)\\ &= \frac{-1}{2}\frac{1}{\beta + \gamma}\sigma^2 && (36) \end{aligned}$$

When the regulator sets the tax at its ex-ante optimal level $p^P = 0$ quantity follows, $q^P = \frac{\theta}{\beta}$. Welfare losses are given by

$$\mathbb{E}W^{SO} - \mathbb{E}W^P = \mathbb{E}\left[\frac{1}{2}(p^{SO} + p^P)(q^{SO} - q^P) - \frac{1}{2}\gamma(q^{SO} + q^P)(q^{SO} - q^P)\right] \tag{37}$$

$$= \frac{-1}{2}\frac{\gamma^2}{\beta^2(\beta+\gamma)}\sigma^2 \tag{38}$$

To check consistency, note that we reproduced the result by Weitzman (1974) that $\mathbb{E}W^Q > \mathbb{E}W^P$ iff $\beta < \gamma$.

B.2 Optimal Response

Optimal Response is defined through the competitive equilibrium condition $\forall m = 1, ..., N$: $p_m = \theta_m - \beta q_m$, rational expectations $\forall 1 \leq m \leq n \leq N : \mathbb{E}_m p_n = p_m$, and expected efficiency $p_m = \gamma \mathbb{E}_m Q$. These properties enable us to construct the dynamics for prices and quantities. Optimal Response solves the set of equations $B'_m = \mathbb{E}_m B'_{m+1} = \mathbb{E}_m B'_{m+2} = ... = \mathbb{E}_m B'_N = \mathbb{E}_m C'$, where competitive markets ensure $B'_n = p_n$:

$$\mathbb{E}_m B'_n = \alpha^{\frac{n-m}{N-1}}\theta_m - \beta \mathbb{E}_m q_n \tag{39}$$

$$\mathbb{E}_m C' = \gamma Q_m + \frac{\gamma}{N}\sum_{n=m+1}^{N} \mathbb{E}_m q_n \tag{40}$$

for $Q_m = \sum_{n=1}^m q_n/N$. We multiply the first equation by $\frac{\gamma}{N\beta}$, sum over all curent plus future $n = m, ..., N$, add the second equation, write $\mathbb{E}_m p_N$ for expected marginal benefits and costs, to get:

$$\left(1 + \frac{\gamma}{\beta}x_m\right)\mathbb{E}_m p_N = \gamma Q_m + \frac{\gamma}{\beta}A_m\theta_m \tag{41}$$

where $x_m = (N - m + 1)/N$ is the share of remaining periods (including period m), and $A_m = \sum_{k=m}^N \alpha^{\frac{k-m}{N-1}}/N$ is the cumulative increase in current plus future marginal productivity induced by θ_m. Combining with the price equation (5), is rewritten as

$$p_m = \frac{\beta\gamma}{\beta + \gamma x_m}Q_{m-1} + \frac{\gamma}{\beta + \gamma x_m}A_m\theta_m \tag{42}$$

$$q_m = \frac{-\gamma}{\beta + \gamma x_m}Q_{m-1} + \frac{1}{\beta}\left(1 - \frac{\gamma}{\beta + \gamma x_m}A_m\right)\theta_m, \tag{43}$$

which gives the recursive solution, $p_m^{OR}(Q_{m-1}, \theta_m), q_m^{OR}(Q_{m-1}, \theta_m)$.

Note that because the OR does not equalize prices over periods, and thus the prices follow a random walk, there is a non-vanishing welfare loss relative to the Social Optimum: $\mathbb{E}W^{OR} - \mathbb{E}W^{CQ} = O(1)$.

B.3 Responsive Quota

Responsive quota has the same allocation as OR for $m = 1, ..., N-1$, but for the last period

$$p_N = p_{N-1} + \mu_N, \tag{44}$$

$$q_N = \frac{\alpha^{\frac{1}{N-1}} \theta_{N-1} - p_{N-1}}{\beta}. \tag{45}$$

At the last period, there is history of cumulative emissions Q_{N-1}, an expected demand increase $\alpha^{1/(N-1)}\theta_{N-1}$, and a final demand shock $(1 - \alpha^{2/(N-1)})^{1/2}\mu_N$. The Optimal Response fully adapts quantities and prices (q_N, p_N) to the new information μ_N. The Responsive Quota fixes last-period quantities q_N to the expected level. Thus, in the multi-period model, welfare losses of Responsive Quota compared to the Optimal Response can only arise in the last period. Formally, as $(Q^{OR} - Q) = (q_N^{OR} - q_N)/N$ the welfare gap becomes (where we leave the RQ superscripts):

$$\mathbb{E}W^{OR} - \mathbb{E}W = \mathbb{E}\left[\frac{1}{2N}(p_N^{OR} + p_N)(q_N^{OR} - q_N) - \frac{\gamma}{2N}(Q^{OR} + Q)(q_N^{OR} - q_N)\right] \tag{46}$$

Since $\mathbb{E}(q_N^{OR} - q_N) = \mathbb{E}_{N-1}(q_N^{OR} - q_N) = 0$, we can multiply it by a constant $\mathbb{E}_{N-1}p_N^{OR}$, $\mathbb{E}_{N-1}p_N$, $\mathbb{E}_{N-1}q_N^{OR}$ or $\mathbb{E}_{N-1}q_N$, keeping zero. Also, when it is multiplied by Q, the part multiplied by Q_{N-1} is zero and only the interaction with q_N remains. Thus, the above equation transforms into

$$\mathbb{E}W^{OR} - \mathbb{E}W = \mathbb{E}\left[\frac{1}{2N}\left((p_N^{OR} - \mathbb{E}_{N-1}p_N^{OR}) + (p_N - \mathbb{E}_{N-1}p_N)\right)(q_N^{OR} - q_N)\right] \tag{47}$$

$$- \mathbb{E}\left[\frac{\gamma}{2N^2}\left((q_N^{OR} - \mathbb{E}_{N-1}q_N^{OR}) + (q_N - \mathbb{E}_{N-1}q_N)\right)(q_N^{OR} - q_N)\right], \tag{48}$$

where the division by N^2 in the second line appears because of production aggregation specified in (19). On closer inspection, the above resembles exactly the first line of (35). That is, welfare losses of Responsive Quota relative to the Optimal Response equal those of Quota relative to the Social Optimum in a one-period model with noise μ_N, marginal costs slope γ/N, and divided by N to correct for the shorter length of period. Thus, we can take the second line of (35) and transform it into

$$\mathbb{E}W^{OR} - \mathbb{E}W^{RQ} = \frac{-1}{2N(\beta+\gamma/N)}(1-\alpha^{2/(N-1)})\sigma^2 \tag{49}$$

In the limit, we have $N(1-\alpha^{2/(N-1)}) \to -2\ln(\alpha)$, so that

$$\lim_{N\to\infty} N^2(\mathbb{E}W^{OR} - \mathbb{E}W^{RQ}) = \frac{-\ln(\alpha)}{2\beta}\sigma^2. \tag{50}$$

Another way to write this is $\mathbb{E}W^{OR} - \mathbb{E}W^{RQ} = O(N^{-2})$, the second result of Theorem 1.

B.4 Endogenous Taxes

Endogenous Taxes has the same allocation as OR for $m = 1, ..., N-1$, but for the last period we have $p_N = p_{N-1}$ and

$$q_N = \frac{\theta_N - p_{N-1}}{N\beta}. \tag{51}$$

To determine the welfare losses relative to OR, we follow the same argument as for RQ. OR optimally determines last-period allocation q_N, p_N, while the Endogenous Taxes fixes last-period prices p_N to the expected level. Thus, welfare of Endogenous Taxes compared to the Optimal Response has the same welfare losses as Prices in the one-period model (37), but with γ replaced by γ/N, and divided by N to account for the period length:

$$\mathbb{E}W^{OR} - \mathbb{E}W^{ET} = \frac{-\gamma^2}{2N^3\beta^2(\beta+\gamma/N)}(1-\alpha^{2/(N-1)})\sigma^2. \tag{52}$$

In the limit, this gives us

$$\lim_{N\to\infty} N^4(\mathbb{E}W^{OR} - \mathbb{E}W^{ET}) = \frac{-\gamma^2\ln(\alpha)}{2\beta^3}\sigma^2, \tag{53}$$

restated as the first result of Theorem 1: $\mathbb{E}W^{OR} - \mathbb{E}W^{ET} = O(N^{-4})$.[7]

[7] Confirmed by numerical simulations suggest it must be N^{-4}

B.5 Banking with fixed cumulative supply

Cumulative Quota solves the set of equations $\mathbb{E}_m B'_{m+1} = \mathbb{E}_m B'_{m+2} = ... = \mathbb{E}_m B'_N$, with

$$\mathbb{E}_m B'_n = \alpha^{\frac{n-m}{N-1}} \theta_m - \beta \mathbb{E}_m q_n \tag{54}$$

We sum all equations, divide by N, exploit $\frac{1}{N}\sum_{n=m+1}^{N} q_n = -Q_m$, write $\mathbb{E}_m p_N$ for expected marginal benefits. Combining with the price equation (5), we find

$$p_m = \frac{\beta}{x_m} Q_{m-1} + \frac{A_m}{x_m} \theta_m \tag{55}$$

$$q_m = \frac{-1}{x_m} Q_{m-1} + \frac{1}{\beta}\left(1 - \frac{A_m}{x_m}\right)\theta_m. \tag{56}$$

Note that the allocation characterization for Cumulative Quota converged to Optimal Response for $\gamma \to \infty$. As CQ does not adapt cumulative production to observed demand changes, it is uniformly more costly than Optimal Response: $\mathbb{E}W^{OR} - \mathbb{E}W^{Banking} = O(1)$.