

Prices versus Quantities with Policy Updating*

William A. Pizer[†] and Brian Prest[‡]

March 19, 2016

Abstract

This paper considers the effect of uncertainty on the comparative advantage of price-based regulation over intertemporally tradable quantity regulation when policies are updated. In particular, we examine the case where uncertainty about costs and benefits unfolds over time, and an information asymmetry implies that information known to the firm in one period does not influence government policy until the next period. Under this setup quantity regulation can achieve the first-best outcome and is always preferred to a price instrument when the government sets both policies to maximize expected net benefits. With quantity trading over time, the firm's intertemporal optimization implies current prices equal discounted expected future prices. This allows the government to circumvent the information asymmetry through policy updates that influence future expected prices and, in turn, current prices. No such opportunity exists with price regulation. We also consider the possibility that policy updates are driven in part by political "noise" rather than true values of costs and benefits. Here, we find that price regulation will be preferred if the variance of shocks in the political noise process exceeds the variance of the true shocks to marginal costs and benefits. Applied to climate change or other applications where marginal benefits are flat, this condition for preferring price regulation simplifies to whether noise shocks have higher variance than benefit shocks, and cost uncertainty does not matter. All of these results sharply contrast with the Weitzman (1974) results that the relative slopes of marginal costs and benefits determine the comparative advantage of prices versus quantities and that benefit uncertainty by itself does not matter.

1 Introduction

For more than forty years, economic thinking about the relative welfare advantage of alternative price and quantity regulation under uncertainty has gone something like this: When the regulator is uncertain about costs known to firms, the difference in expected net benefits between otherwise equivalent price and quantity policies hinges on the difference between the slopes of the marginal benefit and cost schedules, multiplied by the variance of the cost shock (Weitzman 1974). One intuition is that government policy is attempting to replicate society's expected marginal benefits in the form of a demand schedule in the regulated market. A flat schedule (prices) works better when marginal benefits are relatively flat (e.g., Pizer 2003); a vertical schedule (quantities) works

*The authors appreciate comments from participants in the Duke Environmental Economics Workshop and Roger von Haefen.

[†]Sanford School of Public Policy, Duke University. billy.pizer@duke.edu

[‡]University Program in Environmental Policy, Duke University. brian.prest@duke.edu

better when marginal benefits are relatively steep. A corollary to the Weitzman result is that uncertainty about marginal benefits (unless correlated with costs) does *not* affect comparative advantage of the two instruments (Stavins 1996).

As we show in this paper, however, the results change significantly in a dynamic setting when policies are updated and quantity regulation is intertemporally tradable. In contrast to the Weitzman result, in such a setting we find that quantity regulation is preferred regardless of the relative slopes and that uncertainty about benefits matters for the magnitude of the preference. In response to any new information about costs and benefits, there is now an expectation that future policies will be updated to reflect this information. The important difference between price and quantity regulation is no longer which instrument better mimics the marginal benefit schedule, but which instrument better allows future price expectations to influence current prices. This strongly favors intertemporally tradable quantity regulation. That is, with newfound expectations of a high future price for a tradable pollution permit, current tradable permits that can be saved for future use will have a high price now. The same is not true with newfound expectations of a high future pollution tax; there is no simply way to arbitrage current and future obligations. Current and expected future prices under a tax will diverge.

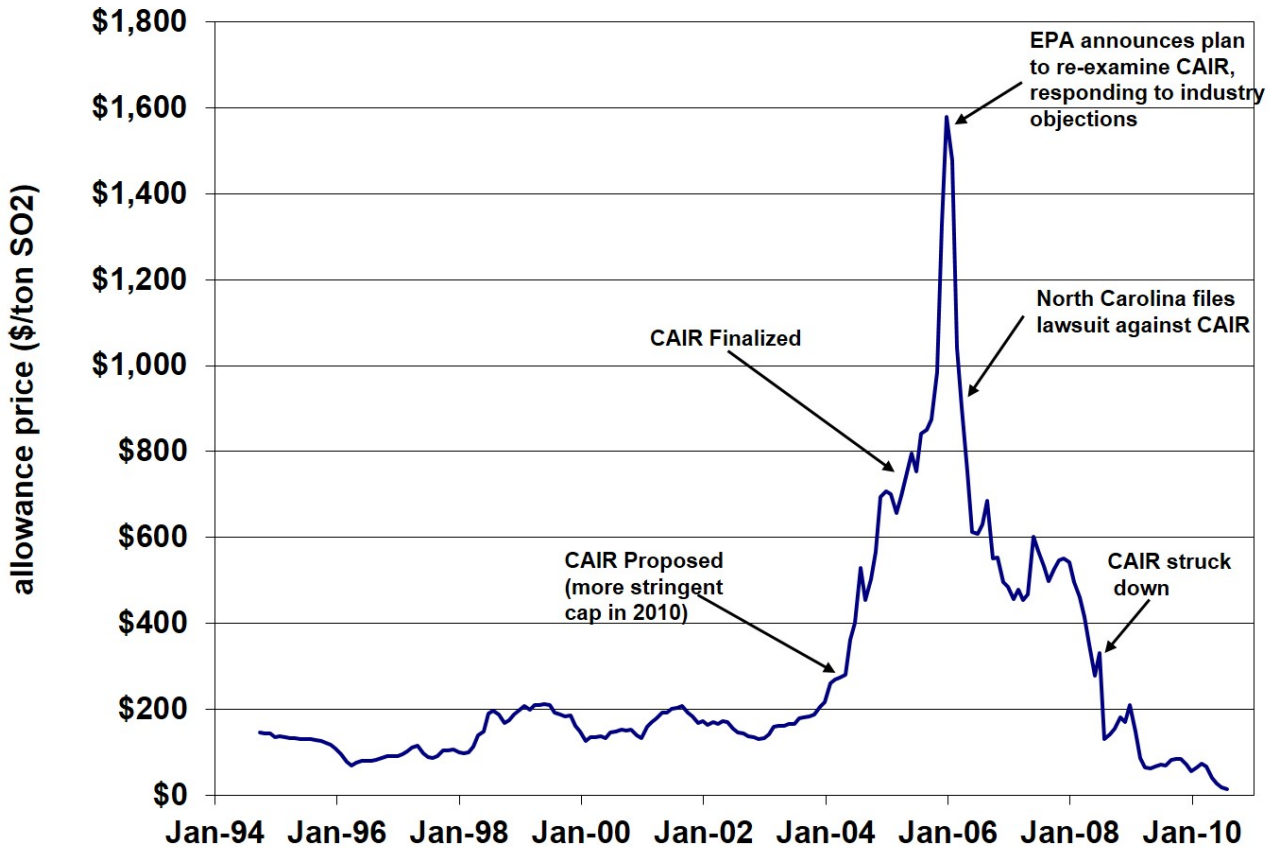


Figure 1: History of SO₂ allowance spot market prices as 2010 updates were developed

In Figure 1, we see exactly this kind of behavior in the SO₂ trading program from 2004 until 2006. During that time, future reforms to the program — a policy update in the form of the Clean Air Interstate Rule or CAIR — were being debated. This new rule would have updated the existing policy to reflect newer evidence of improved health benefits from lower emission levels and higher

marginal benefits. While none of the changes to the program would have come into effect prior to 2010, the current (spot) price moved in parallel to shifting expectations of future price because of the ability to save or “bank” current permits for future use. Higher net benefits are achieved almost immediately as current prices adjust to reflect the improved marginal benefit estimates that will be implemented in the future. Had an SO₂ tax been in place, however, an expectation of higher tax in 2010 would not have influenced the current tax. This creates a seemingly unambiguous preference for quantity regulation.

Starting in 2006, however, the price fell as the EPA reconsidered the proposed rules and later court challenges succeeded in vacating the regulation. These developments call into question whether policy updates to a particular regulation are always driven by aggregate welfare maximization. With this in mind, our paper considers the possibility that the regulator’s objective deviates from the true marginal benefits by some amount. This could reflect special interest politics, as environmental advocates or business interests hold too much sway over the government, or legal or other constraints, as arguably occurred here. When such welfare deviations arise, price regulation can again be preferred if these deviations — what we might call uncertain political “noise” — are expected to be sufficiently large.

The SO₂ program is far from unique in terms of intertemporal trading and policy updating. The linkage over time of markets for regulated quantities through permit banking is consistent with virtually all existing tradable permit programs.¹ Meanwhile, most real-world environmental regulations are updated over time in response to new information. The 1970 Clean Air Act makes specific provisions for the National Ambient Air Quality Standards to be updated every 5 years in response to new information (see 42 U.S. Code §7409). More recently, the European Union announced a package of changes in their Emissions Trading Scheme, in large part owing to low prices and (presumably) low compliance costs.² California is similarly contemplating revised targets for the next phases of their program, hopefully based on improved information.³ At the U.S. federal level, in 2010 an interagency working group estimated the social cost of carbon (SCC)⁴ to be \$24 per metric ton in 2010. That number was revised upward to \$37 based on new modeling estimates of damages in 2013.⁵

There is a similar list of policy updates where one might question whether aggregate welfare maximization is the main goal. These include New Jersey’s decision to exit the Regional Greenhouse Gas Initiative (RGGI) in May 2011, Australia’s decision to terminate their carbon pricing program in July 2014, and the U.S. Supreme Court decision to stay implementation of the Clean Power Plan in February 2016.

In the remainder of the paper, we first review the basic Weitzman result and other relevant literature. We then show how the Weitzman result changes if we allow policy updates. To illustrate our main result, we first present a simple two-period model with a common shock across both periods and no discounting. We then show how the model generalizes to multiple periods with AR(1) error processes and discounting. Next, we introduce the possibility that updates might be influenced by a stochastic process other than new information about costs and benefits, such as

¹This includes lead phase-down program in the 1980s, the acid rain program in the 1990s and 2000s, the European Emissions Trading Scheme, California’s cap-and-trade program, among others. Most programs allow limited borrowing.

²See http://ec.europa.eu/clima/policies/2030/index_en.htm.

³See <http://www.arb.ca.gov/cc/scopingplan/scopingplan.htm>.

⁴The SCC represents the marginal benefit of reducing CO₂ emissions by one metric ton.

⁵Both figures are estimated for emissions reductions occurring in 2015 and presented in 2007 dollars. See Interagency Working Group on Social Cost of Carbon (2010) and (2013) for details.

shifting special interests or evolving legal or other constraints. Finally, we apply our model to the case of global climate change, estimating that intertemporally tradable quantities provide an expected welfare gain over price regulation of \$175 million in one year absent any political noise. If policies remain fixed for five years, the expected welfare gain would be more than \$2 billion.

2 The Weitzman Result and Subsequent Work

In his seminal 1974 article, Martin Weitzman explored the difference between price and quantity regulation in a simple linear model. Assuming costs and benefits given by

$$C(q, \theta) = c_0 + (c_1 + \theta)(q - \hat{q}) + \frac{c_2}{2}(q - \hat{q})^2 \quad (1)$$

$$B(q, \eta) = b_0 + (b_1 + \eta)(q - \hat{q}) - \frac{b_2}{2}(q - \hat{q})^2, \quad (2)$$

linear marginal costs and benefits can be written as

$$MC(q, \theta) = c_1 + \theta + c_2(q - \hat{q}) \quad (3)$$

$$MB(q, \eta) = b_1 + \eta - b_2(q - \hat{q}). \quad (4)$$

He then derives his main result. Namely, the welfare advantage of prices over quantities is

$$\Delta = \frac{\sigma_\theta^2}{2c_2^2}(c_2 - b_2). \quad (5)$$

Here, θ represents cost uncertainty while benefit uncertainty is captured in a similar way by η . The parameters c_1 and c_2 describe expected marginal costs and the slope of marginal costs at a point \hat{q} , and similarly so for b_1 and b_2 with regard to marginal benefits. Generally, $c_2 > 0$ and $b_2 > 0$, reflecting upward-sloping marginal costs and downward-sloping marginal benefits. Marginal costs and benefits themselves can be positive or negative, depending on whether q is a bad or a good (e.g., $c_1 \leq 0$ and $b_1 \leq 0$). In his framework, the linear model is viewed as an approximation to more general functions about the point \hat{q} where marginal costs equal marginal benefits on average. That is, he (and we) assume $MC(\hat{q}, 0) = MB(\hat{q}, 0)$ and $c_1 = b_1$. While this assumption is not required to derive (5), the optimal government policies simplify to $q = \hat{q}$ for quantities and $p = c_1$ for prices.

His appealingly simple result (5) shows that direction of the welfare advantage Δ depends only on the relative slopes of marginal costs and marginal benefits, c_2 and b_2 . We previously mentioned the intuition that a price policy, filling in the otherwise absent market demand role, better matches the marginal benefit schedule when it is relatively flat (and vice-versa for quantity policies when the marginal benefit schedule is relatively steep). Another intuition is that price regulation offers welfare savings because such a policy does more when costs are low and less when costs are high. To the extent costs are convex (large c_2), this improves welfare. At the same time, a price policy introduces variability into q , which is fixed under quantity regulation. To the extent benefits are concave (large b_2), variable q reduces welfare. The net effect depends on the sign of $c_2 - b_2$. Put simply, if marginal benefits are relatively steep compared to marginal costs, then getting the quantity “right” is important in welfare terms, implying that quantity instruments are expected to perform better. Otherwise, price instruments are to be preferred.

The magnitude of both the expected cost saving and benefit loss depends on variability in costs — here, the variation in θ . Benefit uncertainty — variability in η — does not appear in the

expression. Benefit uncertainty does lower expected welfare, but it does so equally for price and quantity regulation.

Weitzman's findings have been extended to many other contexts in the literature. This includes alternative forms of uncertainty (Fishelson, 1976; Yohe, 1978; Stranlund and Ben-Haim, 2008) and non-linear marginal costs and benefits (Yohe, 1978; Kelly, 2005). In all of these cases, the basic result generally remains that flatter marginal benefits favor prices and steeper marginal benefits favor quantities. Other work has considered the correlation between marginal cost and benefit uncertainty (Yohe 1978; Stavins, 1996). Such correlation can overwhelm the basic Weitzman intuition in theory, but it has never been demonstrated to be relevant in practice. Yet other work has examined whether the choice of price or quantity controls affect other outcomes (e.g., investment in Chao and Wilson, 1993, and Zhao, 2003), but none of this latter work speaks to the normative question of welfare impacts.

Most relevant for our paper are extensions to Weitzman that look at uncertainty and instrument choice in a dynamic policy context. Newell and Pizer (2003), Hoel and Karp (2002), and Fell, MacKenzie, and Pizer (2012), all extend the original Weitzman framework to multiple periods where benefits can depend on the accumulated level of the pollutant, rather than the annual flow. Like earlier work, they find broadly similar results: flatter marginal benefits and steeper marginal costs favor prices, and vice-versa for quantity regulation. Importantly, none of these papers consider the possibility that policies might be updated. They both fix the policy trajectory once and for all in the initial period.

Kaplow and Shavell (2002) consider non-linear price policies under which the optimal price is simply set as the marginal benefit schedule. Implicit in that study is the idea that prices or quotas can be updated (or automatically adjust) in real time based on realizations of cost and benefit uncertainty, something that we have not yet seen in practice. They do not consider the implications of intertemporally tradable quantities.

Newell, Pizer, and Zhang (2005) are entirely focused on policy updating, but only in the limited sense of whether a quantity policy could respond to cost shocks in order to to mimic a price policy over time. The idea of benefit uncertainty and that policy updating might be focused on new information about benefits is absent.

In an unrelated paper focused on the idea of an allowance reserve, Murray, Newell, and Pizer (2009) argue that emissions trading with an allowance reserve could provide a higher level of cost-effectiveness than either emission taxes or emissions trading alone. Part of their argument is related to our idea that emissions trading has an advantage when policies are updated in response to meaningful new information. However, they abstract from the notion of welfare and benefits, and they instead assume there is a cumulative emission goal that is updated in the future and only seek to minimize costs.

This paper can be viewed as extending and clarifying the notion in Newell, Pizer, and Zhang (2005) and Murray, Newell, and Pizer (2009) that updating seems to provide some sort of advantage for quantity controls. In contrast to both of these papers, benefits and benefit uncertainty is explicitly considered. This allows us to show that quantity controls are, in fact, unambiguously preferred with truthful updating in response to learning about uncertainty. This then raises the question of whether policy updating is, in fact, based on true information or political "noise."

3 Policy Updating

We now expand the earlier Weitzman model to two periods with the same, but uncertain, costs and benefits in both periods. We consider the possibility that policies are updated between periods. That is, we imagine a regulator who sets a policy for an initial period prior to learning about costs and benefit uncertainty, θ and η in (1) and (2). As in the original Weitzman model, we assume an information asymmetry whereby the regulated firm learns about θ and η before taking action. In our two-period model, the government learns about θ and η alongside the firm, but is unable to act quickly enough to change the first-period policy directly. The government is able to change second-period policy and does so according to a pre-announced rule. Meanwhile, the firm can deduce the government’s second-period policy before taking action in the first period, through knowledge of θ and η and the government’s updating rule.⁶ The information structure and decisions over time are pictured in Figure 2.

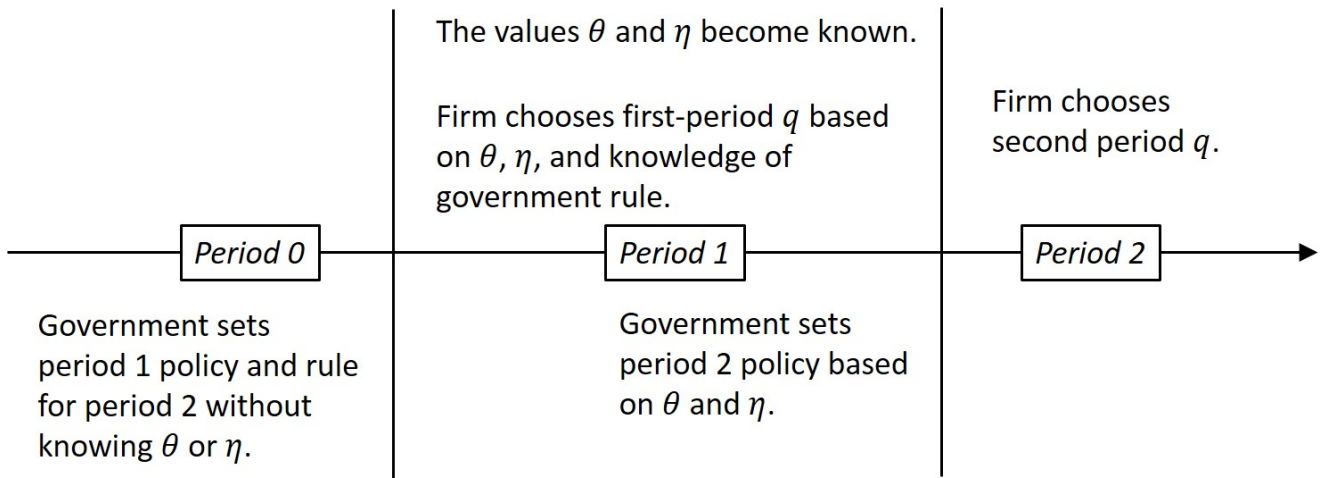


Figure 2: Timeline of events with updating

In addition to the assumed information structure and decision sequence, we make the further assumption that the quantity instrument can be traded between the two periods. As we have noted, this is consistent with virtually all observed tradable permit systems that allow current period permits to be banked for use in the future and many that allow some volume of future period permits to be used early or borrowed.

We now characterize firm behavior, the government’s welfare-maximizing price and quantity policies, and the relative advantage of prices versus quantities, for both fixed and updated policies.

3.1 Firm behavior

In the traditional Weitzman framework, firm behavior under a generic quantity policy \tilde{q} was to trivially set $q = \tilde{q}$. Here, it is not so simple in the first period as the firm has the option to use the assigned quantities \tilde{q}_1 or to deviate, creating a “bank” $B = \tilde{q}_1 - q_1$.⁷ This bank can be in surplus

⁶In this case, we are assuming the government is a Stackelberg leader and the firm does not attempt to strategically influence the government’s second-period behavior. This is consistent with the idea of the regulated “firm” representing an aggregation of multiple firms in a competitive market.

⁷Here, we adopt the convention that q is a “bad,” and therefore a positive value for B defined this way is helpful in the second period. Were q assumed to be a good, we could instead define $B = q_1 - \tilde{q}_1$ in order for positive values

or deficit $B \geq 0$ to be made up in the second period. For that reason, we work backwards: In the second period, there is no option except $q_2 = \tilde{q}_2 + B$. Then, in the first period we can write the firm's problem as:

$$\min_B C(\tilde{q}_1 - B, \theta) + C(\tilde{q}_2 + B, \theta).$$

We ignore discounting for the moment. We also take advantage of the fact that there is no uncertainty from the perspective of the firm. It knows the cost and benefit shocks before it makes any decisions, and it can deduce the government's optimal second-period policy based on θ and η . This leads to an arbitrage condition for the firm: $MC(q_1, \theta) = MC(q_2, \theta)$. That is, equalize marginal costs across periods. Given the same cost function in both periods, we have a simple solution:

$$q_1 = q_2 = \frac{\tilde{q}_1 + \tilde{q}_2}{2}. \quad (6)$$

This is true regardless of what kind of quantity policy (fixed or updating) is used to define \tilde{q}_2 .

Firm behavior under a price policy is the same as derived by Weitzman. Facing prices \tilde{p}_t in either period for output q_t , firms will choose q_t to maximize profits. That is,

$$\max_{q_t} \tilde{p}_t q_t - C(q_t, \theta).$$

This leads to a simple rule for the firm: $MC(q_t, \theta) = \tilde{p}_t$. That is, set marginal costs each period to the regulated price (noting marginal costs will be negative for a "bad"). This yields Weitzman's response function,

$$h(\tilde{p}_t, \theta) = \hat{q} + \frac{\tilde{p}_t - c_1 - \theta}{c_2}.$$

3.2 Optimal Price and Quantity Policies

For comparison, we first consider what happens when the government chooses its policies to maximize expected net benefits over both periods *without* knowing the values of η and θ . Absent updating, this replicates the traditional Weitzman result over two periods, which we denote with superscript W . In the case of quantities, it is clear from equation (6) that — with our assumption of banking / borrowing between periods — only the total quantity volume $\tilde{q} = \tilde{q}_1 + \tilde{q}_2$ matters as a policy choice variable. Following the Weitzman approach that costs and benefits are approximated around \hat{q} where $b_1 = c_1$, it is trivial to show that expected net benefits are maximized when $\tilde{q}^W = 2\hat{q}$. The allocation between periods does not matter.

In the case of prices, the symmetry of the problem leads to the original Weitzman solution $\tilde{p}_t^W = c_1 = b_1$, resulting in equilibrium quantities of $\hat{q} - \frac{\theta}{c_2}$ in each period.

With updating, the government has the ability to pick \tilde{q}_2 and \tilde{p}_2 after observing θ and η . In the case of prices, this is relatively straightforward because there is no behavioral link between periods and the government faces two distinct optimizations. In period one, $\tilde{p}_1^U = c_1$ as before, where superscript U indicates the government's welfare maximizing solution to the updating policy problem. In period two, we have a first-order condition for the optimal updating policy \tilde{p}_2^U :

$$(MB(h(\tilde{p}_2^U, \theta), \eta) - MC(h(\tilde{p}_2^U, \theta), \theta)) h_p(\tilde{p}_2^U, \theta) = 0.$$

Substituting and rearranging, we have

$$c_2(\theta - \eta) + (b_2 + c_2)(\tilde{p}_2^U - c_1 - \theta) = 0,$$

for B to imply greater flexibility in period 2.

or

$$\tilde{p}_2^U = c_1 + \frac{b_2\theta + c_2\eta}{b_2 + c_2}.$$

This is the expression for the optimal, updated price in the second period. Note that in the second period, the regulator is able to achieve the first-best outcome as there is no longer an information asymmetry.

Optimal quantities are also straightforward. Let $\tilde{q} = \tilde{q}_1 + \tilde{q}_2$. From the previous discussion, we know that the firm chooses $q_1 = q_2 = \tilde{q}/2$. This is true regardless of how the government divides \tilde{q} between periods one and two. Therefore, the government's problem amounts to figuring out the cumulative, two-period quantity policy \tilde{q}^U to maximize net benefits:

$$\max_{\tilde{q}} [2B(\tilde{q}/2, \eta) - 2C(\tilde{q}/2, \theta)],$$

where we let \tilde{q}^U be this optimized policy. Note that by determining the cumulative quantity target in the second period, once it observes the uncertain outcomes, the government circumvents the information asymmetry!

Working through the first-order condition

$$2(b_1 + \eta - b_2(\tilde{q}^U/2 - \hat{q}) - c_1 - \theta - c_2(\tilde{q}^U/2 - \hat{q})) = 0,$$

or, rearranging,

$$\tilde{q}^U/2 = \hat{q} + \frac{\eta - \theta}{b_2 + c_2}.$$

Without loss of generality, we can set $\tilde{q}_1^U = \hat{q}$ and $\tilde{q}_2^U = \hat{q} + 2\frac{\eta - \theta}{b_2 + c_2}$ as the optimal, updated quantity policy. Moreover, this achieves the first-best outcome in both periods (in contrast to the price policy, which only achieves the first-best outcome in the second period). That is, knowing the firm will equalize marginal costs and divide the cumulative allocation across both periods, the regulator adjusts the quantity policy in the second period to accommodate the first-best outcome in both periods.

3.3 Comparative advantage

Without updating, the comparative advantage of prices versus quantities, Δ^W , is simply the Weitzman result over two periods, or

$$\Delta^W = 2\Delta = \frac{\sigma_\theta^2}{c_2^2}(c_2 - b_2).$$

With updating, both policies achieve the first-best outcome in the second period. The comparative advantage stems from differences in the first-period outcome, namely the loss associated with the standard price instrument,

$$\begin{aligned} \Delta^U &= E \left[(B(h(\tilde{p}_1^U), \theta), \eta) - C(h(\tilde{p}_1^U), \theta), \theta)) - (B(\tilde{q}^U/2, \eta) - C(\tilde{q}^U/2, \theta)) \right] \\ &= \frac{-(\sigma_\eta^2 + (b_2/c_2)^2\sigma_\theta^2)}{2(b_2 + c_2)}. \end{aligned} \tag{7}$$

We provide an algebraic proof of (7) in the appendix, but a simple graphical proof is as follows. In the first period, the updating price policy is simply the standard price policy, $\tilde{p}_1^U = c_1$ and leads

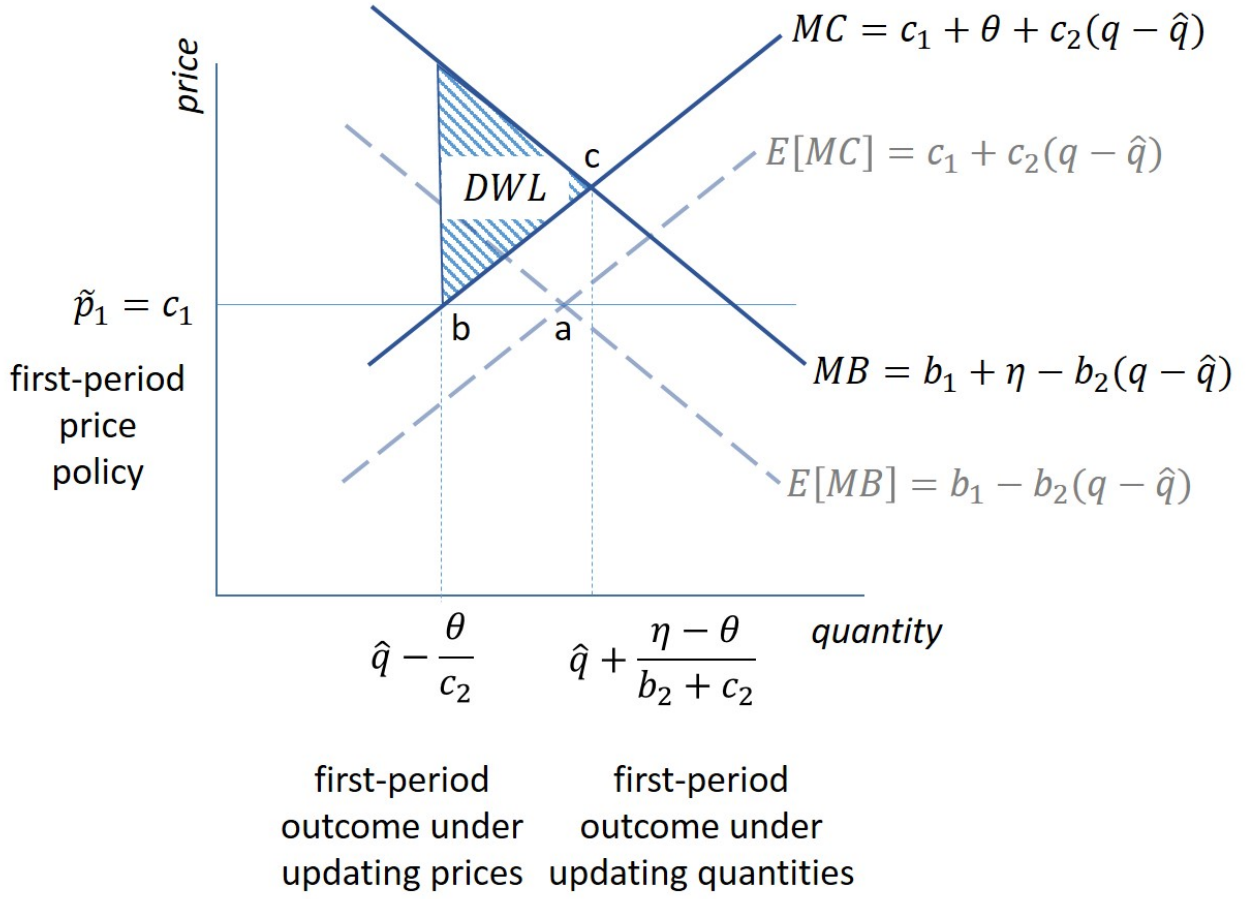


Figure 3: Illustration of deadweight loss from price policy

to the outcome at point b in Figure 3. Meanwhile the updating quantity policy obtains the first best $q = \tilde{q}^U/2 = \hat{q} + \frac{\eta - \theta}{b_2 + c_2}$ at point c .

We can see from the figure that the price policy leads to the shaded deadweight loss, DWL . This is a triangle where the “height” in the quantity dimension is $\frac{\theta}{c_2} + \frac{\eta - \theta}{b_2 + c_2} = \frac{\eta + (b_2/c_2)\theta}{b_2 + c_2}$. The “base” along the price dimension equals the height times $b_2 + c_2$. The area is then one-half base times height, or $\frac{(\eta + (b_2/c_2)\theta)^2}{2(b_2 + c_2)}$. Taking expectations of this area yields the expressed loss in (7)

3.4 Discussion

It is clear that the price policy is never preferred with updating. The quantity policy is first best in both periods while the price policy is generally only first best in the second period. In the special case where marginal benefits are flat ($b_2 = 0$) and there is no benefit uncertainty ($\sigma_\eta^2 = 0$), the relative advantage is zero. In such a case, prices achieve the first best. *But otherwise, quantities with updating is superior to prices. Quantities always achieves the first-best outcome and the welfare loss from price policies depends on both cost and benefit uncertainty.*

This is quite different from the Weitzman result, without updating, where the difference between the marginal benefit and cost slopes determines an otherwise ambiguous preference between price and quantity controls. Moreover, both price and quantity controls are second best in the Weitzman framework due to the information asymmetry faced by the government. Benefit uncertainty does

not appear in the comparative advantage expression (5) because it equally affects the expected welfare loss from both prices and quantities.

Three assumptions underpin these different results vis-a-vis Weitzman. First, we need to assume the quantity policy gives the firm the flexibility to arbitrage its regulatory obligation between periods while the price policy does not. For example, through intertemporal trading of a quantity control, overcompliance in one period creates a permit bank and reduces the obligated action in the subsequent period. While the price policy does allow the firm to vary its quantity choice, doing more in one period does not change the incentives or obligations in the other. Second, we need an information and decision structure that has the regulator setting policy in two periods, with new information revealed about first-period costs and benefits (a) after the first-period policy is fixed but (b) before the firm chooses its first-period action and the regulator sets second-period policy. Finally, we assume that (uncertain) costs and benefits are the same in both periods and there is no discounting, so the firm seeks to exactly equalize marginal costs and this is also the first-best optimum. How reasonable are these assumptions?

The first assumption is generally consistent with observed policies. That is, most quantity-based policies allow intertemporal flexibility. Even where there are restrictions, these restrictions are typically not binding. For example, most programs have some limits on borrowing — yet in virtually all of these programs large banks exist so the borrowing limit is non-binding.

The second assumption is the standard information asymmetry assumed in the traditional Weitzman framework — firms are able to act on better information than the government. Sometimes, this is framed as the idea that government simply does not have the same quality of information, particularly cost information, as the firm. A more relevant reason for this assumption is that policies are typically established well before compliance obligations begin. This gives firms a certain amount of lead time to plan their compliance — but also means government policies are always based on lagged information relative to what is available when policies are binding.

Either way, new information *is* available to the firm by the time compliance is required. The information is either simultaneously or subsequently revealed to the government and the firm can anticipate the government will eventually update the regulation in response to this information. As we noted in the introduction, such policy updating is a common phenomena and to be expected. It is less clear whether updating is always optimal in practice, which is a question that we consider in the next section.

The third assumption implies that the first-best price is the same in both periods. With non-zero discounting and additional uncertain shocks in the future, the firm will instead equate marginal cost today with expected discounted marginal cost in the future. However, it is relatively straightforward to extend the simple two-period model to a more general multi-period model. If we allow the government to specify trading ratios as they periodically update policies in the future, our main result continues to hold as we now demonstrate.

4 Extensions

The preceding discussion focused on expanding the Weitzman model to two periods with policy updating and intertemporally tradable quantities, and then showing how his original result regarding the relative advantage of prices over quantities changes. Namely, the relative advantage is no longer ambiguous; quantity regulation is always preferred.

In this section we extend the model further. We first consider multiple periods with discounting where uncertainty follows an arbitrary AR(1) process. Government policy is based on an announced

rule for setting policy in each period t that uses period $t - 1$ information. Meanwhile, firms are able to use period t information when they choose their action in period t . This maintains the information asymmetry assumed in the previous section. We then consider the possibility that policy updating is not set to maximize net benefits but is distorted by some amount of political noise, such as special interests or uncertain legal or other constraints.

Under the assumption that the regulator can also set trading ratios between periods in addition to setting allocations, the first extension does not change our main result: Quantity regulation still achieves the first-best outcome and is always preferred on the basis of expected net benefits. Adding noise that has a constant variance each period does not change the unambiguous preference for quantity controls, but it is no longer first best. If noise variance is increasing, however, and the increase is greater than the error associated with price regulation, price regulation will instead be preferred.

4.1 Multiple periods

In order to extend the model to multiple periods, we assume that η and θ now follow an AR(1) process up to some final period T rather than taking on a single, permanent value across two periods. That is,

$$\begin{aligned}\eta_t &= \rho_\eta \eta_{t-1} + \mu_t \\ \theta_t &= \rho_\theta \theta_{t-1} + \nu_t.\end{aligned}$$

Here we assume μ_t and ν_t are *i.i.d.* random errors. We modify our intertemporal trading to allow for a non-uniform trading ratios over time. That is,

$$B_t = \tilde{R}_t B_{t-1} + \tilde{q}_t - q_t,$$

where B_t is the bank at the end of period t , \tilde{q}_t is the allocation in period t , and the trading ratio \tilde{R}_t converts a one unit deficit or surplus of period $t - 1$ quantities into those available period t . We assume the problem ends in period T with $\mu_T = \nu_T = B_T = 0$ (after which there is no uncertainty and no need for banking). We show in the appendix that, if the government follows the allocation rule

$$\tilde{q}_t = \hat{q} + \frac{\rho_\eta \eta_{t-1} - \rho_\theta \theta_{t-1}}{b_2 + c_2} + \tilde{R}_t \frac{\mu_{t-1} - \nu_{t-1}}{b_2 + c_2},$$

and sets the trading ratio \tilde{R}_t according to the rule

$$\tilde{R}_t = \frac{p_{t-1}^*}{\beta E[p_t^* | t-1]},$$

then the first-best outcome $q_t^* = \hat{q} + \frac{\eta_t - \theta_t}{b_2 + c_2}$ continues to be achieved each period, based on the assumption that the firm chooses to maximize profit each period. Here, β is the discount rate used by the firm and (as before) $p_t^* = c_1 + \frac{b_2 \theta_t + c_2 \eta_t}{b_2 + c_2}$ is the first-best price outcome. Note that both \tilde{q}_t and \tilde{R}_t only depend on $t - 1$ information, and therefore meet the information asymmetry that we assume.

The idea of trading ratios is not as theoretical as it might seem. Such an approach has been proposed, though to our knowledge never implemented, several times. For example, the American Clean Energy and Security Act in the 111th Congress (H.R. 2454, a.k.a. the ‘‘Waxman-Markey Bill’’) proposed to exchange allowances in existing state programs for new federal allowances at

a ratio that preserved the original value of the state allowances when they were first issued.⁸ Similarly, the Clean Air Interstate Rule increased the stringency of the original Acid Rain Trading Program by effectively applying a trading ratio of 0.5 from 2009 to 2010 and 1/1.43 from 2014 to 2015.⁹

Given that the quantity policy continues to deliver the first-best outcome, the comparative advantage of prices each period is the welfare loss from prices that are set with lagged information:

$$\Delta^U = \frac{-(\sigma_\mu^2 + (b_2/c_2)^2\sigma_\nu^2)}{2(b_2 + c_2)}. \quad (8)$$

Expression (8) closely resembles expression (7), with the variances of η and θ replaced with those of the innovation terms μ_t and ν_t . That is, the error in the price policy each period is determined by the innovation terms μ_t and ν_t rather than η_t and θ_t . A short proof is provided in the appendix.

Note that we can rewrite the allocation rule as $\tilde{q}_t = E[q_t^*|t-1] + \tilde{R}_t(q_{t-1}^* - E[q_{t-1}^*|t-2])$. That is, the allocation in period t equals the expected first-best quantity level for period t as viewed in period $t-1$, plus a “top up” that offsets the (now observed) error in the earlier period $t-1$ allocation. It should not be too surprising that the simple two-period model can be extended to multiple periods. This follows the underlying intuition in Newell, Pizer, and Zhang (2005), where the government constantly updates a quantity policy to mimic a price policy. In both cases, firms anticipate the correction and, thanks to the arbitrage condition, choose the current quantity accordingly.

4.2 Policy Updating with Noise

While expectations of improved cost-benefit information in future policy decisions can improve current-period regulatory outcomes under intertemporally tradable quantity controls, what if we expect something other than aggregate costs and benefits to influence future policy? It is easy to imagine policy decisions being driven in part by special interests and more narrowly defined benefits rather than aggregate, societal costs and benefits. If we expect such changes to an intertemporally tradable quantity target in the future, we should be worried that these changes will be transmitted back to the present.

In particular, suppose that the government acts *as if* marginal benefits in (3) are subject to an additive disturbance ϵ reflecting uncertain special-interest political pressure rather than true information about societal benefits. One could also interpret this term as the Lagrangian multiplier on uncertain legal, political, or other constraints that bound the policy choice. The government now makes use of a “noisy” marginal benefit function given by

$$MB_{\text{noisy}}(q_t, \eta_t, \epsilon_t) = MB(q_t, \eta_t) + \epsilon_t = b_1 + \eta_t - b_2(q - \hat{q}_t) + \epsilon_t \quad (9)$$

that it wishes to equate with marginal costs in order to set policy. The function MB without any subscript continues to refer to true, societal marginal benefits, while MB_{noisy} is the divergent, special-interest or constrained version of marginal benefits used to make policy. We assume the noise is not static, and instead ϵ_t follows a random walk $\epsilon_t = \epsilon_{t-1} + \omega_t$.

In adding this notion of political noise, we make the same asymmetric information assumption about noise as we do about the true shocks to costs and benefits. Namely, the firm has more up-to-date information about political noise when the firm acts in period t than is reflected in the

⁸See Title VII, Subtitle B, Section 790.

⁹See 70 Fed. Reg. 91 (May 12, 2005), page 25291.

period t policy that was set some time before. Moreover, this new period t information is used to update policies that come into effect in period $t + 1$. More specifically, the firm chooses q_t with knowledge of θ_t , η_t , and now ϵ_t , while the government policy \tilde{p}_t (under a price policy) or \tilde{q}_t and \tilde{R}_t (under a quantity policy) is set when only θ_{t-1} and η_{t-1} are known, and when ϵ_{t-1} describes the effect of political noise. Looking ahead, an updated policy will take effect in period $t + 1$ based on period t information. This process repeats until an arbitrary terminal period T .

Returning for a moment to our SO₂ trading program example, firms were well aware of evolving information, political forces, and legal opinions that were shaping the regulatory environment for 2010 and beyond. This awareness influenced pollution behavior and permit prices beginning as early as 2004. From 2004 to 2006, these were arguably driven by better information about welfare-improving, societal marginal benefits reflected in the initially proposed CAIR, which suggested such benefits were well in excess of \$1600 per ton (U.S. EPA 2005). During 2006-2009, the influences were arguably non-welfare-improving political forces and legal decisions, as the market price diverged from marginal benefits. In our model terminology, consider the period 2004-2009 (and earlier years going back to 1995) as an elongated period t . This policy was established under the 1990 Clean Air Act Amendments, which we would call period $t - 1$. During 2004-2009, firms expected the policy to be updated based on this period t information for 2010 and beyond, which we would call period $t + 1$.

Solving for the optimal price policy is relatively straightforward because the choices each period are not linked by the firm. Choosing \tilde{p}_t such that $E[MB_{\text{noisy}}(h(\tilde{p}_t, \theta_t), \eta_t, \epsilon_{t-1}) - MC(h(\tilde{p}_t, \theta_t), \theta_t) | t - 1] = 0$, we have the government's "optimal" price policy when period $t - 1$ information is available and period t policy is being set. Note that in setting period t policy when available information and political noise are given by $t - 1$ indexed variables (e.g., we are actually in period $t - 1$ when t policy is set), the government is seeking to take expectations of θ_t and η_t but not ϵ_t . The value ϵ_{t-1} enters directly as the government's current period $t - 1$ special interests or current legal constraints are what matter for setting policy. As shown in the appendix, this leads to

$$\tilde{p}_t = E[p_t^* | t - 1] + (c_2 / (b_2 + c_2)) \epsilon_{t-1}.$$

The expressions p_t^* and q_t^* continue to represent the first-best outcome in terms of true aggregate net benefits, but the government's objective now deviates from the goal of maximizing expected, aggregate net benefits. Turning to optimal quantity policies, we show in the appendix that choosing

$$\tilde{q}_t = E[q_t^* | t - 1] + (\epsilon_{t-1} / (b_2 + c_2)) + \tilde{R}_t \left(q_{t-1}^* - E[q_{t-1}^* | t - 2] + \frac{\epsilon_{t-1} - \epsilon_{t-2}}{b_2 + c_2} \right)$$

and

$$\tilde{R}_t = \frac{p_{t-1}^* + (c_2 / (b_2 + c_2)) \epsilon_{t-1}}{\beta (E[p_t^* | t - 1] + (c_2 / (b_2 + c_2)) \epsilon_{t-1})}$$

as the quantity policy yields $q_t = q_t^* + \frac{\epsilon_t}{b_2 + c_2}$ each period. More precisely, the government acting in period t , with knowledge of θ_t and η_t and influenced by political noise ϵ_t , will exactly achieve $MB_{\text{noisy}}(q_t, \eta_t, \epsilon_t) - MC(q_t, \theta_t) = 0$ by setting period $t + 1$ policy. In each successive period, the government achieves its "first-best" outcome by setting policy for the next period based on its own noisy view of marginal benefits MB_{noisy} . But this outcome is clearly not first best from a aggregate welfare perspective.

What happens to our comparative advantage expression? As we show in the appendix, the noisy price policy obtains $q_t = q_t^* + (\epsilon_{t-1} - (\mu_t + (b_2/c_2)\nu_t) / (b_2 + c_2))$ and the noisy quantity policy obtains $q_t = q_t^* + \epsilon_t / (b_2 + c_2)$. Both deviate from the first-best q_t^* , and the comparative welfare

advantage of a noisy price policy compared to a noisy quantity policy simply compares the relative losses arising from these two deviations,

$$\Delta^N = \frac{1}{2(b_2 + c_2)} (\sigma_\omega^2 - (\sigma_\mu^2 + (b_2/c_2)^2 \sigma_\nu^2)).$$

Here, Δ^N will be positive only if the variance from the added noise distortion, $\omega_t = \epsilon_t - \epsilon_{t-1}$, under a quantity policy is larger than the variance from not getting the optimal price right $p_t^* - E[p_t^* | t-1]$ under a price policy. Rewriting $\epsilon_t = \epsilon_{t-1} + \omega_t$, we note that ϵ_{t-1} contributes equally to the quantity deviation in both periods and, therefore, does not appear in the comparative advantage expression as it equally disadvantages both policies. This is analogous to the disappearance of benefit uncertainty in the comparison of price and quantity outcomes in the traditional Weitzman framework.

5 Application to climate change policy

Climate change policy has been lurking in the background as an important application of this paper and has motivated examples of stylized facts and assumptions. But what happens when we put quantitative estimates of climate change mitigation costs and benefits, and their associated uncertainties, into our expressions for comparative advantage?

Ignoring political noise, our expression for Δ^U depends on the variance of cost and benefit shocks, as well as the slopes of marginal costs and benefits. A standard approximation in the climate change literature is that marginal benefits are flat and $b_2 = 0$ (e.g., Newell and Pizer, 2008). With that assumption, both b_2 and σ_ν^2 vanish from the expression and we only need estimates σ_μ^2 and c_2 . Newell and Pizer (2003) provide an estimate of $c_2 = 1.6 \times 10^{-7} \$/\text{ton}^2$. We use the recent update to the U.S. Government's Social Cost of Carbon (Interagency 2013) to calibrate the error in marginal benefits. Namely, the central estimate changed from \$24/ton to \$37/ton, or by \$13/ton, in 3 years. Assuming η follows a random walk, that gives 3-year variance of $169 (\$/\text{ton})^2$ or $\sigma_\mu^2 = 56 (\$/\text{ton})^2$ for the 1-year variation. We then have

$$\begin{aligned} \Delta^U &= \frac{-(\sigma_\mu^2 + (b_2/c_2)^2 \sigma_\nu^2)}{2(b_2 + c_2)} \\ &= \frac{-\sigma_\mu^2}{2c_2} \quad (\text{assuming } b_2 = 0) \\ &= \frac{-56 (\$/\text{ton})^2}{2(1.6 \times 10^{-7} \$/\text{ton}^2)} \\ &= -\$175 \text{ million.} \end{aligned}$$

That is, if policy is set for each year t based on year $t-1$ information (e.g., the policy is set one year in advance), the global loss from price policies is \$175 million. Of course, policies are not set every year. The recent Paris agreement suggests policies might be set every five years. In the first year, the loss would be \$175 million, but over five years, the loss would be \$2.6 billion.¹⁰ Were policies updated every twenty years, which seems more in line with U.S. policy adjustments (e.g., the 1990

¹⁰Because we assume benefits follow a random walk, the variance in the second year will be twice the variance in the first year; the variance in the third year will be three times the variance in the first year (and so on). The sum over t years is therefore $t(t+1)/2$ times the value in a single year.

Clean Air Amendments followed 20 years after the 1970 Clean Air Act; the 2010 adjustments to the SO₂ regulations came 20 years after 1990 amendments), the loss would be \$36.8 billion.

A tougher question is to estimate the variance of the political noise innovation ω for the expression Δ^N that allows for such noise. The 2015 Paris Agreement calls for limiting climate change to 2°, or possibly 1.5°, which would imply considerably higher views about marginal benefits. Meanwhile, the position of many climate skeptics would be that marginal benefits are close to zero (if not negative).

Rather than attempt to quantify how this may affect policy over time, we make the observation that when marginal benefits are flat, the condition for price regulation to be preferred simplifies to $\sigma_\omega^2 > \sigma_\mu^2$. That is, prices are preferred if changes driven by political noise are larger than changes driven by true changes in estimated marginal benefits. Though a subjective judgment, it certainly seems plausible that this condition would be satisfied.

6 Conclusion

After governments set their policies, new information often arises (or is revealed) about the benefits and costs of those policies. Much of the previous work comparing price and quantity regulation has focused on the importance of this information asymmetry in a one-period world or when policy remains fixed indefinitely. In this framework, Weitzman (1974) demonstrated that price-based regulation is preferred when marginal benefits are relatively flat compared to marginal costs and that the magnitude of the preference also depends on the variance of cost uncertainty.

Our paper argues that governments eventually, if not regularly, update key regulatory policies. In this paper, we have shown that when such policy updates seek to maximize welfare, and quantity regulation is tradable over time, this updating policy framework unequivocally favors quantity regulation. Trading over time creates an arbitrage condition whereby expected future price changes will change prices today. This can be used by the regulator to overcome the information asymmetry. Such direct anticipatory influences cannot arise under price-based regulation, which will only change prices at the moment the policy update comes into force or through indirect means (such as changes in long-term investments that also link costs over time).

When we consider the possibility of noisy policymaking, where something other than welfare maximization drives the policy updates, a trade-off between price and quantity regulation re-emerges. Quantity controls allow both new information about true costs and benefits *and new noise* to enter the arbitrage condition. The relative advantage expression now depends on the difference between these variances. When marginal benefits are flat (as in the case of carbon dioxide emissions and global climate change), the expression simplifies to depend only on the difference between the added noise variance and the added marginal benefit variance. This is quite different from the original Weitzman result and most of the existing literature.

The important lesson here is that dynamics matter, as does the policymaking process. The traditional (economic) debate over price versus quantity regulation has emphasized relative slopes. When we consider the reality of policymaking over time and intertemporally flexible policy, however, the emphasis is whether policymaking is noisy or not. Quantity-based regulation is arguably the best policy when regulation is set and updated truthfully over time. But if the policy process is noisy, price-based policies may be preferred. Our specific discussion of climate change suggests that the latter could be the case, as do a number of events over the past decade highlighting sources of political noise. Such events include the collapse of SO₂ markets in the wake of legal challenges to the Clean Air Interstate Rule, the exit of New Jersey from the Regional Greenhouse

Gas Initiative, debate over changes to the European Union Emissions Trading Scheme, and the repeal of Australia's carbon price scheme, among others. In many of these situations, the swing in prices (frequently to zero) would seem to exceed plausible changes to the true optimal price, recommending prices.

Another perspective on this result is simply that expectations matter in markets with intertemporal trading. That can be a good thing when the policy process is rational and well-behaved, but it can be a bad thing when the policy process is noisy and subject to significant swings in policy objectives. If policymaking is expected to be noisy and erratic, so too will be allowance markets.

References

- Adar, Zvi and James M. Griffin. "Uncertainty and the Choice of Pollution Control Instruments", *Journal of Environmental Economics and Management*, 3(1976):178-188.
- Biglaiser, Gary, John K. Horowitz, John Quiggin. "Dynamic pollution regulation". *Journal of Regulatory Economics*, 8(1995): 33-44.
- Chao, Hung-Po and Robert Wilson. "Option Value of Emission Allowances", *Journal of Regulatory Economics*, 5(1993): 233-249.
- Fell, Harrison, Ian A. Mackenzie, and William A. Pizer. "Prices versus Quantities versus Bankable Quantities." *Resource and Energy Economics*, 34.4 (2012): 607-23.
- Fishelson, Gideon. "Emission Control Policies under Uncertainty", *Journal of Environmental Economics and Management*, 3(1976): 189-197.
- Hoel, Michael, and Larry Karp. "Taxes versus Quotas for a Stock Pollutant." *Resource and Energy Economics*, 24.4 (2002): 367-84.
- Interagency Working Group on the Social Cost of Carbon, United States Government, Technical Support Document: Social Cost of Carbon for Regulatory Impact Analysis, (2010).
- Interagency Working Group on the Social Cost of Carbon, United States Government, Technical Support Document: Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis, (2013).
- Kaplow, Louis, and Steven Shavell. "On the superiority of corrective taxes to quantity regulation." *American Law and Economics Review* 4.1 (2002): 1-17.
- Kelly, David L. "Price and quantity regulation in general equilibrium", *Journal of Economic Theory*, 125(2005): 36-60.
- Murray, Brian C., Richard G. Newell, and William A. Pizer. "Balancing cost and emission certainty: An allowance reserve for cap-and-trade" *Review of Environmental Economics and Policy* 3(2009): 84-103.
- Newell, Richard and William Pizer. "Regulating stock externalities under uncertainty," *Journal of Environmental Economics and Management*, 45(2003): 416-432.
- Newell, Richard G. and William A. Pizer. "Indexed regulation" *Journal of Environmental Economics and Management* 56(2008): 221-233.

- Newell, Richard, William Pizer, and Jiangfeng Zhang. "Managing permit markets to stabilize prices", *Environment and Resource Economics*, 31(2005):133-157.
- Pizer, William A. "Combining price and quantity controls to mitigate global climate change", *Journal of Public Economics*. 85(2002): 409-434.
- Stavins, Robert. "Correlated uncertainty and the choice of pollution control instruments", *Journal of Environmental Economics and Management*, 30(1996): 218-252.
- Stranlund, John K. and Yakov Ben-Haim. "Price-based vs. quantity-based environmental regulation under Knightian uncertainty: An info-gap robust satisficing perspective", *Journal of Environmental Management*, 87(2008): 443-449.
- United States Environmental Protection Agency. Regulatory Impact Analysis for the Final Clean Air Interstate Rule, (2005).
- Weitzman, Martin. "Prices versus quantities", *Review of Economic Studies*, (1974): 477-491.
- Yohe, G. W. "Towards a general comparison of price controls and quantity controls under uncertainty", *Rev. Econom. Stud.* 45(1978): 229-238.
- Zhao, Jinhua. "Irreversible abatement investment under cost uncertainties: tradable emission permits and emissions charges", *Journal of Public Economics*, 87(2003): 2765-2789.

Appendix

Comparative Advantage in Two Periods

In the body we showed that the firm's reaction function is $h(\tilde{p}_t, \theta) = \hat{q} + \frac{\tilde{p}_t - c_1 - \theta}{c_2}$. The government's optimal policy in the first period is to tax at $\tilde{p}_1 = c_1$, meaning $h(\tilde{p}_1, \theta) = \hat{q} - \frac{\theta}{c_2}$. We also showed that under a quantity instrument, the firm produces $\tilde{q}^U/2 = \hat{q} + \frac{\eta - \theta}{b_2 + c_2}$.

Now define Δ^U , the difference between the expected net benefits under a price versus a quantity instrument. If $\Delta^U < 0$, quantity regulation is preferred. Since there is no more uncertainty in the second period, the net benefits in that period are the same across instruments, meaning we need only consider the difference in period 1. We now derive our primary result.

$$\begin{aligned}
\Delta^U &= E \left[(B(h(\tilde{p}_1^U, \theta), \eta) - C(h(\tilde{p}_1^U, \theta), \theta)) - (B(\tilde{q}^U/2, \eta) - C(\tilde{q}^U/2, \theta)) \right] \\
&= E \left[(b_1 - c_1 + \eta - \theta)(h(\tilde{p}_1^U, \theta) - \tilde{q}^U/2) - \frac{b_2 + c_2}{2} ((h(\tilde{p}_1^U, \theta) - \hat{q})^2 - (\tilde{q}^U/2 - \hat{q})^2) \right] \\
&= E \left[(\eta - \theta) \left(\frac{-\theta}{c_2} - \frac{\eta - \theta}{b_2 + c_2} \right) - \frac{b_2 + c_2}{2} \left(\left(\frac{-\theta}{c_2} \right)^2 - \left(\frac{\eta - \theta}{b_2 + c_2} \right)^2 \right) \right] \\
&= -E \left[(\eta - \theta) \left(\frac{\theta}{c_2} + \frac{\eta - \theta}{b_2 + c_2} \right) + \frac{b_2 + c_2}{2} \left(\left(\frac{-\theta}{c_2} \right)^2 - \left(\frac{\eta - \theta}{b_2 + c_2} \right)^2 \right) \right] \\
&= -E \left[\left(\frac{\frac{b_2}{c_2} \theta \eta - \frac{b_2}{c_2} \theta^2 + \eta^2 - \eta \theta}{(b_2 + c_2)} \right) + \frac{(b_2 + c_2)^2 \theta^2 - c_2^2 (\eta - \theta)^2}{2c_2^2 (b_2 + c_2)} \right] \\
&= -E \left[\left(\frac{\frac{b_2}{c_2} \theta \eta - \frac{b_2}{c_2} \theta^2 + \eta^2 - \eta \theta}{(b_2 + c_2)} \right) + \frac{((\frac{b_2}{c_2})^2 + 2\frac{b_2}{c_2}) \theta^2 - \eta^2 + 2\eta \theta}{2(b_2 + c_2)} \right] \\
&= -E \left[\frac{2\frac{b_2}{c_2} \theta \eta - 2\frac{b_2}{c_2} \theta^2 + 2\eta^2 - 2\eta \theta + ((\frac{b_2}{c_2})^2 + 2\frac{b_2}{c_2}) \theta^2 - \eta^2 + 2\eta \theta}{2(b_2 + c_2)} \right] \\
&= -E \left[\frac{\eta^2 + (\frac{b_2}{c_2})^2 \theta^2 + 2\frac{b_2}{c_2} \theta \eta}{2(b_2 + c_2)} \right] \\
&= \frac{-(\sigma_\eta^2 + (b_2/c_2)^2 \sigma_\theta^2)}{2(b_2 + c_2)}.
\end{aligned}$$

QED.

Multiple Periods

We first define the general auto-regressive error structure:

$$\begin{aligned}\eta_t &= \rho_\eta \eta_{t-1} + \mu_t \\ \theta_t &= \rho_\theta \theta_{t-1} + \nu_t,\end{aligned}$$

with $\eta_0 = \theta_0 = 0$ and $\mu_t = \nu_t = 0$ for $t \geq T$. The banking constraint with a dynamic trading ratio is given by

$$B_t = \tilde{R}_t B_{t-1} + \tilde{q}_t - q_t,$$

with $B_0 = 0$ and $B_t = 0$ for $t \geq T$. That is, we assume the problem effectively ends at period T , beyond which there is no uncertainty. The government can set the optimal policy and there is no need for banking.

Equilibrium

Our strategy here is to guess the government's rule for setting the allocation and dynamic trading ratio each period. We then show that, with these rules, the first-best quantity level is both feasible for the firm and satisfies the firm's first order conditions for profit maximization.

Our proposed allocation rule in every period is given by:

$$\tilde{q}_t = \hat{q} + \frac{\rho_\eta \eta_{t-1} - \rho_\theta \theta_{t-1}}{b_2 + c_2} + \tilde{R}_t \frac{\mu_{t-1} - \nu_{t-1}}{b_2 + c_2}.$$

The proposed dynamic trading ratio is given by

$$\begin{aligned}\tilde{R}_t &= \frac{p_{t-1}^*}{\beta E[p_t^* | t-1]} \\ &= \frac{c_1 + \frac{b_2 \theta_{t-1} + c_2 \eta_{t-1}}{b_2 + c_2}}{\beta \left(c_1 + \frac{b_2 \rho_\theta \theta_{t-1} + c_2 \rho_\eta \eta_{t-1}}{b_2 + c_2} \right)}\end{aligned}$$

for $t > 1$ where p_t^* is the first-best optimal price corresponding to q_t^* . $\tilde{R}_1 = 0$.

First we show that the first-best optimal quantity level $q_t^* = \hat{q} + \frac{\eta_t - \theta_t}{b_2 + c_2}$ is feasible for the firm. That is, it meets the terminal banking condition $B_T = 0$. Note that choosing this quantity level leads to a banking rule

$$B_t = \tilde{R}_t B_{t-1} + \tilde{q}_t - q_t^* = \tilde{R}_t \left(B_{t-1} + \frac{\mu_{t-1} - \nu_{t-1}}{b_2 + c_2} \right) - \frac{\mu_t - \nu_t}{b_2 + c_2},$$

where we have substituted the rule \tilde{q}_t . With $B_0 = \mu_0 = \nu_0 = 0$, we have $B_1 = -\frac{\mu_1 - \nu_1}{b_2 + c_2}$, $B_2 = -\frac{\mu_2 - \nu_2}{b_2 + c_2}$, ... $B_{T-1} = -\frac{\mu_{T-1} - \nu_{T-1}}{b_2 + c_2}$, $B_T = -\frac{\mu_T - \nu_T}{b_2 + c_2}$. With $\mu_T = \eta_T = 0$, we then have $B_T = 0$. Therefore, a choice of $q_t = q_t^*$ is feasible for the firm.

We now show that picking $q_t = q_t^*$ satisfies the first-order condition for the firm. Namely, each period the firm faces:

$$V_t(B_{t-1}, \theta_{t-1}, \eta_{t-1}, \nu_t, \mu_t) = \max_{B_t} -C(\tilde{R}_t B_{t-1} + \tilde{q}_t - B_t, \theta_t) + \beta E[V_{t+1}(B_t, \theta_t, \eta_t, \nu_{t+1}, \mu_{t+1}) | t],$$

where we have replaced q_t with B_t as the choice variable to simplify taking derivatives. The first-order condition is then

$$MC(q_t, \theta_t) = -\beta E \left[\frac{\partial V_{t+1}(B_t, \theta_t, \eta_t, \nu_{t+1}, \mu_{t+1})}{\partial B_t} \middle| t \right].$$

We also have:

$$\frac{\partial V_t(B_{t-1}, \theta_{t-1}, \eta_{t-1}, \nu_t, \mu_t)}{\partial B_{t-1}} = -\tilde{R}_t MC(q_t, \theta_t).$$

Thus, the first order condition is that $MC(q_t, \theta_t) = \beta E[\tilde{R}_{t+1} MC(q_{t+1}, \theta_{t+1})|t]$. Given R_{t+1} only depends on information at t , we can further simplify to

$$MC(q_t, \theta_t) = \beta \tilde{R}_{t+1} E[MC(q_{t+1}, \theta_{t+1})|t].$$

This is simply the no arbitrage condition with discounting, trading ratios, and uncertainty about future shocks. Namely, today's marginal cost must equal the discounted, trading-ratio-adjusted, expected marginal cost next period. Given the definition of $\tilde{R}_{t+1} = p_t^*/\beta E[p_{t+1}^*|t]$, this condition is satisfied when $q_t = q_t^*$. QED.

Comparative Advantage

As shown above, the quantity policy achieves the first best outcome in every period. Therefore, the comparative advantage is simply the deadweight loss associated with the price instrument. Under a price instrument, the government sets a price in order to maximize welfare conditional on $t - 1$ information, namely,

$$E[MC(h(\tilde{p}_t, \theta_t), \theta_t)|t - 1] = E[MB(h(\tilde{p}_t, \theta_t), \eta_t)|t - 1],$$

where $h(\tilde{p}_t, \theta_t)$ is the previously-defined reaction function,

$$h(\tilde{p}_t, \theta_t) = \hat{q} + \frac{\tilde{p}_t - c_1 - \theta_t}{c_2}.$$

Substituting (and canceling $b_1 = c_1$), we have

$$\rho_\theta \theta_{t-1} + c_2 \left(\frac{\tilde{p}_t - c_1 - \rho_\theta \theta_{t-1}}{c_2} \right) = \rho_\eta \eta_{t-1} - b_2 \left(\frac{\tilde{p}_t - c_1 - \rho_\theta \theta_{t-1}}{c_2} \right),$$

or, rearranging,

$$\tilde{p}_t = c_1 + \frac{b_2 \rho_\theta \theta_{t-1} + c_2 \rho_\eta \eta_{t-1}}{b_2 + c_2}.$$

Using the same reaction function, this results in quantity,

$$h(\tilde{p}_t, \theta_t) = \hat{q} + \frac{\rho_\eta \eta_{t-1} + (b_2/c_2) \rho_\theta \theta_{t-1}}{b_2 + c_2} - \frac{\theta_t}{c_2}.$$

The optimal quantity is

$$q_t^* = \hat{q} + \frac{\eta_t - \theta_t}{b_2 + c_2}.$$

The difference between the quantities is

$$\begin{aligned} q_t^* - h(\tilde{p}_t, \theta_t) &= \frac{\eta_t - \rho_\eta \eta_{t-1}}{b_2 + c_2} + \frac{\theta_t}{c_2} - \frac{\theta_t + (b_2/c_2)\rho_\theta \theta_{t-1}}{b_2 + c_2} \\ &= \frac{\mu_t + (b_2/c_2)\nu_t}{b_2 + c_2}. \end{aligned}$$

This is analogous to the figure in section 3.3, where the quantity deviation was $\frac{\eta + (b_2/c_2)\theta}{c_2 + b_2}$, with η and θ replaced by their innovations. Following a similar graphical proof, equation (7) becomes

$$\Delta U = \frac{-(\sigma_\mu^2 + (b_2/c_2)^2 \sigma_\nu^2)}{2(b_2 + c_2)}.$$

QED.

Noise

Prices

Given the condition

$$E[MB_{\text{noisy}}(h(\tilde{p}_t, \theta_t), \eta_t, \epsilon_{t-1}) - MC(h(\tilde{p}_t, \theta_t), \theta_t)|t-1] = 0$$

using the reaction function above, the definitions of noisy marginal benefits (9) and marginal costs (4), and canceling $b_1 = c_1$ we have

$$\rho_\eta \eta_{t-1} - b_2 \left(\frac{\tilde{p}_t - c_1 - \rho_\theta \theta_{t-1}}{c_2} \right) + \epsilon_{t-1} = \rho_\theta \theta_{t-1} + c_2 \left(\frac{\tilde{p}_t - c_1 - \rho_\theta \theta_{t-1}}{c_2} \right).$$

Re-arranging, we have

$$\tilde{p} = c_1 + \frac{c_2(\rho_\eta \eta_{t-1} + \epsilon_{t-1}) + b_2 \rho_\theta \theta_{t-1}}{b_2 + c_2} = E[p_t^*|t-1] + \frac{c_2 \epsilon_{t-1}}{b_2 + c_2}.$$

Quantity Equilibrium

We now have a new policy rule

$$\tilde{q}_t = E[q_t^*|t-1] + (\epsilon_{t-1}/(b_2 + c_2)) + \tilde{R}_t \left(q_{t-1}^* - E[q_{t-1}^*|t-2] + \frac{\epsilon_{t-1} - \epsilon_{t-2}}{b_2 + c_2} \right)$$

and

$$\tilde{R}_t = \frac{p_{t-1}^* + (c_2/(b_2 + c_2))\epsilon_{t-1}}{\beta(E[p_t^*|t-1] + (c_2/(b_2 + c_2))\epsilon_{t-1})}.$$

As before, we show that the firm's choice $q_t = q_t^* + \epsilon_t/(b_2 + c_2)$ (implying $p_t = p_t^* + (c_2/(b_2 + c_2))\epsilon_t$) is both feasible given the terminal banking condition $B_T = 0$ and satisfies the firm's first-order condition for profit maximization (the arbitrage condition $p_t = \beta \tilde{R}_t E[p_{t+1}|t]$).

To see feasibility note that this choice of q_t yields $B_1 = -(q_1^* - E[q_1^*|0] + \epsilon_1/(b_2 + c_2))$ assuming $\epsilon_0 = \theta_0 = \eta_0 = 0$ and more generally

$$B_t = - \left(q_t^* - E[q_t^*|t-1] + \frac{\epsilon_t - \epsilon_{t-1}}{b_2 + c_2} \right)$$

for $t < T$. Assuming $\omega_T = \mu_T = \nu_T = 0$, we have $B_T = 0$. Thus, $q_t = q_t^* + \epsilon_t/(b_2 + c_2)$ is feasible for the firm.

As before, the no arbitrage condition will imply $p_t = MC(q_t, \theta_t) = \beta \tilde{R}_{t+1} E[MC(q_{t+1}, \theta_{t+1})|t] = \beta \tilde{R}_{t+1} E[p_{t+1}|t]$. Assuming $\tilde{R}_t = (p_{t-1}^* + (c_2/(b_2 + c_2))\epsilon_{t-1})/(\beta(E[p_t^*|t-1] + (c_2/(b_2 + c_2))\epsilon_{t-1}))$, this is satisfied. QED.

Comparative Advantage

In the previous comparative advantage derivation we showed that a quantity deviation of $-(\mu_t + (b_2/c_2)\nu_t)/(b_2 + c_2)$ from the first-best q_t^* , which corresponds to a price deviation of $-(b_2\nu_t + c_2\mu_t)/(b_2 + c_2)$ each period from p_t^* , lead to a deadweight loss of

$$\Delta^U = \frac{-(\sigma_\mu^2 + (b_2/c_2)^2\sigma_\nu^2)}{2(b_2 + c_2)}$$

each period.

We now have a deviation of $-(b_2\nu_t + c_2(\mu_t + \epsilon_{t-1}))/ (b_2 + c_2)$ from p_t^* under the noisy price policy and $c_2\epsilon_t/(b_2 + c_2)$ under the noisy quantity policy each period. These correspond to deviations from q_t^* of $(\epsilon_{t-1} - (\mu_t + (b_2/c_2)\nu_t))/ (b_2 + c_2)$ for price regulation and $\epsilon_t/(b_2 + c_2)$ for quantity regulation. Comparing to the previous result, we can compute welfare changes relative to the first best as

$$\frac{-(\sigma_{\epsilon_{t-1}}^2 + \sigma_\mu^2 + (b_2/c_2)^2\sigma_\nu^2)}{2(b_2 + c_2)}$$

for the price policy and

$$\frac{-\sigma_{\epsilon_t}^2}{2(b_2 + c_2)}$$

for the quantity policy each period. We can take the difference to find the comparative advantage of noisy prices compared to noisy quantities,

$$\Delta^N = \frac{\sigma_\omega^2 - (\sigma_\mu^2 + (b_2/c_2)^2\sigma_\nu^2)}{2(b_2 + c_2)},$$

where $\sigma_\omega^2 = \sigma_{\epsilon_t}^2 - \sigma_{\epsilon_{t-1}}^2$ under the assumption $\epsilon_t = \epsilon_{t-1} + \omega_t$. QED.