

NBER WORKING PAPER SERIES

PRICES VERSUS QUANTITIES WITH POLICY UPDATING

William A. Pizer
Brian Prest

Working Paper 22379
<http://www.nber.org/papers/w22379>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2016

The authors appreciate comments from seminar participants at the 2016 AERE Summer Meeting, the Beijing Institute of Technology, the 2016 NBER Spring Meeting, and the Duke Environmental Economics Workshop, as well as Gilbert Metcalf, Terry Dinan, and Roger von Haefen. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by William A. Pizer and Brian Prest. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Prices versus Quantities with Policy Updating
William A. Pizer and Brian Prest
NBER Working Paper No. 22379
June 2016
JEL No. D62,D92,H41,Q58

ABSTRACT

This paper considers how policy updates and trading of regulated quantities over time changes the traditional comparative advantage of prices versus quantities. Quantity regulation that can be traded over time leads firms to set current prices equal to expected future prices. A government seeking to maximize net societal benefits can take advantage of this behavior with a sequence of quantity policy updates that achieves the first best in all periods. Under price regulation where current prices remain fixed until future policy changes occur, no such opportunity exists to achieve the first best, and prices are never preferred. However, if we assume policy updates are driven in part by political "noise" rather than maximizing net societal benefits, the result changes and prices can again be preferred. The comparative advantage now depends the relative variance of noise shocks compared to true cost and benefit shocks. This contrasts sharply with the traditional comparative advantage that depends on the relative slopes of marginal costs and benefits. Applied to climate change, we estimate the comparative advantage of intertemporally tradable quantities (over prices) to be \$2 billion over five years. This estimate grows if updates occur less frequently or could be made negative by political noise.

William A. Pizer
Sanford School of Public Policy
Duke University
Box 90312
Durham, NC 27708
and NBER
billy.pizer@duke.edu

Brian Prest
36 Intuition Circle
Durham, NC 27705
brian.prest@duke.edu

Prices versus Quantities with Policy Updating*

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

June 21, 2016

Abstract

This paper considers how policy updates and trading of regulated quantities over time changes the traditional comparative advantage of prices versus quantities. Quantity regulation that can be traded over time leads firms to set current prices equal to expected future prices. A government seeking to maximize net societal benefits can take advantage of this behavior with a sequence of quantity policy updates that achieves the first best in all periods. Under price regulation where current prices remain fixed until future policy changes occur, no such opportunity exists to achieve the first best, and prices are never preferred. However, if we assume policy updates are driven in part by political “noise” rather than maximizing net societal benefits, the result changes and prices can again be preferred. The comparative advantage now depends the relative variance of noise shocks compared to true cost and benefit shocks. This contrasts sharply with the traditional comparative advantage that depends on the relative slopes of marginal costs and benefits. Applied to climate change, we estimate the comparative advantage of intertemporally tradable quantities (over prices) to be \$2 billion over five years. This estimate grows if updates occur less frequently or could be made negative by political noise.

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 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).

*The authors appreciate comments from seminar participants at the 2016 AERE Summer Meeting, the Beijing Institute of Technology, the 2016 NBER Spring Meeting, and the Duke Environmental Economics Workshop, as well as Gilbert Metcalf, Terry Dinan, 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

As we show in this paper, however, the outcomes can change significantly in a dynamic setting. In a model that extends Weitzman to multiple periods with policy updating, we show that intertemporally tradable quantity regulation can always achieve the first-best outcome. Price regulation can only achieve the first best in special cases and is never strictly preferred in this setting. Moreover, the comparative advantage now depends on uncertainty about both costs and benefits. We obtain this result despite the maintained assumption that the government sets its policy for each period based on previous-period information, while the firm chooses its action based on current-period information.

This result arises because of the interaction between policy updating and the arbitrage condition created by an intertemporally tradable quantity regulation. In the static model, cost shocks lead price and quantity regulation to diverge in different ways from the first best. The same outcomes would occur in our dynamic setting if the quantity instrument could not be traded across periods, regardless of whether policies are updated or not. With intertemporal trading, however, the opportunity to save (or borrow) regulated quantities for (or from) the next period implies an arbitrage condition between current and expected next-period prices as firms cost minimize. A benefit-maximizing government can take advantage of this behavior with a particular strategy of updating the quantity regulation in order to achieve the first best. It can do this even with its constraint of setting policy each period based on previous-period information.

By contrast, under price regulation there is no comparable arbitrage condition. The government's strategy of updating the price regulation can influence expectations about next-period prices and policies, but cost-minimizing behavior by firms does not change the current-period price that they face. That is, if firms shift polluting activity in response to expectations about future pollution taxes, they do not affect the current pollution tax. Coupled with the assumption that the current-period policy is set based on previous-period information, the best that a benefit-maximizing government can do with a price policy in this setting is to achieve a second-best outcome similar to the result in Weitzman (1974).

Viewed another way, a dynamic policy setting can relax the information requirements needed to achieve the first-best policy. An intertemporally tradable quantity policy need not be set to achieve the first-best outcome in any particular period so long as firms expect the quantity allocation to be adjusted appropriately in the future. Our proof requires a final period where uncertainty is finally resolved and, in general, the regulator may need to use trading ratios to manipulate the arbitrage condition. However, in a simple two-period model without discounting, we show how the government merely sets the first-best allocation in the second period plus a "top-up" for the first period.

Of course it far from obvious whether policy updates to a particular regulation are always driven by society's welfare maximization. With this in mind, our paper considers the possibility that the regulator's objective deviates from society's marginal net 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. This possibility — what we might call uncertain political "noise" — favors price regulation as the adverse influence is at least delayed until the policy takes effect.

The motivation for this line of thinking stems from an attempt to explain observed outcomes under real-world quantity regulation and to relate these outcomes to the question of comparative advantage. Figure 1 shows the price behavior in the SO₂ trading program from 1995 through 2010. Starting in 2004, future reforms to the program — a policy update in the form of the Clean Air Interstate Rule or CAIR — began to be debated. This new rule would have updated the existing policy to reflect newer evidence of improved health benefits from lower emission levels




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

(that is, higher marginal benefits of the policy). While none of the changes to the program would have come into effect prior to 2010, the current (spot) price of SO₂ emission permits moved in parallel to shifting expectations of future prices because of the ability to save or “bank” current permits for future use. Current prices adjusted to reflect the improved marginal benefit estimates that were to be implemented in the future, thereby achieving these higher net benefits in practice almost immediately. Had an SO₂ tax been in place, however, an expectation of a higher tax in 2010 would not have influenced the current tax. This feature arguably favors quantity regulation, as the planned improvement in the policy target is realized immediately, before the change has actually been implemented.

Starting in 2006, however, the price fell as the EPA reconsidered the proposed rules and later court challenges succeeded in vacating the regulation. In this case, particularly as the price fell to zero well before the new policy came into effect, an SO₂ tax would have at least maintained the previous level of net benefits until any policy update occurred. Thinking through events in this market, one sees that the distinction between price and quantity regulation relates to the equalization of current and expected future prices under quantity regulation. The comparative advantage then depends on the degree to which future policy updates and prices reflect welfare maximization or something else.

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 five 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 (1974) result and other relevant literature. We then show how the Weitzman result changes if we allow policy updates and intertemporal quantity trading. To illustrate our main results, we present a simple two-period model with a common shock across both periods and no discounting. After showing how intertemporally tradable quantity regulation can achieve the first best outcome, we introduce the possibility that policy updates might be driven by influences other than new information about costs and benefits, such as shifting special interests or evolving legal or other constraints. This reintroduces the possibility that price regulation is strictly preferred, with the comparative advantage determined by the magnitude of this policy noise relative to new information about costs and benefits. All of these results can be generalized to multiple periods with AR(1) error processes and discounting, as we demonstrate in section 5. 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 \$2 billion per five-year period, assuming policies are updated every five years, a number that grows if updates occur less frequently.

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 represented by

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

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

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.

linear marginal costs and benefits can be written as

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

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

Here, θ represents cost uncertainty while benefit uncertainty is captured in a similar way by η . For an arbitrary \hat{q} , the parameters c_1 and c_2 describe marginal costs in a flexible way, and similarly so for b_1 and b_2 . As in Weitzman, we assume $c_2 > 0$ and $b_2 \geq 0$, $MC(0, \theta) < MB(0, \eta)$ and $MC(q, \theta) > MB(q, \eta)$ for large enough q . This implies a single first-best outcome, regardless of whether q is a good or a bad.⁶

In his paper, the linear model is described 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 the key results, it simplifies the exposition and the expected optimal government policies. These become $\tilde{q}^O = \hat{q}$ for quantities and $\tilde{p}^O = c_1 = b_1$ for prices. Here and throughout, we use a tilde over a variable to indicate a chosen government policy; the superscript O more specifically refers to the *Original* Weitzman framework for policymaking.

With this setup, Weitzman derives his main result. When price and quantity regulations are set to maximize expected net benefits, the expected welfare advantage of prices is given by

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

Weitzman’s appealingly simple result (5) shows that the direction of the welfare advantage Δ^O 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

⁶Weitzman focuses on q as a good (e.g., clean air) for simplicity. However, our motivating example is pollution where marginal benefits may be flat and where quantity controls are more easily viewed as regulating a bad (e.g., emissions). There is nothing in the model or notation that favors one or the other, but the reader should keep in mind that for bads, marginal benefits and hence prices are negative; firms must pay rather than being paid for producing bads.

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. Kaplow and Shavell (2002) take a different tack. They propose to fix the price only after firms report their total emissions based on the expected marginal benefit schedule. This allows them to achieve welfare improvements over either fixed prices or fixed quantities, but raises questions as such policies have not been seen 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. Yates and Cronshaw (2001) consider the value of intertemporal permit trading in a deterministic setting, noting that allowing one-for-one banking may not be optimal and introduce the idea of trading ratios (something we return to in our multi-period model). Newell and Pizer (2003) and Fell, MacKenzie, and Pizer (2012) both extend the original Weitzman framework to multiple periods where benefits can depend on the accumulated level of the pollutant, rather than the annual flow. Despite the key difference that one allows intertemporal trading and the other does not, they find results broadly similar to Weitzman. Flatter marginal benefits and steeper marginal costs still favor prices, and vice-versa for quantity regulation. Importantly, these papers do not consider the possibility that policies might be updated.

Several papers do consider policy updates in a multi-period setting. Newell, Pizer, and Zhang (2005) are entirely focused on policy updating, but only in the limited sense of whether an intertemporally tradable 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 (a hybrid price-quantity policy), 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 intertemporal emissions trading has an advantage when policies are updated in response to meaningful new information. Like the previous paper, this abstracts from the notion of welfare and benefits. They assume a cumulative emission goal that is updated in the future and only seek to minimize costs.

Finally, Hoel and Karp (2002) compare welfare outcomes under price and quantity regulation with policy updates but without intertemporal trading of the quantity instrument. They focus on a stock pollutant, whereby the fluctuations in the accumulated stock that occur under a price policy can be attenuated with policy updates. This feature coupled with the lack of intertemporal trading of quantities gives prices an extra welfare advantage in this setting.

These papers focused on a multi-period setting can be categorized by whether they include four features: uncertainty, intertemporal trading, policy updating, and explicit welfare analysis. None of them contain all four, which leads us to the current paper. This paper can be viewed as extending and clarifying the notion that policy updates coupled with intertemporal trading seems to provide some sort of advantage for quantity controls. In contrast to previous work that included both of these features, however, we consider benefits and benefit uncertainty and focus on welfare. This allows us to show that policy updates seeking to maximize societal benefits lead to a clear preference for quantity regulation that can be traded over time.

3 Policy Updating

We now expand the earlier Weitzman model to two periods with the same, but uncertain, costs and benefits in both periods. The regulator sets the first-period price or quantity policy prior to learning about costs and benefit uncertainty, θ and η in (1) and (2), but we consider the possibility that s/he sets the second-period policy once the uncertainty is known. That is, we change the information asymmetry to represent a lag in policy implementation in the context of evolving knowledge rather than a perpetual difference in what may be known or unknown.

We pause for a moment to highlight this slight shift in interpretation. Weitzman focuses on the difficulty of the government to know true costs that are more naturally known to the firms that experience them. In the same way, one could view benefit uncertainty as a deep uncertainty about the true marginal benefits that the government simply cannot know. From such a vantage point, the government only can infer θ after firms respond to the policy and may never learn η at all. We take a slightly different perspective. We think of θ and η as an evolution of society’s best estimate of marginal costs and benefits over time. That is, everyone’s best guess is $\theta = \eta = 0$ in period zero. Information improves in period one. This may be due to firms’ responses to the government policy but is also a function of improved science and experience that arises from the steady march of time. In our view the government learns about θ and η alongside the firm but is unable or unwilling to act quickly enough to change the first-period policy directly.⁷ The government is able to set second-period policy based on this observed information 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 this pre-announced rule.⁸ The information structure and decisions over time are depicted in Figure 2.

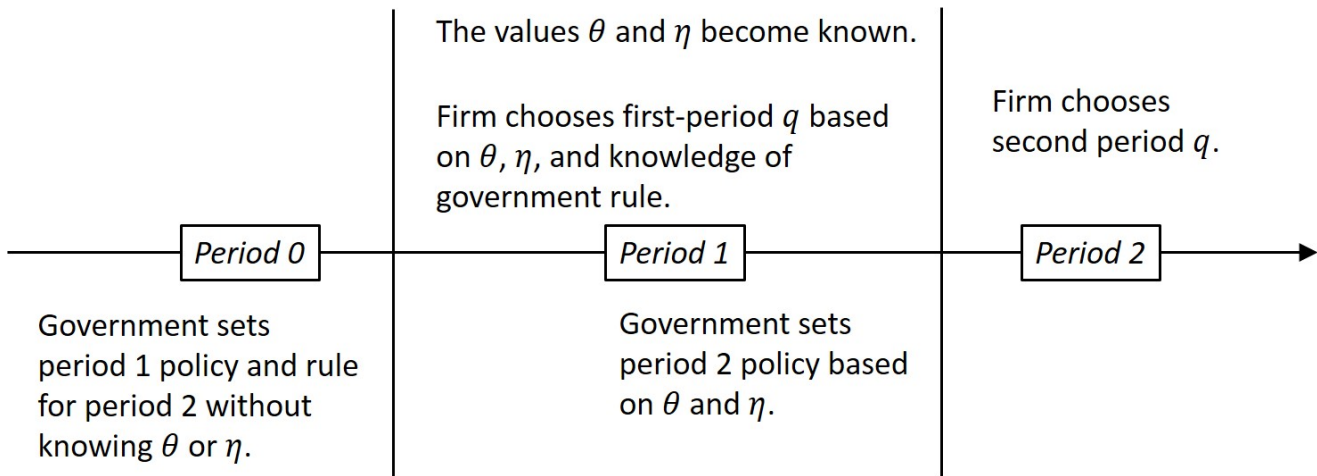


Figure 2: Timeline of events with updating in a simple two-period model

⁷We intentionally do not specify the length of each period. As in Hoel and Karp (2002), the length of each period primarily determines how often policies are updated. To the extent policymaking is a slow process, it is perhaps best to imagine a period being several or even many years long. Under this assumption, it makes sense to assume first-period policy will not be changed directly. While there are examples of tax changes enacted late in a given year being applied retroactively to liabilities incurred earlier in that year, using a retroactive mechanism over a much longer period seems infeasible in practice.

⁸In 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.

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 saved, or “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 $B = \tilde{q}_1 - q_1$. This bank can be in surplus 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:

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

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 knows 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:

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

That is, divide the cumulative quantity regulation equally over the two periods. 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,

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

recalling that \tilde{p}_t will be positive or negative depending on whether q_t is a good or a bad. 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. Given the definition of marginal costs in (3), 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 (that is, without updating). Absent updating, this replicates the earlier Weitzman policy result but now extended over two periods, which we denote with superscript W . In the case of intertemporally tradable quantities, it is clear from equation (6) that only the total quantity volume $\tilde{q} = \tilde{q}_1 + \tilde{q}_2$ matters as a policy choice variable and the allocation between periods does not matter. 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 the cumulative quantity is given by $\tilde{q}^W = 2\hat{q}$, which could be implemented with $\tilde{q}_t^W = \hat{q}$ in each period. 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, $\tilde{p}_1^U = c_1$ leading to equilibrium quantities of $\hat{q} - \frac{\theta}{c_2}$ 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 when all the uncertain outcomes are now known:

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

Substituting and rearranging, we have

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

or

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

This is the expression for the first-best, 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 denote this optimized policy. Note that because the government can choose the second-period allocation after it observes the uncertain outcomes, it can determine the cumulative quantity target \tilde{q} using the true values of θ and η , thereby circumventing the information asymmetry!

Working through the first-order condition

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

or, rearranging,

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

Letting \tilde{q}_1^U be arbitrary, we have

$$\tilde{q}_2^U = 2\hat{q} + 2\frac{\eta - \theta}{b_2 + c_2} - \tilde{q}_1^U = q^* + (q^* - \tilde{q}_1^U) \quad (7)$$

as the optimal, updated quantity policy. Intuitively, the second-period allocation is the first-best outcome for the second period $q^* = \hat{q} + \frac{\eta - \theta}{b_2 + c_2}$, plus a "top-up" ($q^* - \tilde{q}_1^U$) to adjust the first-period

allocation to the first-best level. 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. This suggests that the comparative advantage for the updating policies will simply reflect the first-period loss under the price policy, which we now show.

3.3 Comparative Advantage

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

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

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,

$$\begin{aligned} \Delta^U &= E [(B(h(\tilde{p}_1^U), \theta), \eta) - C(h(\tilde{p}_1^U), \theta), \theta)) - (B(\tilde{q}^U/2, \eta) - C(\tilde{q}^U/2, \theta))] \\ &= \frac{-(\sigma_\eta^2 + (b_2/c_2)^2\sigma_\theta^2)}{2(b_2 + c_2)}. \end{aligned} \quad (8)$$

We provide an algebraic proof of (8) in the appendix, but a simple graphical proof is as follows.⁹ In the first period, the updating price policy is simply the standard price policy, $\tilde{p}_1^U = c_1$ and leads 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 .

Whereas the quantity outcome is first best, 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 (8).

3.4 Discussion of Key Assumptions

It is clear that the price policy is never preferred in this model. 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. *Thus, quantities with updating is weakly 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

⁹The point where the x -axis intercepts the y -axis in Figure 3 should not be interpreted as the origin. It is zero for the x -axis but will be a negative number for the y -axis in the case of a bad. That is, this figure arises in the first quadrant for a good but the fourth quadrant for a bad.




Figure 3: Illustration of deadweight loss from price policy

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.

What's going on? It is clear that there is an important interaction between policy updating and intertemporal trading, both of which tend to be present in practice. We have demonstrated that without updating and with intertemporal trading, the original Weitzman result remains in both periods. It is similarly easy to see that without intertemporal trading and with updating, the original Weitzman result remains in the first period (because the first period is just the Weitzman setup). But these two features together lead to a different result. In this setting, if firms have expectations about second-period policy updates *before* they make choices in the first period, the potential to trade across periods leads them to equalize current and expected future prices driven by the expected policy update. The result must be different. But must it achieve first-best?

Several additional implicit or explicit assumptions lead to the first-best outcome. First, we make a rational expectation assumption that firms are, in fact, expecting what the government ultimately does. This seems reasonable, as anything else should eventually lead to revised expectations. Second, we need to identify a quantity-policy updating rule for the government such that the expected next-period prices indeed achieve the first-best outcome. In the simple two-period model, we made that easy by assuming the periods were identical with a single common shock and no discounting. We will see in section 5 that more generally, with different shocks in different periods

and more periods, there is still an updating rule that achieves the first-best in every period. Finally, we assume the government regularly updates policy based on societal net benefits and not some other objective. This is perhaps the most questionable assumption, to which we now turn.

4 Policy Updating with Noise

While improved cost-benefit information in future policy decisions can improve current-period regulatory outcomes under intertemporally tradable quantity controls, what if something other than improved information influences future policy? It is easy to imagine policy decisions being driven in part by special interests and more narrowly defined benefits rather than broad, 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 when the government updates its policy it acts *as if* marginal benefits in (4) 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. Put simply, suppose the government now makes use of a “noisy” marginal benefit function given by

$$MB_{\text{noisy}}(q, \eta, \epsilon) = MB(q, \eta) + \epsilon = b_1 + \eta - b_2(q - \hat{q}) + \epsilon \quad (9)$$

that it wishes to equate with marginal costs in order to update its 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 update policy. For simplicity, we assume there is no noise when policy is initially set for the first period, and ϵ refers to the distortion introduced in period 1 along with the other true shocks to costs and benefits. In section 5 we extend this to an evolving noise process.

In adding this notion of political noise, we make the same 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 the evolution of political noise than is necessarily reflected in current policy. More specifically, the government sets policy for period 1, \tilde{p}_1^N or \tilde{q}_1^N , before any political noise or shocks to costs and benefits occur. Here, superscript N refers to the policy chosen in our setup with noise. By the time it acts in period 1, the firm knows the value of ϵ as well as θ and η that will be used to update policies coming into effect in period 2. At the same time or soon after, the government sets \tilde{p}_2^N or \tilde{q}_2^N based on that same realization of θ , η , and ϵ .

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 changes 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 1. This policy was established under the 1990 Clean Air Act Amendments, which we would call period 0. During 2004-2009, firms expected the policy to be updated based on this period 1 information for 2010 and beyond, which we would call period 2.

Solving for the optimal price policy is relatively straightforward because the choices each period are not linked by the firm. Choosing \tilde{p}_1^N such that $E[MB_{\text{noisy}}(h(\tilde{p}_1^N, \theta), \eta, 0) - MC(h(\tilde{p}_1^N, \theta), \theta)] = 0$, we have the government's "optimal" price policy for period 1 as viewed in period 0. Note that the government is seeking to take expectations of θ and η but not ϵ . The value of ϵ — zero before the first period — enters directly, as the government's current special interests or current legal constraints are what matter for setting policy. This leads to

$$\tilde{p}_1^N = E[p^*]. \quad (10)$$

For period 2, uncertainty is revealed and noise is realized, so the government chooses \tilde{p}_2^N such that $MB_{\text{noisy}}(h(\tilde{p}_2^N, \theta), \eta, \epsilon) - MC(h(\tilde{p}_2^N, \theta), \theta) = 0$, or

$$\tilde{p}_2^N = p^* + \frac{c_2 \epsilon}{b_2 + c_2}. \quad (11)$$

The expression p^* represents the first-best price outcome in terms of true aggregate net benefits,¹⁰ but the government's objective now deviates from the goal of maximizing expected, aggregate net benefits. Using the firm's reaction function, which is unchanged from our basic model, the equilibrium quantity outcomes are:

$$\begin{aligned} q_1 &= h(\tilde{p}_1^N, \theta) = \hat{q} - \frac{\theta}{c_2} \\ &= q^* - \frac{\eta + (b_2/c_2)\theta}{b_2 + c_2} \\ q_2 &= h(\tilde{p}_2^N, \theta) = \hat{q} + \frac{\eta - \theta}{b_2 + c_2} + \frac{\epsilon}{b_2 + c_2} \\ &= q^* + \frac{\epsilon}{b_2 + c_2}. \end{aligned}$$

where q^* is the first-best outcome in each period. These quantity outcomes are the same as those in our basic model plus the noise distortions in the second period.

Under a quantity policy with updating and intertemporal trading, we know the firm will simply set $q_t = \frac{\tilde{q}_1^N + \tilde{q}_2^N}{2}$ and what matters is this cumulative allocation as before. With the revealed values of θ and η and the realized value of ϵ , the government can choose \tilde{q}_2^N based on an arbitrary choice of \tilde{q}_1^N such that $MB_{\text{noisy}} = MC$ each period, or

$$\begin{aligned} b_1 + \eta - b_2 \left(\frac{\tilde{q}_1^N + \tilde{q}_2^N}{2} - \hat{q} \right) + \epsilon &= c_1 + \theta + c_2 \left(\frac{\tilde{q}_1^N + \tilde{q}_2^N}{2} - \hat{q} \right) \\ \tilde{q}_2^N &= 2\hat{q} + 2\frac{\eta - \theta}{b_2 + c_2} + 2\frac{\epsilon}{b_2 + c_2} + 2\frac{b_1 - c_1}{b_2 + c_2} - \tilde{q}_1^N \\ &= q^* + \frac{\epsilon}{b_2 + c_2} + \left(q^* + \frac{\epsilon}{b_2 + c_2} - \tilde{q}_1^N \right), \end{aligned} \quad (12)$$

which is again the same as the allocations in our basic model plus adjustments for the noise distortion. Given that firm behavior divides the cumulative allocation over both periods, we have

$$q_1 = q_2 = \frac{\tilde{q}_1^N + \tilde{q}_2^N}{2} = q^* + \frac{\epsilon}{b_2 + c_2} \quad (13)$$

¹⁰As noted in section 3.2, the first-best price is $p^* = c_1 + \frac{b_2\theta + c_2\eta}{b_2 + c_2}$.

each period, which is the government’s “optimal” outcome given its own noisy view of marginal benefits, $MB_{\text{noisy}}(q_t, \eta, \epsilon)$, which includes the ϵ distortion. This outcome is clearly not first best from an aggregate welfare perspective.

What happens to our comparative advantage expression? In the second period, both policies produce the same outcome (namely $q_2 = q^* + \frac{\epsilon}{b_2+c_2}$), so the comparative advantage derives solely from how the policies differ in the first period. In period 1, the noisy price policy obtains $q_1 = q^* - \frac{\eta+(b_2/c_2)\theta}{b_2+c_2}$ and the noisy quantity policy obtains $q_1 = q^* + \frac{\epsilon}{b_2+c_2}$. Both deviate from the first-best q^* , 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,

$$\Delta^N = \frac{\sigma_\epsilon^2 - (\sigma_\eta^2 + (b_2/c_2)^2\sigma_\theta^2)}{2(b_2 + c_2)}, \quad (14)$$

Here, Δ^N will be positive only if the variance from the added noise distortion, ϵ , under a quantity policy is larger than the variance from not getting the optimal price right $p^* - E[p^*]$ under a price policy. If we think about the improved information about costs and benefits as a *signal*, this amounts to asking whether the signal to noise ratio is greater than 1. Note that if $b_2 = 0$ and marginal benefits are flat, this simplifies to $\Delta^N = \frac{\sigma_\epsilon^2 - \sigma_\eta^2}{2(b_2+c_2)}$, and the signal simply refers to improved information about benefits.

In our SO₂ example, it would seem to be an open question whether the signal to noise ratio is ≥ 1 . The period from 2004 to 2006 was a large signal. The period from 2006 to 2009 was an equally large noise. Therefore, in this case the question of whether prices should be preferred to quantities remains open.

The important point is that when there is policy updating and intertemporal quantity trading, so long as firms have rational expectations, what matters for welfare is the signal to noise ratio associated with policy updates, not the relative slopes of marginal costs and benefits. This is what governs the relative advantage of price regulation compared to quantity regulation in this simple model. And that is because the key feature of quantity-based regulation is not that it fixes the quantity, but because it allows expected future prices to govern the market.

5 Multiple Periods

We now extend our model and results to a multi-period setting. The purpose is to show that there are general price and quantity updating rules for the government that almost exactly replicate the previous two-period results comparing prices and quantities with noise (and without noise, if $\epsilon = 0$). In expanding to multiple periods, we now assume cost and benefit uncertainty follow arbitrary AR(1) processes. That is, we keep the same expressions for marginal benefits and costs each period, (3) and (4), but assume¹¹

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

¹¹One may be concerned that the linear-quadratic approximation of marginal costs and marginal benefits may break down over long periods of time. It can be shown that our results all continue to hold if the b and c terms in the marginal cost benefit functions change arbitrarily over time. In particular, if those terms are additionally indexed by t , the result continues to hold with appropriate adjustments to the allocation policy and trading ratios, but now the b_2 and c_2 values in the comparative advantage expression are similarly indexed by t , and hence the comparative advantage varies over time.

Similarly, we keep the same expression for noisy marginal benefits (9) but now assume the noise process follows a random walk,

$$\epsilon_t = \epsilon_{t-1} + \omega_t,$$

In contrast to the cost-benefit errors that may be mean reverting, we assume the government does not believe its current political position or constraints will revert back towards zero. Put another way, the current political situation is expected to persist.

We assume $\eta_0 = \theta_0 = \epsilon_0 = 0$ and μ_t, ν_t , and ω_t are i.i.d. random errors for $t > 0$ up to some final period T . After that point, for $t > T$ we assume $\mu_t = \nu_t = \omega_t = 0$ and q_t can be fixed forever at the desired rate by the government. This assumption of a final period analogous to the idea that at some point the regulatory problem is fully solved.

The government's policy choice for prices and quantities amounts to a dynamic programming problem each period, choosing today's policy to maximize a given objective subject to future policies doing the same and firm profit-maximization. We assume that only information up to period $t - 1$ is used to set period t policy. Meanwhile, firms are able to use period t information when they choose their (profit-maximizing) action in period t , including knowledge of how the government will behave in the future. This maintains the information structure assumed in the previous section as shown in Figure 4.

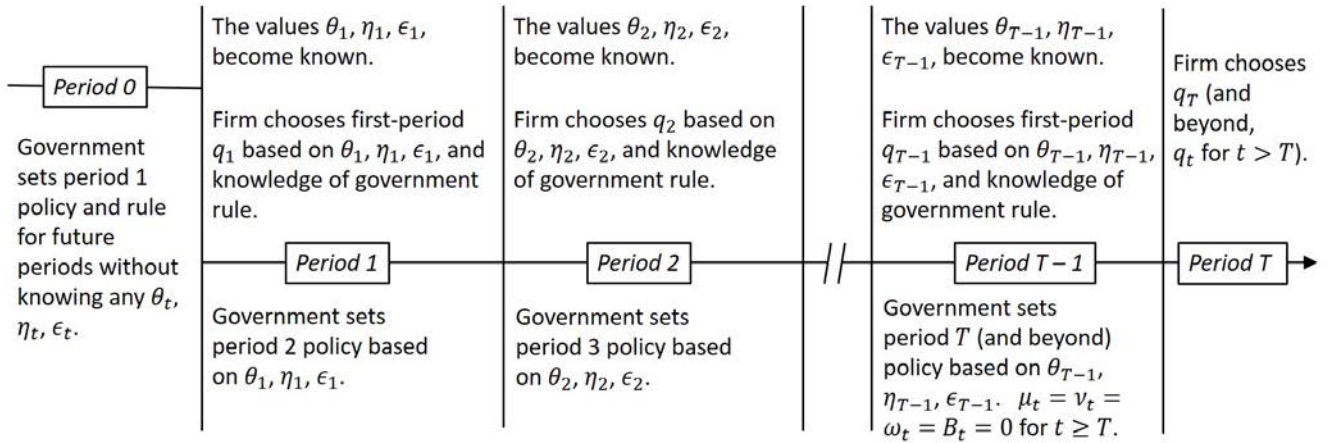


Figure 4: Timeline of events with updating in a T -period model

As before, firms facing quantity regulation with intertemporal trading have the flexibility to over- or under-comply each period. That is, the chosen quantity q_t can be above or below the allocation \tilde{q}_t . With only two periods, we defined a bank that existed for a single period as the surplus or deficit given by $B = \tilde{q}_1 - q_1$, which was then applied to the second period allocation. With multiple periods, we might imagine a bank that accumulates based on $B_t = B_{t-1} + \tilde{q}_t - q_t$ until the final period T when the bank in that period is applied to the final period allocation, similar to the two-period model. Here B_t is the bank at the end of period t ,

Instead, we introduce the one substantive difference between the multiple-period and two-period models: In the multiple-period model, we assume the existence of non-unitary trading ratios between periods under quantity regulation and assume the regulator is able to update those ratios alongside updates to the quantity allocation. Specifically, we modify the aforementioned bank accumulation rule in the following way:

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

The trading ratio is given by \tilde{R}_t and converts a one-unit deficit or surplus of period $t - 1$ quantities at the end of period $t - 1$ into those available at the beginning of period t .

The idea of non-unitary trading ratios, which was previously proposed in Yates and Cronshaw (2001), is not as theoretical as it might seem. Such an approach has been proposed several times and used at least once. For example, 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 2015.¹² Similarly, 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.¹³ Further, New Zealand’s emissions trading scheme has applied a ‘‘two-for-one’’ rule under which one permit allows for two tons of emissions; that rule is expected to change, effectively diminishing the future value of banked allowances. Indeed, New Zealand’s Climate Change Minister has said of the rule, ‘‘It was always a temporary measure...carbon must cost more than it does right now.’’¹⁴ Finally, as China plans to transition from its regional carbon market pilots to a national-level program, it plans to allow permits from the pilots to be carried over to the national one at a discounted factor that depends on the degree of over-allocation in the pilots.¹⁵

Having defined the multiple-period problem, we now need to find government policy rules that satisfy the first-order conditions for their objective function and mimic our simple two-period results. For a price policy, the dynamic programming problem simplifies to a series of one-period optimizations. We show in the appendix that setting

$$\tilde{p}_t^M = E[p_t^*|t - 1] + \frac{c_2\epsilon_{t-1}}{b_2 + c_2} \quad (15)$$

satisfies the government’s first-order condition,

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

The superscript M indicates the multi-period result and $p_t^* = c_1 + \frac{b_2\theta_t + c_2\eta_t}{b_2 + c_2}$ is (as before) the first-best price outcome from a pure welfare perspective. That is, given the government’s position of setting policy one period ahead for period t , (15) represents its best effort to do so based on the available information η_{t-1} and θ_{t-1} and subject to political noise ϵ_{t-1} .

For a quantity policy, we also show in the appendix that, if the government follows the allocation rule

$$\tilde{q}_t^M = E[q_t^*|t - 1] + \frac{\epsilon_{t-1}}{b_2 + c_2} + \tilde{R}_t \left(q_{t-1}^* + \frac{\epsilon_{t-1}}{b_2 + c_2} - \left(E[q_{t-1}^*|t - 2] + \frac{\epsilon_{t-2}}{b_2 + c_2} \right) \right) \quad (16)$$

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

$$\tilde{R}_t^M = \frac{p_{t-1}^* + (c_2/(b_2 + c_2))\epsilon_{t-1}}{\beta(E[p_t^*|t - 1] + (c_2/(b_2 + c_2))\epsilon_{t-1})}, \quad (17)$$

then the choice

$$q_t = q_t^* + \frac{\epsilon_t}{b_2 + c_2} \quad (18)$$

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

¹³See Title VII, Subtitle B, Section 790.

¹⁴See <http://carbon-pulse.com/17084/>

¹⁵http://carbon-pulse.com/18710/?utm_source=CP+Daily&utm_campaign=ec84ec8ee3-CPdaily22042016&utm_medium=email&utm_term=0_a9d8834f72-ec84ec8ee3-37668485

is a feasible choice for the firm each period. In the above expressions, β is the discount rate used by the firm and $q_t^* = \hat{q} + \frac{\eta - \theta}{b_2 + c_2}$ is (as before) the first-best quantity outcome. Note that both \tilde{q}_t^M and \tilde{R}_t^M only depend on $t - 1$ information, and therefore meet the information structure that we assume. After demonstrating (18) is feasible, we show this choice also satisfies the firm's first-order condition for profit maximization. Finally, we note that (18) is exactly the government's choice to set $MB_{\text{noisy}}(q_t, \eta_t, \epsilon_t) = MC(q_t, \theta_t)$ each period. That is, when acting with period t information (setting period $t + 1$ policy), the government exactly achieves its period t optimum. Absent any noise ($\epsilon_t = 0 \forall t$), we have $MB = MC \forall t$ and the government's optimum exactly corresponds to the welfare-maximizing first-best outcome.

These multi-period results are logical extensions of our two-period results. The two-period price policy had different expressions for the first-period (10) and second-period (11) prices. However, these are just special cases of the multi-period price outcome (15) corresponding to the first and last period. When $t = 1$, we have $\epsilon_{t-1} = \epsilon_0 = 0$ as we have assumed $\epsilon_0 = 0$. That is, $\tilde{p}_1^M = E[p_1^*]$. This mimics the first-period price policy in the two-period model, which includes an expected first-best price term but no noise. And when $t = T$, $E[p_t^* | t - 1] = p_T^* - (\mu_T + (b_2/c_2)\nu_T) = p_T^*$ as we have assumed $\mu_t = \nu_t = 0$ for $t \geq T$. That is, $\tilde{p}_T^M = p_T^* + (c_2\epsilon_T)/(c_2 + b_2)$. This mimics the second-period price policy in the two-period model, where there is a noise term but expectations are unnecessary because uncertainty is resolved. More generally, when $T > 2$ and $1 < t < T$, the price policy expression in (15) includes both the expected first-best price and the noise distortion.

Compared to the two-period quantity outcome (13), the multi-period quantity outcome (18) now has a time index on both the quantity optimum q_t^* and noise ϵ_t . It is otherwise identical. The government's multi-period quantity rule (16) includes both a current-period allocation and a "top up" for the previous period, as does the two-period quantity rules (7) and (12).

The main difference is the introduction of the trading ratio \tilde{R}_t^M defined in (17). This is necessary to allow the flexible control of prices over time, as one-for-one trading would necessarily imply expected prices rising at $1/\beta$ compared to current prices. However, trading ratios will be close to unity when the discount rate is small and when ρ_θ and ρ_η are close to 1 (i.e., the cost and benefit processes are approximately random walks). In these cases, $\beta E[p_t^* | t - 1] \approx p_{t-1}^*$, so $\tilde{R}_t^M \approx 1$ and non-unitary trading ratios may not be necessary to approximate the government's optimal policy.

Turning to the comparative advantage of prices versus quantities in a multi-period setting, we have

$$\Delta^M = \frac{\sigma_\omega^2 - (\sigma_\mu^2 + (b_2/c_2)^2\sigma_\nu^2)}{2(b_2 + c_2)}. \quad (19)$$

for each period. Expression (19) again closely resembles the two-period expression (14), with the variances of η , θ , and ϵ replaced with those of the innovation terms μ_t , ν_t , and ω_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 . Meanwhile, under the quantity policy, the equilibrium outcomes incorporate both improved cost and benefit information *and* future noise. Owing to the assumption that noise follows a random walk, the variance of future noise (affecting quantities) is larger than current noise (affecting prices), with their difference being simply σ_ω^2 . This expression is derived in the appendix.

6 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 for the moment (e.g., letting $\sigma_\omega^2 = 0$), our expression for Δ^M depends on the variance of cost and benefit shocks and 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

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

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 per year. 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 \$2.4 billion in present value.¹⁶ Were policies updated every twenty years, which seems more in line with major U.S. policy adjustments (e.g., the 1990 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 \$25.5 billion in present value for each 20 year period.

A tougher question is to estimate the variance of the political noise innovation σ_ω^2 . The 2015 Paris Agreement calls for limiting climate change warming to 2°C, or possibly 1.5°C. Meanwhile, the position of many climate skeptics would be that marginal benefits are close to zero (if not negative). While both of these positions might be viewed as the basis for a legitimate MB estimate, they could also be viewed as a basis for political noise.

Rather than attempt to quantify σ_ω^2 with such information, 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.

¹⁶Because 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. This sum must also be discounted; we use same the 3% discount rate used to generate the estimates for the social cost of carbon.

7 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.

We argue that governments eventually, if not regularly, update key regulatory policies. Moreover, quantity regulation is typically tradable over time. In this case, we show that what matters about quantity policies is not that they fix quantities but that they allow current prices to respond to expected future policy updates. When such policy updates seek to maximize welfare, we show that quantity regulation generally achieves the first-best welfare outcome and prices are never preferred. This highlights an important interaction between policy updates and intertemporal trading: 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 near-term anticipatory influences cannot arise under price-based regulation, which will only change prices at the moment the future policy update comes into force and only influence near-term outcomes to the extent marginal costs are linked over time (through indirect means such as long-term investments).

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 with intertemporal trading and policy updates allow both new information about true costs and benefits *and new noise* to drive current prices. The relative advantage expression then 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 sign of 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.

There are important caveats to our results. We have already noted the possible need for trading ratios between periods depending on the nature of cost and benefit shocks and discounting. There is only limited experience with such ratios in practice, though that may be changing. Price and quantity policies are also rarely enacted in isolation; most exist alongside myriad other regulations. California, for example, has at least four major regulations operating in tandem with its emission trading program (Borenstein et al. 2015). We have not considered how such policies affect our results.

We assume there is a significant lag in policy updates between when the change is well-established and when it occurs. This is relatively standard in regulatory policy, where regulated firms are given considerable lead time and the process itself involves many steps. Tax policy changes, however, typically happen quickly (if not retroactively) to prevent avoidance behavior. It is unclear whether updates to price-based regulation would follow more of a tax policy precedent, or a regulatory policy precedent.

Finally, we assume that once new information is known, particularly persistent cost information, there is confidence that this will be figured into a policy update. If this is not the case, we return to the traditional Weitzman framework until such confidence arises. That is, imagine a cost shock permanently reduces the cost of achieving an emission cap, but there is no expectation of an update to tighten the cap. Without such expectations to raise prices, the Weitzman setup (and results)

would return.

Regardless of these caveats, the underlying intuition remains. The traditional (economic) debate over price versus quantity regulation has emphasized relative slopes. When we consider the reality of policy updates over time and intertemporally tradable quantity regulation, however, it is clear price expectations matter. That can be a good thing for quantity regulation 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. These concerns have a necessary role in comparing real-world price and quantity policies.

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.
- Borenstein, Severin, James Bushnell, Frank A Wolak, Matthew Zaragoza-Watkins. “Expecting the Unexpected: Emissions Uncertainty and Environmental Market Design.” NBER Working Paper 20999, (2015).
- 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.
- Yates, Andrew J. and Mark B. Cronshaw. "Pollution Permit Markets with Intertemporal Trading and Asymmetric Information", *Journal of Environmental Economics and Management*, 42(2001): 104-118.
- Zhao, Jinhua. "Irreversible abatement investment under cost uncertainties: tradable emission permits and emissions charges", *Journal of Public Economics*, 87(2003): 2765-2789.

A Appendix

A.1 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.

$$\begin{aligned}
\Delta^U &= E \left[(B(h(\tilde{p}_1^U, \theta), \eta) - C(h(\tilde{p}_1^U, \theta), \theta)) - (B(\tilde{q}^U/2, \eta) - C(\tilde{q}^U/2, \theta)) \right] \\
&= E \left[(b_1 - c_1 + \eta - \theta)(h(\tilde{p}_1^U, \theta) - \tilde{q}^U/2) - \frac{b_2 + c_2}{2} ((h(\tilde{p}_2^U, \theta) - \hat{q})^2 - (\tilde{q}^U/2 - \hat{q})^2) \right] \\
&= E \left[(\eta - \theta) \left(\frac{-\theta}{c_2} - \frac{\eta - \theta}{b_2 + c_2} \right) - \frac{b_2 + c_2}{2} \left(\left(\frac{-\theta}{c_2} \right)^2 - \left(\frac{\eta - \theta}{b_2 + c_2} \right)^2 \right) \right] \\
&= -E \left[(\eta - \theta) \left(\frac{\theta}{c_2} + \frac{\eta - \theta}{b_2 + c_2} \right) + \frac{b_2 + c_2}{2} \left(\left(\frac{-\theta}{c_2} \right)^2 - \left(\frac{\eta - \theta}{b_2 + c_2} \right)^2 \right) \right] \\
&= -E \left[\left(\frac{\frac{b_2}{c_2} \theta \eta - \frac{b_2}{c_2} \theta^2 + \eta^2 - \eta \theta}{(b_2 + c_2)} \right) + \frac{(b_2 + c_2)^2 \theta^2 - c_2^2 (\eta - \theta)^2}{2c_2^2 (b_2 + c_2)} \right] \\
&= -E \left[\left(\frac{\frac{b_2}{c_2} \theta \eta - \frac{b_2}{c_2} \theta^2 + \eta^2 - \eta \theta}{(b_2 + c_2)} \right) + \frac{((\frac{b_2}{c_2})^2 + 2\frac{b_2}{c_2}) \theta^2 - \eta^2 + 2\eta \theta}{2(b_2 + c_2)} \right] \\
&= -E \left[\frac{2\frac{b_2}{c_2} \theta \eta - 2\frac{b_2}{c_2} \theta^2 + 2\eta^2 - 2\eta \theta + ((\frac{b_2}{c_2})^2 + 2\frac{b_2}{c_2}) \theta^2 - \eta^2 + 2\eta \theta}{2(b_2 + c_2)} \right] \\
&= -E \left[\frac{\eta^2 + (\frac{b_2}{c_2})^2 \theta^2 + 2\frac{b_2}{c_2} \theta \eta}{2(b_2 + c_2)} \right] \\
&= \frac{-(\sigma_\eta^2 + (b_2/c_2)^2 \sigma_\theta^2)}{2(b_2 + c_2)}.
\end{aligned}$$

QED.

A.2 Multiple Periods with Noise

In this section, we prove our assertion from section 5 that the main result holds in multiple periods with noise. The general auto-regressive error structure for costs and benefits is given by

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

and the random walk structure for the noise term:

$$\epsilon_t = \epsilon_{t-1} + \omega_t$$

with $\eta_0 = \theta_0 = \omega_0 = 0$ and $\mu_t = \nu_t = \omega_t = 0$ for $t \geq T$.

A.2.1 Equilibrium under a Price Policy

In each period $t - 1$, the government sets a price for period t to maximize its noisy view of the expected net benefits function. The first order condition is

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

Using the reaction function above, the definitions of noisy marginal benefits (9) and marginal costs (3), 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 an equilibrium price of

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

Using the firm reaction function, $h(\tilde{p}_t, \theta_t) = \hat{q} + \frac{\tilde{p}_t - c_1 - \theta_t}{c_2}$, we can find the equilibrium outcomes under the price policy:

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

That is, under a price policy, the realized outcome in each period deviates from the societal first-best level by the amount $\frac{\epsilon_{t-1} - (\mu_t + (b_2/c_2)\nu_t)}{b_2 + c_2}$ each period.

A.2.2 Equilibrium under a Quantity Policy

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 its optimal policy and there is no need for banking.

Our strategy is to guess the government's rule for setting the allocation and dynamic trading ratio each period. We then show that, with these rules, this optimal 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

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

and the proposed dynamic trading ratio is given by

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

for $t > 1$ where $p_t^* = c_1 + \frac{b_2\theta_t + c_2\eta_t}{b_2 + c_2}$ is the first-best (i.e., excluding noise) price corresponding to q_t^* . We set $\tilde{R}_1 = 0$. We now 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 $p_t = \beta\tilde{R}_{t+1}E[p_{t+1}|t]$).

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 $\epsilon_0 = \theta_0 = \eta_0 = B_0 = 0$ 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 $t < T$. This can be shown by induction. That is, given this value for B_t , $q_t = q_t^* + \frac{\epsilon_t}{b_2 + c_2}$, and the allocation and banking rules shown above, we can show that the relationship holds for $t + 1$:

$$\begin{aligned} B_{t+1} &= \tilde{R}_{t+1}B_t + \tilde{q}_{t+1} - q_{t+1} \\ &= \tilde{R}_{t+1}B_t + \tilde{R}_{t+1} \left(q_t^* - E[q_t^*|t-1] + \frac{\epsilon_t - \epsilon_{t-1}}{b_2 + c_2} \right) + E[q_{t+1}^*|t] - q_{t+1}^* - \frac{\epsilon_{t+1} - \epsilon_t}{b_2 + c_2} \\ &= - \left(q_{t+1}^* - E[q_{t+1}^*|t] + \frac{\epsilon_{t+1} - \epsilon_t}{b_2 + c_2} \right) \end{aligned}$$

as desired. 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.

We now show that picking $q_t = q_t^* + \epsilon_t/(b_2 + c_2)$ satisfies the first-order condition for the firm. Namely, each period the firm faces:

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

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

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

We also have:

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

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 $\tilde{R}_t = (p_{t-1}^* + (c_2/(b_2 + c_2))\epsilon_{t-1}) / (\beta(E[p_t^* | t-1] + (c_2/(b_2 + c_2))\epsilon_{t-1}))$, and using the definition of $MC(q, \theta)$ from (1), this condition is satisfied when $q_t = q_t^* + \epsilon_t / (b_2 + c_2)$.

Thus, $q_t = q_t^* + \frac{\epsilon_t}{b_2 + c_2}$ is feasible for the firm and satisfies the first-order conditions of both the firm.

QED.

Note that this realized outcome deviates from the societal first-best outcome by the amount $\frac{\epsilon_t}{b_2 + c_2}$ in each period.

A.2.3 Optimality of the Proposed Quantity Policy to Government

In the previous section we showed that the proposed quantity policy was feasible and optimal for the firm. Now we show that the proposed policy is optimal for the government (i.e., it leads to the the government's optimal outcome). At period t the government wishes to maximize noisy net benefits based on the contemporaneous values θ_t , η_t , and ϵ_t . That is, at period t the optimal outcome q_t must satisfy the first order condition:

$$\begin{aligned} MB_{\text{noisy}}(q_t, \eta_t, \epsilon_t) &= MC(q_t, \theta_t) \\ b_1 + \eta_t - b_2(q_t - \hat{q}) + \epsilon_t &= c_1 + \theta_t + c_2(q_t - \hat{q}) \\ q_t &= \hat{q} + \frac{\eta_t - \theta_t + \epsilon_t}{b_2 + c_2} \\ &= q_t^* + \frac{\epsilon_t}{b_2 + c_2} \end{aligned}$$

As shown in the previous section, this is exactly the outcome realized under the proposed quantity policy. Therefore the proposed quantity policy is optimal for the government.

QED.

A.2.4 Comparative Advantage

In the previous sections, we showed that in each period we have a deviation from q_t^* under the price policy of $\frac{\epsilon_{t-1} - (\mu_t + (b_2/c_2)\nu_t)}{b_2 + c_2}$ and a deviation under the quantity policy of $\frac{\epsilon_t}{b_2 + c_2}$. These correspond to deadweight losses of $\frac{(\epsilon_{t-1} - (\mu_t + (b_2/c_2)\nu_t))^2}{2(b_2 + c_2)}$ and $\frac{\epsilon_t^2}{2(b_2 + c_2)}$, respectively. We can compute the expected welfare loss per period for each policy, 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. 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.