Policy Instrument Choice with Coasean Provision of Public Goods*

Christopher Costello and Matthew Kotchen

April 2, 2021

Abstract

We examine the interplay between environmental policy instrument choice (i.e., prices *vs.* quantities) and private provision of public goods, which in this context we denote 'Coasean provision.' Coasean provision captures private provision of environmental public goods due to consumer preferences for environmentally friendly goods and services, incentives for corporate environmental management, environmental philanthropy, and even overlapping jurisdictions of policy. We show theoretically that even in a world of perfect certainty, the presence of Coasean provision distinctly affects instrument choice based on the efficiency criterion. We generalize the analysis to account for uncertainty using the classic Weitzman (1974) framework, showing that Coasean provision results in a favoring of prices over quantities with uncertainty over either the marginal benefits or costs of pollution. Our findings suggest that the increasing prevalence of Coasean provision motivates a need in many settings to rethink the design of effective and efficient environmental policy instruments.

1 Introduction

The study of externality problems and solutions provides the foundation for much of environmental economics and policy. The seminal work of A. C. Pigou (1932) developed the basic theory of externalities and proposed a solution by means of Pigouvian taxes. His contribution provides the first foray into what many now refer to as the centralized approaches to environmental policy.[1] An extensive literature has evolved to examine the advantages and disadvantages of various centralized policy instruments, including taxes and subsidies,

**Costello:* University of California, Santa Barbara and NBER (costello@bren.ucsb.edu). *Kotchen:* Yale University and NBER (matthew.kotchen@yale.edu). We are greatful to Bryan Leonard, Andrew Plantinga, Spencer Banzhaf, and seminar participants at Yale University for helpful insights. Costello acknowledges the Bren Chair in Environmental Management and the Environmental Markets Lab.

[1] See Banzhaf (2020) for an interesting historical analysis of whether initial applications of Pigouvian taxes to pollution problems were driven by Pigou's original contribution or whether they emerged separately in other natural resource settings and were later attributed back.

direct standards, and systems of tradeable permits. There is, however, another stream of environmental economics that focuses on decentralized approaches to solving externality problems.[2] This literature has its foundation in the seminal contribution of Ronald Coase (1960). The basic idea is that if property rights are well-defined and there are no transaction costs, then parties can engage is decentralized bargaining to solve externality problems, thereby obviating the need for a centralized, top-down approach.[3]

This paper moves beyond the dichotomy between centralized and decentralized approaches to environmental policy to consider how the presence of incomplete Coasean bargaining affects the choice among centralized policy instruments. We assume two conditions as starting points for analysis. The first, which is quite standard, is that fully resolving some externality problems requires a centralized form of policy, and these are the environmental problems upon which we focus. Deryugina et al. (2020) provide a recent review of real-world applications of the Coase theorem to environmental problems and find examples that include polluters purchasing nearby lands, payments for ecosystem services, and land acquisitions to protect the supply of drinking water. Despite these selected examples, their conclusion echos that in most textbook treatments of Coasean solutions: bargaining is likely to efficiently resolve externality problems in a quite limited set of circumstances where the number of parties involved is exceedingly small, thus underscoring the need, at least in many applications, for centralized policies.

Our second starting point assumption is that implementing a centralized environmental policy alters the institutional setting in ways that can induce Coasean-type bargaining. When a policy is implemented in an otherwise unregulated setting, it establishes rights and responsibilities that can serve as *de facto* property rights, a mechanism for reducing transaction costs, or both. We refer to the possibility for such policy-induced bargaining as Coasean provision. Whereas Coasean bargaining is often discussed in contexts where side-payments can support first-best, efficient outcomes, our notion of Coasean provision captures what are more generally suboptimal outcomes consistent with private provision of a public good (Cornes and Sandler 1985; Bergstrom, Blume, and Varian 1986). In other words, rather than assume away the conditions that give rise to free riding, we acknowledge that free riding occurs, yet consider the potential importance that Coasean provision of public goods may still have on policy instrument choice.

In practice, does such Coasean provision occur after policies are implemented? We argue that the possibility is more than a theoretical curiosity; it is increasingly at play, often

[2]Some of this work is collectively referred to as free market environmentalism, and Anderson and Leal (2001) provide a general introduction.

[3]Mas-Colell et al. (1995) provide a textbook proof of the Coase theorem showing the additional assumptions needed: no income effects, perfect information, rationality, and no endowment effects.

at a large scale, across a range of environmental and natural resource concerns, including climate change, biodiversity conservation, pollution control, and fisheries management. For example, despite the fact that California has a cap and trade program on carbon dioxide emissions, we still observe California companies making unilateral commitments to privately reduce emissions. While land development is commonly regulated and taxed, we regularly see the private purchase of land for conservation purposes. And although fisheries are often regulated with catch limits, seafood supply chains are increasingly committing to procure only sustainably caught seafood. Underlying these examples, and many others, is Coasean provision of environmental public goods motivated by consumer preferences for environmentally friendly goods and services, incentives for corporate environmental management, and direct philanthropy—all of which occur under the backdrop of centralized policy.

Also consistent with our framework are environmental or natural resource policies that take place at different levels of governance or jurisdictions. There exists a literature on nested state and federal environmental regulations (Goulder and Stavins 2011; Goulder, Jacobsen, and van Benthem 2012; Levinson 2012), and our analysis illuminates ways in which policy interactions will depend on the policy instrument choice and level of stringency. For example, many states and cities in the United States have climate policies in place that are independent of, yet contribute to, emission targets at higher levels of government. And outside the United States, for example, the city of Copenhagen has made a public commitment to carbon neutrality, despite the fact that Denmark has a nation-wide carbon tax.[4] Moreover, there are circumstances where one country seeks environmental or natural resource protection in another country (e.g., developed countries seeking to prevent deforestation in developing countries) and our results show how the efficacy and efficiency of these efforts will depend on characteristics of the environmental policies that a country has in place.

The fundamental question that we consider is how the presence of Coasean provision might affect policy instrument choice. We develop a theoretical model with an industry that benefits from pollution and citizens that experience the costs of pollution. A regulator chooses between policy instruments (an emissions tax or a cap-and-trade program) and determines the level policy stringency. Polluters respond by maximizing profits, but are also influenced financially by citizens, who may desire greater abatement than that targeted by the regulation. Any such privately provided abatement on the part of citizens amounts to Coasean provision, whereby citizens explicitly pay polluters to emit less, make purchasing decisions based on a company's environmental commitments, participate in the cap-and-

[4] While a key question in the existing literature on policy interactions for climate change is "leakage," this is not the topic of concern here. Instead, our results are driven by the different incentives for private provision of pubic goods that policy instrument choice creates.

trade program by retiring permits, or combinations of these. Key questions are then: how does instrument choice affect Coasean provision, and conversely, how does the presence of Coasean provision affect instrument choice?

Few bells and whistles are required to generate novel and policy-relevant results. While Weitzman (1974) focuses on the role of uncertainty in his seminal contribution, we start with a deterministic setup. Our first main finding is that the well-known symmetry between price and quantity instruments no longer applies in settings where Coasean provision is possible. An underlying reason is that taxes provide an implicit subsidy to Coasean provision. Surprisingly, the same is not true of auctioned permits in a cap and trade program, even when the auction price and tax are equivalent. This implies, for example, that implementing the seemingly first-best price or quantity instrument (ignoring the potential for Coasean provision) is efficient for the cap, but never for the tax. More generally, we show that the level of regulatory stringency (where seemingly first-best is just one case) affects the comparison between such myopically equivalent policies. And here the results differ: caps continue to dominate when the policy stringency is sufficiently strong, otherwise taxes are more efficient

We then focus on comparisons between conditionally optimal policies, under both certainty and uncertainty. That is, we assume a regulator takes account of the possibility for Coasean provision and thereby seeks to compare policies where the level of stringency is chosen optimally for either the price or quantity instrument. In this case, with certainty, both instruments can implement the first-best level of pollution, and the level of overall efficiency is the same as that which would arise with a prohibition on Coasean provision. Accounting for uncertainty in this framework using the classic Weitzman (1974) approach produces further insights. In contrast to Weitzman (1974), which is a special case of our analysis, we find that uncertainty in the marginal benefits or costs of pollution can affect the ex ante, optimally chosen level of a quantity instrument. The main result from adding uncertainty to our model, however, is that, compared to the standard results, Coasean provision tends to favor prices over quantities, because Coasean provision plays a more prominent role in the former and helps to offset welfare losses from getting the policy "wrong" *ex post*.

Taken as a whole, our findings suggest that the increasing prevalence of Coasean provision in many settings calls for a rethinking of the standard framework for evaluating environmental policies. Policy instrument choice can have a significant impact on the environmental commitments of individuals, companies, and states, and vice-versa, with clear implications for economic welfare.

Although we are not aware of any other research that focuses on the same set of questions, there are a few related contributions upon which we build. The first is the so-called

Buchanan-Stubblebine-Turvey Theorem, which considers how the simultaneous presence of a Pigouvian tax and Coasean bargaining will result in inefficiency (Buchanan and Stubblebine 1962; Turvey 1963). While the basic mechanism underlying that result is at play in our analysis, the framework here is more general because we do not assume the limiting case of perfectly efficient Coasean bargaining. Baumol (1972) argued that the Buchanan-Stubblebine-Turvey setup is implausible in more realistic settings because Coasean bargaining becomes impossible with a large number of actors, where the transaction costs are simply too high. Instead, Baumol (1972) assumes no bargaining at all, which gives rise to the standard framework for comparing centralized policy instruments. Our analysis can thus be viewed as a haromonization of the "Coasean-bargaining-only" and "centralized-policy-only" approaches that is motivated by more contemporary, real-world observations about the presence of privately provided environmental public goods, even when existing policies are in place. Accordingly, our contribution falls in line with the recommendation of Banzhaf et al. (2013) for more research that seeks to bridge the useful insights of both Pigouvian and Coasean approaches to environmental management.[5] In doing so, we also provide a generalization of the canonical Weitzman (1974) framework for policy instrument choice under uncertainty.

In the next section, we describe the model setup, making explicit our definition of Coasean provision. Section 3 defines the policy instruments that we consider, along with the equilibrium conditions that emerge in the presence of Coasean provision. Section 4 considers instrument choice between myopically equivalent policies, where the aim is to illustrate how ignoring Coasean provision distorts standard conclusions about policy instrument choice. Section 5 focuses on a comparison of conditionally optimal instruments. Sections 6 through 9 generalizes the analysis to account for uncertainty in the marginal benefits or costs of pollution. Section 10 concludes with a summary and discussion.

2 Model Setup

We develop the simplest possible model that illustrates the key ideas. "Industry" has demand for pollution, which is a public bad and imposes costs on "citizens." That is, industry benefits from pollution, and citizens benefit from abatement. The initial setup is static and deterministic, although we consider generalizations with uncertainty later in the paper. The aggregate level of pollution is denoted Q. Industry benefits according to a strictly

[5] Another paper by MacKenzie and Ohndorf (2016) analyzes how the inefficiency brought about by the Buchanan-Stubblebine-Turvey Theorem can be offset by a reduction in the costs of establishing property rights, thereby providing a potential argument in favor of Pigouvian taxes. Their analysis does not, however, draw comparisons with other policy instruments or uncertainty.

increasing function $B(Q)$, and we assume the marginal benefits $MB(Q)$ are decreasing. The costly damages of pollution are given by a strictly increasing function $D(Q)$, and we assume the marginal damages $MD(Q)$ are increasing. The aggregate marginal damages are the sum of marginal damages across $j = 1, ..., N$ citizens such that we can write $MD(Q) = \sum_{j=1}^{N} MD_j(Q)$.[6]

We focus on the two classical policy instruments. One is a tax of rate τ applied to each unit of Q. The other is a cap Ω such that the quantity of pollution must satisfy $Q \leq \Omega$. When a cap is used, Ω permits are auctioned off at the highest bid price that clears the market. We assume, following standard approaches, that all revenue from either the tax or auctioned permits is used to provide lump-sum social benefits. Standard analyses in environmental economics are based on the relation $\tau = MB(\Omega)$, which under certainty defines two equivalent instruments with respect to the implied level of Q and overall efficiency. One particular level of policy stringency, which is often the focal point of economic analysis, is that of first-best, defined as the tax and quantity instruments that maximize efficiency by satisfying $\tau = MB(\Omega) = MD(\Omega)$.

The seminal contribution of Ronald Coase (1960) is often considered a reinterpretation of the preceding framework to analyze circumstances where neither of the policy instruments are needed to obtain the efficient, first-best outcome. The Coase Theorem holds that under certain conditions, negotiated bargaining will take place between the two sides (i.e., industry and citizens), and the optimal level of pollution will arise as a result of compensating side payments. An entire literature has emerged to add precision to the conditions that give rise to the Coase Theorem and its potential applications, but the most salient and policy relevant tend to be the establishment of clearly defined property rights and the need for zero transaction costs.[7]

While Coasean bargaining is often viewed as an alternative to other policy interventions (e.g., taxes and caps), our focus here is not on comparing centralized *vs.* decentralized approaches. Instead, we consider the efficiency and distributional implications of policy-induced, Coasean bargaining on the choice between centralized policy instruments. Our starting point is one where an environmental externality exists (creating an environmental public bad), which means that any preexisting Coasean bargaining (if it occurred at all) did not completely resolve the market inefficiency. In particular, the initial unregulated level of

[6]Note that this adding up condition assumes that Q is exogenously given and not the result of private provision of abatement (i.e., less pollution) on the part of citizens, which is a topic that we discuss below and is central to the paper. The distinction here is that when Q is exogenously given, citizens do not have to pay for its provision.

[7]See Medema (2019) for a recent and comprehensive review of the literature related to the Coase Theorem in honor of its sixty year anniversary.

aggregate pollution satisfies $MB(Q_{max}) = 0$. Nevertheless, imposing a centralized, environmental policy alters the institutional setting such that policy-induced, Coasean bargaining might subsequently occur, for reasons that we now discuss related to both property rights and transaction costs.

While environmental policy instruments in the framework presented above are typically viewed as a way for governments to modify polluter incentives in more socially efficient ways, implementing a policy also plays another, perhaps overlooked, role. When a policy is implemented in an otherwise unregulated setting, it establishes the rights and responsibilities of industry and citizens, thereby acting as a *de facto* designator of property rights. That is, it delineates a property right to the polluting industry (e.g., each firm is allowed to pollute X tons) and to the citizens (e.g., they are entitled to environmental quality of Y). Prior to implementation of such policies, it is often unclear whether industry has the right to pollute, citizens have the right to a clean environment, or some combination of both. In these cases, Coasean bargaining may be suppressed because of ambiguities about baseline conditions that establish who is supposed to compensate whom. Indeed, Coase (1960) himself addressed this possibility in his famous confectioner and doctor example. There, it was not until the court decided in favor of the doctor that property rights were clearly delineated and private bargaining could commence. Thus, at least in some cases, implementing an environmental policy may help set the stage for subsequent Coasean bargaining.

Implementing a centralized environmental policy can also reduce transaction costs, providing a second reason for policy-induced, Coasean bargaining. With emissions trading programs, for example, citizens have the ability to purchase and retire pollution rights from a centralized platform rather than needing to engage in costly negotiations with individual firms to reduce pollution.[8] Similarly, with individual transferable quotas for natural resource extraction (e.g., fishing or water rights), citizens sometimes have the ability to participate in these markets and promote conservation. In tandem with policies themselves, changes in technology and information provision can also promote Coasean bargaining after policy implementation. Whether or not explicitly intended for compliance purposes, changes in technology and data availability are dramatically reducing the cost of monitoring and verifying the stocks and flows of many environmental goods. For example, recent advances in satellite and sensor technology means that forests, fishing activity, air pollutants, and water

[8] Banzhaf (2010) makes this point as an explicit argument in favor of cap-and-trade programs. He argues in favor of such policies not only because of the reduction in transaction costs, but also because cap-and-trade programs help solve the additionality problem that may arise with Coasean provision. A recently formed organization, named Carbon Vault, is intended to capitalize on precisely this idea. Institutions and individuals seeking to reduce or offset emissions of carbon dioxide can purchase and retire cap-and-trade emission allowances from several different markets (see https://carbonvault.org/).

are now monitored in real-time around the globe, enabling those seeking to privately provide greater environmental protection to do so in a more efficient and targeted manner.

We use the term "Coasean provision" to capture the idea of policy-induced, Coasean bargaining. Whereas Coasean bargaining is often discussed in contexts where negotiations can support first-best, efficient outcomes, our notion of Coasean provision captures what are more generally suboptimal outcomes consistent with private provision of public goods. From the citizens' perspective we are interested in the potentially market-revealed, marginal willingness to pay to avoid pollution. In the special case of a single citizen and no income effects, this is simply an alternative interpretation of the $MD(Q)$ function defined above.[9] More generally, because of free riding, the market demand for reducing pollution (i.e., abatement) will be based on the private marginal damages to individuals rather than the greater social marginal damages. We denote this function $PMD(Q)$, and it holds by definition that $PMD(Q) \leq MD(Q)$, with the difference including the free riding effect. We also assume that $PMD'(Q) > 0$.

As discussed in the introduction, many different factors can give rise to $PMD(Q)$, including preferences of wealthy individuals driven to environmental causes, corporate environmental management, or both. But to fix ideas it may be helpful to consider a fully micro-founded motivation. Assume citizens have quasilinear preferences of the form $U(x, Q) = x_j - f_i(Q_{max} - A)$, where x_j is private consumption, and A is the aggregate level of privately provided abatement among citizens. It follows that $A = \sum_{j=1}^{N} a_j$ and $Q = Q_{max} - A$. In this case, it is straightforward to verify that $MD(Q) = \sum_{j=1}^{N} f_j'(Q)$ and $PMD(Q) = \max\{f_1'(Q), ..., f_N'(Q)\}$. The latter equation represents the potentially market-revealed, marginal willingness to pay to avoid pollution on the part of the citizens. It is the upper envelope of individual marginal damages, and it is equivalent to the market demand function for private provision of abatement, which is by definition a public good. That is, for any price of abatement p, the aggregate quantity demanded will satisfy $p = PMD(Q_{max} - A)$. An important feature of our setup, however, is that the demand becomes operational after a centralized policy instrument is implemented.[10]

[9]It is necessary to assume no income effects in order for the marginal willingness to pay for abatement to be the same at all levels of Q regardless of whether it is exogenously given or privately provided. Once could, of course, also think in terms of the marginal willingness to accept for pollution.

[10]As mentioned earlier, this does not mean that Coasean bargaining is not possible prior to policy implementation. It means that any pre-policy bargaining has already taken place and is accounted for in the aggregate marginal benefit and damage functions. Our focus here is on the implications of policy-induced, Coasean provision.

3 Instruments

We now consider how the potential for policy-induced, Coasean provision affects equilibrium outcomes with the classical policy instruments of either the pollution tax or cap. Our analysis in this section is positive (i.e., descriptive) and applies to any level of exogenously given policy stringency. In subsequent sections, we turn to normative (i.e., efficiency) concerns related to specific and endogenously chosen levels of stringency.

3.1 A Tax

Consider an exogenously set tax of τ on each unit of pollution. The standard result, without Coasean provision, is that pollution will continue up to the point where industry's marginal net benefits are zero, so the resulting level of pollution will satisfy $MB(\bar{Q}) - \tau = 0$, as shown in Figure 1. This, however, is no longer an equilibrium with $PMD(Q)$ defining the scope for Coasean provision. The logic is standard Coasean bargaining, but based on $PMD(Q)$ rather than the full social marginal damages. Once industry has responded to τ in the usual manner, reducing pollution to $\bar{Q}$, citizens have a private willingness to pay up to $PMD(\bar{Q})$ for the next unit of abatement, even after taking account of free riding. Then, because the private willingness to pay for more abatement exceeds the industry's willingness to accept all the way down to $\hat{Q}$, this is the level of pollution that becomes the equilibrium with Coasean provision. That is, under a tax of τ, Coasean provision reduces pollution from $\bar{Q}$ all the way down to $\hat{Q}$.

Beyond gains from trade that underlie all Coasean solutions, what is the intuition for this somewhat surprising result that conflicts with standard approaches for teaching about pollution taxes? The answer is that imposing the tax provides an implicit subsidy by reducing industry's marginal benefit (i.e., demand) for pollution.[11] We illustrate this with the $MB(Q) - \tau$ curve in Figure 1. The effect is that any side payments that citizens are willing to pay to reduce pollution go even further because polluting firms can avoid paying the tax, in addition to collecting the side payments.

The special case of this setup with $N = 1$ is implicitly considered in Buchanan and Stubblebine (1962). While this case helps to illuminate the potential for Coasean bargaining to occur after implementing a tax, it overlooks a critical feature of the setup in a more general and realistic setting: abatement provides a public good, rather than reducing an

[11] Consider a simple example. Suppose the price of gasoline is \$3.20 per gallon and you drive exactly 227 miles in a week at that price. How much would you be willing to accept to drive one mile less? The answer is that you would accept any price above zero because it must be the case that the final mile earned you nothing in net benefit, else you would have driven more miles.

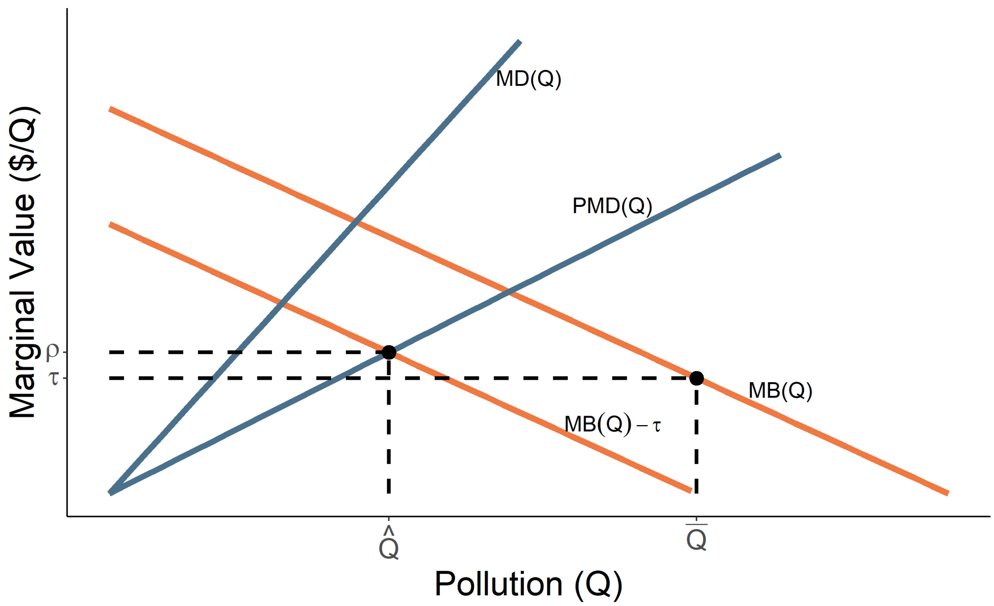

Figure 1: Marginal benefits and damages of pollution and equilibrium with Coasean provision and a tax of τ.

externality imposed on a single agent. This is important, for reasons we consider below, because with $N > 1$ the difference between private marginal willingness to pay and social marginal damages creates distinct welfare implications in the context of policy instrument choice.[12]

Finally, we turn to the side payments themselves. Does such private provision occur in practice in the context of a pollution tax? While our model is highly stylized, it is easy to find real-world examples capturing the same idea. Most municipalities have a landfill tax, yet major companies such as 3M, Coca-Cola, and Johnson & Johnson have recently announced commitments to a new recycling fund.[13] Moreover, as noted previously, the city of Copenhagen has made a public commitment to carbon neutrality despite Denmark having a national carbon tax, and similar arrangements are taking place in many other cities and countries.[14] In both cases, it is reasonable to assume that consumers (or citizens) have exerted market (in the first case) or political (in the second case) pressure to bring about these commitments. As with all Coasean bargaining, the magnitude of the side payments will depend on the outcome of a bargaining game.[15] While the outcome will have clear

[12]Moreover, even in the case of $N = 1$, the standard Coasean argument is support of efficiency will not hold if there are income effects. For example, if Q is a strictly normal good, then $PMD(Q) < MD(Q)$ even for a single citizen.

[13]See https://www.closedlooppartners.com.

[14]See https://international.kk.dk/artikel/carbon-neutral-capital.

[15]One possible case to consider is side payments based on a fixed price that clears all transactions. This would imply a price of $\rho = MB(\hat{Q}) - \tau = PMD(\hat{Q})$, and a total transfer from citizens to industry of

distributional consequences, it will not affect overall efficiency, which is the focus of analysis here.

3.2 A cap and trade

We now consider an exogenously given cap of Ω units of pollution that takes the form of tradeable allowances. In the standard setup, without Coasean provision, the market clearing auction price would be equal to $MB(\Omega)$. Whether any subsequent Coasean provision would occur in this case hinges on whether citizens are willing to pay more than polluters are willing to accept. At the cap of Ω, citizens' demand for abatement implies a marginal willingness to pay of $PMD(\Omega)$ for the first unit of additional abatement, and polluters are willing to accept $MB(\Omega)$. Thus, two cases emerge corresponding to whether the cap is sufficiently stringent or weak, where the threshold level of stringency that distinguishes the two cases satisfies $PMD(\tilde{Q}) = MB(\tilde{Q})$, as shown in Figure 2.

We begin with the case of a sufficiently stringent cap such that $\Omega^L \leq \tilde{Q}$. It is straightforward to see that Coasean provision will play no role in this case. The reason is that when polluters comply with the stringent cap, citizens are simply not willing to pay the permit price to achieve additional abatement. Their willingness to pay falls short by the red dashed line shown in Figure 2. It follows that the equilibrium level of pollution remains at Ω^L, and the market clearing permit price is equal to $MB(\Omega^L)$, both of which are consistent with the standard textbook analysis of a cap-and-trade program.

The more interesting case is one with a sufficiently weak cap such that $\Omega^H > \tilde{Q}$. Here, Coasean provision will play a role because, even after accounting for free riding, the citizens' demand for abatement indicates a marginal willingness to pay that exceeds $MB(\Omega^H)$. This is illustrated as the blue dashed line in Figure 2. In order to reveal this willingness to pay, we assume that citizens are able purchase permits, either directly from the initial allocation auction, or subsequently from firms themselves in a secondary market.[16] This implies a combined industry and citizen inverse demand function for government issued permits that can be written as

$$p(\Omega) = \begin{cases} MB(\Omega) & \text{if } \Omega \leq \tilde{Q} \\ MB(\tilde{Q}) & \text{if } \Omega > \tilde{Q} \end{cases}.$$

$\rho(\bar{Q} - \hat{Q})$. While Figure 1 illustrates a case where $\rho > \tau$, nothing rules the possibility for ρ to be less than or equal to τ.

[16] Other papers have examined various aspects of citizen participation in cap-and-trade markets for air pollution, including questions about what it implies about efficiency of the cap (Israel 2007), the interaction with incentives for lobbying (Malueg and Yates 2006), and the potential for compounding inefficiencies due to market power (Eshel and Sexton 2009)

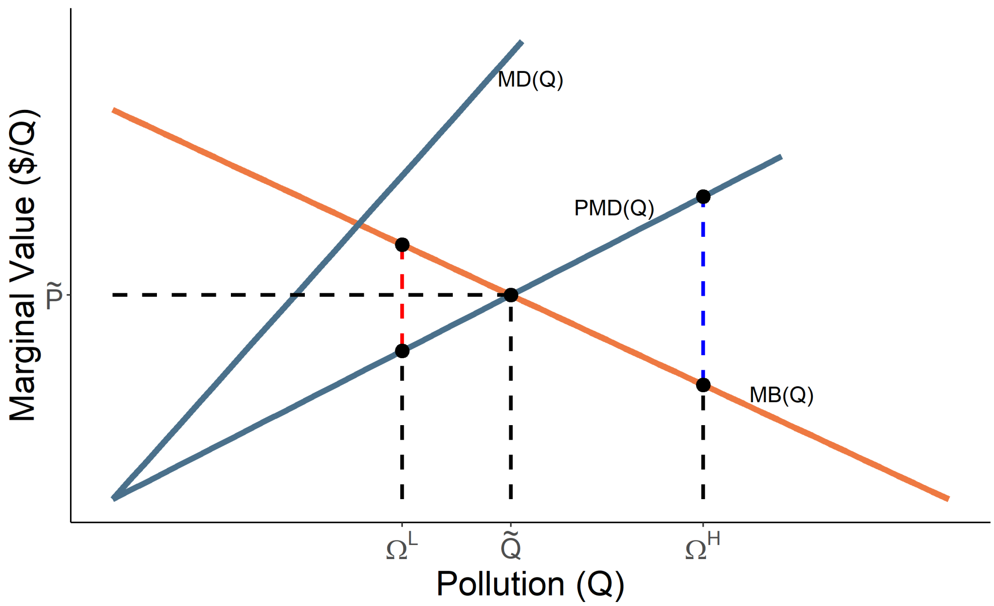

Figure 2: Marginal benefits and damages of pollution and equilibrium with Coasean provision and a cap and trade program defined by Ω.

With a sufficiently weak cap, it follows that the equilibrium permit price is $\tilde{P} = MB(\tilde{Q})$, and the equilibrium level of pollution is $\tilde{Q} < \Omega^H$, as depicted in Figure 2. An important observation for the moment is that equilibrium pollution is less than that targeted by the policy, as abatement of the amount $\Omega^H - \tilde{Q}$ occurs because of Coasean provision.[17]

4 Myopic Policies

We now turn to comparisons of instrument choice in the presence of potential Coasean provision. We begin with a comparison of tax and cap policies that satisfy $\tau^* = MB(\Omega^*) = MD(\Omega^*)$. We refer to these as myopically optimal (M-Optimal) policies because they represent the equivalent, first-best instruments that would be chosen if a regulator overlooked the possibility for Coasean provision; these are also the textbook levels of stringency for efficient environmental policy. We then generalize the comparative analysis to account for all myopically equivalent policies (i.e., not only the myopically first-best). In the next section, we define and evaluate conditionally optimal policies, where the regulator chooses the first-best tax or cap taking account of Coasean provision. In all cases, we consider overall efficiency based on standard welfare measures.

[17]Note that we have implicitly assumed that with either a tax or cap, Coasean provision is a function of $PMD(Q)$, which itself does not depend on the instrument choice. This implies, for example, that any transaction costs associated with Coasean provision are the same with either the tax or cap. While this is the starting point assumption we make throughout, it is one that we discuss again later in the paper.

It is straightforward to see that the welfare maximizing level of pollution will satisfy $MB(Q) = MD(Q)$ regardless of whether or not there is Coasean provision. Drawing on the results above, we know that in the presence of Coasean provision, the M-Optimal cap implements precisely this level of pollution, but the M-Optimal tax does not. This establishes the first result, which begins to show how the standard equivalence between price and quantity instruments breaks down in the presence of Coasean provision.

Proposition 1. *When comparing M-Optimal policies, the cap Ω^* implements the first-best level of social welfare, but the tax τ^* does not.*

The basic intuition for this result is that with M-Optimal levels of policy stringency, Coasean provision does not occur with the cap, but it does with the tax because of the implicit subsidy it confers to bargaining. That is, Coasean provision under the M-Optimal tax leads to inefficiently low levels of pollution.

We now generalize our analysis to consider *any* level of policy stringency, where the comparison of instruments is based on taxes and caps that we refer to as myopically equivalent (M-Equivalent). That is, M-Equivalent policies must satisfy $\tau = MB(\Omega)$, where M-Optimal policies are a special case that accounts for marginal damages as well (i.e., the condition is also equal to $MD(\Omega)$). We find that the results differ in interesting and important ways at different levels of stringency.

But first, because we have already established that at least some M-Equivalent policies do not implement the same level of equilibrium pollution, we need a definition of policy stringency to compare the tax and cap without relying on the same quantities of pollution. We have chosen to normalize stringency based on the level of the tax, such that stringency is defined as $S = \tau$, which implies that the correspondingly stringent M-Equivalent cap must satisfy $S = MB(\Omega)$.[18] This implies that S denotes the level of the tax and the permit price that is consistent with an M-equivalent cap.

Proposition 2. *When comparing M-Equivalent policies, there exists a particular level of stringency $\mathring{S}$ such that the two instruments produce the same level of welfare, which is less than efficient. Moreover, welfare with a tax is greater for all $S < \mathring{S}$, whereas welfare with a cap is greater for all $S > \mathring{S}$.*

Proof. The equilibrium condition for the tax defines an implict function $\hat{Q}(S)$ that is strictly decreasing. It follows that welfare with the tax, $W_\tau(S) = B(\hat{Q}(S)) - D(\hat{Q}(S))$, is a strictly

[18]This choice of stringency measure is without loss of generality. One could alternatively define stringency based on the cap Ω, derive the M-Equivalent tax, and prove all of the same results. One advantage of normalizing based on the tax is that a higher level of S corresponds to greater stringency (i.e., less pollution).

concave function. To characterize the cap equilibrium, define a particular level of stringency $\tilde{S} = MB(\tilde{Q}) = PMD(\tilde{Q})$ so that we have the function

$$Q(S) = \begin{cases} \tilde{Q} & \text{if } S \leq \tilde{S} \\ Q : MB(Q) = S & \text{if } S > \tilde{S} \end{cases}.$$

This function, which defines the equilibrium level of pollution given a level of cap stringency, is weakly decreasing in S. Welfare with the cap, $W_c(S) = B(Q(S)) - D(Q(S))$, is therefore constant for $S \leq \tilde{S}$ and strictly concave for $S > \tilde{S}$. Welfare for the cap is maximized at $S^* = MB(Q) = MD(Q)$, which is first-best, and it is straightforward to verify that $W_c(S) > W_\tau(S)$ for all $S > S^*$, because the tax implements an even lower level of pollution that compounds what is already a (weakly) overly stringent policy. We also know that for all $S < S^*$, $W_\tau(S)$ is strictly concave and $W_c(S)$ is non-decreasing. Moreover, $W_\tau(0) = W_c(0)$ and $S = MB(Q_{max}) > S^*$. Together, these conditions imply that $W_\tau(S) = W_c(S)$ at one and only one level of stringency for $0 < S < Q_{max}$. This defines $\mathring{S}$, and it follows that $W_\tau(S) > (<) W_c(S)$ for all $S < (>) \mathring{S}$.[19] $\square$

Proposition 2 shows that the results of Proposition 1 do not hold for all M-Equivalent policies. While the cap is always more efficient (and first best) when comparing M-Optimal instruments, we find that the more efficient M-Equivalent instrument depends on the level of stringency. Underlying the result is the observation that a tax always induces a lower level of pollution than a M-Equivalent cap, and this explains why the tax is more efficient when the policy is sufficiently weak, whereas the cap is more efficient when the policy is sufficiently stringent.

5 Conditionally Optimal Policies

The previous section considered exogenously set policies that aim to achieve the same level of pollution without recognizing the potential for Coasean provision. We referred to these as myopic polices. What happens if, instead, the regulator explicitly accounts for Coasean provision and chooses the policy stringency to maximize welfare? In this section we derive and compare conditionally optimal (C-Optimal) policies, where the level of stringency is chosen optimally for either the tax or cap. Additionally, we compare the efficiency results to the case where Coasean provision is not permitted to take place. This provides a useful bench-

[19] Although discussed in the next section, a graphical illustration of this proof is shown in Figure 3 for the special case of linear marginal benefits and damages of pollution.

mark for comparison and illuminates the trade-offs that regulators face over questions of whether to allow citizen participation in cap-and-trade programs, side payments to industry to reduce pollution with a tax, or both.

We know that regardless of instrument choice, the efficient level of pollution will satisfy $MB(Q^*) = MD(Q^*)$. We also established previously that the M-Optimal cap implements precisely this level of pollution as the equilibrium with $\Omega^* = Q^* = \Omega^+$, where we use the plus to denote C-Optimal policies. The M-Optimal tax at $\tau^* = MB(Q^*)$ is not first best, however, because the equilibrium condition is $MB(Q) - \tau = PMD(Q)$. This means that τ^* implements an inefficiently low level of pollution, that is, the tax is effectively too stringent. Nevertheless, lowering the tax to $\tau^+ = MD(Q^*) - PMD(Q^*)$ implements Q^* as the equilibrium level of pollution and is therefore C-Optimal. Together, these results prove the following proposition.

Proposition 3. *When comparing C-Optimal policies, both τ^+ and Ω^+ can implement the first-best level of social welfare, and it is the same level that arises through welfare maximization with a prohibition on Coasean provision.*

The intuition underlying Proposition 3 hinges on appropriately calibrating the C-Optimal tax. Rather than reflecting marginal damages at the optimal level of pollution, τ^+ reflects marginal damages net of the citizens' private marginal willingness to pay for abatement. Anticipating the extent of Coasean provision, the regulator lowers the tax and lets citizens contribute to lowering pollution down to the optimal level. A further insight of Proposition 3 is that the level of maximized social welfare is not only invariant to the policy instrument, but is the same as what could be achieved with a ban on Coasean provision (and concomitant re-adjustment of policy stringency).

Figure 3 illustrates these results graphically, along with a summary of some previous results. Assuming linear marginal benefit and damage functions, the figure plots welfare against policy stringency under three different regulatory scenarios: a tax or cap with Coasean provision, and no Coasean provision, in which case the tax and cap are equivalent. As a starting point, note that without Coasean provision, S^* denotes the stringency of the equivalent tax and cap that maximizes welfare. This level of stringency is equal to the M-Optimal and C-Optimal cap with Coasean provision. A lower stringency of τ^+ is required for the C-Optimal tax, which also implements the same level of maximized welfare.

Additional results about M-Equivalent policies that were proved earlier can be seen in the figure by recognizing that each point on the horizontal axis corresponds to a particular pair of M-Equivalent instruments. $\mathring{S}$ denotes the level of stringency identified in Proposition 2 where the M-Equivalent tax and cap implement the same level of welfare, and the tax is

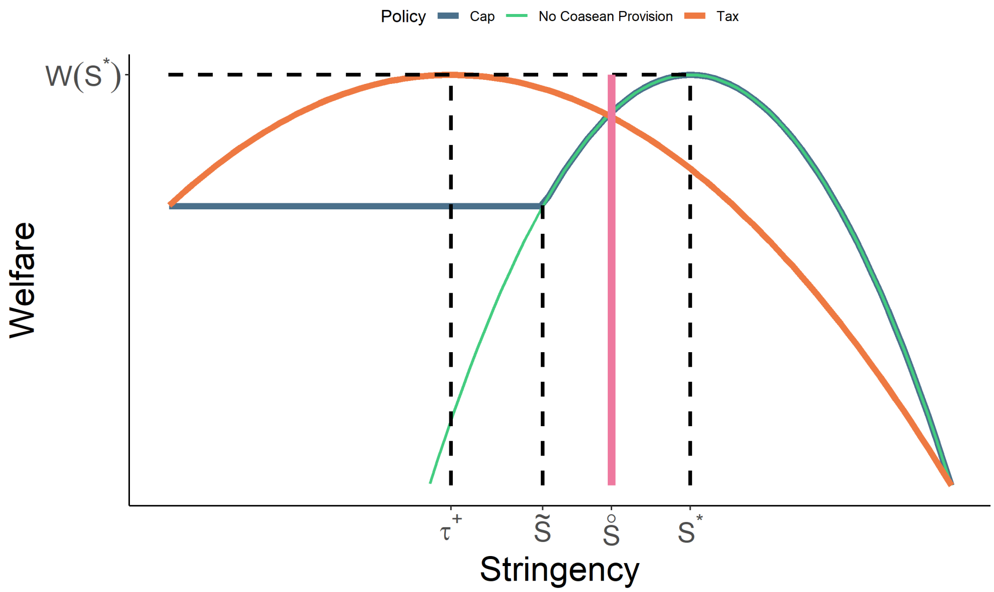

Figure 3: Welfare as a function of policy stringency for a tax and cap, and a policy scenario with no Coasean provision, in which case the tax and cap are equivalent. The figure is based on the simplifying assumption of linear benefit and damage functions, and the case where $PMD(Q) \leq MD(Q)$ holds strictly.

preferred on an efficiency basis at all lower levels of stringency, whereas the cap is preferred at all higher levels, including the M-Optimal levels at S^*. Moreover, note that welfare remains constant with the cap at all levels of stringency less than $\tilde{S}$, because with sufficiently low levels of stringency, citizens always engage in enough Coasean provision to bring pollution back to the point that satisfies $\tilde{S} = PMD(\tilde{Q}) = MB(\tilde{Q})$.

6 Introducing Uncertainty

Our analysis of C-Optimal policies has thus far assumed the regulator has perfect knowledge about the benefits of pollution, the damages of pollution, and the scope for Coasean provision. In this section, we show how incorporating uncertainty about the marginal benefits or the marginal damages affects our conclusions about C-Optimal policies and the choice between them. Our approach adheres closely to the Weitzman (1974) setup, thereby establishing new results as a generalization of those already familiar in the literature. We therefore begin by focusing on uncertainty over the marginal benefits of pollution before turning to uncertainty about the marginal damages.

The policymaker seeks to maximize overall welfare and does so by choosing the optimal stringency of the C-Optimal policies and then choosing the more efficient of the two instruments. To keep things tractable, we adopt linear functional forms, where all parameters are positive. Expected marginal benefits of pollution are given by $MB(Q) = \alpha - \kappa Q$, and realized marginal benefits are $MB(Q) \pm \delta$, where δ captures the uncertainty. In the high state of the world, the marginal benefit is shifted up by δ, which occurs with probability .5, and in the low state of the world, it is shifted down by δ with probability .5. The marginal damages of pollution are given by $MD(Q) = \gamma Q$, which we assume are known with certainty until Section 9. Allowing for Coasean provision, the demand for abatement is given by $PMD(Q) = \beta MD(Q)$, where $0 \leq \beta \leq 1$. The parameter β therefore governs the scope for Coasean provision: $\beta = 0$ implies no scope, and $\beta = 1$ is consistent with one citizen and no income effects, which implicitly matches the standard Coasean assumption.

The central result of Weitzman (1974), using our notation, is that the welfare advantage of a tax compared to a cap with uncertainty over the marginal benefits of pollution is

$$\Delta^W = \delta^2 \left(\frac{\kappa - \gamma}{2\kappa^2} \right), \tag{1}$$

where the superscript W stands for "Weitzman." The equation makes clear that taxes and caps deliver equivalent welfare in the absence of uncertainty (i.e., $\delta = 0$) or if the slopes of the marginal benefit and damage functions are the same (i.e., $\gamma = \kappa$). More generally, taxes (or caps) are preferred if the marginal damage (benefit) function is flatter, that is, if $\gamma < (>) \kappa$.[20]

Our aim in the next three sections is to consider the ways in which the standard results, including equation (1), change in the presence of Coasean provision. We also show in Section 9 that the Weitzman (1974) result about welfare invariance between policy instruments under uncertainty about the marginal damages of pollution no longer holds, and we find a clear and general result of always favoring taxes over caps in the presence of Coasean provision.

[20] An implicit assumption of Weitzman (1974) is that the level of uncertainty is sufficiently small to ensure that his welfare measures underlying equation (1) do not hit corner solutions. In particular, when solving for candidate, ex ante optimal policies, he implicitly assumes the deadweight loss triangles do not run into the vertical or horizontal axes. The condition can be written as

$$\delta \leq \alpha \min \left(\frac{\kappa}{2\gamma + \kappa}, \frac{\gamma}{2\kappa + \gamma} \right)$$

If $\kappa < (>) \gamma$, the condition implies that a tax (cap) optimized to the high state of the world weakly binds in the low state of the world. We will use this condition later in the paper as part of the proof to Proposition 5.

7 C-Optimal Policies with $MB(Q)$ Uncertainty

Before comparing the instruments, we must first consider how introducing uncertainty affects the stringency of the C-Optimal tax and cap. In the standard Weitzman (1974) setup, with no Coasean provision, the levels of stringency for both the tax and cap that maximize expected welfare are invariant to the introduction and level of uncertainty.[21] In our setup, with Coasean provision, we show that this result continues to hold for the tax but not for the cap.

Two observations help to motivate our formal results. First, regarding taxes, we showed in the previous section that greater scope for Coasean provision results in a lowering of the the C-Optimal tax, because the planner anticipates Coasean provision and calibrates the tax to maintain the first-best level of pollution. Recall that the equilibrium condition is $\tau^{+} = MD(Q^{*}) - PMD(Q^{*})$. The same logic is preserved with uncertainty, and as we prove below, there is no effect of uncertainty on the C-Optimal tax. Second, we showed previously that the deterministic C-Optimal cap is unaffected by Coasean provision. We show below that this result continues to hold with uncertainty, provided that the scope for Coasean provision is modest (i.e., β is sufficiently small). However, if β is large enough, we find that Coasean provision will occur in the low-MB state of the world but not the high, and this implies that the C-Optimal cap must be adjusted to account for the backstop that Coasean provision offers under uncertainty.

We begin by establishing the expected deadweight loss of any, arbitrary policy in the presence of Coasean provision.[22] Let Q_i^{*} denote the welfare-maximizing level of pollution in state of the world $i \in \{L, H\}$, which is invariant to the policy instrument choice and whether or not Coasean provision takes place. These solutions are shown in the first row of Table 1. Now let $Q_{i\mathcal{P}}$ denote the equilibrium level of pollution in state of the world i given the use of any, arbitrary tax or cap policy $\mathcal{P} \in \{\tau, \Omega\}$. These quantities and the equilibrium conditions that give rise to them are summarized in the other rows of Table 1. It follows that the difference between the first-best and equilibrium levels of pollution for either policy and state of the world can be written as $D_{i\mathcal{P}} \equiv |Q_i^{*} - Q_{i\mathcal{P}}|$. Then, conditional on policy $\mathcal{P}$, the deadweight loss in state i is given by integrating between the marginal benefit and marginal damage curves, which is an area equal to $\frac{1}{2}D_{i\mathcal{P}}^{2}(\gamma + \kappa)$. Finally, recognizing that

[21] Without Coasean provision, the policies that maximize expected welfare are a tax of $\tau^{W} = \frac{\alpha\gamma}{\gamma+\kappa}$ and a cap of $\Omega^{W} = \frac{\alpha}{\gamma+\kappa}$, and both implement the same level of pollution without uncertainty. These results implicitly rely on the assumption of no corner solutions as described in footnote 20.

[22] As will become clear, it is convenient to establish results based on minimizing deadweight loss rather than maximizing welfare. This is innocuous because deadweight loss of any policy in any state of the world (high or low) is just the loss in welfare under that policy relative to the first-best policy in that state of the world.

Table 1: Pollution levels under different policies in each state of the world.

Variable	Low MB ($i = L$)	High MB ($i = H$)	Condition
Q_i^*	$\frac{\alpha-\delta}{\gamma+\kappa}$	$\frac{\alpha+\delta}{\gamma+\kappa}$	$MB(Q) \pm \delta = MD(Q)$
$Q_{i\tau}$	$\frac{\alpha-\tau-\delta}{\beta\gamma+\kappa}$	$\frac{\alpha-\tau+\delta}{\beta\gamma+\kappa}$	$MB(Q) - \tau \pm \delta = PMD(Q)$
$Q_{i\Omega}$	$\min\left(\Omega, \frac{\alpha-\delta}{\beta\gamma+\kappa}\right)$	$\min\left(\Omega, \frac{\alpha+\delta}{\beta\gamma+\kappa}\right)$	$Q = \Omega$ or $MB(Q) \pm \delta = PMD(Q)$

state i occurs with probability 0.5, the expected deadweight loss under policy $\mathcal{P}$ is

$$E[DWL_{\mathcal{P}}] = \frac{\gamma+\kappa}{4}\left(D_{L\mathcal{P}}^2 + D_{H\mathcal{P}}^2\right), \tag{2}$$

which is a helpful expression for proving several of the subsequent results.

We require one more intermediate step. Lemma 1 below shows that the C-Optimal cap with uncertainty is always one of two possible solutions, depending on whether uncertainty is sufficiently large to trigger Coasean provision.

Lemma 1. *In the presence of uncertainty over the marginal benefits of pollution, the C-Optimal cap is either Ω^+ or $\Omega^{++} = \Omega^+ + \frac{\delta}{\gamma+\kappa}$, where the latter is the efficient quantity of pollution conditional on the high-MB state (i.e., Q_H^*). If it is Ω^+, there is no Coasean provision. If it is Ω^{++}, there is Coasean provision in the low-MB state only.*

Proof. We define a threshold cap set at $\Omega = \tilde{Q}_L \equiv \frac{\alpha-\delta}{\beta\gamma+\kappa}$, which solves $MB(\tilde{Q}_L) - \delta = \beta MD(\tilde{Q}_L)$ and is decreasing in the level of uncertainty. Any given pollution cap is either weakly smaller or larger than $\tilde{Q}_L$. First consider caps that are smaller, so $\Omega \leq \tilde{Q}_L$. Such caps would induce no Coasean provision in either the low or high states, so $Q_{H\Omega} = Q_{L\Omega} = \Omega$ and therefore $E[DWL_\Omega] = \frac{\gamma+\kappa}{4}\left(\left(\frac{\alpha-\delta}{\gamma+\kappa} - \Omega\right)^2 + \left(\frac{\alpha+\delta}{\gamma+\kappa} - \Omega\right)^2\right)$. Minimizing the expected deadweight loss thus requires setting the cap at $\Omega^+ = \frac{\alpha}{\gamma+\kappa}$, so any cap $\Omega \leq \tilde{Q}_L$ is welfare-dominated by Ω^+. Now consider caps that are larger than the threshold, so $\Omega > \tilde{Q}_L$. Such caps would induce Coasean provision in the low state and therefore $Q_{L\Omega}$ is independent of Ω. This implies further that $D_{L\Omega}$ is independent of Ω and minimizing Equation (2) is equivalent to minimizing $D_{H\Omega}^2$, which is solved by the cap $\Omega^{++} = \frac{\alpha+\delta}{\gamma+\kappa}$, so any cap $\Omega > \tilde{Q}_L$ is welfare-dominated by Ω^{++}. Finally, at this solution, there is no Coasean provision in the high state because $Q_{H\Omega} = \Omega^{++}$. □

Invoking Lemma 1, our findings regarding the C-Optimal policies under uncertainty are summarized as follows:

Proposition 4. *In the presence of uncertainty over the marginal benefits of pollution, the optimal tax is equal to the C-Optimal tax without uncertainty, τ^+. The optimal cap is equal* *to the C-Optimal cap without uncertainty, Ω^+, if $\beta \leq \beta_c(\delta)$, where $\beta_c(\delta)$ is a unique critical threshold that is decreasing in δ. Otherwise, the optimal cap rises to Ω^{++}.*

Proof. We begin with the tax. Using the definitions in Table 1, we can solve for $D_{L\tau} = A(\tau) - \delta B$ and $D_{H\tau} = A(\tau) + \delta B$, where $A(\tau) \equiv \frac{\alpha-\tau}{\beta\gamma+\kappa} - \frac{\alpha}{\gamma+\kappa}$ and $B \equiv \frac{1}{\beta\gamma+\kappa} - \frac{1}{\gamma+\kappa}$. Substituting these expressions into Equation (2) and rearranging yields $E[DWL_\tau] = \frac{\gamma+\kappa}{2}(A(\tau)^2 + (\delta B)^2)$. Because B is independent of τ, minimizing the expected deadweight loss with respect to the tax is equivalent to minimizing $A(\tau)^2$, which yields $\tau^+ = \frac{\alpha\gamma(1-\beta)}{\gamma+\kappa} = MD(Q^*) - \beta MD(Q^*)$.

Turning to the cap, Lemma 1 establishes that the only two candidate solutions are Ω^+ and Ω^{++}, and it is sufficient for us the determine which has the lower deadweight loss. Substituting the candidate policies into Equation (2) yields

$$E[DWL_{\Omega^+}] = \frac{\delta^2}{2(\gamma+\kappa)} \tag{3}$$

$$E[DWL_{\Omega^{++}}] = \frac{(1-\beta)^2\gamma^2(\alpha-\delta)^2}{4(\beta\gamma+\kappa)^2(\gamma+\kappa)} \tag{4}$$

Setting these equations equal to each other and solving for β yields a unique critical threshold:

$$\beta_c(\delta) = \frac{\alpha\gamma - \delta(\gamma + \kappa\sqrt{2})}{\gamma(\alpha + \delta(\sqrt{2}-1))}, \tag{5}$$

where we have made explicit the dependence of β_c on uncertainty, δ. We know the threshold is unique because the ratio $\frac{DWL_{\Omega^+}}{DWL_{\Omega^{++}}}$ is monotonically increasing in β and thus crosses 1 only once. That $\beta_c(\delta)$ is decreasing in δ follows immediately from Equation (5). Because the ratio is less than (equal to, greater than) 1 for all $\beta < (=, >)\beta_c(\delta)$, it follows that when $\beta \leq \beta_c$ the C-Optimal cap is Ω^+, and when $\beta \geq \beta_c$ the C-Optimal cap is Ω^{++}. $\square$

Figure 4 illustrates different possibilities for the C-Optimal cap. Ω^+ is the efficient level of pollution without uncertainty. Ω^{++} is the efficient level of pollution conditional on the high-MB state of the world. The figure depicts values of β and δ such that Coasean provision establishes a lower bound on pollution $\tilde{Q}_L > \Omega^+$ when the cap is set at Ω^{++}.[23] The question, then, is: Which cap is preferred? The expected deadweight loss of choosing Ω^+ is the standard Weitzman (1974) result and equal to area $(a+b)/2 = a$; shown in Figure 4 as the lower orange triangle. In contrast, the deadweight loss of choosing Ω^{++} is area $(a+c)/2$, because there is no deadweight loss in the high state. Hence the optimal cap is Ω^{++} if and only if area a is greater than area c (which is the case under the parameters shown in Figure

[23] This is a necessary but not sufficient condition for Ω^{++} to be the C-Optimal cap.

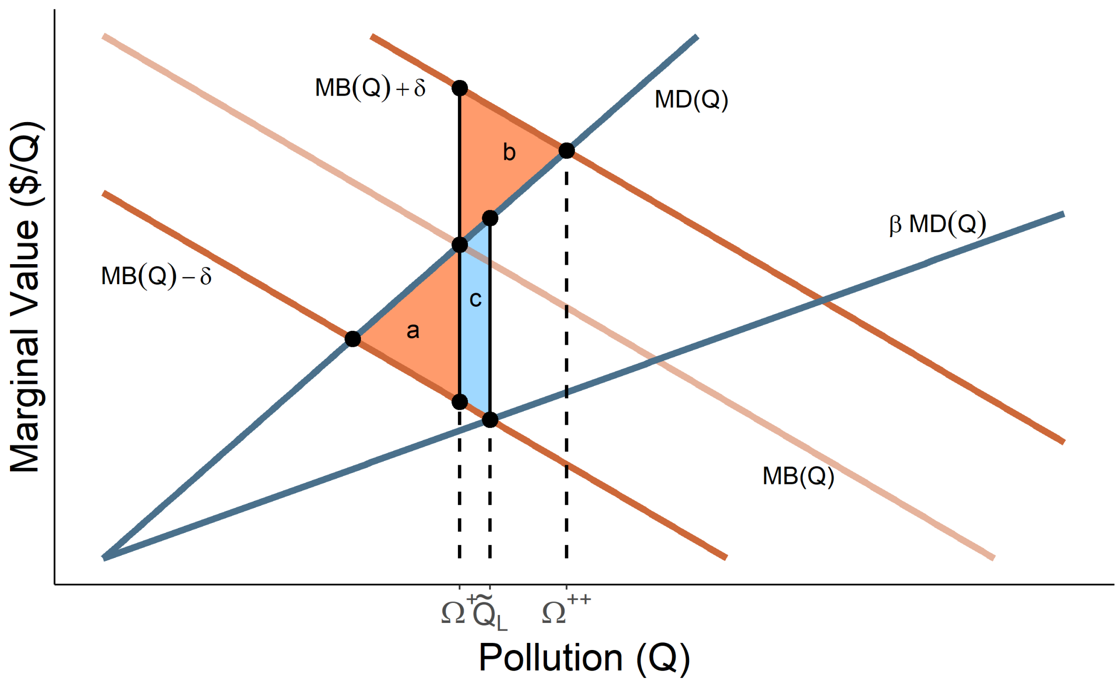

Figure 4: Graphical representation of caps Ω^+ and Ω^{++} and their deadweight losses for an arbitrary value of β.

4), and this is more likely to occur with greater uncertainty (δ) and greater scope for Coasean provision (β).

In sum, Proposition 4 reveals the effect of uncertainty on policy stringency. It has no effect on the stringency of the tax, though the same is true for the cap only if β is sufficiently small. However, if the scope for Coasean provision is sufficiently large, then the optimal cap is slackened, knowing that Coasean provision will serve as a lower bound on the welfare loss in the low-MB state of the world.[24]

8 Instrument Choice with $MB(Q)$ Uncertainty

Having established the C-Optimal policies under uncertainty about the marginal benefits of pollution in Proposition 4, we now consider the question of policy instrument choice: prices *vs.* quantities? Our approach continues to rely on a comparison of expected deadweight losses, where the preferred instrument is the one with a lower expected loss.

Substituting the C-Optimal tax policy into equation (2) yields the expected deadweight

[24] This result is related to the optimal cap set by a regulator seeking to learn about $MB(Q)$ over time by overtly setting a slack cap and observing the resulting pollution level (Costello and Karp 2004).

loss under the C-Optimal tax:

$$E[DWL_{\tau^+}] = \frac{(1-\beta)^2\gamma^2\delta^2}{2(\beta\gamma+\kappa)^2(\gamma+\kappa)}. \tag{6}$$

With respect to the C-Optimal cap, we have already derived the expected deadweight losses for the two possible cases of Ω^+ and Ω^{++} in equations (3) and (4), respectively. As shown in Proposition 4, these two cases also correspond with whether β is less than or greater than $\beta_c(\delta)$. Subtracting equation (6) from equation (3) yields the welfare advantage of the tax over the cap when $\beta \leq \beta_c(\delta)$:

$$\Delta|_{\beta\leq\beta_c(\delta)} = \delta^2\left(\frac{2\beta\gamma+\kappa-\gamma}{2(\beta\gamma+\kappa)^2}\right) \tag{7}$$

Note that Weitzman's result in equation (1) is a special case of equation (7) when $\beta = 0$. Now, subtracting equation (6) from equation (4) yields the welfare advantage of the tax over the cap when $\beta \geq \beta_c(\delta)$:

$$\Delta|_{\beta\geq\beta_c(\delta)} = \left(\frac{(1-\beta)^2\gamma^2}{4(\beta\gamma+\kappa)^2(\gamma+\kappa)}\right)\left((\alpha-\delta)^2-2\delta^2\right) \tag{8}$$

Using these deadweight loss expressions, our next proposition focuses on the question of instrument choice, given different levels of the scope for Coasean provision.

Proposition 5. *In the presence of uncertainty about the marginal benefits of pollution, expected welfare with the tax is greater than that for the the cap if and only if* $\beta > \beta^* \equiv \frac{\gamma-\kappa}{2\gamma}$.

Proof. It is sufficient to prove that for a given β, equations (7) and (8) are greater than zero if and only if $\beta > \beta^*$. Equation (7) evaluated at β^* is equal to zero, and the expression is clearly positive or negative for all values of β that are smaller or bigger, respectively. Turning to equation (8), note that the sign is the same as that of the second term in parentheses. It follows that equation (8) is positive if and only if $\delta < \frac{\alpha}{1+\sqrt{2}}$, which holds by the implicit assumption in Weitzman (1974) that we made explicit in footnote 20. In particular, it is straightforward to verify that $\alpha \min\left(\frac{\kappa}{2\gamma+\kappa}, \frac{\gamma}{2\kappa+\gamma}\right) < \frac{\alpha}{1+\sqrt{2}}$, and because the left-hand side is weakly greater than δ, this completes the proof. $\square$

The fundamental insight of Proposition 5 is that a greater β—i.e., scope for Coasean provision—tends to imply an advantage to taxes over caps. The reason is that greater β lowers the tax-induced spread between equilibrium pollution levels in the low- and high-MB states of the world. Then, because these pollution levels are both closer to those that are

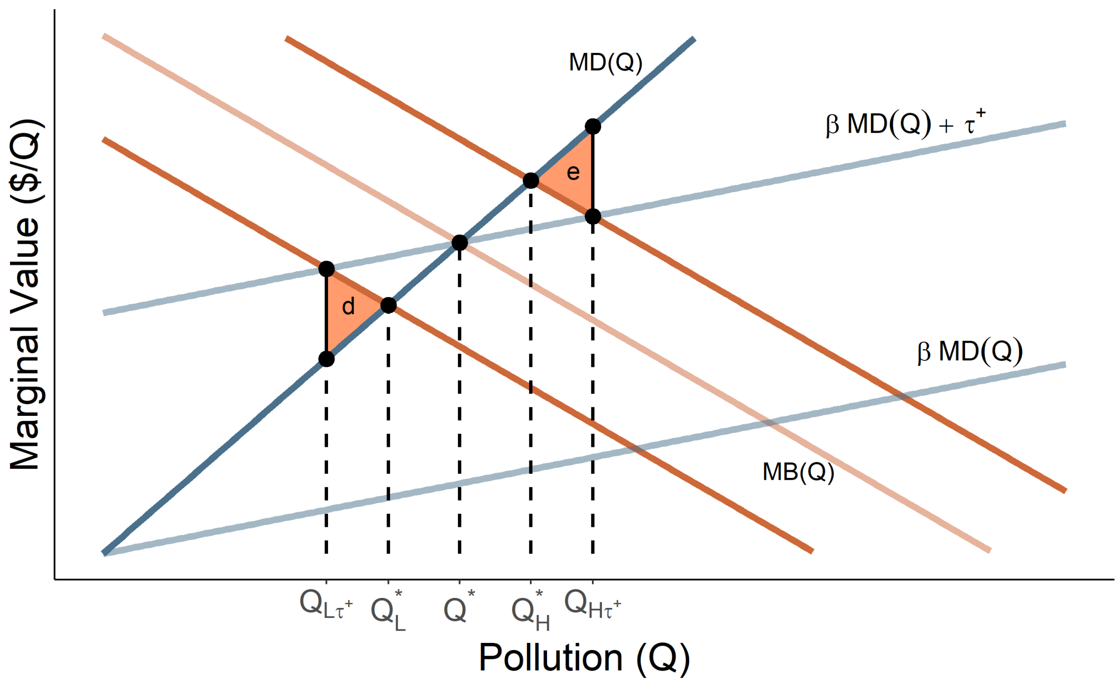

Figure 5: Expected deadweight loss with a tax, uncertainty, and Coasean provision is equal to $(d+e)/2$, which is decreasing in β.

ex-post optimal, expected welfare with the tax is greater for reasons that do not similarly affect the cap.

Figure 5 illustrates the mechanism at work. While the efficient quantities of pollution in the low and high states (Q_L^* and Q_H^*) are determined by $MB(Q) \pm \delta = MD(Q)$, the equilibrium quantities ($Q_{L\tau^+}$ and $Q_{H\tau^+}$) are determined by $MB(Q) \pm \delta = \beta MD(Q) + \tau^+$ (see Table 1). The deadweight loss in the low and high states are thus areas d and e, respectively, with the expected deadweight loss equal to $(d+e)/2$. The figure makes clear how the expected deadweight loss with the tax is decreasing in β; for an increase in β makes $\beta MD(Q) + \tau^+$ steeper while maintaining the same intersection with $MB(Q)$ at Q^*.[25] In the extreme case of $\beta = 1$, the deadweight loss is zero, and in the case of $\beta = 0$, we have the case considered in Weitzman (1974).

Finally, it is useful to compare our results on instrument choice to the familiar baseline of Weitzman (1974). We can show that the presence of Coasean provision expands the parameter space over which taxes dominate caps. To see this, set equation (7) equal to zero and solve for the condition where taxes are strictly preferred to caps: $\frac{\kappa}{\gamma} > 1 - 2\beta$. Without Coasean provission (i.e., $\beta = 0$), we recover precisely Weitzman's result in equation (1).

[25]To see why the intersection is the same, recall that the tax is set such that $\tau^+ + \beta MD(Q^*) = MB(Q^*)$, so $d\tau^+/d\beta = -MD(Q^*)$ to maintain the result.

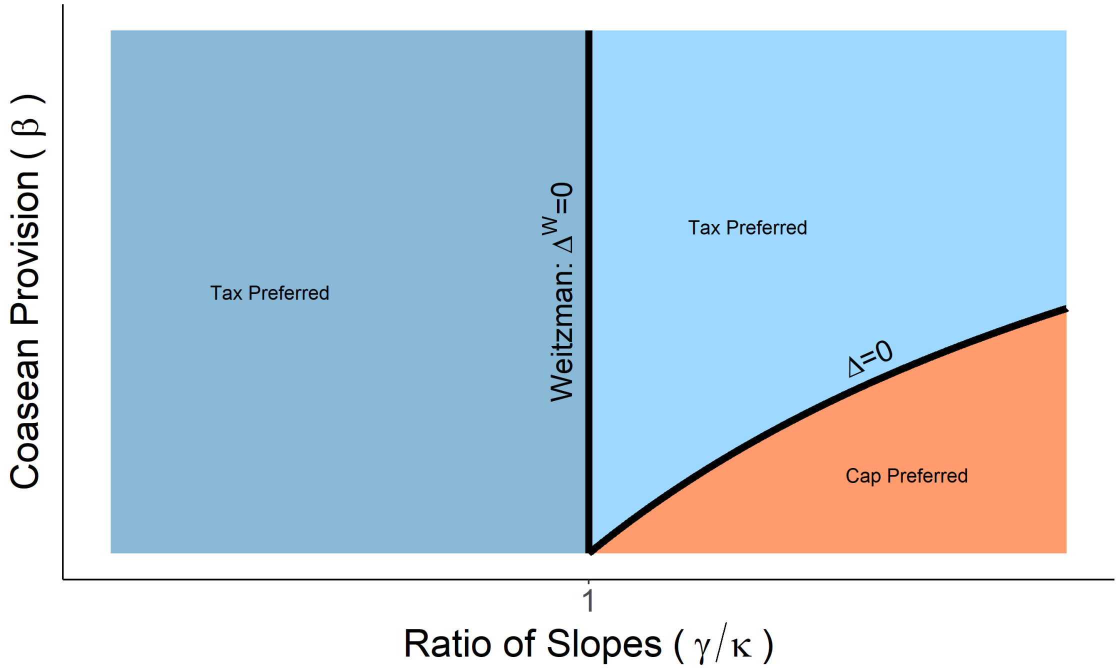

Figure 6: Parameter space over which the tax or cap delivers higher welfare. Blue areas indicate a preference for the tax; red area indicates a preference for the cap. Without Coasean provision the cap equals the tax when $\Delta^W = 0$.

Moreover generally, if taxes are preferred in Weitzman's setup, they are always preferred with Coasean provision; however, certain caps that are preferred with Weitzman's setup are in fact dominated by taxes in the presence of Coasean provision.

Figure 6 illustrates these results. The horizontal axis is the ratio $\frac{\gamma}{\kappa}$, and the vertical axis is β. Weitzman's result, applicable at $\beta = 0$, is that the taxes or caps are always preferred to the left or right of a ratio equal to one, respectively. With Coasean provision (i.e., $\beta > 0$), however, the dividing threshold is represented by the $\Delta = 0$ curve (satisfying $\frac{\kappa}{\gamma} = 1 - 2\beta$), which flips the region above from preferring caps to preferring taxes. Thus, Coasean provision expands the κ and γ parameter space over which taxes are preferred to caps, as indicated by the light blue shaded area in Figure 6.

9 Instrument Choice with $MD(Q)$ Uncertainty

The classic Weitzman (1974) analysis pays little attention to uncertainty about the marginal damages of pollution because he shows that it "affects price and quantity modes *equally* adversely" (p. 485). That is, he finds a welfare invariance between policy instruments with uncertainty about marginal damages, which does not (on its own) affect the level of pollution with either the tax or cap. In what follows, we show that this result no longer holds in the presence of Coasean provision, and moreover, a clear welfare preference emerges for taxes over caps.[26]

We continue to define the functions with certainty as $MB(Q) = \alpha - \kappa Q$ and $MD(Q) = \gamma Q$, along with $PMD(Q) = \beta MD(Q)$. In parallel fashion, we capture uncertainty in the marginal damages as $MD(Q) = \gamma Q \pm \delta$, while retaining the assumption that private demand for polution reduction is a fixed proportion β of marginal damages. As before, Coasean Provision can be turned off by setting $\beta = 0$, which conforms to the standard Weitzman (1974) setup.

Our first set of results is that the C-Optimal policies are nearly identical to those derived earlier in the case of uncertainty about the marginal benefits of pollution:

Proposition 6. *In the presence of uncertainty about the marginal damages of pollution, the optimal tax is equal to the C-Optimal tax without uncertainty, τ^+. The optimal cap is equal to the C-Optimal cap without uncertainty, Ω^+, if $\beta \leq \beta_d(\delta)$, where $\beta_d(\delta)$ is a unique critical threshold that is decreasing in δ. Otherwise, the optimal cap rises to Ω^{++}.*

Proof. The ex-post welfare maximizing policy in the high- and low-MD states are $Q_H^* = \frac{\alpha-\delta}{\gamma+\kappa}$ and $Q_L^* = \frac{\alpha+\delta}{\gamma+\kappa}$. For any tax τ, the pollution levels that result are $Q_\tau = \frac{\alpha-\tau\pm\beta\delta}{\beta\gamma+\kappa}$. Thus, the difference in pollution levels with any tax can be written as $D_{L\tau} = A(\tau) + \delta B$ and $D_{H\tau} = -A(\tau) + \delta B$, where $A(\tau) \equiv \frac{\alpha-\tau}{\beta\gamma+\kappa} - \frac{\alpha}{\gamma+\kappa}$ and $B \equiv \frac{1}{\gamma+\kappa} - \frac{\beta}{\beta\gamma+\kappa}$. Substituting these expressions into Equation (2) and rearranging yields $E[DWL_\tau] = \frac{\gamma+\kappa}{2}(A(\tau)^2 + (\delta B)^2)$. Because B is independent of τ, minimizing the expected deadweight loss with respect to the tax is equivalent to minimizing $A(\tau)^2$, which yields $\tau^+ = \frac{\alpha\gamma(1-\beta)}{\gamma+\kappa} = MD(Q^*) - \beta MD(Q^*)$.

Any cap Ω gives rise to a pollution level in the low- and high-MD states of $Q_{\Omega,L} = \min\left(\Omega, \frac{\alpha+\beta\delta}{\beta\gamma+\kappa}\right)$ and $Q_{\Omega,H} = \min\left(\Omega, \frac{\alpha-\beta\delta}{\beta\gamma+\kappa}\right)$. A C-Optimal cap will always bind in the low-MD state of the world, but owing to Coasean provision, it may or may not bind in the high-MD state. If it binds in the high-MD state, then the distances from welfare-maximizing pollution levels are $D_L = E + F - \Omega$ and $D_H = F - E + \Omega$, where $E \equiv \frac{\alpha}{\gamma+\kappa}$, and $F \equiv \frac{\delta}{\gamma+\kappa}$. Invoking equation (2) and simplifying yields an expected deadweight loss of $E[DWL_{bind}] = \frac{\gamma+\kappa}{2}(E^2 + F^2 + \Omega^2 - 2E\Omega)$. Applying the first-order condition and solving reveals that $\Omega^* = E = \frac{\alpha}{\gamma+\kappa} = \Omega^+$.

If the cap fails to bind in the high-MD state, then the distances from welfare-maximizing pollution levels are $D_L = J - \Omega$ and $D_H = I$, where $J \equiv \frac{\alpha+\delta}{\gamma+\kappa}$ and $I = \frac{\alpha-\beta\delta}{\beta\gamma+\kappa} - \frac{\alpha-\delta}{\gamma+\kappa}$.

[26] We consider uncertainty over marginal damages without simultaneous uncertainty over marginal benefits. While others (Stavins 1996) have considered simultaneous and correlated uncertainties, we leave such analyses in the presence of Coasean provision for future research.

In this case, invoking equation (2) and simplifying yields an expected deadweight loss of $E[DWL_{slack}] = \frac{\gamma+\kappa}{4}(J^2 - 2J\Omega + \Omega^2 + I^2)$. Applying the first-order condition and solving reveals that $\Omega^* = J = \frac{\alpha+\delta}{\gamma+\kappa} = \Omega^{++}$.

Thus, the two candidate caps are Ω^+ and Ω^{++}, with the associated deadweight losses:

$$E[DWL_{\Omega^+}] = \frac{\delta^2}{2(\gamma+\kappa)} \tag{9}$$

$$E[DWL_{\Omega^{++}}] = \frac{(\alpha\gamma+\delta\kappa)^2(1-\beta)^2}{4(\beta\gamma+\kappa)^2(\gamma+\kappa)}, \tag{10}$$

where (9) binds the the high-MD state, and (10) does not. To complete the proof, however, we must find the cutoff value, $\beta_d(\delta)$, above which it is optimal to switch from Ω^+ to Ω^{++}. Setting Equations 9 and 10 equal and solving, gives

$$\beta_d = \frac{\delta\kappa(1-\sqrt{2}) + \alpha\gamma}{\delta\gamma\sqrt{2} + \delta\kappa + \alpha\gamma}, \tag{11}$$

which is clearly decreasing in δ. Finally, if $\beta < \beta_d(\delta)$, then Ω^+ is optimal, and we must show that it is binding in the high-MD state. So we must show $\frac{\alpha}{\gamma+\kappa} < \frac{\alpha-\beta\delta}{\beta\gamma+\kappa}$, or, rearranging, we must show that $\beta < \frac{\alpha\gamma}{\alpha\gamma+\delta(\gamma+\kappa)} \equiv \beta^+$. Comparing β^+ to $\beta_d(\delta)$ in equation (11), we see that $\beta < \beta_d(\delta) < \beta^+$. Instead, if $\beta > \beta_d(\delta)$, then Ω^{++} is optimal, and we must show that it is non-binding in the high-MD state of the world. So we must show $\frac{\alpha-\beta\delta}{\beta\gamma+\kappa} < \frac{\alpha+\delta}{\gamma+\kappa}$, or, rearranging, we must show that $\beta > \frac{\alpha\gamma-\kappa\delta}{\alpha\gamma+2\gamma\delta+\delta\kappa} \equiv \beta^{++}$. Comparing β^{++} to $\beta_d(\delta)$ in equation (11), we see that $\beta > \beta_d(\delta) > \beta^{++}$, which concludes the proof. $\square$

The only difference between Propositoins 4 and 6 is the critical threshold that determines whether Ω^+ or Ω^{++} is the C-Optimal cap. Contributing to this difference is the fact that uncertainty with marginal damages also affects the private demand for pollution reductions (i.e., $PMD(Q)$).[27] In this case, the intuition behind the switch from Ω^+ to Ω^{++} is nevertheless very similar to what we described previously. When β is small (i.e. less than $\beta_d(\delta)$), setting the myopically optimal cap (Ω^+) will not induce Coasean provision, so the deadweight loss calculations are as if Coasean provision plays no role. But when β is large, Coasean provision will serve as a backstop in the high-MD state. In other words, if marginal damages turn out to be high, Coasean provision occurs and renders the cap non-binding. Recognizing this, the policymaker can set a weaker cap in order to lower the deadweight loss that would occur in the low-MD state. Indeed, in this case, the optimal cap is set at Ω^{++}, which completely eliminates deadweight loss in the low-MD state.

[27] In particular, note by equations (5) and (11) that $\beta_c(\delta) \neq \beta_d(\delta)$, except in special cases.

Turning now to the main question of instrument choice, we have the following clear result always favoring taxes over caps:

Proposition 7. *In the presence of uncertainty about the marginal damages of pollution, the expected welfare with the tax is greater than that for the cap for all values of* $\beta > 0$ *and* $\delta > 0$.

Proof. We begin by calculating the expected deadweight loss of the C-Optimal policies and then show that expected deadweight loss under the tax is always smaller. Expected deadweight loss under the tax is

$$E[DWL_{\tau^+}] = A^2\left(\frac{\gamma+\kappa}{2}\right), \tag{12}$$

where $A = \frac{\delta\kappa(1-\beta)}{(\gamma+\kappa)(\beta\gamma+\kappa)}$ is the amount of pollution in excess of the socially optimal amount in the high-MD state. Expected deadweight loss under the cap depends on which cap is optimal. If Ω^+ is C-Optimal, then no Coasean provision occurs in either state of the world. In that case, we have

$$E[DWL_{\Omega^+}] = G^2\left(\frac{\gamma+\kappa}{2}\right), \tag{13}$$

where $G = \frac{\delta}{\gamma+\kappa}$. Algebraic manipulation reveals that $G > A$, so $E[DWL_{\tau^+}] < E[DWL_{\Omega^+}]$. Instead, if Ω^{++} is the C-Optimal cap, there is no deadweight loss in the low-MD state and we have

$$E[DWL_{\Omega^{++}}] = (A+B)^2\left(\frac{\gamma+\kappa}{4}\right), \tag{14}$$

where $B = \frac{\alpha\gamma(1-\beta)}{(\gamma+\kappa)(\beta\gamma+\kappa)}$. Thus, we seek to compare $\frac{(A+B)^2}{4} \lesseqgtr \frac{A^2}{2}$. Recognizing that A and B share the same denominator, let $A = \delta\kappa c$ and $B = \alpha\gamma c$, for a constant $c = \frac{1-\beta}{(\gamma+\kappa)(\beta\gamma+\kappa)}$. Algebraic manipulation shows that the inequality of interest is $-\delta^2\kappa^2 + 2\delta\kappa\alpha\gamma + \alpha^2\gamma^2 \lesseqgtr 0$, which is a quadratic expression in δ. The positive root of this expression is $\delta^0 = \frac{\alpha\gamma}{\kappa}(1+\sqrt{2})$; so $E[DWL_{\tau^+}] < E[DWL_{\Omega^{++}}]$ if and only if $\delta < \delta^0$.

Now recall the implicit assumption from Weitzman's original analysis whereby $\delta \leq \alpha\min\left(\frac{\kappa}{2\gamma+\kappa}, \frac{\gamma}{2\kappa+\gamma}\right)$ (see Footnote 20). Denote this threshold value of δ as δ^W, and we will show that $\delta < \delta^W < \delta^0$. The first inequality is implied by Weitzman, and upon proving the second, we will have shown that $\delta < \delta^0$, which is the desired result. There are two cases to examine: $\kappa > \gamma$ and $\kappa \leq \gamma$. If $\kappa > \gamma$, then $\min\left(\frac{\kappa}{2\gamma+\kappa}, \frac{\gamma}{2\kappa+\gamma}\right)$ is given by $\frac{\gamma}{2\kappa+\gamma}$, so we compare $\frac{\gamma(1+\sqrt{2})}{\kappa} \lesseqgtr \frac{\gamma}{2\kappa+\gamma}$. Cross multiplying gives the inequality $(1+\sqrt{2})(2\kappa+\gamma) \lesseqgtr \kappa$, so $\delta^0 > \delta^W$ (when $\kappa > \gamma$). Instead, if $\kappa \leq \gamma$, then $\min\left(\frac{\kappa}{2\gamma+\kappa}, \frac{\gamma}{2\kappa+\gamma}\right)$ is given by $\frac{\kappa}{2\gamma+\kappa}$. Cross multiplying, and making use of $\kappa < \gamma$, reveals that $\delta^0 > \delta^W$. We conclude that $\delta^W < \delta^0$ and therefore that $\delta < \delta^0$, which completes the proof.

□

Under uncertainty over marginal damages, the C-Optimal tax and cap both give rise to pollution levels that are, in some sense, intermediate. They are higher than ex-post efficient (in the high-MD state, when optimal pollution levels are low) and lower than ex-post efficient (in the low-MD state, when optimal pollution levels are high). The welfare advantage of taxes arises because Coasean provision, which occurs with the tax, but generally not with the cap, always nudges that pollution level closer to the ex-post efficient level. In contrast, under uncertainty over marginal benefits, the C-Optimal tax always overshoots – it leads to less pollution than is ex-post optimal (in the low-MB state, when optimal pollution levels are low) and more pollution than is ex-post optimal (in the high-MB state, when optimal pollution levels are high). This induces a horse race between taxes and caps vis-à-vis welfare in the case when uncertainty concerns the benefits of pollution.

10 Summary and Discussion

We have shown throughout this paper that the presence of Coasean provision affects policy instrument choice. In a world of certainty, if policies are set without regard to Coasean provision, then the standard equivalence between price and quantity instruments breaks down. It turns out that between the myopic, first-best instruments, caps are more efficient. More generally, between myopically equivalent policies, caps are more efficient than taxes only when the level of policy stringency is sufficiently strong. When each of the policies is chosen optimally, they can both implement the first-best level of pollution, but the tax is lowered from the textbook Pigouvian level in anticipation of Coasean provision. Such an adjustment is not warranted with an optimally chosen cap, and under such a cap, we would expect no Coasean provision. These results are summarized in the first column of Table 2.

The other columns in Table 2 summarize our results in the presence of uncertainty with respect to either the marginal benefits or marginal costs of pollution. While Weitzman (1974) analyzes uncertainty in both $MB(Q)$ and $MD(Q)$, he concludes that only the first case has consequences for policy instrument choice. Our analysis is thus a generalization of Weitzman's to account for Coasean provision, and we find that the results differ in significant and policy-relevant ways. First, with uncertainty over $MB(Q)$, we find that the prospect of Coasean provision expands the scenarios in which price instruments are preferred to quantity instruments. This result obtains because Coasean provision interacts with a pollution tax in a manner that tends to reduce deadweight loss in any state of the world. Second, with uncertainty over $MD(Q)$, we find that the equivalence result derived by Weitzman (i.e. that prices and quantities deliver equivalent pollution levels and welfare) no longer holds. Instead, we find that with Coasean provision, pollution levels differ across policies and prices always

Table 2: Summary of C-Optimal results, instrument choice, and Coasean provision, with and without two types of uncertainty.

Outcome of Interest	No Uncertainty	$MB(Q)$ Uncertainty	$MD(Q)$ Uncertainty
C-Optimal Tax	τ^+	τ^+	τ^+
C-Optimal Cap	Ω^+	Ω^+ if $\beta < \beta_c$ Ω^{++} otherwise	Ω^+ if $\beta < \beta_d$ Ω^{++} otherwise
Optimal Policy Choice	Equivalent	Tax if $\beta > \beta^*$ Cap otherwise	Tax
Coasean Provision at Optimal Choice	Only with tax	If $\beta > \beta^*$	Always

welfare dominate quantities. The reason is that an optimally chosen tax can account for Coasean provision in ways that bring the equilibrium level of pollution closer to what is ex-post optimal. Overall, we find that compared to Wetizman's classic results, the presence of Coasean provision tips the balance towards favoring prices over quantities.

More generally, we hope the analysis contributes to a new area of research that seeks to bridge useful insights from both Pigouvian and Coasean approaches to environmental and natural resource management (Banzhaf, Fitzgerald, and Schnier 2013). Rather than view the approaches as either/or substitutes, we consider settings where both simultaneously operate. While the analysis produces novel and policy-relevant results—calling for a rethinking of policy instrument choice in the presence of Coasean provision—it also raises questions that warrant further consideration. We briefly discuss three in particular to conclude the paper.

Is Coasean provision likely to be important in the real world? While our analysis is purely theoretical, it is motivated by the increasing prevalence of what can be reasonably deemed Coasean provision. Despite the existence of wide-ranging environmental and natural resource policies, the private provision of environmental public goods is on the rise. It occurs through direct philanthropy, corporate environmental management, and consumer preferences for environmentally friendly goods and services. We nevertheless recognize that for some environmental problems, the extent to which voluntary provision will have a significant impact can be limited. These might be considered relatively low-β scenarios. But relatively high-β scenarios consistent with our model certainly exist, as evidenced by the extent of provision observed above and beyond regulatory requirements. Examples include the large-scale impact of Walmart's sourcing of sustainably harvested seafood despite fisheries regulations, climate change policies at the state level that exceed federal requirements, and international efforts to promote conservation in other countries viewed as having insufficient protections.

What about alternative motives for Coasean provision? We have assumed throughout that Coasean provision is motivated by the benefit of providing a public good (i.e., abatement), where public and private provision are perfect substitutes. But the literature on privately provided public goods considers alternative motives that include signaling (Glazer and Konrad 1996), reputation (Harbaugh 1998), and warm-glow altruism (Andreoni 1989; Andreoni 1990). A key feature of these motives is that utility from provision comes from the act of giving rather than the incremental change to the level of the public good. While such motives may underlie Coasean provision in some circumstances, we leave it to future research to examine how different motivational assumptions may operate in this setting. One reason is that behavior motivated in this way is distinct from Coasean-type bargaining, because demand for reputation benefits and warm glow is effectively demand for a private good. We might, however, expect some of the differences between taxes and caps to be attenuated because the extent of Coasean provision would not depend on the direct effect of the policies on levels of pollution. This line of inquiry also adds a wrinkle to the analysis vis-à-vis welfare measures. In settings with both public and private provision, where the later is driven by warm glow, one must contend with an additional set of questions related to non-neutrality between the mechanisms of provision and whether warm-glow benefits should be included in welfare calculations (Chilton and Hutchinson 1999; van 't Veld 2020).

Asymmetries in transaction costs are also worthy of further inquiry. An implicit assumption throughout our analysis is that transaction costs associated with Coasean provision are invariant to the choice of policy instrument. But this assumption may be unrealistic in some settings. For example, cap-and-trade programs create centralized markets to facilitate transactions that may include citizens purchasing and retiring permits, in addition to trades among regulated firms. With taxes, however, how Coasean provision takes place may be less clear, perhaps relying on bilateral negotiations, and more susceptible to standard critiques about the limitations of Coasean bargaining. To the extent such differences do arise, extensions to the analysis are possible, where, for example, β could differ depending on the policy instrument being employed. While this would alter the precise conditions that we derive, many of the qualitative findings about the potential importance of Coasean provision would remain.

References

Anderson, T. L. and D. R. Leal (2001). *Free Market Environmentalism, Revised Edition.* New York: Palgrave.

Andreoni, J. (1989). Giving with impure altruism: Applications to charity and ricardian equivalence. *Journal of Political Economy 97*(6), 1447–1458.

Andreoni, J. (1990). Impure altruism and donations to public goods: A theory of warm-glow giving. *Economic Journal 100*(401), 464–477.

Banzhaf, H. S. (2010). The free-market environmentalist case for cap-and-trade). *Paper Prepared for the Workshop on Tough Questions for Free Market Environmentalism, Property and Environment Research Center (PERC) Bozeman, Montana*.

Banzhaf, H. S. (2020). A history of pricing pollution (or, why Pigouvian taxes are not necessarily Pigouvian). *NBER Working Paper 27683*.

Banzhaf, H. S., T. Fitzgerald, and K. Schnier (2013). Nonregulatory approaches to the environment: Coasean and Pigouvian perspectives. *Review of Environmental Economics and Policy 7*(2), 238–258.

Baumol, W. J. (1972). On taxation and the control of externalities. *The American Economic Review 62*(3), 307–322.

Bergstrom, T., L. Blume, and H. Varian (1986). On the private provision of public goods. *Journal of Public Economics 29*(1), 25–49.

Buchanan, J. and W. Stubblebine (1962). Externality. *Economica 29*, 371–384.

Chilton, S. M. and W. G. Hutchinson (1999). Some further implications of incorporating the warm glow of giving into welfare measures: A comment on the use of donation mechanisms by champet al. *Journal of Environmental Economics and Management 37*(2), 202–209.

Coase, R. H. (1960). The problem of social cost. *Journal of Law and Economics 3*, 1–44.

Cornes, R. and T. Sandler (1985). The simple analytics of pure public good provision. *Economica 52*(205), 103–116.

Costello, C. and L. Karp (2004). Dynamic taxes and quotas with learning. *Journal of Economic Dynamics and Control 28*(8), 1661–1680.

Deryugina, T., F. Moore, and R. S. J. Tol (2020). Applications of the Coase Theorem. *University of Sussex, Working Paper Series No. 08-2020*.

Eshel, D. M. D. and R. J. Sexton (2009). Allowing communities to trade in imperfectly competitive pollution-permit markets. *Journal of Regulatory Economics 36*, 60–82.

Glazer, A. and K. A. Konrad (1996). A signaling explanation for charity. *American Economic Review 86*(4), 1019–1028.

Goulder, L. H., M. R. Jacobsen, and A. A. van Benthem (2012). Unintended consequences from nested state and federal regulations: The case of the pavley greenhouse-gas-per-mile limits. *Journal of Environmental Economics and Management 63*(2), 187–207.

Goulder, L. H. and R. N. Stavins (2011). Challenges from state-federal interactions in us climate change policy. *American Economics Review 101*, 253–257.

Harbaugh, W. T. (1998). The prestige motive for making charitable transfers. *American Economic Review 88*(2), 277–282.

Israel, D. (2007). Environmental participation in the U.S. sulfur allowance auctions. *Environmental and Resource Economics 38*, 373–390.

Levinson, A. (2012). Belts and suspenders: Interactions among climate policy regulations. *in Design and Implementation of U.S. Climate Change Policies, D. Fullerton and C. Wolfram eds.*.

MacKenzie, I. A. and M. Ohndorf (2016). Coasean bargaining in the presence of Pigouvian taxation. *Journal of Environmental Economics and Management 75*, 1–11.

Malueg, D. A. and A. J. Yates (2006). Citizen participation in pollution permit markets. *Journal of Environmental Economics and Management 51*(2), 205–217.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995). *Microeconomc Theory*. London: Oxford University Press.

Medema, S. G. (2019). The Coase Theorem at sixty. *Journal of Economic Literature*.

Pigou, A. C. (1932). *The Economics of Welfare, Fourth Edition*. London, UK: Macmillan and CO., Limited.

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

Turvey, R. (1963). On divergences between social cost and private cost. *Economica 30*(119), 309–313.

van 't Veld, K. (2020). Eco-labels: Modeling the consumer side. *Annual Review of Resource Economics 12*, 187–207.

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