Policy Instrument Choice with Coasean Provision of Public Goods*

Christopher Costello and Matthew Kotchen

September 13, 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

*Costello: University of California, Santa Barbara and NBER (costello@bren.ucsb.edu). Kotchen: Yale University and NBER (matthew.kotchen@yale.edu). We are grateful for helpful insights from Spencer Banzhaf, Ken Bao, Bryan Leonard, Andrew Plantinga, Arthur van Benthem, and seminar participants at Yale University and the 2021 Spring meeting of the NBER EEE group. 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.

disadvantages of various centralized policy instruments, including taxes and subsidies, direct standards, and systems of tradeable permits. There is, however, another stream of environmental economics that focuses on decentralized approaches to solving externality problems. This literature is partially founded on the seminal contribution of Ronald Coase (1960), where he finds that if property rights are well-defined and there are no transaction costs, then parties can engage in decentralized bargaining to solve externality problems.[2] And under these specific circumstances, there is no 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 alone is likely to efficiently resolve externality problems in a quite limited set of circumstances where the number of parties involved is exceedingly small. This finding underscores the need, at least in many applications, for centralized policies.

Our second starting point assumption is that Coasean-type bargaining can occur alongside centralized environmental policy; indeed, there may even be circumstances under which implementation of a centralized policy can *induce* Coasean-type bargaining. When a policy is implemented in an otherwise unregulated setting, it not only increases salience of an issue, it establishes rights and responsibilities that can serve as *de facto* property rights, a mechanism for reducing transaction costs, or both. When Coasean-type bargaining occurs in the presence of centralized policy, we refer to it as Coasean provision.[4] Whereas Coasean bargaining is often discussed in contexts where bargaining 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,

[2]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.

[3]A more nuanced interpretation of the Coase Theorem is that transaction costs are legitimate costs that should be accounted for in attempts to resolve externality problems. This idea in central to more decentralized approaches such as those based on information provision (Tietenberg 1998) and broad notions of free market environmentalism (Anderson and Hill 1975; Anderson and Leal 2001).

[4]We use the term Coasean provision whether or not the Coasean-type bargains are strictly induced by the centralized policy; for as we discuss later, the incentives for Coasean provision might also be preexisting.

and Varian 1986). In other words, rather than assume away the conditions that give rise to free riding or other impediments to efficient bargaining, we acknowledge that they occur and consider the potential importance of Coasean provision on policy instrument choice.

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

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. 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.[5] 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 de-

[5] 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 analysis considers how taking account of incentives for private provision of pubic goods affects policy instrument choice.

termines the level policy stringency. Polluters respond by maximizing profits, part of which stems from direct or indirect interaction with 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. Examples are diverse and include citizens explicitly paying polluters to emit less, consumers making purchasing decisions based on a company's environmental performance, non-polluters participating in a cap-and-trade program by retiring permits, or any combination of voluntary behaviors that reduce pollution in a regulated sector. 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 on an efficiency basis when the policy is sufficiently stringent, otherwise, for weaker policies, taxes are more efficient.

We then show how to optimally calibrate a policy's stringency to account for Coasean provision; we refer to such policies as 'conditionally optimal', and we conduct this analysis with and without Weitzman-style uncertainty. Once the policy stringency has been adjusted for Coasean provision, in the 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 no scope for 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. However, the main result from adding uncertainty is that Coasean provision tends to favor prices over quantities, compared to Weitzman's standard result. The reason is that Coasean provision plays a more prominent role with price instruments that 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 real world applications 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 important 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). One consequence of that result is that a tax, set at the standard Pigouvian level, would induce over-abatement. While the basic mechanism underlying this 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, nor do we restrict attention to the standard Pigouvian tax stringency. 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 generalization and synthesis of the "Coasean-bargaining-only" and "centralized-policy-only" approaches that is motivated by 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.[6] 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 make 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 with no uncertainty, where we analyze myopically equivalent and conditionally optimal levels of stringency. Sections 5 through 8 generalize the analysis to account for uncertainty in the marginal benefits or costs of pollution. Section 9 concludes with a summary and discussion.

[6] 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.

2 Setup with Coasean Provision

"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. The aggregate level of pollution is denoted Q. Industry benefits according to a strictly 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 $MD(Q) = \sum_{j=1}^{N} MD_j(Q)$.

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 two centralized 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 versus decentralized approaches. Instead, we consider the efficiency implications of Coasean-type bargaining that may occur concurrently or as the result of implementing a centralized policy instrument. Our starting point is one where an environmental externality exists (creating an environmental

[7]See Medema (2019) for a recent and comprehensive review of the literature related to the Coase Theorem in honor of its 60 year anniversary.

public bad), which means that any preexisting Coasean-type bargaining (if it occurred at all) did not completely resolve the market inefficiency.

We use the term "Coasean provision" to capture the behavior on which we focus. 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 sub-optimal 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.[8] 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 private marginal damage function in a reduced form way as $PMD(Q)$, and it holds by definition that $PMD(Q) \leq MD(Q)$, with the difference including the free riding effect and perhaps others factors such as transaction costs. We also assume that $PMD'(Q) > 0$.[9]

Different factors can give rise to $PMD(Q)$, including the preferences of wealthy individuals driven to environmental causes or the willingness to pay of citizens for more environmentally friendly goods and services. The function itself simply represents citizen demand for private provision of an environmental public good. To fix ideas, it is nevertheless helpful to consider an example that is fully micro-founded. Assume citizens have quasilinear preferences of the form $U(x, Q) = x_j - f_j(Q_{max} - A)$, where x_j is private consumption, $A = \sum_{j=1}^{N} a_j$ is the aggregate level of privately provided abatement, and Q_{max} satisfies $MB(Q_{max}) = 0$. This implies that $Q = Q_{max} - A$. 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 $p = PMD(Q_{max} - A)$, and this defines Q as a function of p.

Figure 1 shows the relationship between these different functions. Coasean provision would reduce the level of pollution from the quantity Q_{max} down to $\tilde{Q}$. Over this range, the

[8]It is necessary to assume no income effects in order for the marginal willingness to pay for abatement to equal the marginal damages 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.

[9]Malueg and Yates (2006) employ a similar approach to model the potential for voluntary participation of individuals in a cap-and-trade program. While this aspect of their paper is similar to ours, they focus on incentives for lobbying and not on questions about instrument choice based on prices versus quantities.

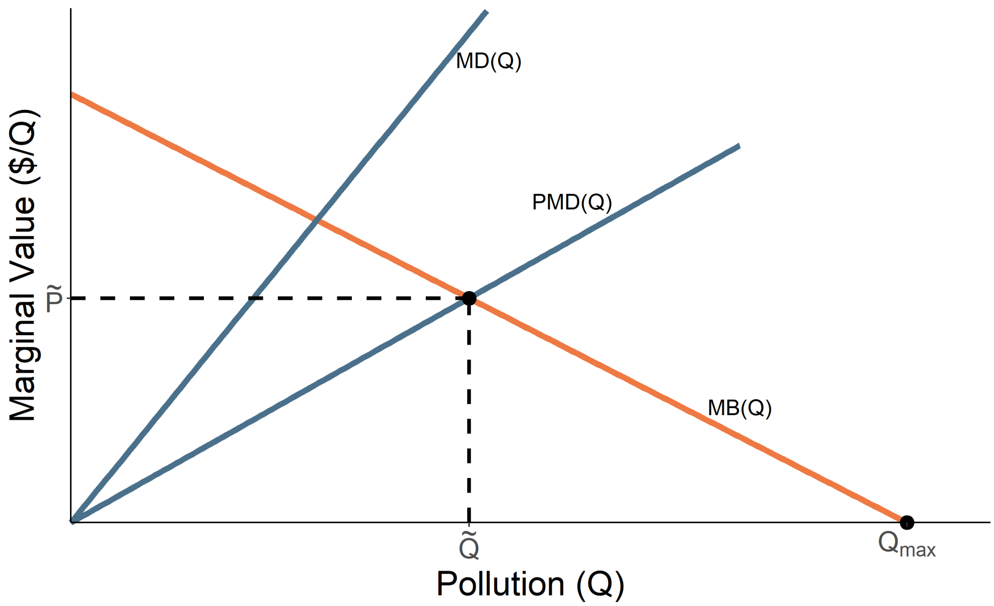

Figure 1: Coasean provision in the context of the marginal benefits and damages of pollution.

citizens' marginal willingness to pay for abatement is greater than industries' marginal willingness to accept. While the actual side payments that facilitate these transactions would depend on the relative bargaining power, one such outcome of interest might be the fixed price $\tilde{P}$ that clears all transactions. Deryugina et al. (2020) discuss a number of real-world examples of such Coasean provision in environmental and natural resource settings. We note further that the basic setup is consistent with all forms of private provision of environmental public goods—including philanthropy aimed at improving environmental quality, environmental improvements through citizen-pressured corporate environmental management, the exchange of environmentally friendly goods and services, and even transfers among countries in the form of payments for carbon offsets or forest conservation.

Also consistent with our analysis is an alternative starting point of Q_{max} (see Figure 1). Under this alternative representation, we capture situations where the $PMD(Q)$ might only become operational after a centralized policy is implemented. This is motivated by circumstances where imposing a centralized, environmental policy alters the institutional setting such that policy-induced, Coasean provision might subsequently occur, with reasons related to both property rights and transaction costs. When a policy is implemented in an otherwise unregulated setting, it often 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 provision may be suppressed prior to policy implementation because of ambiguities about baseline conditions that establish who is supposed to compensate whom.[10] Thus, at least in some cases, implementing an environmental policy may help set the stage for subsequent Coasean provision.

Implementing a centralized environmental policy can also increase salience and reduce transaction costs, providing additional reasons for policy-induced, Coasean provision. The simple fact that a policy is put in place may signal to citizens the importance of an environmental issue that alerts them to questions about whether the policymaker is doing enough and thus whether private provision is called for. And 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.[11] 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.

Finally, in tandem with policies themselves, changes in technology and information provision can also promote Coasean provision 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 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. Such trends can explain why the $PMD(Q)$ curve, and therefore the scope for Coasean provision, might look different before and after implementation of a centralized environmental policy.

[10]Indeed, Coase (1960) himself described a similar situation 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.

[11]Banzhaf (2010) makes this point as an explicit argument in favor of cap-and-trade programs. He argues in support 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 allowances from several different markets (see https://carbonvault.org/).

3 Policy Instruments

We now consider how the potential for Coasean provision affects equilibrium outcomes with implementation of the classical policy instruments of either a 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, where we assume that all revenue is used to produce public benefits of equal value in a way that does not affect marginal incentives of either industry or citizens. 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 2 for an arbitrary level of the tax. This, however, is no longer an equilibrium with $PMD(Q) > 0$ defining the scope for Coasean provision. The logic is standard Coasean bargaining, but based on the $PMD(Q)$ rather than the full social marginal damages, and we assume throughout the paper, unless otherwise indicated, that $PMD(Q) < MD(Q)$.[12] 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 equilibrium pollution level with Coasean provision. That is, under a tax of τ, Coasean provision reduces pollution from the standard policy target of $\bar{Q}$ all the way down to $\hat{Q}$. Although not shown in the figure, setting a tax rate lower at $\tau \leq MB(\tilde{Q})$ but greater than zero will still induce an equilibrium level of pollution $\hat{Q} < \tilde{Q}$, regardless of whether the Coasean provision is preexisting or policy induced. That is, regardless of whether the no-policy level of pollution were $\tilde{Q}$ or Q_{max}, the level of pollution that results from a tax of τ with Coasean provision will always satisfy $MB(\hat{Q}) - \tau = PMD(\hat{Q})$.

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. We illustrate this with the $MB(Q)-$

[12]This assumes the existence of a market failure so that the question of policy instrument choice has relevance. Otherwise, if $PMD(Q) = MD(Q)$, we have the full Coasean solution at the efficient level of pollution without any policy instrument.

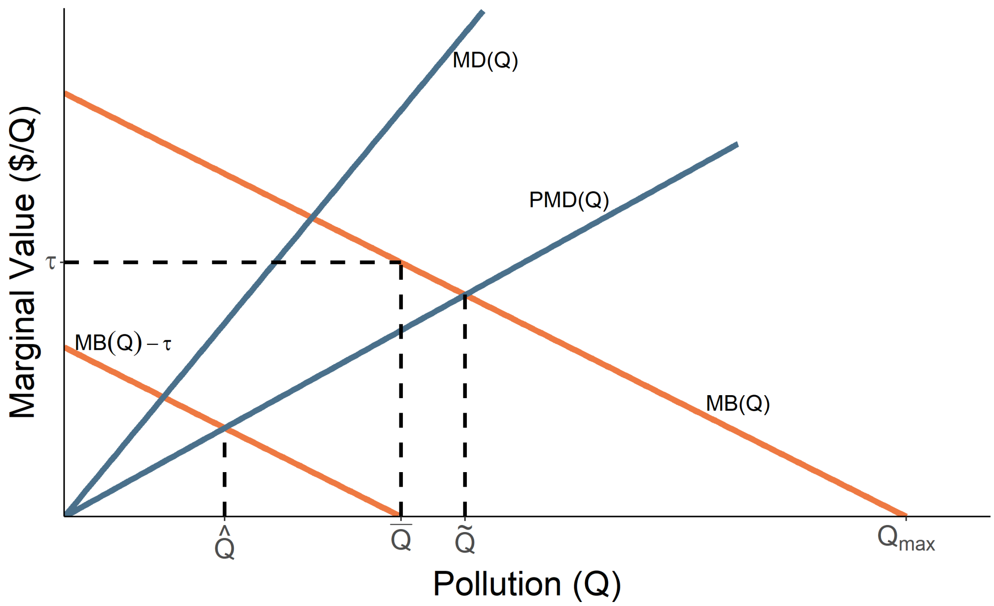

Figure 2: Equilibrium with Coasean provision and a tax of τ.

τ curve in Figure 2. The effect is that any side payments (either explicit or implicit) 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 a single polluter and a single citizen is implicitly considered in Buchanan and Stubblebine (1962). While that special 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 externality imposed on a single agent. This is important, for reasons we consider below, because with any degree of free riding or transaction costs, the difference between private marginal willingness to pay and social marginal damages creates distinct welfare implications in the context of policy instrument choice.[13]

Finally, we turn to the question of whether such Coasean provision occurs in practice alongside implementation of a pollution tax. Real-world examples matching the setup simply require the existence of a pollution tax along with additional citizen (or consumer) driven reductions in pollution on the part of the regulated industry. For example, despite the fact that Denmark has a national carbon tax, the city of Copenhagen has made a public commitment to carbon neutrality, and similar arrangements are taking place in many other cities

[13]Moreover, even with a single citizen, the standard Coasean argument is support of efficiency does 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.

and countries.[14] Also consistent with the setup are markets for environmentally friendly goods and services, where consumers pay a price premium for green products form an industry subject to a tax, e.g., purchases of green electricity in a power sector subject to a carbon tax. Another example is that 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 is partnership with non-governmental organizations that aims to reduce solid waste even further.[15]

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 3.

We begin with the case of a sufficiently stringent cap such that $\Omega^L \leq \tilde{Q}$, which is also shown in Figure 3. 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. Citizens' willingness to pay falls short by the red dashed line shown in Figure 3. 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 the one shown in Figure 3 with a sufficiently weak cap such that $\Omega^H > \tilde{Q}$. First consider the case where Coasean provision is policy induced, that this, with no policy the level of pollution would be Q_{max}. In that case, the cap at Ω^H causes abatement by industry (from Q_{max} down to Ω^H), but it also triggers additional Coasean provision, because 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 3. In order to express this willingness to pay in a market, we assume that citizens are able to purchase permits, either directly from the initial allocation auction, or subsequently in a

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

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

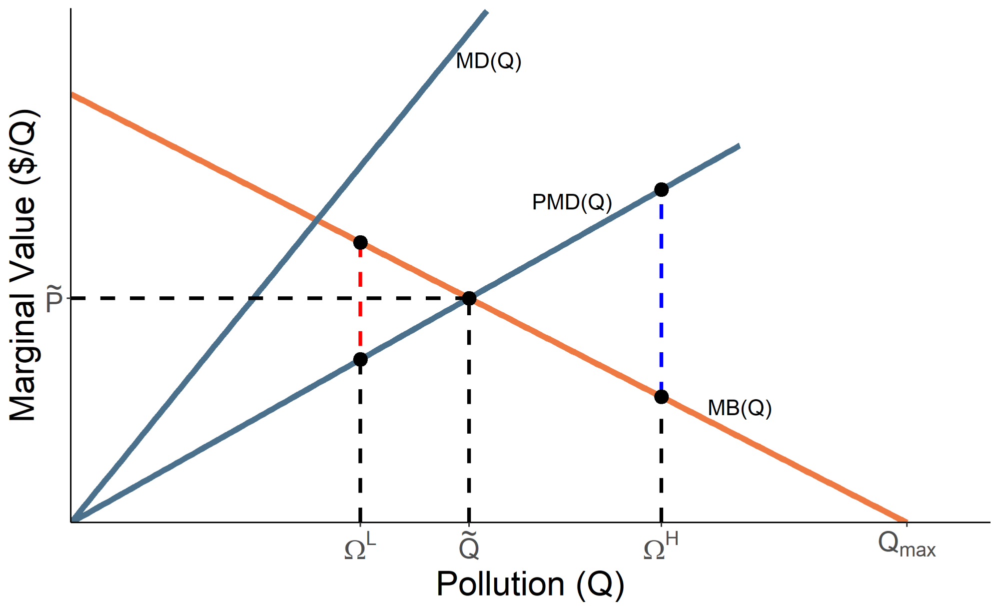

Figure 3: Equilibrium with Coasean provision and a cap-and-trade program set at Ω.

secondary market.[16] This implies a combined industry and citizen inverse demand function for government issued permits that can be written as

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

Hence, 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 3. An important observation for the moment is that equilibrium pollution is less than that targeted by the policy, as abatement of the amount $\Omega^H - \tilde{Q}$ arises from Coasean provision.

How would the preceding logic and equilibrium differ if Coasean provision where preexisting, that is, if the no-policy level of pollution were $\tilde{Q}$? One possibility is that the policy is simply not binding and the equilibrium level of pollution remains at $\tilde{Q}$. The other possibility is that citizens shift from what they had been doing to incentivize pollution reductions to purchasing and retiring permits with exactly the same result—an equilibrium level of pollution at $\tilde{Q}$, along with a market-clearing permit price of $\tilde{P}$. We have thus shown, as we did previously with the tax, that the equilibrium condition with implementation of a cap is invariant to whether the Coasean provision is preexisting or policy induced. Throughout the

[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)

remainder of the paper we therefore make no further mention of the distinct possibilities for preexisting or policy-induced Coasean provision, for while it matters for distribution, it has no bearing on efficiency or on the ultimate level of pollution.[17]

Does Coasean provision occur in practice along side implementation of cap-and-trade programs? Asproudis and Weyman-Jones (2020) provide a recent overview of the many instances in which communities, citizens, and environmental organizations participate in tradeable allowance markets. These include the European Union Emissions Trading Program, the U.S. Sulfur Allowance Trading Program, and the Southern California Regional Clean Air Incentives Market. Indeed, several environmental organizations have formed around the idea of promoting such transactions, and as mentioned previously, third parties (e.g., Carbon Vault) are being formed to facilitate the purchase and retirement of allowances in a range of markets.

4 Instrument Choice with No Uncertainty

We now turn to comparisons of instrument choice in the presence of Coasean provision. We begin with a comparison of tax versus cap policies assuming that a regulator overlooks the possibility for Coasean provision when choosing policy stringencies. We then consider 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.

4.1 M-Optimal Policies

Let us first consider tax of τ^* and cap of Ω^* 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 does not take account of the possibility for Coasean provision. They are also the textbook levels of stringency for efficient environmental policy.

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, because $\Omega^* < \tilde{Q}$. However, the M-Optimal tax

[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 assumption we make throughout, it is one that we discuss again later in the paper.

does not, because the equilibrium conditions implies $MB(\hat{Q}) - \tau^* = PMD(\hat{Q})$ and thus $\hat{Q} < \Omega^*$. 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 because the equilibrium level of pollution is inefficiently low.*

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.

4.2 M-Equivalent Policies

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). In particular, M-Equivalent policies 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.

Based on this definition of stringency, we now assert the next proposition, which we prove formally and illustrate graphically.

Proposition 2. *Consider M-Equivalent policies that aim to be binding but still allow some level of pollution (i.e., $0 < S < MB(0)$). 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 the tax is greater for all $S < \mathring{S}$, whereas welfare with the cap is greater for all $S > \mathring{S}$.*

[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).

Proof. We first define welfare as a function of stringency for each instrument. The equilibrium condition for a tax of stringency S is $MB(Q) - S = PMD(Q)$. This implicit function implies pollution as a function of tax stringency, $\hat{Q}(S)$, which is strictly decreasing. It follows that welfare with the tax, $W_\tau(S) = B(\hat{Q}(S)) - D(\hat{Q}(S))$, is a strictly 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

$$Q(S) = \begin{cases} \tilde{Q} & \text{if } S \leq \tilde{S} \\ MB^{-1}(S) & \text{if } S > \tilde{S} \end{cases}.$$

This function, which defines the equilibrium level of pollution given a level of cap stringency, is weakly decreasing in S. It follows that welfare with the cap, $W_c(S) = B(Q(S)) - D(Q(S))$, is constant for $S \leq \tilde{S}$ and strictly concave for $S > \tilde{S}$. Note that because the equilibrium conditions do not depend on whether the Coasian provision is preexisting or policy induced, the welfare functions apply in either case.

We now show that these two continuous welfare functions can cross only once. Welfare under either policy instrument is maximized at a quantity of pollution that satisfies $MB(Q^*) = MD(Q^*)$. Using the equilibrium conditions, this implies an optimal level of stringency for the tax of $S^+ = MB(Q^*) - PMD(Q^*)$ and the cap of $S^* = MB(Q^*)$. It follows that $S^+ < S^*$ and $W_\tau(S^+) = W_c(S^*)$. Then, because $W_\tau(S)$ is strictly decreasing for all $S > S^+$ and $W_c(S)$ is weakly increasing for all $S < S^*$, there exists a unique level of stringency $\mathring{S} \in (S^+, S^*)$ that satisfies $W_\tau(\mathring{S}) = W_c(\mathring{S}) < W_\tau(S^+) = W_c(S^*)$.

There are two steps remaining to complete the proof. The first is to show that $W_\tau(S) > W_c(S)$ for all $S < S^+$. At these low levels of stringency, the level of pollution is inefficiently high with the tax, but it is even higher with the cap, so welfare must be greater with the tax than the cap. The second step is to show that $W_\tau(S) < W_c(S)$ for all $S > S^*$. At these high levels of stringency, the level of pollution is inefficiently low with the cap, but it is even lower with the tax, so welfare must be greater with the cap than the tax. □

Proposition 2 shows that the results of Proposition 1 do not generalize to 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 an 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.

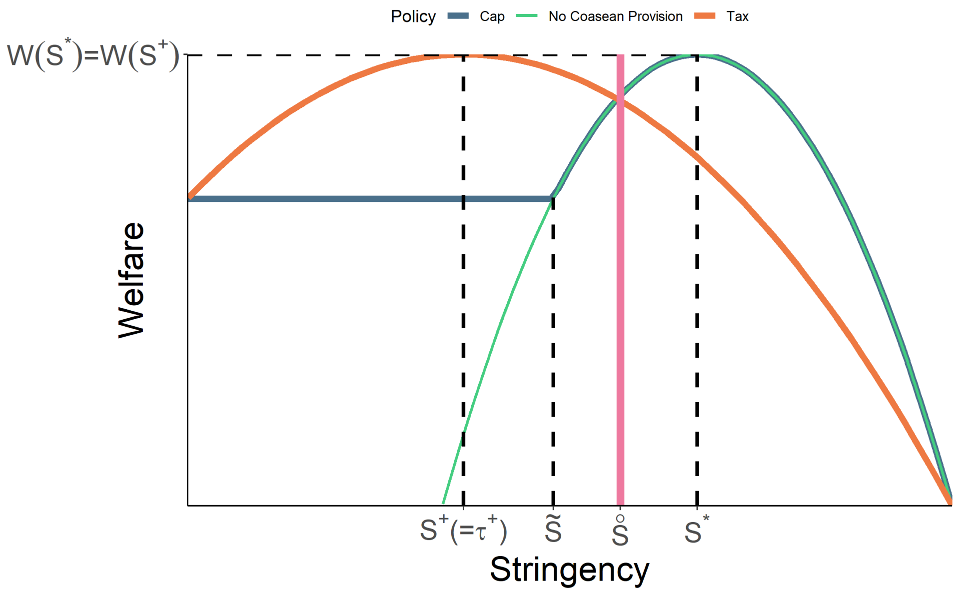

Figure 4: Welfare as a function of policy stringency for a tax and cap, and a policy scenario with no Coasean provision, in which case the tax and cap are equivalent. The figure is based on the simplifying assumption of linear marginal benefit and damage functions, and the case where $PMD(Q) \leq MD(Q)$ holds strictly.

Figure 4 illustrates graphically the results of Proposition 2, assuming linear marginal benefit and damage functions. The figure plots welfare against policy stringency for both the tax and cap scenarios with Coasean provision. It also includes a third scenario of no Coasean provision as a familiar starting point. Consider how stringency affects welfare in the absence of Coasean provision, as shown with the thin green curve. This applies when $PMD(Q)$ is effectively zero, and S^* denotes the optimal stringency of both the tax and cap that maximizes welfare. Now consider use of a cap in the presence of Coasean provision (blue curve). This curve coincides with the previous curve, so welfare is unaffected by Coasean provision, except at sufficiently low levels of stringency where $S < \tilde{S}$. In those cases, Coasean provision establishes an upper bound on pollution and therefore a lower bound on welfare. Thus, if a cap is used, Coasean provision provides a backstop for pollution, but only when the cap is sufficiently weak. Finally, consider use of a tax with Coasean provision (orange curve). In that case, the entire welfare curve is shifted to the left because for any level of M-Equivalent stringency, the tax produces a lower equilibrium level of pollution. This effect means that the tax welfare dominates the cap at sufficiently low levels of stringency, whereas the cap dominates the tax at sufficiently high levels of stringency.

4.3 C-Optimal Policies

With M-Equivalent policies, we considered exogenously set levels of stringency that aim to achieve the same level of pollution without recognizing the potential for Coasean provision. Here we show how, instead, a regulator can explicitly account for Coasean provision and choose the respective policy stringencies to maximize welfare. We refer to these as conditionally optimal (C-Optimal) policies, where the level of stringency is chosen optimally for either the tax or cap.

We know that regardless of instrument choice, the efficient level of pollution will satisfy $MB(Q^*) = MD(Q^*)$. We have in fact used this observation already in the proof of Proposition 2. We also established previously that the M-Optimal cap implements precisely this level of pollution as the equilibrium with $\Omega^* = Q^* = \Omega^+$, where the plus notation denotes C-Optimal policies. The M-Optimal tax at $\tau^* = MB(Q^*)$ does not implement the first best level of pollution, 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 standard Pigouvian tax is effectively too stringent. It is straightforward to see that lowering the tax to $\tau^+ = MD(Q^*) - PMD(Q^*)$ implements Q^* as the equilibrium level of pollution and is therefore C-Optimal. This result is also one we employ in the proof of Proposition 2.

Together, these results for C-Optimal policies prove the following proposition.

Proposition 3. *When comparing C-Optimal policies, both τ^+ and Ω^+ implement the first-best level of social welfare, and it is the same level that arises through welfare maximization without 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; it is the same as that which could be achieved even without Coasean provision. Figure 4 illustrates these results too. The "backing off" of the C-Optimal tax is reflected in the way that $\tau^+ = S^+ < S^*$, and the welfare effects are shown in the way that $W(S^+) = W(S^*)$, which is also equal to the level of maximized welfare with no Coasean provision.

5 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 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 8. 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., $\kappa = \gamma$). More generally, taxes (or caps) are preferred if the marginal damage (benefit) function is flatter, that is, if $\gamma < (>) \kappa$.[19]

[19] 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. Violations of this assumption were examined by Goodkind and Coggins (2015). In our setting, we implement the implicit Weitzman conditions as requiring that the full range of candidate caps or taxes (i.e., those spanning what would be ex-post optimal in the low or high states of the world) must deliver interior welfare measures. Restricting attention to this range ensures that both the cap and tax derived as being ex-ante optimal in Weitzman

Our aim in the next three sections is to consider the ways in which the standard uncertainty results, including equation 1, change in the presence of Coasean provision. We also show in Section 8 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 instead a clear result of always favoring taxes over caps.

6 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.[20] 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.[21] Let Q_i^* denote the welfare-maximizing ex-post level of

(1974) are indeed welfare-maximizing. Applied to our setting, 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.

[20] 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 19.

[21] As will become clear, it is convenient to establish results based on minimizing deadweight loss rather

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 $MB(Q) \pm \delta = PMD(Q)$

pollution in state of the world $i \in \{L, H\}$; this 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 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^*).*

Proof. The cap Ω^{++} ensures zero deadweight loss in the high state and might induce Coasean provision in the low state. Any cap above Ω^{++} would yield a weakly higher level of pollution than either Q_L^* or Q_H^*, and thus weakly increases the deadweight loss in either state. This means we can rule out any cap $\Omega > \Omega^{++}$.

We next rule out any cap $\Omega < \Omega^+$. Let $\mathcal{L}(\Omega)$ denote expected deadweight loss from cap Ω. We know from Weitzman (1974) that if Ω^+ binds in both states, then it welfare-dominates

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.

all other caps that also bind in both states (call this benchmark deadweight loss $\bar{\mathcal{L}}$). The cap Ω^+ is always binding in the high state because $\Omega^+ = \frac{\alpha}{\gamma+\kappa} < \frac{\alpha+\delta}{\beta\gamma+\kappa}$ (where the latter term is the crossing of $PMD(Q)$ and $MB(Q)+\delta$). Suppose it is also binding in the low state. Then, any lower cap $\Omega^- < \Omega^+$ must also be binding in the low state. Since both candidate caps are binding in both states of the world, we know that $\mathcal{L}(\Omega^+) = \bar{\mathcal{L}} < \mathcal{L}(\Omega^-)$, so Ω^+ welfare-dominates any lower cap. Instead, suppose Ω^+ is not binding (i.e. induces Coasean provision) in the low state, then we have to compare it to both marginally lower caps that are also non-binding and to substantially lower caps that are binding. First, a marginally lower cap has no effect on deadweight loss in the low state, but raises deadweight loss in the high state, so it is welfare-decreasing. Now consider a substantially lower cap, $\Omega^- << \Omega^+$, which binds in both states of the world. Since it binds in both states of the world, this cap has $\mathcal{L}(\Omega^-) > \bar{\mathcal{L}}$. And because in this case Ω^+ is not binding, we know that $\mathcal{L}(\Omega^+) < \bar{\mathcal{L}}$ (because the deadweight loss in the high state is the same as in the binding case, but Coasean provision reduces deadweight loss in the low state relative to the binding case). Putting these together reveals $\mathcal{L}(\Omega^+) < \bar{\mathcal{L}} < \mathcal{L}(\Omega^-)$. Thus, we can rule out any cap $\Omega < \Omega^+$.

It follows that the only possible caps that could minimize expected deadweight loss are in the interval $[\Omega^+, \Omega^{++}]$. Consider a strictly interior cap $\Omega' \in (\Omega^+, \Omega^{++})$. If Ω' binds in the low state, then $\Omega^+ < \Omega'$ must also bind in the low state, so Ω' is strictly dominated by Ω^+. Instead, if Ω' does not bind in the low state, then raising the cap above Ω' has no effect on low-state pollution (or deadweight loss), but lowers deadweight loss in the high state. Thus, the cap should be raised. This argument perpetuates until Ω^{++} is reached, so the optimal cap must be either Ω^+ or Ω^{++}. □

Lemma 1 expands the set of caps that can maximize welfare, relative to what was found in Weitzman (1974), where the cap Ω^+ was *always* the optimal cap. We find that the prospect of Coasean provision gives rise to a second possibility that will be optimal under some circumstances (Ω^{++}). A consequence of Lemma 1 is that if the optimal cap is Ω^+, there is no Coasean provision in either state of the world, and if the optimal cap is Ω^{++}, then there is Coasean provision, but only in the low-MB state.

The next proposition summarizes our findings on the C-Optimal policies under uncertainty:

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 5 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^{++}.[22] 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 5 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 based on the parameters used to generate Figure 5). How does the preference of the high cap (Ω^{++}) depend on β? From Figure 5, a higher β pivots the $PMD(Q) = \beta MD(Q)$ upward, which has no consequence for area a, but shrinks area c. Thus, a higher value of β further advantages the high cap. Instead, if β is sufficiently low, area c becomes larger than area a, and the lower cap (Ω^+) is

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

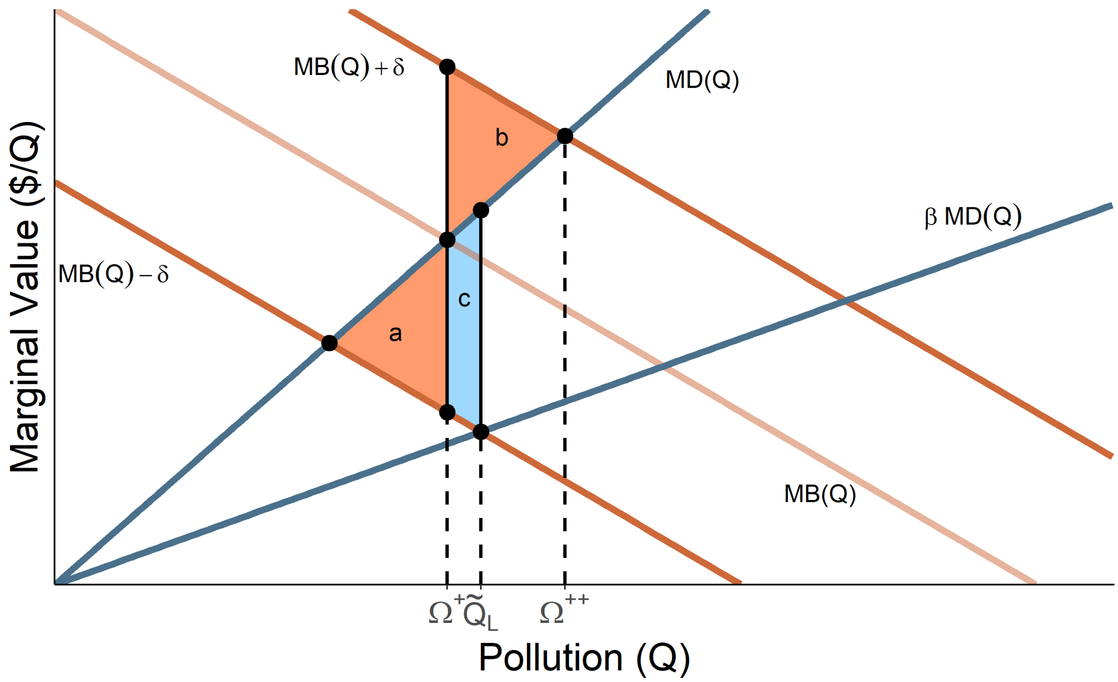

Figure 5: Graphical representation of caps Ω^{+} and Ω^{++} and their deadweight losses for an arbitrary value of β.

optimal. A similar graphical comparison can be made with higher or lower δ. Indifference between the two caps is given algebraically by equation (5).

In sum, Proposition 4 reveals the effect of uncertainty on policy stringency. It has no effect on the stringency of the tax, and 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.[23]

7 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

[23]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 larger or smaller, 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 19. 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. □

The fundamental insight of Proposition 5 is that a greater β—i.e., scope for Coasean provision—tends to imply that taxes are preferred to 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

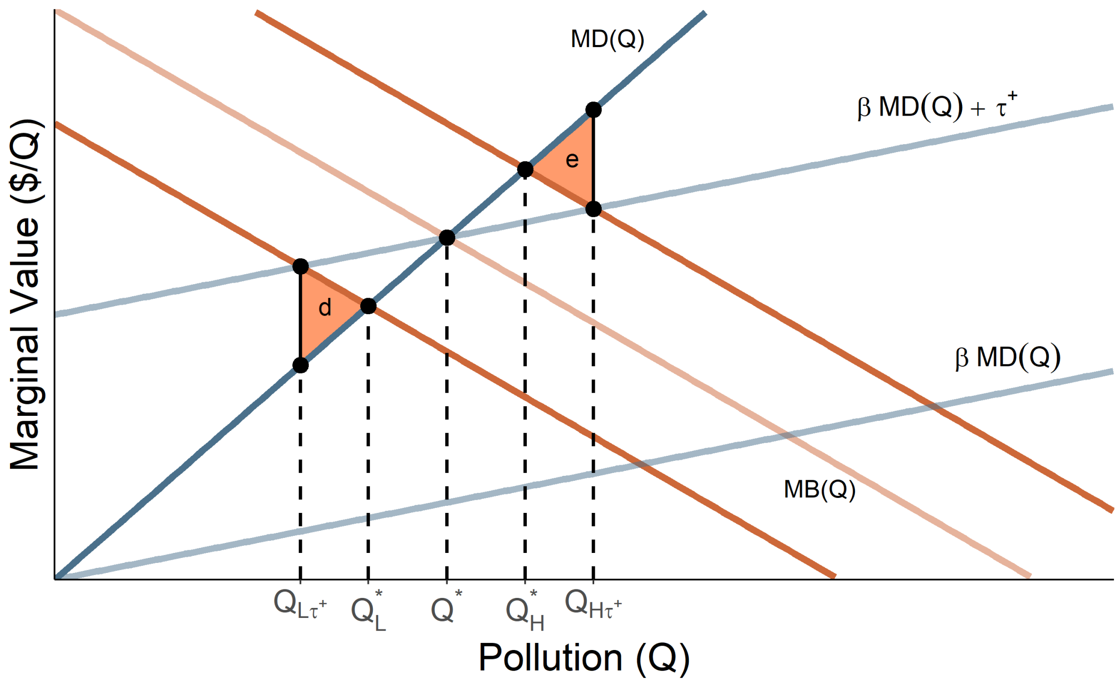

Figure 6: Expected deadweight loss with a tax, uncertainty, and Coasean provision is equal to $(d + e)/2$, which is decreasing in β.

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

Figure 6 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 $(d + e)/2$ (indeed, areas d and e are symmetric, so expected deadweight loss is just area d). How does this deadweight loss depend on β? It is clear from the figure that pivoting $\beta MD(Q)$ up allows greater scope for Coasean provision, and this increase in β makes $\beta MD(Q) + \tau^+$ steeper while maintaining the same intersection with $MB(Q)$ at Q^*.[24] This raises $Q_{L\tau^+}$ and lowers $Q_{H\tau^+}$, implying that in any state of the world, the deadweight loss is lower. 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

[24]To see why the intersection of $\beta MD(Q) + \tau^+$, $MB(Q)$, and $MD(Q)$ is the same for any β, consider the following: First, $MB(Q) = MD(Q)$ only at Q^*, which is independent of β. To see that $\beta MD(Q^*) + \tau^+$ intersects at the same point, recall that the tax is set such that $\tau^+ = MB(Q^*) - \beta MD(Q^*)$, so $\beta MD(Q^*) + \tau^+ = MB(Q^*)$, which, as we just argued, equals $MD(Q^*)$.

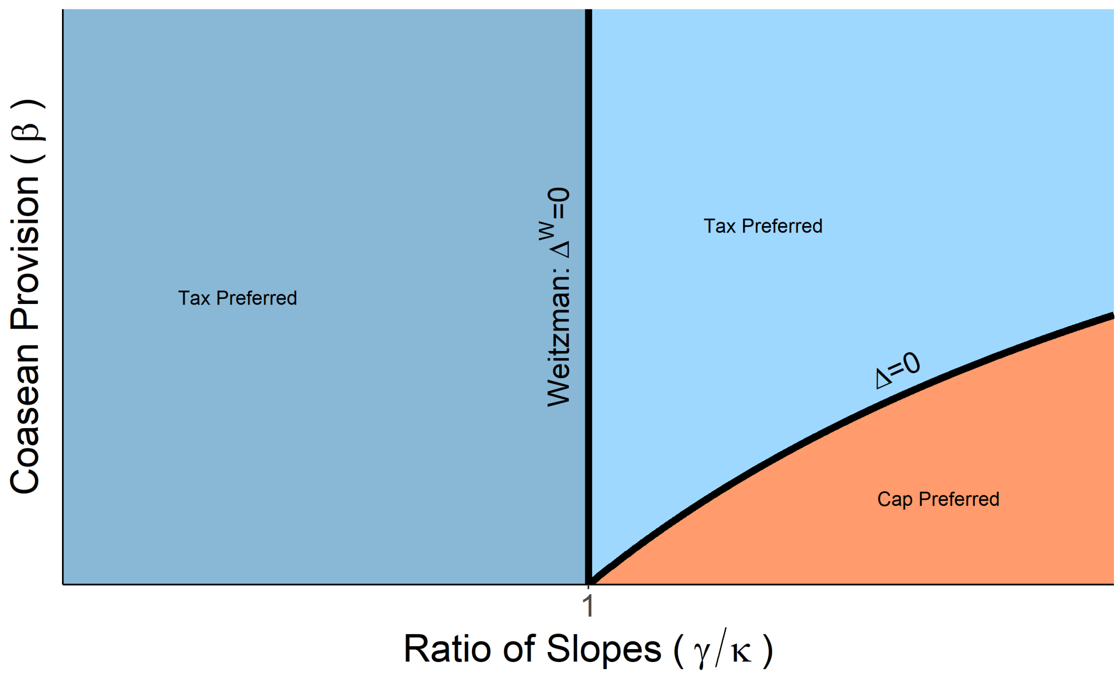

Figure 7: Parameter space over which the tax or cap delivers higher welfare. Blue areas indicate a preference for the tax, and red areas indicate a preference for the cap. Without Coasean provision (i.e., $\beta = 0$), welfare for the cap equals that for the tax when $\Delta^W = 0$.

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 provision (i.e., $\beta = 0$), we recover precisely Weitzman's result in equation (1). 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 7 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 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 7.

8 Instrument Choice with $MD(Q)$ Uncertainty

Uncertainty over marginal damages may arise for numerous reasons, including scientific uncertainty over the physical effects of pollution, sparse data on epidemiological effects of pollution, or uncertainty over ecosystem services that may be affected by pollution. Weitzman (1974) 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.[25]

Marginal benefits are given by $MB(Q) = \alpha - \kappa Q$ and expected marginal damages are $MD(Q) = \gamma Q$. We continue to assume $PMD(Q) = \beta MD(Q)$. Commensurate with our treatment of uncertainty above, let realized marginal damages be given by $MD(Q) = \gamma Q \pm \delta$, where we retain the assumption that private demand for pollution 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-

[25] 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.

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 $I = \frac{\alpha-\beta\delta}{\beta\gamma+\kappa} - \frac{\alpha-\delta}{\gamma+\kappa}$. 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 yields

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

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. □

The only difference between Propositions 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., $PMD(Q)$).[26] 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

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

$\beta_d(\delta)$), 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.

Turning now to the main question of instrument choice, we have the following sharp 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 $\frac{(A+B)^2}{4} \lesseqgtr \frac{A^2}{2}$, or more compactly, $\frac{(A+B)}{2} \lesseqgtr \frac{A}{\sqrt{2}}$. 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 reveals the inequality $\sqrt{2}\alpha\gamma \lesseqgtr \delta\kappa(2-\sqrt{2})$, which implies the cutoff value $\delta^0 = \frac{\alpha\gamma}{\kappa}(1+\sqrt{2})$; so $E[DWL_{\tau^+}] < E[DWL_{\Omega^{++}}]$ if and only if $\delta < \delta^0$.

Recall the implicit assumption from Weitzman's original analysis that corner solutions cannot be hit in any state of the world under the range of candidate policies. We formalized this in footnote 19 for the case of uncertainty over $MB(Q)$. Instead, with uncertainty over $MD(Q)$ the condition is that $\delta < \delta^W = \frac{\alpha\gamma}{2\gamma+\kappa}$. We will show that $\delta < \delta^W < \delta^0$. The first inequality is implied from Weitzman. To evaluate the second, we wish to show $\frac{\alpha\gamma}{2\gamma+\kappa} < \frac{\alpha\gamma}{\kappa}(1+\sqrt{2})$, which is easily verified by cross multiplying and simplifying. This proves that $\delta^W < \delta^0$, which implies that $\delta < \delta^0$ and 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 this way, Coasean provision is harnessed, along with the policy itself, to produce "reasonably good" outcomes under any state of the world.[27] In contrast, we showed above that under uncertainty over marginal benefits, the C-Optimal tax always overshoots—that is, it leads to less pollution than is ex-post optimal in the low-MB state (when optimal pollution levels are low) or 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.

9 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

[27] In a somewhat related manner, Shavell (1984) examined the relative merits of liability rules versus direct regulation to manage accident risk. Like our result that Coasean provision can work in tandem with an optimal policy, he found that both liability rules and regulation could be jointly applied, with each providing a kind of backstop for the other's shortcomings.

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

(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 over 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 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 Weitzman'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 is reasonably characterized as 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, Coasean provision must take place in a more decentralized manner. 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 P. J. Hill (1975). The evolution of property rights: a study of the american west. *The Journal of Law and Economics 18*(1), 163–179.

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.

Asproudis, E. and T. Weyman-Jones (2020). How the ENGOS can fight the industrial/business lobby with their tools from their own field? ENGOS participation in emissions trading market. *Sustainability 12*(8553).

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 Champ et 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. 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.

Goodkind, A. L. and J. S. Coggins (2015). The Weitzman price corner. *Journal of Environmental Economics and Management 73*, 1–12.

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.

Shavell, S. (1984). A model of the optimal use of liability and safety regulation. *The Rand Journal of Economics 15*(2), 271–280.

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

Tietenberg, T. (1998). Disclosure strategies for pollution control. *Environmental and Resource Economics 11*(3-4), 587–602.

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.