

NBER WORKING PAPER SERIES

INFORMATIONALLY EFFICIENT CLIMATE POLICY:
DESIGNING MARKETS TO MEASURE AND PRICE EXTERNALITIES

Derek Lemoine

Working Paper 30535
<http://www.nber.org/papers/w30535>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2022

This work benefited from comments at the 2020 Centre for Economic Policy Research Climate Change Workshop, the 2021 National Bureau of Economic Research Summer Institute, the 2022 Conference on Efficient Climate Policies in an Uncertain World (Collège de France), the 2022 Encuentro Anual de la Sociedad Economía de Chile, and several conferences and seminars. This paper's ideas originated in the applied policy perspective of “Incentivizing Negative Emissions Through Carbon Shares”. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Derek Lemoine. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Informationally Efficient Climate Policy: Designing Markets to Measure and Price Externalities
Derek Lemoine
NBER Working Paper No. 30535
October 2022
JEL No. D82,G14,H23,Q54,Q58

ABSTRACT

I study how policymakers can access and act on the information about climate change damages that is dispersed throughout the economy. I analyze a new dynamic deposit-refund instrument (called “carbon shares”) that I show can: i) efficiently price emissions conditional on information, ii) efficiently incentivize removal of past emissions conditional on information, and iii) efficiently aggregate dispersed information about the social cost of emissions. Conventional emission taxes generally succeed at only the first of these objectives. Rather than projecting damages in all future periods and all possible states of the world in order to calculate the optimal tax, the regulator here estimates damages as they are realized and empowers markets to perform price discovery about future damages.

Derek Lemoine
Department of Economics
University of Arizona
McClelland Hall 401EE
Tucson, AZ 85721
and CEPR
and also NBER
dlemoine@email.arizona.edu

A data appendix is available at <http://www.nber.org/data-appendix/w30535>

Despite the validity in principle of the tax-subsidy approach in the Pigouvian tradition, in practice it suffers from serious difficulties. For we do not know how to estimate the magnitudes of the social costs, the data needed to implement the Pigouvian tax-subsidy proposals.

– Baumol (1972, p. 316)

The practical problem, however, arises precisely because these facts are never so given to a single mind, and because, in consequence, it is necessary that in the solution of the problem knowledge should be used that is dispersed among many people.

– Hayek (1945, p. 530)

1 Introduction

Economists have long emphasized the informational advantages of market-based policies for controlling pollution. Market-based policies only require the regulator to measure the social cost of pollution (i.e., the externality), whereas command-and-control policies also require the regulator to measure each firm or agent’s cost of reducing pollution. Because actual firms and agents know more about their own options to reduce pollution than does the regulator, market-based policies increase efficiency by empowering them to employ their most cost-effective options. However, even market-based policies may demand a lot of information from a regulator: it is often quite challenging to measure the social cost of pollution, as Baumol (1972) recognizes in the opening quotation. In particular, economists have struggled to measure the social cost of the carbon emissions that drive global climate change.¹ As a result, economists have still not converged on an emission price to recommend to policymakers interested in market-based solutions.²

To date, measurement of social costs has been centralized among academics and regulators, but much information about the cost of climate change is dispersed throughout society.

¹Nordhaus (2019, p. 1998) acknowledges, “In reality, projecting impacts is the most difficult task and has the greatest uncertainties of all the processes associated with global warming.” Some economists even criticize the social cost estimates underlying the most prominent climate-economy models as “completely made up, with no theoretical or empirical foundation” (Pindyck, 2013, p. 868). Recent work estimates the effects of weather shocks and uses these effects to project the consequences of future climate change (Deschnes and Greenstone, 2007; Carleton and Hsiang, 2016), but despite many advances in this literature, fundamental questions remain about how to map consequences of weather shocks to consequences of climate change (Dell et al., 2014; Lemoine, 2021b).

²In a survey of 289 economists with expertise on climate change (Holladay et al., 2009), estimates of the social cost of carbon had a standard deviation of \$339 per tCO₂, around three times larger than the average estimate. The social costs of carbon calculated in Pindyck (2019) from another survey of 113 economists also shows substantial dispersion, with the standard deviation again of comparable magnitude to the mean.

Every person and firm on this planet is exposed to climate change. Each may have some information about their own particular exposure and about their own particular ability to adapt. Since at least Hayek (1945), a rich tradition in economics views markets as an algorithm that aggregates dispersed information about the costs and benefits of the many goods produced in society. Yet environmental economists have not studied how to design markets to aggregate information about externalities, and thereby to perform price discovery for social costs.

I design and analyze a new market-based policy instrument that simultaneously measures and controls externalities. In my setting, as is customary, firms in different sectors of the economy trade off the benefit of emitting carbon against the cost of complying with current policies. Carbon emissions generate warming that impacts each sector of the economy in an uncertain, stochastic, and potentially correlated fashion. I introduce two novel features to this environment. First, firms can pay to remove old emissions from the atmosphere.³ Second, information about climate change damages is heterogeneous. Agents measure climate impacts in their own sectors, and a regulator measures the aggregate effect of climate change from data on final good production and temperature. Both types of measurements may be arbitrarily imperfect. Agents' measurements are private information, whereas the regulator reports its measurement to all actors in the economy and can use it to set policy.

I show that a regulator who uses an emission tax policy is unable to optimally use new information about social costs or to observe all of the information about social costs that is dispersed throughout society.⁴ First, if the regulator did have access to all agents' information about social costs, then it could use an emission tax to efficiently control new emissions (as has long been known). However, an emission tax on its own cannot incentivize the removal of past emissions other than as means to offset ongoing emissions: once an emitter has paid the penalty, subsequent information is irrelevant even if it dramatically alters the estimated social cost of previously emitted carbon and warrants paying for its removal. An emission tax ultimately places the risk of needing to fund large-scale carbon removal on taxpayers, which could impose an impractical fiscal burden (Bednar et al., 2019).⁵ Second, the regulator will not actually have access to all agents' information. I show that this regulator's emission tax is informationally inefficient in the plausible case that the regulator measures aggregate

³Carbon dioxide removal, or negative emission, strategies include chemically separating carbon dioxide from air ("direct air capture"), capturing emissions from power plants that burn biomass ("bioenergy with carbon capture and storage"), accelerating the weathering of rocks, enhancing uptake of carbon by forests or oceans, and more. See National Research Council (2015), Fuss et al. (2018), and National Academies of Sciences, Engineering, and Medicine (2018) for recent reviews. Recently, Microsoft and Stripe each received bids to undertake carbon removal for around \$150 per tCO₂ on average (Joppa et al., 2021).

⁴Similar critiques apply to cap-and-trade programs, the quantity version of an emission tax.

⁵A forward-thinking regulator could save the revenue collected by an emission tax for the procurement of future carbon removal, but as analyzed in Section 4.1, even that revenue would be insufficient if carbon removal becomes desirable because the regulator learns climate change is more damaging than they had believed at the time of emission.

consequences only imperfectly and also in the plausible case that sectors have heterogeneous value shares and correlated exposure to climate change. In either case, an all-seeing regulator with access to all of the information dispersed throughout society would choose a different emission tax than would a more realistic, information-constrained regulator.

I analyze a new instrument that I show can incentivize optimal carbon removal and can efficiently aggregate dispersed information. The new instrument is a dynamic deposit-refund scheme. The regulator requires that emitters pay a deposit at the time of emission and in exchange gives emitters a tradeable asset that I call a “carbon share”. In each period, the regulator refunds part of the deposit to current shareholders based on whether its measure of aggregate realized damages was as bad as implied by the deposit. The equilibrium value of the carbon share reflects expected refunds, which are by construction smaller than the value of the deposit. Emitters have an incentive to reduce emissions in order to avoid giving up the deposit for the less valuable carbon share. In later periods, a carbon shareholder may decide to remove the underlying unit of carbon from the atmosphere in order to recover the share’s deposit. And in each period, the equilibrium price of the carbon share reflects market participants’ beliefs about the regulator’s future measurements and thus about future damages from climate change.

I show that there exists a fully revealing rational expectations equilibrium in which the price of a carbon share perfectly aggregates the information dispersed throughout society and in which the incentives to reduce emissions and remove carbon are the same as in the welfare-maximizing, informationally efficient benchmark.⁶ This optimal outcome depends on the deposit being sufficiently large and on the regulator making a good-faith (albeit potentially imperfect) effort to measure and report aggregate recent damages. A large deposit is critical because the private cost of emitting carbon and the private benefit of removing carbon are both defined by traders’ expectation of the difference between the deposit and expected refunds. I call this difference the expected stream of “damage charges” that correspond to the regulator’s future measurements of aggregate damages. If the initial deposit is small, then some periods’ damage charges are likely to be constrained by the deposit and thus be smaller than the regulator’s measured damages. As the deposit becomes large, traders’ expectations of damage charges converge to their expectations of future measurements of damages and thus to their current estimates of the marginal damage from carbon emissions. Thus the private cost of emitting converges to the social cost of carbon emissions as the deposit becomes large. Numerical simulations suggest that a deposit around 2–3 times as large as the estimated social cost of emissions approximates optimal emission and removal incentives.

⁶Following the convention in finance and information economics, I use “rational expectations” to indicate that Bayesian traders use prices to learn about others’ information and trade optimally conditional on their posterior beliefs. The usage common in macroeconomics is slightly different, as that literature does not typically model asymmetric information. See Vives (2008, Chapter 3) and Campbell (2017, Chapter 12) for discussions.

I also study the Bayesian Nash equilibrium of a game in demand functions in order to assess equilibria that are “implementable” via a specific trading mechanism (see, among others, Vives, 2014; Rostek and Yoon, 2020). Each trader submits a demand schedule that accounts for their private information and for what they would infer about other traders’ information from any given carbon share price they might observe. I show that carbon share prices in implementable equilibria use traders’ private information to learn about correlated impacts in other sectors and to reduce the impact of measurement error in aggregate data. Encouragingly, an all-seeing, informationally efficient regulator would use dispersed information in the same way. Discouragingly, traders’ risk aversion, their awareness of their information leaking into market prices, and the existence of noise traders prevent the carbon share price from being fully revealing. Although implementable equilibria do not achieve full informational efficiency, they do aggregate information in ways similar to informationally efficient beliefs but not achievable by a regulator with access only to aggregate data.

This new policy shifts much of the work of projecting possible climate change damages from the regulator to markets. However, the regulator does still play a critical role in measuring realized damages from climate change (after all, climate change is still an externality). Traders use their private information about climate change damages and the observed carbon share price to forecast the regulator’s future measurements because those determine future damage charges and refunds. Whatever the regulator attempts to measure therefore determines the market incentives to reduce emissions and to remove old emissions.⁷ The regulator’s measurements should be unbiased on average, as is also critical to setting the proper emission tax. However, here the regulator’s measurements can also be arbitrarily noisy without undercutting the policy too severely, because carbon emission and removal incentives ultimately depend on the dispersed information about future measurements that is aggregated by the carbon share market.

This paper intersects with several distinct literatures. As mentioned above, a central theme throughout environmental economics is the importance of using market-based instruments to control pollution, whether in the form of emission taxes or cap-and-trade programs (see, among others, Metcalf, 2009; Stavins, 2022).⁸ Instead of focusing exclusively on the role of emission prices in determining the budget sets of firms and households, I also consider how to design markets so that emission prices convey information about damages from climate change.⁹ Where dispersed information about damages is relatively unimportant

⁷The regulator could choose to measure whatever it cares about. For instance, it could apply equity weighting and/or value nonmarket impacts.

⁸The recommendation to address climate change by taxing emissions dates to at least Nordhaus (1977), and attempts to econometrically estimate the consequences of climate change date to at least Mendelsohn et al. (1994). Weitzman (1974) shows that asymmetric information about abatement costs can break the equivalence between an emission tax and cap. I emphasize asymmetric information about the externality, which I show can make an emission tax or cap informationally inefficient.

⁹Other work focuses on the revelation of beliefs about the magnitude of climate change: Schlenker and Taylor (2021) show that weather derivatives are sensitive to climate model projections, and Hsu (2011)

and cleanup of past emissions is irrelevant (as may be true of particulate matter or lead pollution), the new policy instrument performs like an emission tax or cap-and-trade program. But where information about damages is dispersed, the new policy instrument acts like improving the information underlying an emission tax or cap-and-trade program, and where cleanup of past emissions is potentially relevant, the present policy can incentivize such cleanup without requiring the regulator to directly fund it. Climate change clearly demonstrates dispersed information about impacts and the possibility of ex post cleanup, and many other externalities will too.¹⁰

The new instrument that I analyze is a dynamic deposit-refund instrument. Static deposit-refund schemes resolve difficulties monitoring—and thus taxing—improper waste disposal (e.g., Bohm, 1981; Russell, 1987; Fullerton and Kinnaman, 1995; Torsello and Vercelli, 1998).¹¹ I posit no problem monitoring either the act of emission or the act of carbon removal, but I do resolve difficulties in incentivizing carbon removal that arise from the regulator’s imperfect ability to tax past emitters for emerging climate damages.¹² My dynamic deposit-refund scheme resolves the difficulties of taxing past emitters by imposing all costs upfront and offering rewards for subsequently claiming ownership of pollution. From the perspective of emitters, my policy combines a tax (in the form of the deposit) and a subsidy (in the value of the carbon share received), with emission incentives determined by the difference between the deposit and the value of the carbon share. This type of emission

proposes futures on emission taxes set, following McKittrick (2011), according to a predefined function of temperature (see also Aliakbari and McKittrick, 2018). Here the focus is on predicting the aggregation of local damages rather than on predicting the level of global temperature, here price discovery determines emission and carbon removal decisions rather than merely providing a signal useful for long-run investment, and the carbon share market described here is likely to have more liquidity than a prediction market or a futures market because the traded asset is a property right created by the (still commonplace) act of emitting.

¹⁰For instance, consider the externalities produced by orbital debris in space. Satellite owners could post a bond to fund an “orbital-use share” that would incentivize both optimal debris creation and optimal debris cleanup. Fees for launching satellites are the analogue of an emission tax. They fail to incentivize either active measures to avoid creating debris post-launch or cleanup of debris post-impact. Rao et al. (2020) propose orbital-use fees that are the analogue of taxing the stock of pollution, a policy option discussed in footnote 12. Such a policy is vulnerable to judgment-proofness problems induced by market churn. It also does not aggregate private information like the carbon share price does.

¹¹Deposit-refund schemes have also been understood as means to avoid the fiscal costs of subsidies and the distributional costs of taxes (Bohm, 1981). Here one of the motivations is to avoid the fiscal costs of using the public purse to directly fund carbon removal.

¹²If the regulator could tax the stock of carbon in the atmosphere (rather than the flow of emissions into the atmosphere), then the incentive to remove carbon would respond to new information as the stock tax was updated, analogous to how incentives under carbon shares respond to updated refunds. Emitters’ incentives would depend on their expectations of future stock taxes (see Appendix B), analogous to how incentives under carbon shares depend on expected future damage charges. However, the stock tax would be evaded as emitters go out of business over the many decades that carbon persists in the atmosphere, whereas carbon shares are instruments with positive payoffs that investors are willing to pay for. Stock taxes have been proposed in the context of climate change (Lemoine, 2007), mine remediation (White et al., 2012; Yang and Davis, 2018), and space orbits (Rao et al., 2020).

incentive is familiar from the static generalization of deposit-refund schemes in Fullerton and Wolverton (2000). Here, however, the level of the subsidy is not fixed by the regulator but is instead determined in equilibrium (in the form of the share prices) by private actors' information about climate change impacts.¹³

The possibility of removing enough carbon from the atmosphere to make aggregate emissions “net negative” has become a prominent part of the climate policy discourse. The Intergovernmental Panel on Climate Change projects that limiting warming to 1.5°C (2°C) would require up to 700 (250) Gt CO₂ of net negative emissions over this century (IPCC, 2022), the 2021 U.S. Infrastructure Investment and Jobs Act provided \$3.5 billion to establish carbon removal hubs, and the 2022 U.S. Inflation Reduction Act increased tax credits for capturing and storing carbon from the air from \$50 to \$180 per ton of CO₂. Despite the increasingly prominent discussion and promotion of carbon removal, I know of no work on market-based approaches to incentivizing optimal use of these technologies.¹⁴ In the absence of alternative policy instruments, many assume that governments would directly subsidize carbon removal, despite concerns about the fiscal burden such subsidies would impose (see Bednar et al., 2019; Edenhofer et al., 2021).¹⁵ I here describe a policy instrument that is revenue-positive for the government and efficiently adapts the scale of carbon removal to new information about the cost of climate change.

This paper constitutes a novel link between environmental economics and the literature on asymmetric information in financial markets. Since Grossman (1976, 1978), much literature studies financial markets' ability to efficiently aggregate dispersed information. I here study an asset tied to an externality. I consider a setting that mixes elements of large- and small-market models: there are many small traders, but they are attached to a finite number of sectors and thus have a finite number of distinct information sets. Each small trader does not account for how its own trades affect the price (i.e., is nonstrategic) but does account for how the aggregate decisions of similarly informed traders move the price (i.e., uses information rationally). The setting and the conclusion that implementable equilibria are partially revealing share features of the monopolistic competition environment in Kyle (1989), but Kyle (1989) assumes that traders' signals are independent of each other. My focus

¹³I show that the ideal deposit would equal the worst-case social cost of carbon. Others have previously proposed that fees on materials or products be set to their most harmful possible environmental fate, with fees refunded in accord with the harmfulness of actual outcomes (e.g., Solow, 1971; Mills, 1972; Bohm and Russell, 1985; Costanza and Perrings, 1990; Boyd, 2002). These informal proposals employ arguments based on ambiguity aversion, difficulties monitoring pollution, or difficulties posed by judgment-proofness.

¹⁴Conventional emission pricing policies could incentivize use of carbon dioxide removal technologies up to the point at which net emissions are zero. However, the European Union's flagship cap-and-trade program historically has not provided the credits for carbon dioxide removal that could sustain even this limited incentive (Scott and Geden, 2018; Rickels et al., 2020). Bednar et al. (2021) propose “carbon removal obligations” that would extend standard cap-and-trade schemes to allow temporarily overshooting longer-run carbon targets.

¹⁵Bednar et al. (2019) calculate that the subsidies required for carbon removal could exceed even the share of output that the U.S. spends on defense.

on correlated information is similar to studies of “fundamental value” (Vives, 2011; Rostek and Weretka, 2012; Vives, 2014), but here payoffs are common across traders because they depend on damage charges that are applied uniformly to all shareholders (the externality is common to shareholders). My environment therefore mixes payoffs that are a pure common value with an information structure reminiscent of studies of fundamental value.

Dynamics are also central to my analysis. My traders aim to predict the regulator’s future measurements, both because they directly determine a carbon share’s refunds and because they determine capital gains as future traders use those signals to update their own beliefs about subsequent refunds. Much of the literature on asymmetric information in financial markets analyzes static models. I analytically solve a dynamic equilibrium model of information aggregation by assuming that each generation of traders has access to all public information from previous periods (i.e., to all previous prices and regulator reports) but not to any private information from previous periods. This assumption is similar to the critical assumption in Vayanos (1999, 2001) because traders begin each period in a symmetric informational position.¹⁶

The next section describes the economic and informational environment. Section 3 derives outcomes in the informationally efficient, welfare-maximizing benchmark. Section 4 analyzes emission taxes, focusing on the constraints they face in using new information and in observing all available information. Section 5 formally defines carbon shares and establishes the conditions under which this policy resolves the issues demonstrated with emission taxes. The final section concludes. The appendix contains numerical details, analysis of stock taxes, and proofs.

2 Setting

2.1 Production, Consumption, and Emissions

Let there be a unit mass of households and $N > 1$ intermediate-good sectors, each of which is perfectly competitive. Output from sector i in period t is

$$Y_{it} = \exp[-\zeta_{it}T_t] L_{it} Y^{it}(e_{it}).$$

$L_{it} \in [0, 1]$ is labor offered by households to sector i in exchange for wage w_{it} . The representative firm in sector i has gross production function $Y^{it}(e_{it})$, with $e_{it} \geq 0$ indicating emissions. $Y^{it}(\cdot)$ is strictly increasing and concave. Temperature T_t imposes damages ζ_{it} in sector i at

¹⁶Other dynamic models assume that traders are risk-neutral (e.g., Kyle, 1985) or myopic (e.g., Singleton, 1987), or they assume ordered information structures (e.g., Wang, 1993) that do not make sense in the present paper’s multisector economy. In Vives (1995), the asset is liquidated after a finite number of periods, at which time its true payoff is revealed. Here carbon shares can endure forever and the true nature of climate impacts is never revealed.

time t .¹⁷ The multiplicative effect of climate damages follows the DICE model (Nordhaus, 1992, 2013), among others, and the exponential form for damages follows Golosov et al. (2014) and Lemoine (2021a), among others. Firms can condition their emission decisions on ζ_{it} .

The representative firm in sector i can fund the removal of quantity $R_{it} \geq 0$ of emissions from the atmosphere. It purchases this emission removal from a competitive industry whose costs $c(R_t)$ as a share of total output (see below) depend on aggregate removal $R_t \triangleq \sum_{i=1}^N R_{it}$, with $c(\cdot)$ strictly positive, strictly increasing, and strictly convex.¹⁸

Cumulative emissions up to time t are $M_t = M_0 + \sum_{s=0}^{t-1} [\sum_{i=1}^N e_{is} - R_s]$, with pre-policy cumulative emissions $M_0 \geq 0$ given. Time t warming is $T_t = \alpha M_t$, with $\alpha > 0$. This representation recognizes that carbon dioxide is a globally mixed pollutant and follows recent scientific findings that global temperature is approximately a linear function of cumulative emissions (see Dietz and Venmans, 2019, among others). Firms are small, so they ignore the effects of their own emissions on temperature.

Total output is Cobb-Douglas:

$$Y_t = \prod_{i=1}^N (Y_{it})^{\kappa_i},$$

with each $\kappa_i > 0$ and $\sum_{i=1}^N \kappa_i = 1$. Aggregate consumption C_t is no greater than net output:

$$C_t \leq (1 - c_t(R_t))Y_t.$$

The representative household has logarithmic utility:

$$u(C_t) = \ln(C_t).$$

Time t welfare is the present value of expected utility:

$$\sum_{s=t}^{\infty} \frac{1}{(1+r)^{s-t}} E_t[u(C_s)],$$

with per-period discount rate r and with the information set defined in each application below.

In equilibrium, firms maximize the expected present value of profits subject to prices, households maximize utility subject to budget constraints, and all markets clear. Time 0 consumption is the numeraire.

¹⁷The index i could equivalently be interpreted as indicating either regions or sector-region pairs.

¹⁸Convexity in removal costs reflects both the cost of removing carbon from the atmosphere and the potential scarcity of sites for storing carbon after removal.

2.2 Informational Environment

I now describe the informational environment. All agents are Bayesian.

Agents affiliated with sector i observe $\zeta_{it} + \lambda_{it}$, where $\zeta_{it} = \zeta_i + \bar{\zeta}_i + \epsilon_{it}$. The ζ_i are unknown and unobserved. The $\bar{\zeta}_i$ are public knowledge and represent prior expected damages in each sector. Assume only that $\sum_{k=1}^N \kappa_k \bar{\zeta}_k < (r/\alpha) \kappa_i Y^{i0'}(0)/Y^{i0}(0)$ for some $i \in \{1, \dots, N\}$, which will ensure that welfare-maximizing aggregate emissions are strictly positive in the initial period. The ϵ_{it} and λ_{it} are random variables that are each normally distributed, mean-zero, unobserved, and uncorrelated either across sectors or over time. The variance of each ϵ_{it} is $\sigma^2 > 0$, and the variance of each λ_{it} is $\omega^2 \geq 0$. The ϵ_{it} represent random exposure to global temperature. That randomness could result from randomness in the mapping from global temperature to temperatures in locations relevant to sector i and/or from randomness in sector i 's exposure to its locations' temperatures. The λ_{it} represent agents' potentially imperfect ability to measure the effect of temperature on sectoral production.

The regulator and firms have a common jointly normal prior over the ζ_i at time 0. Each ζ_i has a prior mean of zero and has prior variance $\tau_0^2 > 0$.¹⁹ The correlation between any pair ζ_i and ζ_j (for $i \neq j$) is $\Gamma \in [0, 1]$, a known parameter. This correlation determines how signals of damages in one sector provide information about damages in another sector. If $\Gamma = 0$, then the unknown component of damages is independent across sectors. If $\Gamma > 0$, then the unknown component of damages has a common component across sectors, as when impacts in one sector affect other sectors or as when vulnerability to weather is correlated across sectors.

The regulator does not observe sectoral production or input choices. Instead, at the end of time t , the regulator uses observed total output to measure $\tilde{\zeta}_t + \tilde{\lambda}_t$, where $\tilde{\zeta}_t \triangleq \sum_{i=1}^N \kappa_i \zeta_{it}$.²⁰ The $\tilde{\lambda}_t$ are random variables that are normally distributed, mean-zero, and serially uncorrelated. Their variance is $\tilde{\omega}^2 \geq 0$. They reflect the possibility of measurement error in aggregate data and of additional imprecision due to having to estimate $\tilde{\zeta}_t$ from aggregate data. The regulator shares the measured $\tilde{\zeta}_t + \tilde{\lambda}_t$ with all agents in the economy.

The timing within a period t is that intermediate-good firms make emission decisions, markets clear based on realized production, agents observe $\zeta_{it} + \lambda_{it}$, and finally the regulator observes $\tilde{\zeta}_t + \tilde{\lambda}_t$.

¹⁹ Assuming a prior mean of zero is not restrictive, as nonzero means are absorbed into the $\bar{\zeta}_i$.

²⁰ Firms' equilibrium production choices are independent of T_t (see (A-12) through (A-14)), so the regulator can estimate $\tilde{\zeta}_t$ from a time series of Y_t and T_t .

3 Informationally Efficient, Welfare-Maximizing Benchmark

Begin by considering welfare-maximizing emissions and carbon removal. Define \hat{E}_t as the expectation operator based on all information available up to time t , $\hat{\mu}_t$ and $\hat{\Omega}_t$ as the vector of posterior means and the posterior covariance matrix for the ζ_i based on information up to time t , and $\hat{\mu}_t$ and $\hat{\Omega}_t$ as the posterior mean and variance of $\sum_{i=1}^N \kappa_i \zeta_i$ based on information up to time t .²¹

Welfare-maximizing outcomes solve the following Bellman equation:

$$\hat{W}(T_t, \hat{\mu}_t, \hat{\Omega}_t) = \max_{\mathbf{L}_t, \mathbf{e}_t, R_t \geq 0} \hat{E}_t \left[u(C_t) + \frac{1}{1+r} \hat{W}(T_{t+1}, \hat{\mu}_{t+1}, \hat{\Omega}_{t+1}) \right],$$

where \mathbf{L}_t and \mathbf{e}_t indicate vectors of labor and emissions in each sector. Taking first-order conditions and then recursively substituting from the envelope theorem yields the following conditions that must hold for all i :

$$\frac{\kappa_i Y^{it'}(e_{it})}{Y^{it}(e_{it})} \begin{cases} = \frac{1}{r} \alpha \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] & \text{if } e_{it} > 0 \\ \leq \frac{1}{r} \alpha \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] & \text{if } e_{it} = 0 \end{cases}, \quad (1)$$

$$\frac{c'_t(R_t)}{1 - c_t(R_t)} \begin{cases} = \frac{1}{r} \alpha \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] & \text{if } R_t > 0 \\ \geq \frac{1}{r} \alpha \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] & \text{if } R_t = 0 \end{cases}. \quad (2)$$

On the right-hand side, the terms in brackets yield per-period expected damages per unit of warming, the α converts to units of emissions, and the $1/r$ converts to present value. The first condition equates the marginal benefit of emissions to the marginal social cost of emissions when emissions are strictly positive. If $Y^{it'}(0)$ is sufficiently small, then $e_{it} = 0$. The second condition equates the marginal cost of carbon removal to the marginal social cost of emissions (i.e., the marginal benefit of carbon removal) when carbon removal is strictly positive. If $c'_t(0)$ is sufficiently large, then $R_t = 0$. As $\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \hat{\mu}_t$ increases, e_{it} falls and R_t either increases or remains zero. Negative emissions, in which $R_t > \sum_{i=1}^N e_{it}$, become optimal when $\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \hat{\mu}_t$ is sufficiently large.

Now consider beliefs. As in other literature, the benchmark of informational efficiency updates beliefs from the available signals $\zeta_{it} + \lambda_{it}$ and $\tilde{\zeta}_t + \tilde{\lambda}_t$.

²¹Throughout, I use a hat ($\hat{\cdot}$) to indicate outcomes under the informationally efficient welfare-maximizing benchmark, a tilde ($\tilde{\cdot}$) to indicate outcomes under an emission tax policy, and a breve ($\breve{\cdot}$) to indicate outcomes under the carbon share policy.

Proposition 1 (Informationally Efficient Beliefs). *There exists $\hat{Z}_t \in [0, 1)$ such that $\hat{Z}_t \rightarrow 0$ as $\tilde{\omega}^2/\omega^2 \rightarrow \infty$ and*

$$\begin{aligned} \hat{\mu}_t = & \hat{Z}_t \left[\frac{1}{t} \sum_{j=0}^{t-1} [\tilde{\zeta}_j + \tilde{\lambda}_j] - \sum_{k=1}^N \kappa_k \bar{\zeta}_k \right] \\ & + \frac{(1 - \hat{Z}_t)(1 - \Gamma)\tau_0^2 - \hat{Z}_t\sigma^2/t}{(1 - \Gamma)\tau_0^2 + \sigma^2/t + \omega^2/t} \sum_{k=1}^N \kappa_k \left[\frac{1}{t} \sum_{j=0}^{t-1} [\zeta_{kj} + \lambda_{kj}] - \bar{\zeta}_k \right] \\ & + \frac{\sigma^2/t + (1 - \hat{Z}_t)\omega^2/t}{(1 - \Gamma)\tau_0^2 + \sigma^2/t + \omega^2/t} \frac{N\Gamma\tau_0^2}{(1 - \Gamma)\tau_0^2 + \sigma^2/t + \omega^2/t + N\Gamma\tau_0^2} \frac{1}{N} \sum_{k=1}^N \left[\frac{1}{t} \sum_{j=0}^{t-1} [\zeta_{kj} + \lambda_{kj}] - \bar{\zeta}_k \right]. \end{aligned} \quad (3)$$

Proof. Apply the projection theorem to a random vector formed from $\sum_{k=1}^N \kappa_k \zeta_k$, the N sectoral signals, and the aggregate signal. See Appendix C. \square

Corollary 1 (Special Cases for Informationally Efficient Beliefs).

i If $\omega^2 = 0$, then

$$\begin{aligned} \hat{\mu}_t = & \frac{(1 - \Gamma)\tau_0^2}{(1 - \Gamma)\tau_0^2 + \sigma^2/t} \overbrace{\sum_{k=1}^N \kappa_k \left[\frac{1}{t} \sum_{j=0}^{t-1} \zeta_{kj} - \bar{\zeta}_k \right]}^{=\lim_{\tilde{\omega}^2 \rightarrow 0} \frac{1}{t} \sum_{j=0}^{t-1} [\tilde{\zeta}_j + \tilde{\lambda}_j] - \sum_{k=1}^N \kappa_k \bar{\zeta}_k} \\ & + \frac{\sigma^2/t}{(1 - \Gamma)\tau_0^2 + \sigma^2/t} \frac{N\Gamma\tau_0^2}{(1 - \Gamma)\tau_0^2 + \sigma^2/t + N\Gamma\tau_0^2} \frac{1}{N} \sum_{k=1}^N \left[\frac{1}{t} \sum_{j=0}^{t-1} \zeta_{kj} - \bar{\zeta}_k \right]. \end{aligned}$$

ii If $\tilde{\omega}^2 = 0$ and $\Gamma = 0$, then

$$\hat{\mu}_t = \frac{\tau_0^2}{\tau_0^2 + \sigma^2/t} \left[\frac{1}{t} \sum_{j=0}^{t-1} \tilde{\zeta}_j - \sum_{k=1}^N \kappa_k \bar{\zeta}_k \right].$$

iii If $\tilde{\omega}^2 = 0$ and each $\kappa_i = 1/N$, then

$$\hat{\mu}_t = \frac{(1 - \Gamma)\tau_0^2 + N\Gamma\tau_0^2}{(1 - \Gamma)\tau_0^2 + \sigma^2/t + N\Gamma\tau_0^2} \left[\frac{1}{t} \sum_{j=0}^{t-1} \tilde{\zeta}_j - \sum_{k=1}^N \kappa_k \bar{\zeta}_k \right].$$

Proof. See Appendix D. \square

The informationally efficient benchmark aggregates information from the private signals and the public signals. The first part of the corollary describes mean beliefs when dispersed agents do not suffer measurement error ($\omega^2 = 0$). In this case, the informationally efficient benchmark has no use for the aggregate signals $\tilde{\zeta}_j + \tilde{\lambda}_j$. Instead, it weights the sectoral signals as if a perfectly measured version of the aggregate signal were available (first line) and, when $\Gamma > 0$, it also uses the disentangled (i.e., unweighted) sectoral signals directly because signals in relatively unimportant sectors with small κ_i provide information about damages in all other sectors (second line). In fact, if $\Gamma = 1$, informationally efficient beliefs do not weight the sectoral signals by the κ_i at all (the first line vanishes) because each sector's signal provides the same information about aggregate damages as any other sector's signal, whether or not a sector receives only a small weight in aggregate output.

The second and third parts of Corollary 1 describe mean beliefs when there is no measurement error in the aggregate signal ($\tilde{\omega}^2 = 0$). The second part shows that the informationally efficient benchmark has no use for the disaggregated sectoral signals $\zeta_{kj} + \lambda_{kj}$ when sectoral signals are independent ($\Gamma = 0$). And the third part shows that the informationally efficient benchmark has no use for the disaggregated sectoral signals $\zeta_{kj} + \lambda_{kj}$ when sectors have identical value shares of output and thus receive identical weights in the aggregate signal (i.e., each $\kappa_i = 1/N$). In either case, mean beliefs are what one would expect from simple applications of the familiar univariate normal-normal updating rule to the aggregate signal $\tilde{\zeta}_j + \tilde{\lambda}_j$.

Proposition 1 shows that informationally efficient beliefs in general use both types of information. The aggregate signal (first line in (3)) provides information that can mitigate the consequences of measurement error in sectoral signals, and the sectoral signals (second line) provide information used to construct an alternate version of the aggregate signal that mitigates the consequences of measurement error in the aggregate signal. Efficient updating also leverages correlation across sectoral effects (third line) to learn from sectors whose small κ_i mean they do not directly matter much for aggregate outcomes.

4 Regulation by Emission Taxes

Now consider a regulator who maximizes welfare by taxing firms' period t net emissions at rate ν_t . Firms can avoid the tax either by reducing emissions or by simultaneously contracting for removal to offset ongoing emissions. The regulator returns any tax revenue to households as lump-sum transfers. I initially assume that the regulator's tax revenue must be weakly positive in each period before weakening that assumption in Section 4.1.

The regulator sets the time t tax at the beginning of the period so as to maximize welfare conditional on its time t beliefs and subject to market equilibrium. The regulator's chosen time t tax is therefore a function of the aggregate measurements from times 0 through $t - 1$. Denote the regulator's mean belief about $\sum_{k=1}^N \kappa_k \zeta_k$ at the time t information set as $\tilde{\mu}_t$.

The following proposition gives the optimal emission tax:

Proposition 2 (Emission Tax). *There exists $\bar{\nu}_t > 0$ such that $\sum_{i=1}^N e_{it} - R_t = 0$ if and only if $\nu_t \geq \bar{\nu}_t$. The regulator maximizes welfare with a tax of*

$$\nu_t = \min \left\{ \bar{\nu}_t, C_0 \frac{\alpha}{r} \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \tilde{\mu}_t \right] \right\}.$$

Proof. See Appendix E. □

When $\nu_t < \bar{\nu}_t$, the tax is determined by the present value of expected aggregate damages.²² Using that tax, firms' first-order conditions (A-13) and (A-14) become

$$\begin{aligned} \frac{\kappa_i Y^{it'}(e_{it})}{Y^{it}(e_{it})} &= \frac{\alpha}{r} \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \tilde{\mu}_t \right], \\ \frac{c'_t(R_t)}{1 - c_t(R_t)} &= \frac{\alpha}{r} \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \tilde{\mu}_t \right]. \end{aligned}$$

Once we account for the possibility of corner solutions, these are the conditions for welfare-maximization given in (1) and (2) as long as $\tilde{\mu}_t = \hat{\mu}_t$.

We see two possible reasons why a tax may not attain the informationally efficient welfare-maximizing benchmark. First, it could be that $C_0 \frac{\alpha}{r} \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] > \bar{\nu}_t$. In this case, the benchmark from Section 3 would have negative net emissions whereas the regulator's feasible equilibrium has zero net emissions. I refer to this possibility as inducing a loss due to inefficiency in using information, as it indicates a failure to implement the information collected from the economy. Second, it could be that $\tilde{\mu}_t \neq \hat{\mu}_t$. In this case, the benchmark from Section 3 would be based on beliefs that are different from the regulator's beliefs. I refer to this possibility as inducing a loss from observing information, as it indicates a failure to collect all of the information in the economy. I explore each in turn.

4.1 Loss Due to Inefficiency in Using Information

Temporarily assume that the regulator observes all signals $\zeta_{it} + \lambda_{it}$, as when firms truthfully communicate to the regulator or the regulator collects the same data as firms. In this case, $\tilde{\mu}_t = \hat{\mu}_t$, so the regulator can set the informationally efficient tax and we have no loss from inefficiency in observing information.

²²The combination of logarithmic utility and the damage specification means that uncertainty is not priced directly, as in Golosov et al. (2014). For a more general constant relative risk aversion utility function, the optimal tax would be sensitive to future consumption and would include a risk premium (see Lemoine, 2021a).

It is immediately obvious from the foregoing analysis that the policy described in Proposition 2 imposes losses relative to the welfare-maximizing benchmark when there is some chance that the constraint $\bar{\nu}_t$ binds (i.e., that negative net emissions become optimal). That chance is driven by the possibility of observing information that makes $\tilde{\mu}_t$ large and by the possibility that technological progress in carbon removal makes R_t large for any given ν_t (i.e., makes $\bar{\nu}_t$ small).

One might object that the modeled revenue constraint is too stringent. A forward-thinking regulator could save the revenue collected from emission taxes and dedicate it to funding carbon removal.²³ In effect, such a policy establishes a lockbox for emission tax revenue that allows the regulator to procure some level of negative net emissions without needing to raise money from taxpayers. It changes the revenue constraint from a static one that must hold in each period to a dynamic one that must hold across periods.

Consider the implications of a dynamic revenue constraint in a world in which the regulator does not learn about damages but in which technological progress in carbon removal can eventually make negative net emissions optimal even under the prior belief. Proposition 2 establishes that the regulator's unconstrained-optimal tax would be

$$\nu_t = C_0 \frac{\alpha}{r} \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \tilde{\mu}_t \right].$$

This tax is unaffected by the possibility of technological progress in carbon removal technologies and is constant over time in the absence of learning about damages (i.e., when $\tilde{\mu}_t$ is constant over time). Because this tax is also the subsidy a regulator would like to offer for carbon removal in a negative net emission scenario, the tax that the regulator collects at the time of emission is exactly equal to the subsidy the regulator would subsequently offer to remove that same unit of emission from the atmosphere. The dynamic revenue constraint would therefore never bind unless it became optimal to bring future carbon stocks below their initial level M_0 .

Now let the regulator learn about damages. In that case, observing unfavorable information about climate damages could increase $\tilde{\mu}_t$ by enough to make negative net emissions optimal whether or not there is progress in carbon removal technology. The tax the regulator collects at the time of emission is then strictly less than the subsidy the regulator would subsequently offer to remove that same unit of emission from the atmosphere. The regulator has sufficient tax revenue in its lockbox to bring carbon only part of the way back to M_0 . The more pessimistic damage estimates become, the more likely this constraint on the regulator's ability to fund negative net emissions becomes binding. And if carbon removal technology simultaneously progresses quickly, then this constraint becomes even more likely to bind. The regulator might again be unable to procure the optimal level of negative net emissions without raising funds from taxpayers.

²³I thank Jim Stock for this suggestion.

Further, when the lockbox might fail to incentivize optimal removal, the regulator distorts emission decisions in anticipation that the dynamic revenue constraint might bind:

Corollary 2 (Emission Tax With A Lockbox).

1. *Consider a period in which net emissions are strictly positive. The tax with a lockbox is strictly greater (strictly less) and emissions are strictly less (strictly greater) than given in Proposition 2 if marginally raising that tax increases (decreases) tax revenue.*
2. *Consider a period in which net emissions are weakly (strictly) negative. The optimal price that the regulator offers for carbon removal is weakly (strictly) less, and net emissions are weakly (strictly) greater, than the welfare-maximizing benchmark.*

Proof. See Appendix F. □

The marginal value of a higher emission tax is comprised of its marginal value in a setting without revenue constraints and its effect on future negative emission constraints via its effect on the revenue that will be stored in the lockbox. The latter effect distorts the emission tax away from its unconstrained-optimal level in order to prepare for the possibility that sufficiently negative net emissions become optimal. When emissions are strictly positive, raising an emission tax increases revenue by charging more per unit of net emissions but reduces revenue by reducing net emissions. The first part of the corollary establishes that the optimal tax in the presence of a lockbox is higher (lower) when the former (latter) dominates. Small distortions in the emission tax do not impose first-order costs today, but they can raise extra revenue that provides first-order benefits by weakening a potential future constraint on negative net emissions.²⁴

Once the regulator is already paying for negative emissions out of the lockbox, reducing the emission tax clearly leaves more revenue in the lockbox: a lower price requires the regulator to pay less per unit of carbon removal and also procures less carbon removal. The second part of the corollary establishes that a regulator already paying for negative emissions prepares for the possibility of future binding constraints by reducing the price offered for carbon removal.

In summary, we have learned that:

1. A regulator who must obey a period-by-period revenue constraint cannot procure negative emissions using an emission tax. An emission tax therefore cannot attain first-best should negative emissions eventually become optimal, even if the regulator has perfect information.

²⁴The possibility of hitting the constraint at some future time did not distort the optimal tax in Proposition 2 because the combination of logarithmic utility and multiplicative-exponential damages makes the optimal tax independent of future emission and removal trajectories. This combination of assumptions also makes the lockbox work perfectly in the absence of learning about damages.

2. A regulator who stores emission tax revenue in a lockbox for funding future carbon removal can fund the optimal level of carbon removal if the regulator does not learn about damages over time. In the current setting, an emission tax-plus-lockbox therefore can attain first-best in the presence of technological progress in carbon removal.
3. A regulator who stores emission tax revenue in a lockbox for funding future carbon removal might not be able to fund the optimal level of carbon removal if the regulator learns about damages over time. An emission tax-plus-lockbox therefore may not attain first-best in the presence of new information about climate damages. And the optimal use of the lockbox distorts emissions and removal decisions in all other periods so as to increase funds in the lockbox.

That final point shows that information is critical to inefficiencies in procuring negative emissions: an emission tax may not be able to optimally use new information about the social cost of emissions, should that new information be sufficiently pessimistic and the regulator not be able to costlessly offer arbitrarily large subsidies from taxpayer funds.

4.2 Loss Due to Inefficiency in Observing Information

Now allow asymmetric information but assume that carbon removal is infeasible, as when $c'_t(0)$ is so large for all t that carbon removal is irrelevant under plausible beliefs. We have no loss from inefficiency in using information and instead analyze a loss from inefficiencies in observing information.

The following result describes the regulator's time t posterior estimate of $\sum_{k=1}^N \kappa_k \zeta_k$ formed from observing the aggregate signals $\tilde{\zeta}_j + \tilde{\lambda}_j$.

Proposition 3 (Regulator's Beliefs).

$$\tilde{\mu}_t = \frac{(1 - \Gamma)\tau_0^2 \sum_{i=1}^N \kappa_i^2 + \Gamma\tau_0^2}{(1 - \Gamma)\tau_0^2 \sum_{i=1}^N \kappa_i^2 + \Gamma\tau_0^2 + \frac{1}{t}[\tilde{\omega}^2 + \sigma^2 \sum_{i=1}^N \kappa_i^2]} \frac{1}{t} \sum_{j=0}^{t-1} \left(\tilde{\zeta}_j + \tilde{\lambda}_j - \sum_{k=1}^N \kappa_k \tilde{\zeta}_k \right). \quad (4)$$

Proof. Follows from application of the conventional univariate normal-normal Bayesian updating formula, observing that the prior variance is $\tau_0^2 \sum_{i=1}^N \kappa_i^2 + 2\Gamma\tau_0^2 \sum_{i=1}^N \sum_{k=i+1}^N \kappa_i \kappa_k$ and using $\sum_{i=1}^N \kappa_i = 1$. \square

The weight placed on the aggregate measurement in (4) increases in Γ : positive correlation increases the regulator's prior uncertainty about aggregate damages and thus mechanically increases the weight placed on the aggregate measurement. Posterior beliefs are exactly the same as those formed by a counterfactual regulator in a world with variance $\tilde{\tau}_0^2 = (1 - \Gamma)\tau_0^2 + \Gamma\tau_0^2 / [\sum_{i=1}^N \kappa_i^2]$ and correlation $\tilde{\Gamma} = 0$. In contrast, Proposition 1 showed that positive correlation among the unknown sector-specific effects ζ_i increases the weight that

informationally efficient beliefs place on the disentangled sectoral measurements, because such beliefs recognize that each sector's measurement contains information about all other sectors.

The following corollary delineates conditions under which the regulator's beliefs are informationally efficient.

Corollary 3 (Informationally Efficient Regulator). *For $t > 0$, $\tilde{\mu}_t = \hat{\mu}_t$ with probability 1 if (i) $\tilde{\omega}^2 = 0$ and either (iia) $\Gamma = 0$ or (iib) each $\kappa_i = 1/N$.*

Proof. Follows straightforwardly from Corollary 1 and Proposition 3. □

The regulator's beliefs are informationally efficient if there is no measurement error at the aggregate level ($\tilde{\omega}^2 = 0$) and either there is no correlation among sectoral effects ($\Gamma = 0$) or sectors have identical weights in production (each $\kappa_i = 1/N$) that render correlation unimportant for learning. Otherwise Proposition 1 showed that informationally beliefs generally use disaggregated signals unavailable to the regulator. Because all three of the conditions in Corollary 3 are plausibly violated in reality, an actual regulator's beliefs are likely to be informationally inefficient.

Figure 1 provides an example that illustrates how the regulator's beliefs (gray) differ from informationally efficient beliefs (black), for the extreme cases when unknown sectoral effects are independent of each other ($\Gamma = 0$, dashed) and are perfectly correlated with each other ($\Gamma = 1$, solid). Based on the calibration described in Appendix A, the emission tax that would be optimal at time 0 beliefs is \$118 per tCO₂. The depicted simulations assume that the initial tax that would be optimal with perfect information about the ζ_i would be twice as large. Each curve averages over 1 million trajectories for $\hat{\mu}_t$ and $\tilde{\mu}_t$.

The left panel assesses how beliefs converge to the truth.²⁵ In these cases, informationally efficient beliefs converge to the truth faster than do the regulator's beliefs. Both types of beliefs converge faster when sectors are perfectly correlated with each other than when sectors are independent of each other: informationally efficient beliefs infer more from each observation in the presence of correlation, and the regulator places more weight on the data when correlation increases the prior variance.

The middle and right panels plot the emission tax chosen, on average, after observing signals from time 0. The middle panel shows that the number of sectors N does not affect beliefs when sectoral effects are independent (the calibration scales σ^2 so that aggregate stochasticity is independent of N). However, correlation makes the average speed of learning increase in the number of sectors N , and informationally efficient beliefs in particular update much faster when the economy has multiple sectors that provide information about each

²⁵The average speed of convergence is used here for illustration, but it is not a measure of the quality of beliefs. Such a measure would also account for the standard deviation of beliefs. For instance, when $\Gamma = 0.5$, the regulator's beliefs converge towards the true value on average faster than do informationally efficient beliefs, but they are more sensitive to randomness in any particular trajectory.




Figure 1: An example of how informationally efficient beliefs and the regulator’s beliefs would each evolve on average (i.e., of $E_0[\hat{\mu}_t|\zeta]$ and $E_0[\tilde{\mu}_t|\zeta]$). Appendix A details the calibration. The emission tax optimal at prior beliefs is \$118 per tCO₂, and simulations assume that the emission tax conditional on true knowledge of the ζ_i would be twice as large (\$236 per tCO₂). Left: Evolution of average beliefs over the first 25 periods, with $N = 10$ and $\tilde{\omega}^2/[\sigma^2 \sum_{i=1}^N \kappa_i^2] = 2$. Middle: The effect of the number of sectors (N) on period 1 beliefs. Right: The effect of aggregate measurement error ($\tilde{\omega}^2$) on period 1 beliefs.

other. The right panel shows that aggregate measurement error $\tilde{\omega}^2$ (which increases to the right) markedly slows learning by the regulator. In contrast, informationally efficient beliefs are less sensitive to aggregate measurement error because they can use the disentangled sectoral signals directly and thereby mitigate that source of error.²⁶

Whereas the drawback in Section 4.1 was inefficiency in using the available information should that information warrant negative emissions, the drawback here is the inefficiency in observing all available information. The best emission tax that a regulator can implement will generally differ from the emission tax that the regulator would choose based on all of the information in the economy.

5 A Policy Framework that Dominates Conventional Emission Pricing

We have seen that conventional emission pricing does not perform ideally at either collecting or using information about the social cost of greenhouse gas emissions. As a result, there is space for a policy to do better than conventional market-based instruments. I now describe such a policy, in the form of a novel dynamic deposit-refund instrument.

²⁶It is hard to detect visually in the figure, but aggregate measurement error does slow learning even for informationally efficient beliefs with $\Gamma = 1$. The convergence of informationally efficient beliefs and the regulator’s beliefs when $\Gamma = 0$ and $\tilde{\omega}^2 = 0$ illustrates Corollary 3.

This new type of policy requires each emitter to post a deposit $D \geq 0$ per unit of emissions. We can express D as

$$D \triangleq \frac{1}{r} C_0 \alpha \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \bar{\mu} \right]. \quad (5)$$

Choosing the deposit is equivalent to choosing a parameter $\bar{\mu} \geq -\sum_{i=1}^N \kappa_i \bar{\zeta}_k$ that defines implied per-period climate damages.²⁷ In exchange for the deposit, the emitter receives a transferable asset that is attached to the unit of carbon emitted. I refer to the asset as a “carbon share” because it reflects a claim on a part of the carbon in the atmosphere.

At the end of each period, the policymaker applies a damage charge Δ_t to each outstanding carbon share. This charge is set equal to the lesser of the period t measured marginal damage from carbon emissions and the per-period damages implied by the deposit:

$$\Delta_t \triangleq C_0 \alpha \min \left\{ \tilde{\zeta}_t + \tilde{\lambda}_t, \sum_{k=1}^N \kappa_k \bar{\zeta}_k + \bar{\mu} \right\}. \quad (6)$$

The damage charge is returned lump sum to consumers.²⁸ The policymaker refunds to carbon shareholders the difference between the damage charge and the per-period damages implied by the deposit:

$$\begin{aligned} d_t &\triangleq r D - \Delta_t \\ &= C_0 \alpha \max \left\{ 0, \sum_{k=1}^N \kappa_k \bar{\zeta}_k + \bar{\mu} - (\tilde{\zeta}_t + \tilde{\lambda}_t) \right\}. \end{aligned} \quad (7)$$

The refunds d_t are weakly positive. No refund is paid in the period of emission. The deposit acts like principal, some of which is returned to agents over time in the form of refunds and some of which is reclaimed by the regulator in the form of damage charges. Over the lifetime of a carbon share, the present value of total refunds and damage charges recovers the deposit:

$$\begin{aligned} \sum_{s=1}^{\infty} \frac{1}{(1+r)^s} [d_{t+s} + \Delta_{t+s}] &= \sum_{s=1}^{\infty} \frac{1}{(1+r)^s} r D \\ &= D. \end{aligned}$$

²⁷The deposit would equal the emission tax from Proposition 2 if $\bar{\mu} = \tilde{\mu}_t$.

²⁸In a second-best setting, revenue from damage charges could be used to offset revenue from distortionary taxes. A full analysis of such a setting should consider how to adapt both the damage charges and the deposit (see Fullerton and Wolverton, 2000, 2005).

	Time t	Time $t+1$	Time $t+2$...	Time $t+s$
Emitter:	Pays D , Receives share worth q_t	Sells share for q_{t+1}			
Shareholder:		Buys share for q_{t+1} , Receives d_{t+1}	Receives d_{t+2}	...	Pays p_{t+s}^R , Receives $(1+r)D - \Delta_{t+s}$

Figure 2: Example of the life of a carbon share. Here the share is attached to a unit of time t emissions, the emitter decides to sell the share at time $t + 1$, and the new shareholder decides to remove the underlying unit of carbon from the atmosphere at time $t + s$.

In each period subsequent to emission, a carbon share's owner decides whether to leave its attached unit of carbon in the atmosphere. If the owner removes the carbon from the atmosphere in time t , they receive $(1+r)D - \Delta_t$ and the share is retired; otherwise they receive refund d_t and can keep or sell the share. The carbon share is therefore an option to claim the deposit by spending on carbon removal. The shareholder receives the refunds d_t whether exercising or holding the option, but the shareholder loses the charges Δ_t as long as the option is unexercised. Shares are clearly valuable, because the worst they do is pay zero refunds. If the owner of a carbon share were to declare bankruptcy or otherwise liquidate, its creditors would want the carbon share so they could receive its refunds and have the option to eventually reclaim the full deposit.

Figure 2 provides an example of payoffs over time under the carbon share policy. At time t , an emitter posts the deposit D and in return receives a carbon share whose market value is q_t (to be analyzed below). At time $t + 1$, the emitter in this example decides to sell the share to a third party for the market price q_{t+1} . That third party claims the time $t + 1$ refunds d_{t+1} and continues to do so until either selling the share or removing the underlying unit of carbon from the atmosphere. At time $t + s$, the third party in this example does decide to remove the underlying unit of carbon from the atmosphere, which costs p_{t+s}^R . At that point, the regulator retires the carbon share and pays the third party $(1+r)D - \Delta_{t+s}$.

I assume all agents discount at rate r . Use \tilde{W} to denote welfare along a realized trajectory under the carbon share policy defined above and use \hat{W} to denote welfare along a realized

trajectory under the welfare-maximizing, informationally efficient benchmark:

$$\begin{aligned}\check{W} &= \sum_{s=0}^{\infty} \frac{1}{(1+r)^s} u(\check{C}_s), \\ \hat{W} &= \sum_{s=0}^{\infty} \frac{1}{(1+r)^s} u(\hat{C}_s).\end{aligned}$$

The full-information expected loss from using the carbon share policy is:

$$\check{L} = E_0 \left[\hat{W} - \check{W} \mid \boldsymbol{\zeta} \right],$$

where $\boldsymbol{\zeta}$ is a vector of the ζ_i .

5.1 Improved Efficiency in Using Information

Temporarily assume that all actors in the economy see all of the $\zeta_{it} + \lambda_{it}$ at each time t . We saw in Section 4.1 that an emission tax may fail to use information efficiently when information justifies negative emissions. I will show that carbon shares can improve outcomes if the deposit D is sufficiently large.

Define \hat{q}_t as the carbon share's value in period t prior to observing the $\zeta_{it} + \lambda_{it}$, $\tilde{\zeta}_t + \tilde{\lambda}_t$, Δ_t , or d_t , where the hat notation reflects that the carbon share price in this section uses informationally efficient beliefs (as opposed to the share price to be defined in Section 5.2). The following lemma establishes the equilibrium value of the carbon share:

Lemma 1 (Carbon Share Value).

$$\hat{q}_t = \sum_{j=0}^{\infty} \frac{1}{(1+r)^j} \hat{E}[d_{t+j}]. \quad (8)$$

Proof. See Appendix G. □

The equilibrium value of the carbon share is the expected present value of the refunds that it claims. The value of holding a carbon share therefore derives from the possibility that damages will not be as bad as implied by $\bar{\mu}$. At the time of emission, the firm's net outlays per unit of non-abated emissions are $D - (\hat{q}_t - \hat{E}[d_t]) \in [0, D]$.²⁹

The following assumption ensures that it would never be optimal to remove enough carbon to bring atmospheric carbon and temperature below their initial levels:

²⁹If future damages were guaranteed to be zero in every period, then the present value of the stream of refunds at the time of emission would be D , and if future damages were guaranteed to exceed the per-period value implied by D in every period, then the present value of the stream of refunds at the time of emission would be zero. Therefore $\hat{q}_t - \hat{E}[d_t] \in [0, D]$.

Assumption 1 (Current Carbon Will Not Be Removed). $\hat{R}_t \leq M_t - M_0$ for all $t \geq 0$.

Even the highest-removal scenarios for the coming century do not project bringing carbon or temperature below current levels (IPCC, 2022), so this assumption is likely to be met by any carbon share policy begun in the next few years. The following proposition relates the period t loss to the deposits required at earlier times:

Proposition 4 (Efficiency Conditional on Information). *Let Assumption 1 hold. Then $\check{L} \rightarrow 0$ as $D \rightarrow \infty$.*

Proof. See Appendix H. □

The proposition establishes that the carbon share policy achieves the welfare-maximizing benchmark as the deposit becomes large. The proof shows that the time t private values for reducing emissions and removing carbon are each equal to

$$\sum_{j=1}^{\infty} \frac{1}{(1+r)^j} \hat{E}_t[\Delta_{t+j}].$$

Emitters lose the difference between the initial deposit and the initial value of the share they receive, and that difference is the present value of expected damage charges. Carbon removal benefits shareholders by preventing the loss of future damage charges. From (6), damage charges are the current period's realized marginal damage when $\bar{\mu}$ (and thus D) is large. Therefore the present value of expected future damage charges under large D is simply the present value of expected marginal damage from emissions, which is the social cost of carbon familiar from much work on the economics of climate change.³⁰

As the deposit becomes large, the carbon share policy maintains the emission reduction incentives of an emission tax but approaches efficiency in using information even in the presence of carbon removal. Comparing to results in Section 4.1, the carbon share policy outperforms an emission tax with a static revenue constraint if net negative emissions might ever become optimal, and the carbon share policy outperforms an emission tax with a dynamic revenue constraint if optimal removal might exhaust the cumulative revenue collected from emission taxes. Compare incentives under the carbon share policy and under an emission tax policy that attaches a distinct lockbox to each unit of emissions, so that the most a regulator can spend to remove old emissions is what the regulator collected at the time of emission. Increasing the emission tax at some time would increase the amount the regulator could later spend on removal, but such a change in the emission tax would overincentivize

³⁰I have used normal distributions for tractability and ease of exposition. If I instead assumed that the distribution of damages had finite support, then the carbon share policy would achieve the welfare-maximizing benchmark as $\bar{\mu}$ (and thus D) approaches some finite value from below. Under this interpretation, the carbon share policy approaches efficiency as the deposit approaches the worst-case social cost of carbon.

emission reductions. In contrast, increasing the carbon share's deposit increases both the efficiency of its emission price and the efficiency of its removal incentives.

The optimal carbon share policy provides the same incentives as would the optimal tax on the stock of carbon previously emitted by a firm (as opposed to the conventional tax on the flow of carbon emissions studied in Section 4).³¹ However, whereas firms could avoid a carbon stock tax by declaring bankruptcy, carbon shares are valuable assets that investors want to hold, financed at the time of emission by the deposit. Carbon shares therefore avoid judgment-proofness problems that can bedevil stock taxes. Moreover, a stock tax would lack the information aggregation benefits to be described in Section 5.2.

One might be concerned about whether the deposit would challenge firms' liquidity (see Shogren et al., 1993). Recall that firms receive a carbon share in return for their deposit and can immediately sell this valuable asset on. From (A-21), their net outlays per unit of emissions are the exact same outlays required by the traditional Pigouvian emission tax. This is why an arbitrarily large deposit does not distort firms' emission incentives. If the market for carbon shares is decently thick, a carbon share policy need not be any more financially challenging than a conventional carbon emission tax.³²

But one might still wonder about the scale of the deposit. If the deposit is not so large, then the highest potential damage charges are truncated by the constraints imposed by the deposits, which reduces the expected damage charges that firms use to guide emission and removal decisions. The possibility of hitting deposits' constraints therefore increases emissions and reduces removal. Ex post, options on carbon shares would reveal whether traders deemed it likely that the value of a carbon share would approach zero, as when damage charges are constrained by the deposit. Ex ante, a numerical exercise detailed in Appendix A provides some indication of how large a deposit may be necessary. This exercise takes damage estimates from the survey in Pindyck (2019) and considers the probability that any given deposit would be insufficient to cover the implied damage charges (i.e., that $\Delta_t < C_0 \alpha [\tilde{\zeta}_t + \tilde{\lambda}_t]$ in equation (6)). In this calibration, expected damages imply a tax of \$118 per tCO₂. Figure 3 shows that a deposit roughly twice as large (\$250 per tCO₂) would suffice in all but the worst 10% of cases, and a deposit roughly three times as large (\$400 per tCO₂) would suffice in all but the worst 5% of cases. An adequate deposit may therefore be well within an order of magnitude of what the carbon tax would have been.

5.2 Improved Efficiency in Observing Information

I now investigate how the market for carbon shares aggregates dispersed information about climate change damages. I therefore allow asymmetric information, as in Section 4.2. I

³¹Appendix B shows that the optimal time t stock tax would be $\lim_{\bar{\mu} \rightarrow \infty} \Delta_t$.

³²Gross outlays are also capped because any firm could avoid posting the deposit by reducing its emissions. The growing number of firms making zero emission pledges and recent cost projections for removal technologies both suggest that even the maximum gross outlays are limited to a reasonable scale.




Figure 3: The probability that deposits of various sizes would bind, based on the calibration in Appendix A and using $\tilde{\omega}^2 = \sigma^2 = 0$.

condition results on arbitrarily large D so as to highlight potential inefficiencies in observing information rather than in using information. From Section 5.1, emissions and carbon removal will be optimal conditional on information.

A continuum of traders normalized to unit mass is attached to sector i . At the beginning of time t , all agents have a symmetric, common prior over $\sum_{k=1}^N \kappa_k \zeta_k$, based on the regulator's measured aggregate damages in earlier periods and the prices of carbon shares in earlier periods. The time 0 prior is as described in Section 2.2, and the prior at the beginning of time t assigns variance τ_t^2 to each ζ_k . The price q_t of carbon shares at the beginning of time t reflects this information. Firms make emission and removal decisions based on this price and the population consumes accordingly. Subsequently, traders attached to sector i measure $\zeta_{it} + \lambda_{it}$.³³ They trade carbon shares based on this differentiated information. The market clears at price \check{q}_t^* , but noise traders make the observed price $\check{q}_t = \check{q}_t^* + \theta_t$, where θ_t is a mean-zero, independently and identically distributed, normal random variable, with variance $\Theta^2 > 0$.³⁴ Between periods t and $t + 1$, the regulator measures $\tilde{\zeta}_t + \tilde{\lambda}_t$ from its data on aggregate output, returns refunds d_t to shareholders based on shareholdings at the end

³³Traders do not need to be only in emitting sectors; they could be in any sector with information about damages. Here, that possibility would be reflected by $Y^{itr}(0) = 0$, in which case sector i would have zero time t emissions but could be affected by damages.

³⁴Noise has long been recognized as critical for the existence of partially revealing equilibria (e.g., Hellwig, 1980; Grossman and Stiglitz, 1980; Diamond and Verrecchia, 1981; Admati, 1985). Noise traders are often also interpreted as stochastic shocks to aggregate supply. Here noisy prices are critical for the existence of equilibrium because traders recognize the consequences of their continuum of companions all bidding with the same information. In Vives (2014), the information structure has fundamental values and small traders, which precludes the need for noise traders, but here the payoffs themselves are pure common values and thus retain the rationale for noise traders despite traders being small (see footnote 37).

of period t , and issues new shares to firms based on period t emissions.³⁵

At the beginning of time t , traders of type i have share holdings y_{it} and wealth w_{it} . The y_{it} include shares issued in all previous periods that are still active (i.e., for which the underlying unit of carbon has not yet been removed). Traders have the ability to invest in a riskless asset with return r . After observing their private signals, traders choose their net demand X_{it} to maximize their expected utility of wealth at the beginning of period $t + 1$:

$$\max_{X_{it}} E_t \left[-\exp \left[-A \left((1+r)(w_{it} + (X_{it} + y_{it})d_t - X_{it}\check{q}_t) + (y_{it} + X_{it})q_{t+1} \right) \right] \middle| \zeta_{it} + \lambda_{it}, \check{q}_t \right],$$

with $A > 0$ the coefficient of absolute risk aversion and E_t indicating expectations based on common information available at the beginning of time t . Traders have exponential utility, as opposed to the logarithmic utility function of the representative household. Exponential utility is critical to the analysis in this section because exponential utility yields linear asset demand functions that are independent of wealth and amenable to aggregation. For these reasons, exponential utility (including its implementation as quadratic payoffs) is used in nearly all literature on asymmetric information in asset markets.³⁶

In a (noisy) rational expectations equilibrium, markets clear with traders inferring from prices whatever information they can and bidding to maximize utility conditional on that information. This equilibrium is defined as fully revealing if the carbon share price reveals the same information about $\sum_{k=1}^N \kappa_k \zeta_k$ as would observing all the disentangled $\zeta_{kt} + \lambda_{kt}$. The following proposition establishes properties of this equilibrium:

Proposition 5 (Efficient Equilibrium). *Let Assumption 1 hold. A fully revealing rational expectations equilibrium with $\check{L} = 0$ exists in the limit as $\Theta^2 \rightarrow 0$ and $D \rightarrow \infty$.*

Proof. See Appendix J. □

There exists a fully revealing rational expectations equilibrium as noise traders lose influence and the deposit becomes large. The price of a carbon share then reflects informationally efficient beliefs and traders hold those same beliefs upon observing their private information and the carbon share price. By aggregating traders' private information about damages, the carbon share market improves on the regulator's ability to estimate damages, and by defining the marginal cost of emitting, the carbon share market simultaneously implements that information to control emissions and incentivize carbon removal. We have therefore designed a decentralized policy instrument that can implement the welfare-maximizing, informationally efficient benchmark that is generally unattainable via an emission tax instrument.

³⁵These new shares can be handled by including time t emitters in the set of time $t + 1$ traders.

³⁶If we give the representative household in Section 2.1 exponential utility over consumption and also let damages be additive rather than multiplicative, then the learning dynamics are unchanged and, as $A \rightarrow 0$ (so as to eliminate risk premia, as with log utility), the regulator's tax in Section 4 and the damage charge in Section 5.1 are altered only by losing the C_0 normalization.

However, it is well-known that a fully revealing rational expectations equilibrium is not always implementable: it may be that no trading mechanism can actually deliver this equilibrium. In particular, if the carbon share price is a sufficient statistic for all information in the economy, then traders should ignore their private information, in which case it is unclear how their private information ends up being summarized by the equilibrium price. I therefore also study an equilibrium in demand functions (e.g., Kyle, 1989). Here traders submit demand functions that account for their observed sectoral signals $\zeta_{it} + \lambda_{it}$ and for the information they would infer from an observed price \check{q}_t . Traders treat \check{q}_t as exogenous (i.e., they are price-takers) but do recognize how their observed signals influence that price through the beliefs of other traders in their sector.³⁷ Following much previous literature, I associate an implementable equilibrium with a Bayesian Nash equilibrium of this game and study linear equilibria.³⁸

Define $\check{\mu}_t$ as traders' posterior estimate of $\sum_{k=1}^N \kappa_k \zeta_k$ at the beginning of time t , after observing \check{q}_s and $\check{\zeta}_s + \check{\lambda}_s$ for all $s \in \{0, \dots, t-1\}$. The following proposition characterizes emissions in an implementable equilibrium.

Proposition 6 (Implementable Equilibrium). *In the limit as $D \rightarrow \infty$, there exists a linear equilibrium in demand functions in which the marginal cost of emissions and marginal benefit of carbon removal are each equal to*

$$\frac{1}{r} C_0 \alpha \left[\sum_{k=1}^N \kappa_k \bar{\zeta}_k + \check{\mu}_t \right]$$

and

$$\begin{aligned} \check{\mu}_t = & \sum_{k=0}^{t-1} \check{\pi}_k \left(\check{\zeta}_k + \check{\lambda}_k - \sum_{j=1}^N \kappa_j \bar{\zeta}_j \right) \\ & + \sum_{k=0}^{t-1} \frac{\check{\pi}_k}{\check{\chi}_k} \underbrace{\left[\frac{(1-\Gamma)\tau_k^2 + \sigma^2}{\tau_k^2 + \sigma^2 + \omega^2} \sum_{i=1}^N \check{\kappa}_{ik} \kappa_i (\zeta_{ik} + \lambda_{ik} - \bar{\zeta}_i) + \frac{N\Gamma\tau_k^2}{\tau_k^2 + \sigma^2 + \omega^2} \frac{1}{N} \sum_{i=1}^N \check{\kappa}_{ik} (\zeta_{ik} + \lambda_{ik} - \bar{\zeta}_i) \right]}_{\text{from } \check{q}_k} \\ & - \sum_{k=0}^{t-1} \frac{\check{\pi}_k}{\check{\chi}_k} \frac{r}{C_0 \alpha (\check{\chi}_k + r)} \theta_k, \end{aligned} \tag{9}$$

where $\check{\kappa}_{kt}$, $\check{\chi}_k$ and $\check{\chi}_k$ are each $\in (0, 1)$. If Γ , ω^2 , $\tilde{\omega}^2$, and Θ^2 are sufficiently small, then $\check{\pi}_k \in (0, 1)$ and $\lim_{\Gamma, \tilde{\omega}^2 \rightarrow 0} \check{\pi}_k$ is arbitrarily close to zero.

³⁷As in Vives (2014), the continuum of traders solves the “schizophrenic” problem of Hellwig (1980) because price-taking behavior is here individually optimal. However, I study an economy with a finite number of sectors and so the signals observed in some sector can affect the price.

³⁸The equilibrium is symmetric in terms of strategies defined over ζ_{it} and \check{q}_t . Of course, the actual demand schedules will not be symmetric as each will depend on the observed ζ_{it} .

Proof. See Appendix K.

Sketch: The proof uses traders' first-order conditions to determine demand for carbon shares in each sector. The equilibrium price \check{q}_t^* sets aggregate net demand to zero given the beliefs traders form from their private signals, the observed carbon share price, and expectations of q_{t+1} . Because normal-normal Bayesian updating implies that $\check{\mu}_{t+1}$ is a linear function of $\check{\mu}_t$, \check{q}_t , and $\check{\zeta}_t$, so too is expected q_{t+1} . The proof then constructs a signal \tilde{q}_t of aggregate damages implied by \check{q}_t . By normal-normal Bayesian updating, each sector i trader's posterior mean for $\check{\zeta}_t + \check{\lambda}_t$ is a linear function of $\check{\mu}_t$, \check{q}_t , and $\check{\zeta}_{it}$. We can thus express the price signal \check{q}_t as an unknown linear function of sectoral signals. The projection theorem yields each type of trader's posterior mean for aggregate damages conditional on observed sectoral information and on the observed price signal. Matching coefficients and applying Brouwer's fixed-point theorem yields posterior beliefs that are self-fulfilling via the price and also yields the market-clearing price, both as functions of the unknown coefficients that determine $\check{\mu}_{t+1}$. Beliefs $\check{\mu}_{t+1}$ follow from multivariate normal-normal updating, matching coefficients to the conjectured form, and recursively substituting backwards for earlier $\check{\mu}_k$ and earlier \check{q}_k . \square

Carbon shares act like imposing an emission tax based on beliefs $\check{\mu}_t$. The first line in (9) determines the weight placed on previous periods' aggregate measurements of damages. The second line describes how agents learn from the past prices \check{q}_k of carbon shares (use (A-44) in (A-32)). Those past prices embed two types of information: a first piece learns from a version of the aggregate signal constructed from sectoral signals, and a second piece takes advantage of the correlation among sectoral effects to learn from the disentangled signals. The ability to construct a version of the aggregate signal that is affected by sectoral measurement error but not by aggregate measurement error and the ability to use the correlation across sectors to learn more efficiently were critical to informationally efficient beliefs but were missing from the regulator's beliefs.³⁹

The most notable differences with respect to the informationally efficient beliefs described in Proposition 1 are the randomness induced by noise traders (third line in (9)) and the presence of the $\check{\kappa}_{ik}$. As the $\kappa_i \rightarrow 1/N$ (i.e., as sectors become symmetric), the $\check{\kappa}_{it}$ approach a constant $\check{\kappa}_t \in (0, 1)$ and the second line approaches

$$\sum_{k=0}^{t-1} \frac{\check{\pi}_k}{\check{\chi}_k} \check{\kappa}_k \left[\frac{(1 - \Gamma)\tau_k^2 + \sigma^2}{\tau_k^2 + \sigma^2 + \omega^2} \sum_{i=1}^N \kappa_i (\zeta_{ik} + \lambda_{ik} - \bar{\zeta}_i) + \frac{N\Gamma\tau_k^2}{\tau_k^2 + \sigma^2 + \omega^2} \frac{1}{N} \sum_{i=1}^N (\zeta_{ik} + \lambda_{ik} - \bar{\zeta}_i) \right].$$

The term in brackets is similar to a term in Proposition 1.⁴⁰ This expression illustrates that the carbon price does aggregate dispersed information in implementable equilibria and

³⁹Under the conditions of the proposition, traders' posterior beliefs do not rely on past share prices ($\check{\pi}_k \approx 0$) when sectoral effects are uncorrelated ($\Gamma = 0$) and the aggregate signal is perfectly measured ($\check{\omega}^2 = 0$). This result should be unsurprising given the analysis of the regulator in Section 4.2.

⁴⁰One difference is the σ^2 in the numerator of the first term, which was missing from Proposition 1. Carbon share traders are trying to predict the aggregate measurement that will be released following the

moreover aggregates that information in a fashion analogous to—albeit not identical to—informationally efficient beliefs.

The downward adjustments due to the $\check{\kappa}$ reflect two forces: traders in each sector shade their bids to reflect their awareness of their own sector's signals leaking into the asset price (from equation (A-43)), and risk-averse traders' demand for carbon shares decreases in the variance of the returns they will earn (from equation (A-23)). The following corollary examines the $\check{\kappa}$ in more detail.

Corollary 4 (Traders' Distortions). *Consider the $\{\check{\kappa}_{1t}, \dots, \check{\kappa}_{Nt}\}$ defined in Proposition 6.*

- i $\lim_{\Theta^2 \rightarrow 0} \check{\kappa}_{it} = 0$*
- ii If $\kappa_i = 1/N$ for all $i \in \{1, \dots, N\}$, then $\check{\kappa}_t$ increases in Θ^2 and $\lim_{\Theta^2 \rightarrow \infty} \check{\kappa}_t = 1/N$.*
- iii Without loss of generality, order sectors by κ_i . As $\Theta^2 \rightarrow \infty$, the sequence $\{\check{\kappa}_{1t}, \dots, \check{\kappa}_{Nt}\}$ is monotone increasing, with $\check{\kappa}_{1t} \leq 1/N$ and $\check{\kappa}_{Nt} \geq 1/N$. The latter two inequalities are strict if $\kappa_1 < 1/N$.*

Proof. See Appendix L. □

The first result implies that carbon share prices fail to aggregate information as noise traders become irrelevant. This is a manifestation of the same force that prevents the fully revealing rational expectations equilibrium from being implementable: traders whose information is fully revealed by the equilibrium price do not trade on that information. In contrast, the proof shows that bid shading vanishes as $\Theta^2 \rightarrow \infty$. Thus there is a tension between minimizing the consequences of the third line in (9) and minimizing the consequences of bid shading for the $\check{\kappa}$. The second part of the corollary shows that risk aversion still matters for the $\check{\kappa}$ as $\Theta^2 \rightarrow \infty$ and that the effects of risk aversion are symmetric when the sectors are symmetric (i.e., with identical κ_i). And the third part of the corollary shows that the effects of risk aversion are more severe for sectors that have smaller value shares in final-good production, because traders in those sectors observe signals that are less informative about returns to carbon shares and thus perceive additional risk that makes them less willing to trade carbon shares.

In sum, a carbon share policy attains informational efficiency in a fully revealing rational expectations equilibrium and does aggregate information in implementable equilibria in a fashion that is imperfect but analogous to the aggregation performed within informationally efficient updating. By using markets to perform price discovery for social cost, a carbon

same period because this measurement determines immediate refunds and subsequent carbon share prices. This aggregate measurement will include their privately observed stochastic sectoral shocks, whose variance is σ^2 . These contemporaneous shocks do not hinder learning about the coming aggregate measurement. In contrast, informationally efficient beliefs formed from time t signals in Proposition 1 predict aggregate measurements in all subsequent periods, in which case the contemporaneous stochastic sectoral shocks are pure noise that hinders learning.

share policy mitigates the difficulties that an emission tax's regulator faces in observing the information available in the economy.

6 Conclusions and Discussion

I have advanced a new perspective on environmental policymaking. Extending the traditional emphasis on asymmetric information about firms' costs of eliminating emissions, I have emphasized asymmetric information about the social cost of emissions. I have shown that conventional emission taxes neither aggregate dispersed information nor enable full use of potential information about the severity of climate change. Instead, I have shown that a new market-based instrument I call "carbon shares" aggregates dispersed information and enables full use of new information without losing the desirable properties of emission taxes and other conventional market-based instruments.

The new policy instrument conceives of a different role for the regulator. Traditionally, the regulator must project the marginal harm from emissions in all possible states of the world and in all future time periods in order to determine an emission price. Here, however, the regulator need only measure damage as it is realized and determine the deposit based on approximate worst-case outcomes. This new instrument shifts the burden of projecting possible future damages from the regulator to market traders. These traders form their own damage estimates from information produced by the regulator, from the observed prices of carbon shares, and from their own private information. This type of belief updating is a common task in markets.

This proposal generates four immediate questions. First, how would a regulator actually estimate the realized aggregate impacts of climate change? This task is different from the task economists have traditionally undertaken in projecting future damages from climate change. It is closer to the attribution studies now regularly undertaken in which climate scientists test how climate change affected the likelihood of recent realized weather events. And governments already do regularly produce measures of recent economic outcomes that are noisy yet are of critical importance for policymaking and directly determine monetary payments. As but one example, the U.S. consumer price index determines social security benefits and other transfer payments and is a prominent input to monetary policy, but it is imperfectly estimated and there is disagreement even about what it should be estimating (National Research Council, 2002; Schultze, 2003). The present challenge may be no greater.

Second, would the regulator have credibility to estimate realized impacts faithfully? If an econometric framework could be developed that became widely accepted, then the estimation may here be institutionalized as with the production of other federal statistics—and to the extent this estimation relies on standard data, it may be less vulnerable to political influence than the U.S. government's estimates of the social cost of carbon have been (see Voosen, 2021). A real-world implementation of the policy might also constrain the change

in damage charges from period to period, which would reduce the flexibility to respond to new information but also reduce vulnerability to transient political influence.

Third, how would this instrument affect incentives to coordinate policy internationally? I have followed a long tradition in analyzing the benchmark of a global regulator. However, climate policy is in practice fragmented among countries. Future work should compare international dimensions of this policy to carbon taxes, cap-and-trade programs, and other policy options. In particular, the ability to institutionalize the damage charge calculations and to explicitly adopt country weights in the damage charge calculations could each affect incentives to coordinate policy: these calculations may have more credibility than a global carbon tax would enjoy and countries may be incentivized to join a coalition in order to have their damages counted.

Finally, would traders have an incentive to collect additional information about climate impacts? To date, the development of better scientific monitoring and modeling systems has primarily been the task of governments and universities. However, such information should have a market value under the policy instrument analyzed herein, as I conjecture that the implementable equilibrium does not suffer the paradox of Grossman and Stiglitz (1976, 1980). One might thus expect traders to invest in information production, so that a carbon share policy may not just aggregate the information already dispersed throughout the economy but also improve that information. Such information might also have a market value under an emission tax (as agents want to understand their own exposure to climate change and to forecast future emission taxes), but this information is plausibly much more valuable under a carbon share policy because it determines the immediate payoffs from trading carbon shares. I leave the analysis of incentives to collect information in various policy environments to future work.

References

- Admati, Anat R. (1985) “A noisy rational expectations equilibrium for multi-asset securities markets,” *Econometrica*, Vol. 53, No. 3, pp. 629–657.
- Aliakbari, Elmira and Ross McKittrick (2018) “Information aggregation in a prediction market for climate outcomes,” *Energy Economics*, Vol. 74, pp. 97–106.
- Baumol, William J. (1972) “On taxation and the control of externalities,” *The American Economic Review*, Vol. 62, No. 3, pp. 307–322.
- Bednar, Johannes, Michael Obersteiner, Artem Baklanov, Marcus Thomson, Fabian Wagner, Oliver Geden, Myles Allen, and Jim W. Hall (2021) “Operationalizing the net-negative carbon economy,” *Nature*, Vol. 596, No. 7872, pp. 377–383.

- Bednar, Johannes, Michael Obersteiner, and Fabian Wagner (2019) “On the financial viability of negative emissions,” *Nature Communications*, Vol. 10, No. 1, p. 1783.
- Bohm, Peter (1981) *Deposit-Refund Systems: Theory and Applications to Environmental, Conservation, and Consumer Policy*, Baltimore: RFF Press.
- Bohm, Peter and Clifford S. Russell (1985) “Comparative analysis of alternative policy instruments,” in Allen V. Kneese and James L. Sweeney eds. *Handbook of Natural Resource and Energy Economics*, Vol. 1: Elsevier Science Publishers B.V. (North-Holland), pp. 395–460.
- Boyd, James (2002) “Financial responsibility for environmental obligations: Are bonding and assurance rules fulfilling their promise?” *Research in Law and Economics*, Vol. 20, pp. 417–486.
- Campbell, John (2017) *Financial Decisions and Markets: A Course in Asset Pricing*, Princeton: Princeton University Press.
- Carleton, Tamma A. and Solomon M. Hsiang (2016) “Social and economic impacts of climate,” *Science*, Vol. 353, No. 6304, p. aad9837.
- Costanza, Robert and Charles Perrings (1990) “A flexible assurance bonding system for improved environmental management,” *Ecological Economics*, Vol. 2, No. 1, pp. 57–75.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2014) “What do we learn from the weather? The new climate-economy literature,” *Journal of Economic Literature*, Vol. 52, No. 3, pp. 740–798.
- Deschnes, Olivier and Michael Greenstone (2007) “The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather,” *American Economic Review*, Vol. 97, No. 1, pp. 354–385.
- Diamond, Douglas W. and Robert E. Verrecchia (1981) “Information aggregation in a noisy rational expectations economy,” *Journal of Financial Economics*, Vol. 9, No. 3, pp. 221–235.
- Dietz, Simon and Frank Venmans (2019) “Cumulative carbon emissions and economic policy: In search of general principles,” *Journal of Environmental Economics and Management*, Vol. 96, pp. 108–129.
- Edenhofer, Ottmar, Max Franks, and Matthias Kalkuhl (2021) “Pigou in the 21st Century: a tribute on the occasion of the 100th anniversary of the publication of The Economics of Welfare,” *International Tax and Public Finance*, Vol. 28, pp. 1090–1121.

- Fullerton, Don and Thomas C. Kinnaman (1995) “Garbage, recycling, and illicit burning or dumping,” *Journal of Environmental Economics and Management*, Vol. 29, No. 1, pp. 78–91.
- Fullerton, Don and Ann Wolverton (2000) “Two generalizations of a deposit-refund system,” *The American Economic Review Papers and Proceedings*, Vol. 90, No. 2, pp. 238–242.
- (2005) “The two-part instrument in a second-best world,” *Journal of Public Economics*, Vol. 89, No. 9, pp. 1961–1975.
- Fuss, Sabine, William F. Lamb, Max W. Callaghan, Jérôme Hilaire, Felix Creutzig, Thorben Amann, Tim Beringer, Wagner de Oliveira Garcia, Jens Hartmann, Tarun Khanna, Gunnar Luderer, Gregory F. Nemet, Joeri Rogelj, Pete Smith, José Luis Vicente Vicente, Jennifer Wilcox, Maria del Mar Zamora Dominguez, and Jan C. Minx (2018) “Negative emissions—Part 2: Costs, potentials and side effects,” *Environmental Research Letters*, Vol. 13, No. 6, p. 063002.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2014) “Optimal taxes on fossil fuel in general equilibrium,” *Econometrica*, Vol. 82, No. 1, pp. 41–88.
- Grossman, Sanford (1976) “On the efficiency of competitive stock markets where trades have diverse information,” *The Journal of Finance*, Vol. 31, No. 2, pp. 573–585.
- (1978) “Further results on the informational efficiency of competitive stock markets,” *Journal of Economic Theory*, Vol. 18, No. 1, pp. 81–101.
- Grossman, Sanford J. and Joseph E. Stiglitz (1976) “Information and competitive price systems,” *The American Economic Review*, Vol. 66, No. 2, pp. 246–253.
- (1980) “On the impossibility of informationally efficient markets,” *The American Economic Review*, Vol. 70, No. 3, pp. 393–408.
- Hayek, F. A. (1945) “The use of knowledge in society,” *The American Economic Review*, Vol. 35, No. 4, pp. 519–530.
- Hellwig, Martin F. (1980) “On the aggregation of information in competitive markets,” *Journal of Economic Theory*, Vol. 22, No. 3, pp. 477–498.
- Holladay, J. Scott, Jonathan Horne, and Jason A. Schwartz (2009) “Economists and climate change: Consensus and open questions,” Policy Brief 5, Institute for Policy Integrity, New York University School of Law.
- Hsu, Shi-Ling (2011) “A prediction market for climate outcomes,” *University of Colorado Law Review*, Vol. 83, No. 1, pp. 179–256.

- IPCC (2022) “Summary for Policymakers,” in P. R. Shukla, J. Skea, R. Slade, A. Al Khourdajie, R. van Diemen, D. McCollum, M. Pathak, S. Some, P. Vyas, R. Fradera, M. Belkacemi, A. Hasija, G. Lisboa, S. Luz, and J. Malley eds. *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge, UK and New York, NY, USA: Cambridge University Press.
- Joppa, Lucas, Amy Luers, Elizabeth Willmott, S. Julio Friedmann, Steven P. Hamburg, and Rafael Broze (2021) “Microsoft’s million-tonne CO₂-removal purchase—lessons for net zero,” *Nature*, Vol. 597, No. 7878, pp. 629–632.
- Kyle, Albert S. (1985) “Continuous auctions and insider trading,” *Econometrica*, Vol. 53, No. 6, pp. 1315–1335.
- (1989) “Informed speculation with imperfect competition,” *The Review of Economic Studies*, Vol. 56, No. 3, pp. 317–355.
- Lemoine, Derek (2007) “Greenhouse gas property: An adaptable climate policy for an uncertain world,” SSRN Working Paper 2256992.
- (2021a) “The climate risk premium: How uncertainty affects the social cost of carbon,” *Journal of the Association of Environmental and Resource Economists*, Vol. 8, No. 1, pp. 27–57.
- (2021b) “Estimating the consequences of climate change from variation in weather,” Working Paper 25008, National Bureau of Economic Research.
- McKittrick, Ross (2011) “A simple state-contingent pricing rule for complex intertemporal externalities,” *Energy Economics*, Vol. 33, No. 1, pp. 111–120.
- Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw (1994) “The impact of global warming on agriculture: A Ricardian analysis,” *The American Economic Review*, Vol. 84, No. 4, pp. 753–771.
- Metcalf, Gilbert E (2009) “Market-based policy options to control U.S. greenhouse gas emissions,” *Journal of Economic Perspectives*, Vol. 23, No. 2, pp. 5–27.
- Mills, Edwin S. (1972) *Urban Economics*.
- National Academies of Sciences, Engineering, and Medicine (2018) *Negative Emissions Technologies and Reliable Sequestration: A Research Agenda*, Washington, DC: The National Academies Press.

- National Research Council (2002) *At What Price?: Conceptualizing and Measuring Cost-of-Living and Price Indexes*, Washington, DC: The National Academies Press.
- (2015) *Climate Intervention: Carbon Dioxide Removal and Reliable Sequestration*, Washington, DC: The National Academies Press.
- Nordhaus, William (2013) “Integrated economic and climate modeling,” in Peter B. Dixon and Dale W. Jorgenson eds. *Handbook of Computable General Equilibrium Modeling*, Vol. 1: North Holland, Elsevier B.V. pp. 1069–1131.
- (2019) “Climate change: The ultimate challenge for economics,” *American Economic Review*, Vol. 109, No. 6, pp. 1991–2014.
- Nordhaus, William D. (1977) “Economic growth and climate: The carbon dioxide problem,” *American Economic Review Papers and Proceedings*, Vol. 67, No. 1, pp. 341–346.
- (1992) “An optimal transition path for controlling greenhouse gases,” *Science*, Vol. 258, No. 5086, pp. 1315–1319.
- Pindyck, Robert S. (2013) “Climate change policy: What do the models tell us?” *Journal of Economic Literature*, Vol. 51, No. 3, pp. 860–872.
- (2019) “The social cost of carbon revisited,” *Journal of Environmental Economics and Management*, Vol. 94, pp. 140–160.
- Rao, Akhil, Matthew G. Burgess, and Daniel Kaffine (2020) “Orbital-use fees could more than quadruple the value of the space industry,” *Proceedings of the National Academy of Sciences*, Vol. 117, No. 23, pp. 12756–12762.
- Rickels, Wilfried, Alexander Proelss, Oliver Geden, Julian Burhenne, and Mathias Fridahl (2020) “The future of (negative) emissions trading in the European Union,” Kiel Working Paper 2164, Kiel Institute for the World Economy.
- Rostek, Marzena J. and Ji Hee Yoon (2020) “Equilibrium theory of financial markets: Recent developments,” prepared for Journal of Economic Literature.
- Rostek, Marzena and Marek Weretka (2012) “Price inference in small markets,” *Econometrica*, Vol. 80, No. 2, pp. 687–711.
- Russell, Clifford S. (1987) “Economic incentives in the management of hazardous wastes,” *Columbia Journal of Environmental Law*, Vol. 13, No. 2, pp. 257–274.
- Schlenker, Wolfram and Charles A. Taylor (2021) “Market expectations of a warming climate,” *Journal of Financial Economics*, Vol. 142, No. 2, pp. 627–640.

- Schultze, Charles L. (2003) “The Consumer Price Index: Conceptual Issues and Practical Suggestions,” *Journal of Economic Perspectives*, Vol. 17, No. 1, pp. 3–22.
- Scott, Vivian and Oliver Geden (2018) “The challenge of carbon dioxide removal for EU policy-making,” *Nature Energy*, Vol. 3, No. 5, pp. 350–352.
- Shogren, Jason F., Joseph A. Herriges, and Ramu Govindasamy (1993) “Limits to environmental bonds,” *Ecological Economics*, Vol. 8, No. 2, pp. 109–133.
- Singleton, Kenneth J. (1987) “Asset prices in a time series model with disparately informed, competitive traders,” in W. Barnett and K. J. Singleton eds. *New Approaches to Monetary Economics*, Cambridge, MA: Cambridge University Press.
- Solow, Robert M. (1971) “The economist’s approach to pollution and its control,” *Science*, Vol. 173, No. 3996, pp. 498–503.
- Stavins, Robert N. (2022) “The relative merits of carbon pricing instruments: Taxes versus trading,” *Review of Environmental Economics and Policy*, Vol. 16, No. 1, pp. 62–82.
- Torsello, Loredana and Alessandro Vercelli (1998) “Environmental bonds: A critical assessment,” in Graciela Chichilnisky, Geoffrey Heal, and Alessandro Vercelli eds. *Sustainability: Dynamics and Uncertainty*, Dordrecht: Springer Netherlands, pp. 243–255.
- Vayanos, Dimitri (1999) “Strategic trading and welfare in a dynamic market,” *The Review of Economic Studies*, Vol. 66, No. 2, pp. 219–254.
- (2001) “Strategic trading in a dynamic noisy market,” *The Journal of Finance*, Vol. 56, No. 1, pp. 131–171.
- Vives, Xavier (1995) “Short-term investment and the informational efficiency of the market,” *The Review of Financial Studies*, Vol. 8, No. 1, pp. 125–160.
- (2008) *Information and Learning in Markets: The Impact of Market Microstructure*, Princeton: Princeton University Press.
- (2011) “Strategic supply function competition with private information,” *Econometrica*, Vol. 79, No. 6, pp. 1919–1966.
- (2014) “On the possibility of informationally efficient markets,” *Journal of the European Economic Association*, Vol. 12, No. 5, pp. 1200–1239.
- Voosen, Paul (2021) “Trump downplayed the cost of carbon. That’s about to change,” *Science*, Vol. 371, No. 6528, pp. 447–448.

- Wang, Jiang (1993) “A model of intertemporal asset prices under asymmetric information,” *The Review of Economic Studies*, Vol. 60, No. 2, pp. 249–282.
- Weitzman, Martin L. (1974) “Prices vs. quantities,” *The Review of Economic Studies*, Vol. 41, No. 4, pp. 477–491.
- White, Ben, Graeme J. Doole, David J. Pannell, and Veronique Florec (2012) “Optimal environmental policy design for mine rehabilitation and pollution with a risk of non-compliance owing to firm insolvency,” *Australian Journal of Agricultural and Resource Economics*, Vol. 56, No. 2, pp. 280–301.
- Yang, Peifang and Graham A. Davis (2018) “Non-renewable resource extraction under financial incentives to reduce and reverse stock pollution,” *Journal of Environmental Economics and Management*, Vol. 92, pp. 282–299.