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, Revised October 2024

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 extends and replaces the 2021 working paper “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.

Informationally Efficient Climate Policy: Designing Markets to Measure and Price Externalities
Derek Lemoine
NBER Working Paper No. 30535
October 2022, Revised October 2024
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 that I label "carbon shares". This instrument empowers markets to perform price discovery about future damages. I show that there exists a limit-case rational expectations equilibrium in which this instrument: i) efficiently prices new emissions, ii) efficiently incentivizes removal of past emissions, and iii) efficiently aggregates dispersed information about the social cost of emissions. Conventional emission taxes achieve only the first of these objectives, for given information. Calibrated numerical simulations suggest benefits from achieving the other two objectives.

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 require the regulator only to measure the social cost of pollution (i.e., the externality) and aggregate abatement costs, whereas command-and-control policies also require the regulator to measure individual firms' costs 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] Even after decades of work, economists have still not converged on

[1]Nordhaus (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 (Deschênes 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).

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. Every person and firm on the 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 economists have not studied how to design markets to aggregate information about externalities, and thereby to perform price discovery for social costs.

I study the informational efficiency of environmental policy. I design and analyze a new market-based policy instrument that simultaneously measures and controls externalities. In my setting, firms in different sectors of the economy trade off the benefit of emitting carbon against the cost of complying with current policies, as is customary. 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 are unbiased but may be arbitrarily noisy. 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 subject to two types of inefficiencies when uncertain about climate damages.[4] First, an emission tax cannot incentivize

[2]In 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 tCO_2, around three times larger than the average estimate. The social costs of carbon calculated in Pindyck (2019) from another survey of 113 economists also show substantial dispersion, with the standard deviation again of comparable magnitude to the mean.

[3]Carbon 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 tCO_2 on average (Joppa et al., 2021).

[4]My analysis of emission taxes extends to an emission cap, which is also set based on a regulator's estimate of social cost and faces similar limitations in achieving negative emissions. Weitzman (1974) shows that asymmetric information about abatement costs can break the equivalence between an emission tax and

the efficient removal of past emissions, which constitutes an *allocative inefficiency*. An emission tax must be paired with a subsidy for removal, but I show that the efficient subsidy cannot be financed by retained emission tax revenue when the regulator learns about damages over time. The regulator is unable to use new information about damages to optimally reallocate economic activity unless it raises new revenue from taxpayers.

The second inefficiency concerns the optimality of that new information. I show that the regulator's emission tax is *informationally inefficient* in the plausible case that the regulator measures aggregate consequences only imperfectly and also in the plausible case that sectors have heterogeneous value shares in final good production 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 design a new instrument that aims to incentivize optimal carbon removal without spending taxpayer funds and to efficiently aggregate dispersed information. The regulator requires that emitters post a deposit at the time of emission and in exchange gives emitters a tradeable security 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 and which vary inversely with market expectations of the regulator's future damage measurements. 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 retire the carbon share and recover its deposit. If, for example, traders become more pessimistic about climate change, then expected refunds fall. Emitting becomes less attractive because the value of the carbon share received falls, and carbon removal becomes more attractive because the opportunity cost of retiring a carbon share falls.

I show that there exists a limit-case, 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 match the allocatively and informationally efficient benchmark.[5] This optimal outcome depends on the

cap. I instead emphasize imperfect and asymmetric information about the externality.

[5]Following the convention in finance and information economics, I use "rational expectations" to indicate

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' expectations 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 damage charges will be constrained by the deposit in periods with high measured climate change damages. As the deposit becomes large, traders' expectations of damage charges converge to their expectations of future damage measurements and thus to their current estimates of the marginal damage from carbon emissions, so that the private cost of emitting converges to their estimates of the social cost of carbon emissions.[6] A numerical calibration suggests that a deposit around 2–3 times as large as the estimated social cost of emissions approximates optimal emission and removal incentives.

The fully revealing equilibrium exists, but, in an instance of the Grossman-Stiglitz paradox, it is unclear how traders' private information could appear in a price that induces them to ignore that same private information. I therefore also study the linear Bayesian equilibrium of a game in traders' demand functions. I show that equilibrium emission incentives are not exactly the same as what an all-seeing regulator would implement, but I also show that, via the equilibrium carbon share price, emission incentives do account for dispersed information about climate change damages that is unavailable to a centralized, tax-wielding planner.

Where dispersed information about damages is relatively unimportant and cleanup of past emissions is irrelevant (as may be true of particulate matter or lead pollution), the present policy performs like an emission tax or cap-and-trade program. But where information about damages is dispersed, the present policy acts like improving the information underlying an emission tax or cap-and-trade program, and where cleanup of past emissions is potentially

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.

[6]The regulator can choose to measure whatever it cares about. Whatever the regulator attempts to measure will determine the market incentives to reduce emissions and to remove old emissions. For instance, the regulator could apply equity weighting, value nonmarket impacts, and/or, as analyzed in Appendix G, induce the market to apply a social discount rate to damages.

relevant, the present policy can incentivize such cleanup without requiring the public to fund it. Climate change clearly demonstrates both dispersed information about impacts and the possibility of ex post cleanup, and many other externalities will too.[7]

I quantitatively assess the importance of allocative inefficiency for climate change policy by extending the DICE-2016R climate-economy model (Nordhaus, 2017).[8] So as to focus on the inefficiency related to carbon removal, this experiment assumes that all actors observe all information in the economy. I find that, relative to the optimal carbon share policy, the expected loss from employing an emission tax by itself (i.e., without a subsidy for carbon removal) is 7–9% of the potential benefits from policy, depending on the damage calibration. The loss is only slightly smaller if the regulator saves the revenue from the emission tax to pay for later carbon removal subsidies. The loss is concentrated in high-damage states of the world, when carbon removal is especially valuable ex post.

I quantitatively assess the importance of informational inefficiency for climate change policy by calibrating a prior distribution for industries' temperature exposure to county-level U.S. data. Simulating the linear Bayesian equilibrium of the carbon share policy, I show that the mean squared error of the implied emission price decreases over time and is closer to informational efficiency than is the mean squared error of the emission tax chosen by a regulator based on public aggregate data. The magnitude of the carbon share policy's informational advantage depends on the fraction of residual year-to-year stochasticity in industries' output that I attribute to variation in exposure to temperature. As that fraction declines, sectoral signals improve and the emission price implied by the carbon share policy converges to the informationally efficient price and to the truth, whereas the emission tax remains constrained by the regulator's inability to access private sectoral information.

A carbon share policy is a dynamic deposit-refund instrument. Static deposit-refund instruments resolve difficulties monitoring—and thus taxing—improper waste disposal (e.g.,

[7]For 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 above.

[8]The DICE-2016R model does not specify how carbon removal is incentivized. Its optimal policy includes carbon removal that would require trillions of dollars of subsidies, which far exceeds the cumulative revenue collected from emission taxes.

Bohm, 1981; Russell, 1987; Fullerton and Kinnaman, 1995; Torsello and Vercelli, 1998).[9,10] I posit no problem monitoring either the act of emission or the act of carbon removal (see footnote 28 below). From the perspective of emitters, my policy combines a tax (via the deposit) and a subsidy (via the value of the carbon share received), with emission incentives determined by the difference between the deposit and the value of the carbon share. Static tax-subsidy schemes have been explored by Fullerton and Wolverton (2000, 2005). The present proposal has at least two critical differences. First, the subsidy here evolves with new information about the damages caused by a stock pollutant, rather than being tied to information at the time of emission. Second, the level of the subsidy, in the form of the carbon share price, is here determined in equilibrium from traders' private information about climate change impacts rather than being fixed by a regulator based on public information.

My proposed instrument can be understood by comparison to two other proposals. First, a regulator might try to resolve the allocative inefficiency by taxing firms for their past contributions to the current stock of atmospheric carbon (as opposed to taxing their current flow of emissions).[11] Under such a policy, emission and removal incentives depend on expectations of future stock taxes, and I show in Appendix F that optimal stock taxes are tied to measurements of realized damages, precisely as the damage charges in a carbon share policy. However, moral hazard would be an issue, as emitters will recognize that there is a high likelihood of going out of business over the many decades that carbon persists in the atmosphere. In addition, different agents could expect different stock taxes and thus set their marginal costs of emissions to different levels, violating the equimarginal principle that is a basic tenet of efficiency. In contrast, carbon shares are valuable securities with positive payoffs funded by an upfront deposit, rendering moral hazard irrelevant, and agents base their marginal cost of emissions on the same observed carbon share price, thereby satisfying

[9] Deposit-refund instruments 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.

[10] I 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.

[11] Stock taxes have been proposed in the context of climate change (Lemoine, 2007; Billette de Villemeur and Leroux, 2019), mine remediation (White et al., 2012; Yang and Davis, 2018), and space orbits (Rao et al., 2020).

the equimarginal principle.

Second, a regulator might try to resolve the informational inefficiency by using prediction markets to aggregate information. Specifically, it could report measurements of realized climate damages, run a prediction market on those measurements, and set an emission tax based on that market's revealed predictions.[12] As with carbon shares, such a policy shifts the work of projecting possible future climate change damages from the regulator to markets. And as with carbon shares, such a policy could in principle aggregate dispersed information. However, many prediction markets suffer from low liquidity, and the tax-cum-prediction-market scheme lacks incentives to remove carbon beyond the point at which net emissions reach zero. In contrast, liquidity in a carbon share market is assured because each unit of emissions creates a new stake in the market that its holder will need to value, and the carbon share market does incentivize negative emissions when pessimistic damage estimates make them optimal.

The allocative challenge of procuring negative emissions is rapidly becoming very salient. The Intergovernmental Panel on Climate Change projects that limiting warming to 1.5°C (2°C) would require up to 700 (250) Gt CO_2 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 CO_2. Despite the increasingly prominent discussion and promotion of carbon removal, I know of no prior work on market-based approaches to incentivizing optimal use of these technologies.[13] 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).[14] I here propose a policy

[12] Some prior work focuses on the revelation of beliefs about the magnitude of climate change rather than about climate change damages: Schlenker and Taylor (2021) show that weather derivatives are sensitive to climate model projections, and Hsu (2011) proposes futures on emission taxes set, following McKitrick (2011), according to a predefined function of temperature (see also Aliakbari and McKitrick, 2018). In the latter papers, price discovery provides a signal useful for long-run investment but does not determine emission and carbon removal decisions.

[13] Conventional emission pricing policies can incentivize use of carbon dioxide removal technologies only up to the point at which net emissions are zero. Bednar et al. (2021) propose "carbon removal obligations" that would extend standard cap-and-trade schemes to allow temporarily overshooting longer-run carbon targets.

[14] Bednar 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.

instrument that is revenue-positive for the government and that efficiently adapts the scale of carbon removal to new (and private) information about the cost of climate change.

The informational challenge of aggregating dispersed signals of social costs is new to the environmental economics literature.[15] Formalizing the insight of Hayek (1945), much work since Grossman (1976, 1978) studies financial markets' ability to efficiently aggregate dispersed information. I here study an asset tied to an externality. Payoffs are common across traders because they depend on damage charges that are applied uniformly to all shareholders. Common value models are typical in the financial market literature. My setting differs from most such work by focusing on the role of correlation among partial signals of the common value.[16] I formally show that correlation can be critical to the importance of dispersed information and to the informativeness of carbon share prices. And I quantitatively analyze that informativeness within a calibrated dynamic simulation.

The next section describes the economic and informational environment. Section 3 derives outcomes in the allocatively and informationally efficient benchmark. Section 4 analyzes the allocative and informational efficiency of emission taxes. Section 5 formally defines the carbon share policy and analyzes its allocative and informational efficiency. Section 6 quantitatively explores each type of inefficiency with calibrated numerical examples. The final section concludes. The appendix contains proofs, extensions, and numerical details.

2 Setting

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

[15] Cantillon and Slechten (2018) analyze how emission permit prices aggregate information about whether the permit market will end up long or short. They do not consider private information about the damages due to climate change (i.e., about the optimal level of emissions).

[16] Here, different agents observe signals of different components of a pure common value. The components may be correlated across agents. Some prior work does permit correlated shocks in cases with pure common values (e.g., Ozsoylev and Walden, 2011; Lambert et al., 2018; Lou et al., 2019; Bergemann et al., 2021). Other work permits agents to have correlated private values (e.g., Vives, 2011, 2014; Rostek and Weretka, 2012, 2015; Rahi and Zigrand, 2018).

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

$$Y_{it} = \exp[-\zeta_{it} T_t] \, L_{it} \, y_{it}(e_{it}). \tag{1}$$

Households offer labor $L_{it} \in [0, 1]$ 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 and can evolve with technology. Temperature T_t imposes damages ζ_{it} in sector i at 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.

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 $c_t(R_t)$ as a share of total output (see below) depend on aggregate removal $R_t \triangleq \sum_{i=1}^{N} R_{it}$. The cost function $c_t(\cdot)$ is strictly positive, strictly increasing, and strictly convex and can evolve over time.[18]

Cumulative emissions up to time t are $M_t = M_0 + \sum_{s=0}^{t-1} \left[\sum_{i=1}^{N} e_{is} - R_s \right]$, 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:

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

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

[18] Convexity in removal costs reflects the energetic and land costs of removing carbon from the atmosphere and also the potential scarcity of sites for storing carbon after removal.

with each $\kappa_i > 0$ and $\sum_{i=1}^N \kappa_i = 1$. Substituting from Y_{it},

$$Y_t = e^{-\tilde{\zeta}_t T_t} \prod_{i=1}^N \left(L_{it}\, y_{it}(e_{it})\right)^{\kappa_i},$$

where $\tilde{\zeta}_t \triangleq \sum_{i=1}^N \kappa_i \zeta_{it}$ is the aggregate damage realization. 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:

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

with per-period discount rate $r > 0$ 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.

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 that $\sum_{k=1}^N \kappa_k \bar{\zeta}_k < (r/\alpha)\kappa_i y'_{i0}(0)/y_{i0}(0)$ for some $i \in \{1, ..., N\}$, 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 $\sigma^2 \geq 0$, which represents 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 variance of each λ_{it} is $\omega^2 \geq 0$, which represents 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 $\tau_0^2 > 0$. The correlation between any pair ζ_i and ζ_j (for $i \neq j$) is $\Gamma \in [0, 1)$, a known parameter.[19] 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 uncertain vulnerability to weather is correlated across sectors.

The regulator does not observe sectoral production or input choices.[20] Instead, at the end of time t, the regulator uses observed total output to measure $\tilde{\zeta}_t + \tilde{\lambda}_t$.[21] 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 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.[22]

Within period t, intermediate-good firms first make emission decisions, then markets clear based on realized production, then 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 $\bar{\zeta}_i$. Assuming that the prior variance τ_0^2 and correlation Γ are constant over sectors is merely for ease of exposition, and assuming that Γ is nonnegative simplifies the discussion by focusing on an especially plausible case. In the calibrated numerical simulation (Section 6.2), I permit the prior variance and correlations to be heterogeneous and the correlations to be negative.

[20] The subsequent analysis would be fundamentally similar if we instead permitted the regulator to observe sectoral production and assumed that agents had private information at the subsectoral level.

[21] Firms' equilibrium production choices are independent of T_t (see (A-17) through (A-19)), so the regulator can estimate $\tilde{\zeta}_t$ from a time series of Y_t and T_t. For this same reason, the assumption below that firms in sector i do not know ζ_{it} when choosing their time t emissions is not critical.

[22] Policy is based on imperfectly measured variables in other prominent contexts. For instance, 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).

3 Allocatively and Informationally Efficient Benchmark

Begin by considering the benchmark: emissions and carbon removal that are both allocatively and informationally efficient. Allocative efficiency refers to achieving the welfare-maximizing emission and carbon removal allocations conditional on information, and informational efficiency refers to forming beliefs about the ζ_i that match the posterior of a fictitious regulator who either saw all of the signals in the economy or had them communicated to her.

Define $\hat{E}_t$ as the expectation operator based on all information available throughout the economy up to time t, $\hat{\boldsymbol{\mu}}_t$ and $\hat{\boldsymbol{\Omega}}_t$ as the $N \times 1$ vector of posterior means and the $N \times N$ posterior covariance matrix for the ζ_i based on all information available throughout the economy 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 all information available throughout the economy up to time t. Welfare-maximizing outcomes solve the following Bellman equation:

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

where $\boldsymbol{L}_t$ and $\boldsymbol{e}_t$ 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:

$$\frac{\kappa_i y_{it}'(e_{it})}{y_{it}(e_{it})} \begin{cases} = \frac{1}{r} \alpha \left[\sum_{k=1}^{N} \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] & \text{if } e_{it} > 0 \\ \leq \frac{1}{r} \alpha \left[\sum_{k=1}^{N} \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] & \text{if } e_{it} = 0 \end{cases}, \tag{2}$$

$$\frac{c_t'(R_t)}{1 - c_t(R_t)} \begin{cases} = \frac{1}{r} \alpha \left[\sum_{k=1}^{N} \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] & \text{if } R_t > 0 \\ \geq \frac{1}{r} \alpha \left[\sum_{k=1}^{N} \kappa_k \bar{\zeta}_k + \hat{\mu}_t \right] & \text{if } R_t = 0 \end{cases}. \tag{3}$$

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} either falls or remains zero 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. Following convention, the benchmark of informational efficiency in period t uses the universe of earlier signals, $\{\zeta_{1j} + \lambda_{1j}, ..., \zeta_{Nj} + \lambda_{Nj}, \tilde{\zeta}_j + \tilde{\lambda}_j\}_{j=0}^{t-1}$.

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

$$\begin{aligned}
\hat{\mu}_t =& \hat{Z}_t \left[\frac{1}{t} \sum_{j=0}^{t-1} [\tilde{\zeta}_j + \tilde{\lambda}_j] - \sum_{k=1}^N \kappa_k \bar{\zeta}_k \right] \\
&+ \frac{(1-\hat{Z}_t)(1-\Gamma)\tau_0^2 - \hat{Z}_t \sigma^2/t}{(1-\Gamma)\tau_0^2 + \sigma^2/t + \omega^2/t} \sum_{k=1}^N \kappa_k \left[\frac{1}{t} \sum_{j=0}^{t-1} [\zeta_{kj} + \lambda_{kj}] - \bar{\zeta}_k \right] \\
&+ \frac{\sigma^2/t + (1-\hat{Z}_t)\omega^2/t}{(1-\Gamma)\tau_0^2 + \sigma^2/t + \omega^2/t} \frac{N\Gamma\tau_0^2}{(1-\Gamma)\tau_0^2 + \sigma^2/t + \omega^2/t + N\Gamma\tau_0^2} \frac{1}{N} \sum_{k=1}^N \left[\frac{1}{t} \sum_{j=0}^{t-1} [\zeta_{kj} + \lambda_{kj}] - \bar{\zeta}_k \right].
\end{aligned} \tag{4}$$

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 H, with $\hat{Z}_t$ defined in equation (A-11). □

Proposition 1 shows that informationally efficient beliefs in general use both aggregate and dispersed information. The aggregate signal (first line in (4)) 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.

The following corollary describes some limit cases:

Corollary 1 (Special Cases for Informationally Efficient Beliefs)**.**

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

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

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

$$\hat{\mu}_t = \frac{1}{t}\sum_{j=0}^{t-1}\tilde{\zeta}_j - \sum_{k=1}^{N}\kappa_k\bar{\zeta}_k.$$

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

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

iv 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 I. □

The first part of Corollary 1 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 to produce a perfectly measured version of the aggregate signal (first line). In the presence of correlation ($\Gamma > 0$), it also uses the unweighted average of the sectoral signals (second line): each sector's signals provide the same information about damages in the other sectors, regardless of whether a given sector has large or small value share κ_i.

The remaining three parts of Corollary 1 describe conditions under which efficient updating has no use for the sectoral signals $\zeta_{kj} + \lambda_{kj}$. All three conditions require that there be no measurement error in the aggregate signal ($\tilde{\omega}^2 = 0$). When there is also no stochasticity ($\sigma^2 = 0$), the aggregate signal is without noise and the other signals are superfluous (part ii). When sectoral signals are independent ($\Gamma = 0$), there is no scope for learning through correlation and thus no scope for improving on the perfectly measured aggregate signal (part iii). And when sectors have identical value shares (each $\kappa_i = 1/N$), the weights inside the perfectly measured aggregate signal do not need to be adjusted in order to optimally use the sectoral signals' correlation (part iv).

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 contracting for removal to offset ongoing emissions. The regulator returns any tax revenue to households as lump-sum transfers.

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 J. □

When $\nu_t < \bar{\nu}_t$, the tax is determined by the present value of expected aggregate damages.[23]

[23]The combination of logarithmic utility and the damage specification means that the optimal tax does not depend on total consumption and uncertainty is not priced directly, both 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).

Using that tax, firms' first-order conditions (A-18) and (A-19) become

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

Once we account for the possibility of corner solutions, these would be the conditions for welfare-maximization given in (2) and (3) if $\tilde{\mu}_t = \hat{\mu}_t$.

There are two reasons why a tax may not attain the allocatively and informationally efficient benchmark.[24] 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 allocation from Section 3 would have negative net emissions whereas the regulator's feasible equilibrium has zero net emissions. In order to attain the optimum, the regulator would have to pair the emission tax with a subsidy for carbon removal. Appendix A shows that if this subsidy must be financed by past or present emission tax revenue, then a forward-looking regulator distorts emission and removal decisions in all periods in order to increase tax revenue—and even then may not be able to fund the optimal level of removal if it learns that damages are especially severe. Appendix A shows that this allocative inefficiency arises only when the regulator learns about damages over time.

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, which derive from only a subset of the information in the economy. The following proposition describes the regulator's posterior estimate:

Proposition 3 (Regulator's Beliefs)**.**

$$\tilde{\mu}_t = \frac{(1-\Gamma)\tau_0^2\sum_{i=1}^{N}\kappa_i^2 + \Gamma\tau_0^2}{(1-\Gamma)\tau_0^2\sum_{i=1}^{N}\kappa_i^2 + \Gamma\tau_0^2 + \frac{1}{t}[\tilde{\omega}^2 + \sigma^2\sum_{i=1}^{N}\kappa_i^2]}\frac{1}{t}\sum_{j=0}^{t-1}\left(\tilde{\zeta}_j + \tilde{\lambda}_j - \sum_{k=1}^{N}\kappa_k\bar{\zeta}_k\right). \quad (5)$$

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

[24] A regulator would design a tradeable emission permit ("cap-and-trade") program so that the expected permit price equaled ν_t from Proposition 2. Such a policy differs from the allocatively and informationally efficient benchmark in the same ways as does the emission tax.

using $\sum_{i=1}^{N} \kappa_i = 1$. □

The weight placed on the aggregate measurement in (5) 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 dispersed sectoral measurements, because such beliefs recognize that each sector's measurement contains information about all other sectors. The regulator cannot access these dispersed measurements and thus sees correlation merely as increasing prior uncertainty.

The following corollary delineates conditions under which the beliefs that determine the emission tax are informationally efficient:

Corollary 2 (Informationally Efficient Emission Taxes). *For $t > 0$, $\tilde{\mu}_t = \hat{\mu}_t$ with probability 1 if (i) $\tilde{\omega}^2 = 0$ and either (iia) $\sigma^2 = 0$, (iib) $\Gamma = 0$, or (iic) 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 stochasticity in climate damages ($\sigma^2 = 0$), there is no correlation among sectoral effects ($\Gamma = 0$), or sectors have identical weights in production (each $\kappa_i = 1/N$). These were the special cases from Corollary 1 in which informationally efficient beliefs do not require the dispersed sectoral signals. Otherwise Proposition 1 showed that informationally beliefs generally do use dispersed signals unavailable to the regulator. Because it is unlikely that any of the conditions in Corollary 2 hold in reality, an actual regulator's beliefs are not likely to be informationally efficient.

5 Regulation by Carbon Shares

We have seen that conventional emission pricing can suffer both allocative and informational inefficiencies when the regulator learns about damages over time. I now describe a policy designed to improve along each margin.

This new type of policy requires each emitter to post a deposit $D \geq 0$ per unit of emissions. Choosing the deposit is equivalent to choosing a parameter $\bar{\mu}$ that defines implied per-period climate damages:[25]

$$D \triangleq \frac{1}{r} C_0 \, \alpha \left[\sum_{k=1}^{N} \kappa_k \bar{\zeta}_k + \bar{\mu} \right]. \tag{6}$$

The regulator invests the deposit and earns interest at rate r. In exchange for the deposit, the emitter receives a transferable security that is attached to the unit of carbon emitted. I refer to the security 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 time t measured marginal damage from carbon emissions and the per-period damages implied by the deposit:

$$\Delta_t \triangleq C_0 \, \alpha \, \min \left\{ \tilde{\zeta}_t + \tilde{\lambda}_t, \, \sum_{k=1}^{N} \kappa_k \bar{\zeta}_k + \bar{\mu} \right\}. \tag{7}$$

Damage charges are returned lump sum to households.[26] The policymaker refunds to carbon shareholders the difference between the damage charge and the per-period damages implied by the deposit:[27]

$$\begin{aligned} d_t \triangleq & \, r \, D - \Delta_t \\ = & \, C_0 \, \alpha \max \left\{ 0, \, \sum_{k=1}^{N} \kappa_k \bar{\zeta}_k + \bar{\mu} - (\tilde{\zeta}_t + \tilde{\lambda}_t) \right\}. \end{aligned} \tag{8}$$

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

[25]The deposit equals the emission tax ν_t from Proposition 2 when $\bar{\mu} = \tilde{\mu}_t$ and $\nu_t < \bar{\nu}_t$.

[26]In a second-best setting, revenue from damage charges could be used to offset revenue from distortionary taxes. Future work could consider how to adapt both the damage charges and the deposit (see Fullerton and Wolverton, 2000, 2005).

[27]The optimal refund will be common to all outstanding shares because marginal damage is the same for all units of carbon (see equation (2)). In other applications, shares could be vintaged by date or location of emission.

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:

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

When deposits are invested at the market interest rate, the carbon share policy is revenue-neutral if climate change damages completely fail to materialize and raises revenue via nonzero damage charges otherwise.

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.[28] Leaving carbon in the atmosphere for one more period means losing the damage charge Δ_{t+1}.

Figure 1 provides an example of cashflows 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 another party for the market price q_{t+1}. That 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 shareholder in this example does decide to remove the underlying unit of carbon from the atmosphere, which costs p^R_{t+s}. As a result, the regulator retires the carbon share in time $t+s$ and pays the shareholder $(1+r)D - \Delta_{t+s}$.

I assume all agents discount at rate r (but see footnote 30 below). Use $\breve{W}$ and $\breve{C}_s$ to denote welfare and consumption along a realized trajectory under the carbon share policy defined above, and use $\hat{W}$ and $\hat{C}_s$ to denote welfare and consumption along a realized trajectory

[28] I have treated removal of carbon as permanent, but the policy can be extended to account for the possibility of imperfect removal (e.g., from a leaking underground reservoir or a forest that may burn): instead of refunding $(1+r)D - \Delta_t$ at the time of removal, the regulator would refund $d_{t+s} + \Delta_{t+s}$ as removal is verified in each subsequent period $t+s$. This type of policy would give shareholders incentives to monitor and maintain removal.

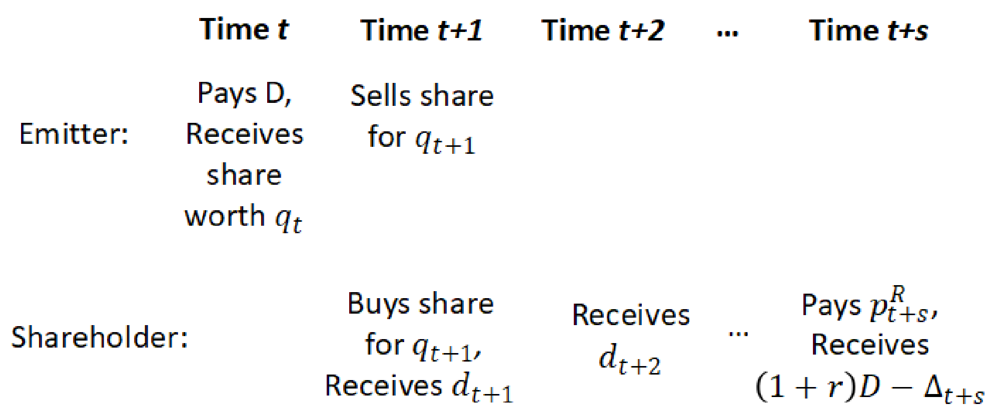

Figure 1: Example of cashflows over the life of a carbon share. Here the share is attached to a unit of time t emissions, the emitter decides to sell the share at time $t+1$, and the new shareholder decides to remove the underlying unit of carbon from the atmosphere at time $t+s$.

under the allocatively and informationally efficient benchmark:

$$\breve{W} = \sum_{s=0}^{\infty} \frac{1}{(1+r)^s} u(\breve{C}_s), \qquad \hat{W} = \sum_{s=0}^{\infty} \frac{1}{(1+r)^s} u(\hat{C}_s).$$

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

$$\breve{L} = E_0 \left[\hat{W} - \breve{W} \middle| \boldsymbol{\zeta} \right],$$

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

5.1 Improved Allocative Efficiency

Temporarily assume that all traders see all of the signals in the economy, so that informational efficiency is not at issue:

Assumption 1 (Common Information)**.** *At the beginning of period t, each trader's information set contains $\{\zeta_{1s} + \lambda_{1s}, ..., \zeta_{Ns} + \lambda_{Ns}, \tilde{\zeta}_s + \tilde{\lambda}_s\}_{s=0}^{t-1}$, and within period t, each trader observes $\{\zeta_{1t} + \lambda_{1t}, ..., \zeta_{Nt} + \lambda_{Nt}\}$.*

We saw in Section 4 that an emission tax can be allocatively inefficient when new information about damages might justify negative emissions. I will show that carbon shares improve outcomes when the deposit D is sufficiently large. Appendix B provides a parallel analysis within a transparent simplified model.

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 (which may not be true of the equilibrium share price to be analyzed in Section 5.2). The following lemma establishes the equilibrium value of the carbon share:

Lemma 1 (Carbon Share Value)**.** *Under Assumption 1,*

$$\hat{q}_t = \sum_{j=0}^{\infty} \frac{1}{(1+r)^j} \hat{E}_t[d_{t+j}] \geq 0. \tag{9}$$

Proof. See Appendix K. □

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 derives from the possibility that damages will not be as bad as implied by $\bar{\mu}$. At the time of emission, a firm's net outlays per unit of non-abated emissions are $D - (\hat{q}_t - \hat{E}_t[d_t]) \in [0, D]$, recognizing that the emitter does not receive refunds d_t.[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:

Assumption 2 (Pre-Policy Carbon Should Not Be Removed)**.** $\hat{R}_t \leq M_t - M_0$ *for all* $t \geq 0$.

This assumption saves us from having to consider how the timing of a carbon share policy's introduction limits the carbon removal it can incentivize. Even the highest-removal scenarios for the coming century do not project bringing either 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.

[29] If future damage measurements 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 damage measurements 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 $\hat{q}_t - \hat{E}_t[d_t] \in [0, D]$.

The following proposition relates the period t loss to the deposits required at earlier times:[30]

Proposition 4 (Efficiency Conditional on Information)**.** *Let Assumptions 1 and 2 hold. Then* $\breve{L} \to 0$ *as* $D \to \infty$*.*

Proof. See Appendix L. □

The proposition establishes that the carbon share policy achieves the allocatively efficient 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 (7), 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.[31] As the deposit becomes large, the carbon share policy attains the emission reduction incentives of an emission tax while improving on the emission tax by approaching allocative efficiency.[32]

[30]One might be concerned that market actors and the regulator would apply different discount rates. Appendix G shows that, because traders care about forecasting whatever damage charge calculation the regulator will use, the regulator can induce traders to apply its own favored discount rate to future climate change damages by rescaling its calculated damage charges.

[31]I 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.

[32]A carbon share policy is different even from an emission tax policy that attaches a distinct lockbox to each unit of emissions from which it later pays for carbon removal. In both cases, the most a regulator can induce firms to spend on removing a particular unit of carbon is constrained by the funds the regulator collected at the time of emission. However, increasing the carbon share's deposit increases the efficiency of its emission price, whereas increasing the emission tax overincentivizes emission reductions.

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).[33] However, whereas firms could avoid a carbon stock tax by declaring bankruptcy, emitters here pay the entire penalty up front. Further, carbon shares are valuable assets that investors want to hold, financed at the time of emission by the deposit. 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. Carbon shares therefore avoid judgment-proofness problems that would bedevil stock taxes. (Moreover, a stock tax would fail to coordinate emission incentives across firms because it would lack the ability to aggregate information analyzed below 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-25), 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 thick, a carbon share policy need not be any more financially challenging than a conventional carbon emission tax.[34]

But one might still wonder about the scale of the deposit. If the deposit is not sufficiently large, then the highest potential damage charges are truncated by the constraints imposed by the deposit, which reduces the expected damage charges that firms use to guide emission and removal decisions. The possibility of hitting a deposit's constraints therefore increases emissions and reduces removal. Ex post, the prices of 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 C 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 (7)). In this calibration, expected damages imply a

[33] Appendix F shows that the optimal time t stock tax would be $\lim_{\bar{\mu}\to\infty} \Delta_t$.

[34] Gross 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.

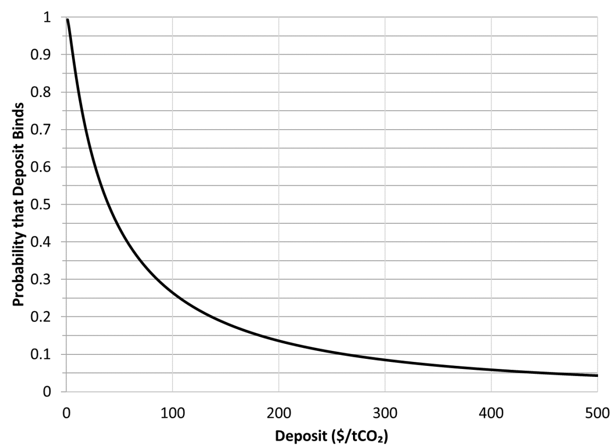

Figure 2: The probability that deposits of various sizes would bind, based on the calibration in Appendix C and using $\tilde{\omega}^2 = \sigma^2 = 0$.

tax of \$118 per tCO_2. Figure 2 shows that a deposit roughly twice as large (\$250 per tCO_2) would suffice in all but the worst 10% of cases, and a deposit just over three times as large (\$400 per tCO_2) 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 Informational Efficiency

I now investigate how the market for carbon shares aggregates dispersed information about climate change damages. I condition results on arbitrarily large D so that, from Section 5.1, allocative efficiency is not at issue.

A continuum of traders of measure $n_i > 0$ 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, the posterior at the beginning of time t assigns variance $\breve{\tau}_t^2$ to each ζ_k, and the posterior mean of $\sum_{k=1}^{N} \kappa_k \zeta_k$ at the beginning of time t is $\breve{\mu}_t$. The price of carbon shares at the beginning of time t is q_t. Firms make emission and removal decisions based on this price and households consume accordingly. Subsequently, traders attached to sector i measure $\zeta_{it} + \lambda_{it}$. They trade carbon shares based on this differentiated information. The market clears at price $\breve{q}_t^*$. 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 of period t, and issues new shares to firms based on period t emissions.[35]

At the beginning of time t, each trader of type i has share holdings z_{it} and wealth w_{it}. The z_{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 $t+1$:

$$\max_{X_{it}} \breve{E}_t \left[-\exp\left\{ -A_i \left((1+r)(w_{it} + (X_{it} + y_{it})d_t - X_{it}\breve{q}_t) + (y_{it} + X_{it})q_{t+1} \right) \right\} \Big| \, \zeta_{it} + \lambda_{it}, \breve{q}_t \right], \tag{10}$$

with $A_i > 0$ the coefficient of absolute risk aversion among traders attached to sector i and $\breve{E}_t$ 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. Log utility is a standard in climate change economics, but exponential utility is critical for tractability within the literature on asymmetric information in asset markets.[36] Foregoing analysis would largely survive if the representative household had exponential utility.[37]

In a rational expectations equilibrium, markets clear with traders inferring from prices whatever information they can and maximizing utility conditional on that information. This equilibrium is "fully revealing" if the carbon share price $\breve{q}_t^*$ reveals the same information about $\tilde{\zeta}_t$ as would observing the N disentangled signals $\{\zeta_{kt} + \lambda_{kt}\}_{k=1}^N$. The following proposition establishes properties of this equilibrium:

Proposition 5 (Efficient Equilibrium)**.** *Let Assumption 2 hold. A fully revealing rational*

[35] Traders 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 $y'_{it}(0) = 0$, in which case sector i would have zero time t emissions but could be affected by damages. The creation of new shares can be handled by including time t emitters in the set of time $t+1$ traders.

[36] Exponential utility yields linear asset demand functions that are independent of wealth and amenable to aggregation.

[37] Specifically, if we give the representative household in Section 2.1 exponential utility over consumption with $A > 0$ the coefficient of absolute risk aversion and also let damages be additive rather than multiplicative, then the learning dynamics are unchanged and, as $A \to 0$ (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.

expectations equilibrium with $\breve{\mu}_t = \hat{\mu}_t$ *and* $\breve{L} = 0$ *exists as* $\sigma^2 \to 0$ *and* $D \to \infty$.

Proof. See Appendix N.
Sketch: Shows that a fully revealing price both clears the market in period 0 and, as $\sigma^2 \to 0$, generates informationally efficient beliefs at the beginning of period 1. Shows analogous results for all t by induction. Result then follows from Proposition 4.

□

A fully revealing rational expectations equilibrium exists and recovers informational efficiency as the fundamental shocks ϵ become small and the deposit becomes large. Informationally efficient updating sees the stochastic shocks to each ζ_{it} as noise to be filtered out for predicting $\sum_{i-1}^{N} \kappa_i \zeta_i$, but it is not noise to time t traders attempting to predict the coming aggregate measurement $\tilde{\zeta}_t + \tilde{\lambda}_t$. These two predictions converge as $\sigma^2 \to 0$ because the remaining shocks in the observed $\zeta_{it} + \lambda_{it}$ become pure noise in either case. 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 approaches the allocatively and informationally efficient benchmark, even though the regulator has the exact same information as in Section 4.

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.[38] In particular, if the carbon share price is a sufficient statistic for all information in the economy, then traders should ignore their private information when updating beliefs, in which case it is unclear how their private information ends up being summarized by the equilibrium price. This is the Grossman-Stiglitz paradox.

I therefore also study a linear Bayesian Nash equilibrium in demand functions. Traders submit demand functions and the market-clearing price $\breve{q}_t^*$ sets aggregate net demand $\sum_{i=1}^{N} X_{it}^*$ to zero. Noise traders make the observed price $\breve{q}_t = \breve{q}_t^* + \theta_t$, where θ_t is a mean-zero, independently and identically distributed, normal random variable, with variance $\Theta^2 > 0$.[39] Traders

[38] The foregoing analysis postulates a price and tests whether it meets market-clearing and rational expectations criteria. The coming analysis postulates a mechanism that converts strategies into a market-clearing price, which meets rational expectations criteria when strategies are in equilibrium. See Kyle (1989) and Vives (2011, 2014).

[39] Noise in price or, equivalently, supply has long been recognized as critical for the existence of partially

treat $\breve{q}_t^*$ 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. Each X_{it}^* solves (10) given Bayesian beliefs updated from all share prices and public measurements prior to time t, from observations of sector i's private time t signals $\zeta_{it}+\lambda_{it}$, and from the time t share price $\breve{q}_t$. In equilibrium, the Bayesian updating process treats the share price as containing exactly the information that the market-clearing price does contain. The equilibrium is symmetric in terms of strategies defined over $\zeta_{it}+\lambda_{it}$ and $\breve{q}_t$.[40]

The following proposition characterizes equilibrium emission incentives.

Proposition 6 (Implementable Equilibrium)**.** *There exist scalars* $\{\tilde{\pi}_{kt}, \breve{\pi}_{kt}, \breve{\chi}_k, \breve{\kappa}_{1k}, ..., \breve{\kappa}_{Nk}\}_{k=0}^{t-1}$ *such that, as* $D \to \infty$*, the linear Bayesian equilibrium sets the marginal cost of emissions and marginal benefit of carbon removal equal to*

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

with

$$\begin{aligned}
\breve{\mu}_t =& \sum_{k=0}^{t-1} \tilde{\pi}_{kt} \left(\tilde{\zeta}_k + \tilde{\lambda}_k - \sum_{j=1}^{N} \kappa_j \bar{\zeta}_j \right) \\
&+ \sum_{k=0}^{t-1} \breve{\pi}_{kt} \underbrace{\left[\frac{(1-\Gamma)\breve{\tau}_k^2 + \sigma^2}{\breve{\tau}_k^2 + \sigma^2 + \omega^2} \sum_{i=1}^{N} \breve{\kappa}_{ik} \kappa_i \left(\zeta_{ik} + \lambda_{ik} - \bar{\zeta}_i \right) + \frac{N \Gamma \breve{\tau}_k^2}{\breve{\tau}_k^2 + \sigma^2 + \omega^2} \frac{1}{N} \sum_{i=1}^{N} \breve{\kappa}_{ik} (\zeta_{ik} + \lambda_{ik} - \bar{\zeta}_i) \right]}_{\text{from } \breve{q}_k} \\
&- \sum_{k=0}^{t-1} \breve{\pi}_{kt} \frac{r}{C_0 \alpha (\breve{\chi}_k + r)} \theta_k. \qquad (11)
\end{aligned}$$

Proof. See Appendix O.

Sketch: Use traders' first-order conditions under a linear conjecture for $\breve{q}_t$ to determine demand for carbon shares in each sector. Because normal-normal Bayesian updating implies that $\breve{\mu}_{t+1}$ is a linear function of $\breve{\mu}_t$, $\breve{q}_t$, and $\tilde{\zeta}_t + \tilde{\lambda}_t$, so too is $\breve{E}_t[q_{t+1} | \zeta_{it} + \lambda_{it}, \breve{q}_t]$. The proof then constructs a signal $\tilde{q}_t$ of aggregate damages implied by $\breve{q}_t$. By normal-normal Bayesian

revealing equilibria when traders are small and payoffs are pure common values (e.g., Hellwig, 1980; Grossman and Stiglitz, 1980; Diamond and Verrecchia, 1981).

[40] I generalize the equilibrium to be asymmetric for the numerical application in Section 6.2.

updating, each sector i trader's posterior mean for $\tilde{\zeta}_t + \tilde{\lambda}_t$ is a linear function of $\breve{\mu}_t$, $\tilde{q}_t$, and $\zeta_{it} + \lambda_{it}$. We can thus express the implied signal $\tilde{q}_t$ as an unknown linear function of sectoral signals. The projection theorem yields each type of trader's posterior mean for aggregate damages conditional on the observed $\zeta_{it} + \lambda_{it}$ and on the $\tilde{q}_t$ implied by the observed $\breve{q}_t$. Matching coefficients and applying Brouwer's fixed-point theorem yields posterior beliefs $\breve{E}_t[\tilde{\zeta}_t + \tilde{\lambda}_t|\zeta_{it} + \lambda_{it}, \breve{q}_t]$ that are self-fulfilling via the price and yields the market-clearing price $\breve{q}_t^*$, both as functions of the unknown coefficients that determine $\breve{\mu}_{t+1}$. Beliefs $\breve{\mu}_{t+1}$ follow from the projection theorem and the solution for earlier carbon share prices. □

A carbon share policy acts like imposing an emission tax based on beliefs $\breve{\mu}_t$. The first line in (11) determines the weight placed by $\breve{\mu}_t$ on previous periods' aggregate measurements of damages. The second line describes how agents learn from the past prices $\breve{q}_k$ of carbon shares (use equation (A-52) in (A-36)). 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 unweighted average of the sectoral 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 in Proposition 1 but were missing from the regulator's beliefs in Proposition 3.

The following corollary establishes the conditions under which past carbon share prices provide useful information:

Corollary 3 (Informative Carbon Share Prices)**.** *Consider the* $\{\breve{\pi}_{0t}, ..., \breve{\pi}_{(t-1)t}\}$ *defined in Proposition 6. These are each equal to zero if (i)* $\tilde{\omega}^2 = 0$ *and either (iia)* $\sigma^2 = 0$*, (iib)* $\Gamma = 0$*, or (iic) each* $\kappa_i = 1/N$*. Otherwise* $\breve{\pi}_{(t-1)t} > 0$*.*

Proof. See Appendix P. □

Traders' posterior beliefs do not rely on the sectoral signals embedded in past share prices ($\breve{\pi}_{kt} = 0$) under the same conditions that informationally efficient beliefs rely only on the aggregate signals (see Corollary 1), which are also the conditions under which an emission tax-wielding regulator's beliefs are informationally efficient (see Corollary 2). In other cases, traders and firms do learn from carbon share prices. Because these beliefs determine emission reduction incentives and use information unavailable to the regulator, the carbon share policy is generally more informationally efficient than an emission tax policy.

However, the carbon share policy does not generally achieve full informational efficiency. There are several differences with respect to the informationally efficient beliefs described in Proposition 1. First, noise traders mechanically induce randomness in the share price that hinders learning (the θ_k in the third line of (11)). Second, the $\breve{\kappa}_{ik}$ downweight information from the carbon share price (in the second line of (11)). The following corollary examines that downward adjustment in more detail:

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

$$\breve{\kappa}_{ik} = \breve{\kappa}_{ik}^{BS} \breve{\kappa}_{ik}^{RA},$$

where $\breve{\kappa}_{ik}^{BS} \in (0, 1)$*,* $\breve{\kappa}_{ik}^{RA} \in (0, 1)$*, and*

i $\lim_{\Theta^2 \to \infty} \breve{\kappa}_{ik}^{BS} = 1$.

ii $\lim_{\Theta^2 \to 0} \breve{\kappa}_{ik}^{BS} = 0$.

iii $\lim_{A_i \to 0} \breve{\kappa}_{ik}^{RA} = 1$*, and* $\lim_{A_i \to 0} \breve{\kappa}_{jk}^{RA} = 0$ *for all* $j \neq i$.

iv $\lim_{A_i \to \infty} \breve{\kappa}_{ik}^{RA} = 0$.

v If $\kappa_i = 1/N$ *for all* $i \in \{1, ..., N\}$*, then* $\breve{\kappa}_{ik}^{RA} = n_i A_i^{-1} / \sum_{j=1}^{N} n_j A_j^{-1}$.

vi Without loss of generality, order sectors by increasing κ_i*. If the sequence* $\{n_1 A_1^{-1}, ..., n_N A_N^{-1}\}$ *is weak monotone increasing, then as* $\Theta^2 \to \infty$*, the sequence* $\{\breve{\kappa}_{1k}^{RA}, ..., \breve{\kappa}_{Nk}^{RA}\}$ *is weak monotone increasing, with* $\breve{\kappa}_{1k}^{RA} \leq n_1 A_1^{-1} / \sum_{j=1}^{N} n_j A_j^{-1}$ *and* $\breve{\kappa}_{Nk}^{RA} \geq n_N A_N^{-1} / \sum_{j=1}^{N} n_j A_j^{-1}$*. The two inequalities are strict if, in addition,* $\kappa_1 < 1/N$.

Proof. See Appendix Q. □

The downward adjustments due to the $\breve{\kappa}_{ik}$ can be decomposed into two terms. The first term, labeled $\breve{\kappa}_{ik}^{BS}$, reflects how traders in each sector shade their bids to reflect the possibility of a winner's curse in this market with a pure common value: from equations (A-42) and (A-51), they underweight their own signals when forming beliefs. This downward adjustment vanishes as noise traders fully determine the price (part i of the corollary) but prevents carbon share prices from aggregating any private information as noise traders become irrelevant (part ii). The latter result 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. Thus there is a tension between minimizing the consequences of the third line in (11) and minimizing the consequences of bid shading for the $\breve{\kappa}_{ik}$.

The second term, labeled $\breve{\kappa}_{ik}^{RA}$, reflects how risk-averse traders' demand for carbon shares decreases in the variance of the returns they will earn: from equations (A-32) and (A-34), the market price is more sensitive to traders who have more precise beliefs and less aversion to risk. As traders in some sector approach risk-neutrality, their demand becomes perfectly elastic and the carbon share price incorporates only their private information (part iii of Corollary 4), but as traders in some sector become infinitely risk-averse, the carbon share price incorporates information only from the other sectors (part iv). When sectors have symmetric roles in final good production, perceived risk is the same in each sector and it is the traders with less risk aversion who reveal more information to the market (part v). When sectors are not symmetric and bid-shading is small, 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 (part vi).

In sum, noise traders, bid shading, and traders' risk aversion each prevent the equilibrium carbon share price from being fully revealing,[41] but that equilibrium does aggregate information in a fashion that is useful to traders and is analogous to the aggregation performed within informationally efficient updating. By using markets to perform price discovery for social cost, a carbon share policy goes some distance towards resolving the difficulties that a regulator wielding an emission tax faces in bringing the economy's dispersed information about social costs into the private cost of emitting.

[41]Moreover, even if it were somehow the case that each $\breve{\kappa}_{ik} = 1$ with $\Theta^2 = 0$, the carbon share price would still fail to aggregate information in the most efficient manner possible. In particular, stochasticity governed by σ^2 makes traders' predictions target a slightly different object than would be necessary for welfare maximization, as described following Proposition 5.

6 Calibrated Numerical Evaluation

I now develop calibrated examples in order to understand the scope that carbon shares have to improve allocative and informational efficiency.

6.1 Allocative Efficiency

I first assess the potential importance of being able to fund negative emissions. As in Section 5.1, I assume all agents can see all signals in the economy, so that there are no informational losses.

I extend the DICE-2016R climate-economy model of Nordhaus (2017) to allow for uncertainty about damages from warming.[42] Prior to 2065, the damage parameter is fixed and negative emissions are not feasible. In 2065, a random component of the damage parameter is realized and negative emissions become feasible. I calibrate the variance of damages to the expert survey of year 2066 losses from climate change in Pindyck (2019), following the implementation in Lemoine (2021a) that adjusts for uncertainty about warming. In one case ("DICE Damages"), I fix the mean of the distribution to match damages in DICE-2016R, and in the other case ("Expert Damages"), I allow the mean to also be determined by the expert survey. The latter case implies more severe losses from warming.[43]

I consider three types of policy environments. First, a simple emission tax cannot incentivize any negative emissions. Second, an emission tax with a lockbox can fund negative emissions by using revenue accumulated from emission taxes over time (see Section 4). Third, a carbon share policy with a sufficiently large deposit can fund negative emissions at the optimal level (see Section 5.1).

The top panel of Table 1 reports the balanced growth equivalent (BGE) increase in consumption from implementing policy (Mirrlees and Stern, 1972), relative to a case without any emission controls. The carbon share policy is far more valuable in the calibration to expert damages, providing expected benefits equivalent to a permanent 40% increase in consumption as opposed to a permanent 1.5% increase in consumption in the calibration

[42] The abatement cost function in DICE-2016R accounts for carbon dioxide removal technologies. I maintain this cost function but remove the exogenous cap on the magnitude of negative emissions. DICE-2016R assumes that negative emissions are funded in a non-distortionary fashion.

[43] I update the carbon and climate modules in accord with recommendations in Dietz et al. (2021). See Lemoine (2020) for the full equations and parameterization and for optimal trajectories in a deterministic version of each calibration.

Table 1: Balanced growth equivalent (BGE) gain from policy options, and reduction in BGE relative to efficient benchmark, in percentage points (pp).

	Expert Damages		DICE Damages	
	BGE (%)	Loss (pp)	BGE (%)	Loss (pp)
In 2015				
Carbon Share	40.3	-	1.51	-
Emission Tax with Lockbox	38.1	2.2	1.40	0.11
Emission Tax	37.7	2.7	1.38	0.14
After high-damage draw in 2065				
Carbon Share	92.6	-	7.34	-
Emission Tax with Lockbox	85.9	6.7	6.96	0.39
Emission Tax	85.6	7.0	6.75	0.59

BGE is relative to a case with abatement fixed at zero (but savings optimized). The BGE translates changes in welfare into the constant relative difference in consumption between two counterfactual consumption trajectories that grow at the same constant rate (Mirrlees and Stern, 1972).

with DICE damages. The second row constrains negative emissions to what can be funded out of cumulative emission tax revenue. BGE falls by 2.2 (0.11) percentage points in the case of expert (DICE) damages. The third row constrains negative emissions to what can be funded out of current emission tax revenue, which means it cannot fund any net negative emissions. Relative to the carbon share policy, BGE is now 2.7 (0.14) percentage points lower in the case of expert (DICE) damages.[44] By enabling negative emissions, a carbon share policy makes adopting climate policy more beneficial.

The lower panel of Table 1 reports BGE when damages happen to be above average.[45] Policy becomes much more valuable when damages are unexpectedly high, with BGE increasing to 93% (7.3%) in the case of expert (DICE) damages. That BGE is also much more sensitive to the policy instrument, as the reduction in BGE from using an emission tax approximately triples. The efficient benchmark increases use of carbon removal when damages are high, which makes the emission tax policy's allocative inefficiency especially

[44]The regulator chooses policy in full knowledge of each constraint. To correct for the chance that the negative emission constraint will bind, the regulator increases the initial period's emission price from \$191 (\$22.7) per tCO_2 to \$233 (\$24.5) per tCO_2 when the regulator knows she will not be able to obtain negative emissions and increases it a bit more to \$235 (\$25.5) per tCO_2 when she knows she can use emission tax revenue to fund later negative emissions, with expert (DICE) damages.

[45]It evaluates welfare after damages are realized at the second-highest quadrature node of the damage distribution.

severe ex post. That allocative inefficiency is only slightly reduced by retaining emission tax revenue to fund a subsidy for carbon removal, because the required subsidies quickly exhaust the saved revenue and also because early periods' emission taxes are set without knowing that damages will turn out to be high.

6.2 Informational Efficiency

I next assess the potential for a carbon share policy to improve informational efficiency relative to an emission tax. As in Section 5.2, I treat D as arbitrarily large, so that there is no allocative inefficiency.

I regress the log change in county-level output for each of 12 U.S. industries on the log change in county temperature over 2002–2019 (see Appendix D). I estimate separate coefficients by industry, I include industry-county and industry-year fixed effects, and I include quadratic time trends at the industry-state level. The estimated covariance matrix for the effects of temperature on each industry becomes the prior covariance matrix for the ζ_i. Generalizing the main analysis, I permit the prior variance τ_{i0}^2 to vary by industry i and the prior correlation Γ_{ik} to vary by industries i and k. I calibrate the κ_i to industry i's average output over the sample divided by average total output over the sample. Setting $\sigma_i^2 = \eta \Xi_i$, I calibrate the Ξ_i to the variance of the residuals for industry i. Fraction η of the residual variation in output growth is assumed due to variation in its relation to temperature (i.e., to variation in the ζ_{it}). I study cases with η equal to 0.1, 0.01, and 0.001.

Appendix D reports that many of the estimated correlations between the ζ_i are large, which will tend to reduce the informational efficiency of regulation by emission taxes. It also reports that the κ_i and σ_i^2 are negatively correlated, which will tend to increase the informational efficiency of regulation by emission taxes.

In the base specification, I set the variance of aggregate measurement error to be ten times the variance in the aggregate measurement induced by regression residuals ($\tilde{\omega}^2 = 10 \sum_{i=1}^{N} (\kappa_i)^2 \Xi_i$) and I set the variance of noise traders to be a small fraction of that aggregate variance ($\Theta^2 = 10^{-12} \sum_{i=1}^{N} (\kappa_i)^2 \Xi_i$). I do not model sectoral measurement error ($\omega^2 = 0$). Following Choi et al. (2007), I set $A_i = 0.043$. I normalize the mass of traders to unity in each sector ($n_i = 1$). Carbon tax parameters follow the deposit calculations in Section 5.1.

Appendix E describes how to solve the linear Bayesian equilibrium, which is here asymmetric due to heterogeneity in prior variances, prior correlations, and the variances of sectoral shocks. Solving for the equilibrium share price in a given period t requires solving a system of 38 equations in 38 unknowns. With a solution for period t in hand, I update variance matrices and then solve the period $t+1$ equilibrium. Once I have the equilibrium parameters in each period, I simulate the evolution of beliefs along 10,000 trajectories of the ϵ_{it}, λ_{it}, $\tilde{\lambda}_t$, and θ_t, with each trajectory using its own draw of the ζ_i from the time 0 prior distribution.

Figure 3 plots the evolution of the mean squared error (MSE) between emission prices of interest and the full-information emission price implied by the draw of the ζ_i. In each panel, the informationally efficient tax uses all sectoral and aggregate signals and therefore has the lowest MSE. A regulator setting a carbon tax sees only the public aggregate information and therefore has the highest MSE. The carbon share price aggregates both public aggregate information and private sectoral information, albeit partially and noisily. It has an intermediate MSE.

The top row plots the base specification described above. The columns vary η. In all panels, a regulator using a carbon tax policy learns very slowly, which mimics the slow progress over the last 20 years in reaching firm agreement about the correct level for a carbon tax. In the left panel, the percentage of variation due to temperature sensitivity (η) is 10%. Informationally efficient beliefs steadily reduce their MSE as time passes and observations accumulate, and MSE falls nearly as quickly under the carbon share policy. But beliefs retain substantial error after 50 periods in either case.

The middle panel in the top row reduces η to 1%. Private information becomes more important, which permits informationally efficient beliefs and the carbon share policy to both substantially reduce MSE even in the first period. The right panel further reduces η to 0.1%. In this case, nearly all noise is introduced at the aggregate level, through measurement error. The carbon share policy is nearly informationally efficient, as the share price nearly perfectly aggregates traders' private information. Moreover, the carbon share price rapidly converges to true beliefs.

The middle row increases the role of noise traders, by increasing Θ^2 by a factor of 10^6 (to $10^{-6}\sum_{i=1}^{N}(\kappa_i)^2\Xi_i$). This change affects only learning under the carbon share policy. It increases the willingness of traders to trade on their information, which makes the carbon share price more informative, but it also directly introduces noise into the carbon share price, which makes it less informative. On net, the latter effect dominates, slowing learning under the carbon share policy. The final column shows that the noise induced by traders imposes

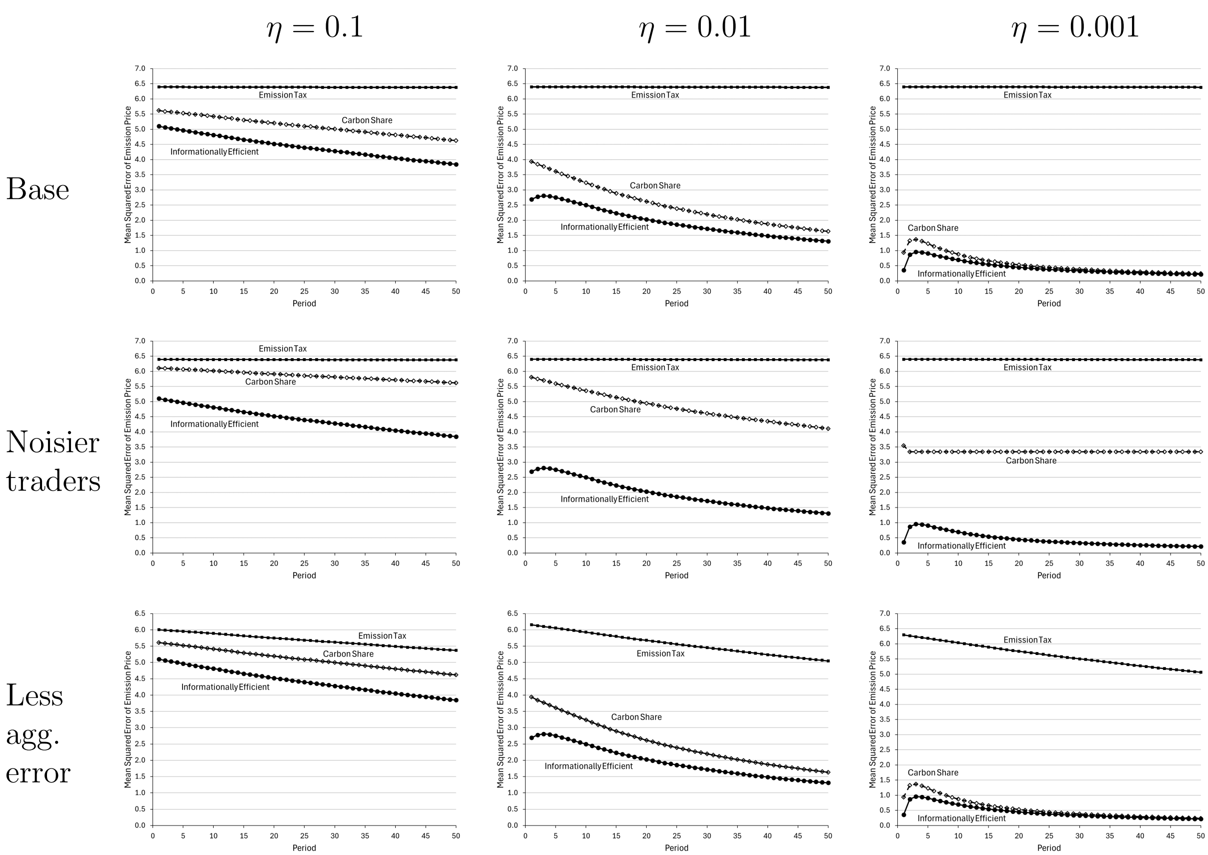

Figure 3: The evolution of the mean squared error in the emission price chosen by a fictitious welfare-maximizing regulator who observes all information available throughout the economy before setting an emission price (circles), chosen by a welfare-maximizing regulator who sets an emission tax on the basis of public aggregate information (squares), and implied by carbon shares that noisily and partially aggregate private sectoral information (diamonds). The columns vary the percentage of the residual variation in an industry's output growth attributed to variation in its relationship to temperature, from 10% (left) to 1% (middle) to 0.1% (right). The rows vary the parameterization of $\tilde{\omega}^2$ and Θ^2 (see text).

a limit on how much traders can learn from sectoral signals under the carbon share policy. Nonetheless, learning is still faster under the carbon share policy than under the carbon tax policy.

The bottom row reduces aggregate measurement error to be comparable to the variance of sectoral shocks, by reducing $\tilde{\omega}^2$ by a factor of 100 (to $0.1\sum_{i=1}^{N}(\kappa_i)^2\Xi_i$). Reducing aggregate measurement error speeds up learning under the carbon tax policy. However, learning under the carbon share policy remains faster, driven by how the share price aggregates private sectoral observations. The carbon share policy retains much scope to improve the informational efficiency of emission pricing.

7 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 generate an allocative inefficiency through their inability to fund optimal cleanup and generate an informational inefficiency through their inability to aggregate dispersed information about social costs. I have shown that a new market-based instrument, which I call "carbon shares" in the context of climate change, achieves allocative and informational efficiency in a limit case and improves on emission taxes in calibrated simulations.

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 or cap. Here, however, the regulator need only measure damage as it is realized and determine the deposit based on approximate worst-case outcomes. A carbon share policy shifts the burden of projecting possible future damages from the regulator to market traders. These traders base their damage estimates on information produced by the regulator, the observed prices of carbon shares, and their own private information. This type of belief updating is a common task in markets.

The benefits of a carbon share policy consist in being able to incentivize the removal of carbon in a revenue-neutral fashion and in learning faster what the true social cost of carbon may be. Both of these benefits are larger when a carbon share policy is begun sooner, before society emits the carbon it may end up either removing or learning to regret. The replacement of an emission tax with a carbon share policy cannot wait until emissions have already been zeroed out.

The present analysis has focused on efficiency in aggregating and acting on the information about climate damages dispersed throughout society, holding that information fixed. Information about climate damages is likely to be more valuable to private actors under a carbon share policy than under an emission tax or a cap-and-trade program. One might thus expect a carbon share policy to spur private investment in information production, so that a carbon share policy not only aggregates the information already dispersed throughout the economy but also improves that information. Future work should compare the incentives to produce information in the various policy environments.

References

Aliakbari, Elmira and Ross McKitrick (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.

Bergemann, Dirk, Tibor Heumann, and Stephen Morris (2021) "Information, market power, and price volatility," *The RAND Journal of Economics*, Vol. 52, No. 1, pp. 125–150.

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.

Cantillon, Estelle and Aurélie Slechten (2018) "Information aggregation in emissions markets with abatement," *Annals of Economics and Statistics*, No. 132, pp. 53–79.

Carleton, Tamma A. and Solomon M. Hsiang (2016) "Social and economic impacts of climate," *Science*, Vol. 353, No. 6304, p. aad9837.

Choi, Syngjoo, Raymond Fisman, Douglas Gale, and Shachar Kariv (2007) "Consistency and heterogeneity of individual behavior under uncertainty," *American Economic Review*, Vol. 97, No. 5, pp. 1921–1938.

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.

Deschênes, 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, Frederick van der Ploeg, Armon Rezai, and Frank Venmans (2021) "Are economists getting climate dynamics right and does it matter?" *Journal of the Association of Environmental and Resource Economists*, Vol. 8, No. 5, pp. 895–921.

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 (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. (1989) "Informed speculation with imperfect competition," *The Review of Economic Studies*, Vol. 56, No. 3, pp. 317–355.

Lambert, Nicolas S., Michael Ostrovsky, and Mikhail Panov (2018) "Strategic trading in informationally complex environments," *Econometrica*, Vol. 86, No. 4, pp. 1119–1157.

Lemoine, Derek (2007) "Greenhouse gas property: An adaptable climate policy for an uncertain world," SSRN Working Paper 2256992.

——— (2020) "Incentivizing negative emissions through carbon shares," Working Paper 27880, National Bureau of Economic Research.

——— (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 economic impact of climate change from weather variation," July.

Lou, Youcheng, Sahar Parsa, Debraj Ray, Duan Li, and Shouyang Wang (2019) "Information aggregation in a financial market with general signal structure," *Journal of Economic Theory*, Vol. 183, pp. 594–624.

McKitrick, Ross (2011) "A simple state-contingent pricing rule for complex intertemporal externalities," *Energy Economics*, Vol. 33, No. 1, pp. 111–120.

Mills, Edwin S. (1972) *Urban Economics*.

Mirrlees, J. A. and N. H. Stern (1972) "Fairly good plans," *Journal of Economic Theory*, Vol. 4, No. 2, pp. 268–288.

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. (1992) "An optimal transition path for controlling greenhouse gases," *Science*, Vol. 258, No. 5086, pp. 1315–1319.

——— (2017) "Revisiting the social cost of carbon," *Proceedings of the National Academy of Sciences*, Vol. 114, No. 7, pp. 1518–1523.

Ozsoylev, Han N. and Johan Walden (2011) "Asset pricing in large information networks," *Journal of Economic Theory*, Vol. 146, No. 6, pp. 2252–2280.

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.

Rahi, Rohit and Jean-Pierre Zigrand (2018) "Information acquisition, price informativeness, and welfare," *Journal of Economic Theory*, Vol. 177, pp. 558–593.

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.

Rostek, Marzena and Marek Weretka (2012) "Price inference in small markets," *Econometrica*, Vol. 80, No. 2, pp. 687–711.

——— (2015) "Information and strategic behavior," *Journal of Economic Theory*, Vol. 158, pp. 536–557.

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.

Shogren, Jason F., Joseph A. Herriges, and Ramu Govindasamy (1993) "Limits to environmental bonds," *Ecological Economics*, Vol. 8, No. 2, pp. 109–133.

Solow, Robert M. (1971) "The economist's approach to pollution and its control," *Science*, Vol. 173, No. 3996, pp. 498–503.

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.

Billette de Villemeur, Etienne and Justin Leroux (2019) "Tradable climate liabilities: A thought experiment," *Ecological Economics*, Vol. 164, p. 106355.

Vives, Xavier (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.

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.