

Asset returns as carbon taxes*

Lautaro Chittaro[†] Monika Piazzesi[‡] Marcelo Sena[§] Martin Schneider[¶]

September 2025

Abstract

In frictionless financial markets, a carbon tax on energy users provides the same incentives as a *replicating asset price schedule* that depends on emissions. In particular, the replicating rate of return on a firm increases linearly in scope 1 emissions relative to enterprise value. We use this result to interpret pollution premia measured by recent empirical studies and conclude that markets currently provide only modest incentives. Replicating a serious carbon tax requires high returns in the right tail of the emission intensity distribution. With heterogeneous investors, such returns are not sustainable unless essentially everyone perceives large nonpecuniary costs from holding dirty capital. Substantial emission reductions can be achieved, however, when even a small share of investors perceive nonpecuniary *benefits* from owning clean electricity capital.

*We thank Hunt Allcott, Hui Chen, Niels Gormsen, Larry Goulder, Lars Hansen, Ben Hebert, Kilian Huber, Per Krusell, Chen Lian, Lasse Pedersen, Jim Poterba, Stephanie Schmitt-Grohe, Kjetil Storesletten, David Thesmar, Annette Vissing-Jorgensen as well as seminar and conference participants at the Banco de Portugal Conference on Monetary Economics, Chicago Booth, Columbia, MIT Sloan, NBER EFG Meeting, Normac, Princeton, SITE, SED, Stanford, and UC Berkeley EEE-fest for helpful comments and suggestions.

[†]Stanford, chittaro@stanford.edu.

[‡]Stanford & NBER, piazzesi@stanford.edu.

[§]Stanford, msena@stanford.edu.

[¶]Stanford & NBER, schneidr@stanford.edu.

1 Introduction

A large fraction of wealth in modern financial markets is invested in green assets.¹ In theory, green investing can affect the real economy by altering firms' costs of capital. If investors demand higher expected returns—pollution premia—on dirtier firms, it becomes more expensive for such firms to raise capital, and they should grow more slowly. A growing empirical literature measures return differentials between dirty and clean firms. While the results are not yet clear-cut, in part because the available time series are short, there is mounting evidence of a positive pollution premium. At the same time, survey evidence shows that a substantial fraction of households is willing to give up some returns to benefit the environment. What is unclear, however, is how strong these preferences for green assets must be for this to have an impact.

This paper asks what it takes for capital markets to mimic a carbon tax, the standard tool economists recommend to fix a climate externality. We first show that, in models with frictionless firm financing, a carbon tax on energy users provides the same incentives as a *replicating asset price schedule* that depends on emissions. Replicating rates of return then increase in a firm's emission intensity, defined as scope 1 emissions relative to enterprise value. By evaluating replicating return schedules at emission intensities currently observed in the U.S. economy, we show that financial markets today provide at most modest incentives, equivalent to a carbon tax on the order of \$10 per ton of CO₂. Computing replicating returns in a quantitative GE model shows that a \$100-per-ton carbon tax requires large returns on dirty production plans. The model allows for heterogeneous preferences: if not all agents are green investors, exact replication is no longer possible, but we still can ask what return schedules can get close to what a carbon tax would accomplish.

The simplest way to derive a replicating asset price schedule is to assume that carbon taxes are collected not directly from polluting firms, but instead from investors who own those firms, in proportion to their ownership stakes. When financial markets are frictionless, how the tax is collected does not matter for allocations. However, collection from investors affects asset prices: the value of a firm to its owners falls exactly by the carbon tax bill. Shareholder value maximization chooses the same production plan whether emissions affect the firm's after-tax income or instead owners' valuation of the firm. In this sense, the asset price schedule replicates the carbon tax. The result is an asset market version of tax incidence equivalence, a principle familiar from goods markets: whether a producer or consumer pays carbon tax also does not matter for allocations, only for prices. In asset markets, it can also be expressed using returns: a carbon tax bill τE from owning a firm worth V calls for compensation by an extra return $\tau E/V$ that is increasing in the emission intensity E/V .

A related equivalence result applies when there is no carbon tax but investors instead dislike dirty assets, as typically assumed in the climate finance literature. Investors who incur nonpecuniary cost τE per share from owning a dirty firm also require compensation by an extra return $\tau E/V$. The return schedule now depends on the preference parameter τ , but we

¹According to Bloomberg Intelligence, global ESG assets are currently roughly 20% of assets under management and are projected to increase to 25% by 2030.

observe the same equilibrium returns and allocations *as if* a carbon tax at rate τ were collected from investors. In both cases, a carbon tax *rate* corresponds to a return *schedule*, whereas the cross-section of replicating returns consistent with a carbon tax rate depends on the endogenous choice of intensities. Like the tax burden, equilibrium returns reflect not only the tax rate, but also features of the economic environment, such as technology. For example, technical progress that makes it cheaper to save energy will lower equilibrium returns consistent with replication of a given carbon tax rate.

To illustrate the magnitude of replicating returns, we conduct two quantitative exercises. First, we evaluate return schedules at the currently observed cross-section of emission intensities. For any given carbon tax rate, the resulting distribution describes returns that we would see today if financial markets were replicating the tax. A \$10 carbon tax requires premia for emissions that are 6 basis points on average. However, the right-skewed distribution of emission intensities implies much higher premia at the top: 60bp at the 99th percentile. For publicly traded electricity companies, the median premium is 1 percent, while the 90th percentile reaches 1.6%. Since the empirical asset pricing literature typically finds pollution premia of only a few percentage points, we conclude that the current financial market incentives are modest.

Our second exercise evaluates return schedules at emission intensities that emerge in equilibrium after firms have adjusted to a high carbon tax. The resulting return distributions reflect the tax rates—or the nonpecuniary costs—that investors would need to bear to enforce strong climate incentives. To predict these adjustments, we use a multisector growth model with extra detail on energy production, calibrated using both sector-level data and firm-level data from the electricity sector. Under a \$100 carbon tax, the return premia become substantial in the right tail of the distribution: the 99th percentile reaches 5%, with high premia in the transport sector, which has limited ability to substitute away from oil. In contrast, electricity firms, which can more easily substitute away from coal, face much lower premia—just 21bp at the median and 1.1% at the 90th percentile. Most adjustment in the electricity sector occurs within firms: the replicating return schedule incentivizes currently dirty firms to clean up their production.

Importantly, incentives are provided by the entire return schedule—not just by equilibrium returns. For example, an electricity firm considering a production plan with today's median emission intensity would face a 10% premium under a \$100 carbon tax. To replicate such a return schedule based on investor preferences, *all* investors must incur large nonpecuniary costs from holding dirty capital and therefore be willing to sacrifice substantial pecuniary returns. This motivates an extension of our model to heterogeneous investors who face short-sale constraints. We assume that green investors interact with traditional investors who do not care about the emission intensity of their portfolios.

When green investors incur nonpecuniary costs, even a small presence of traditional investors greatly weakens the impact of green investing on emissions. Starting from a scenario in which 100% of wealth is held by green investors whose preferences replicate a \$100 carbon tax, the introduction of just 2% of wealth held by traditional investors is enough to undo half of the emission reduction. This fragility arises for two reasons: (i) the distribution of emissions is

heavily right-skewed, and (ii) many of the dirtiest firms also have the greatest potential to reduce their emission intensities. In equilibrium, traditional investors sort into holding high-emission firms that green investors avoid due to high nonpecuniary costs. Moreover, traditional investors are happy to accept the same returns regardless of emission intensity. The resulting return schedule becomes flat at the highest intensities: incentives are no longer provided in a range where they would have the most impact.

Can a small share of green investors still generate meaningful incentives for emission reduction? Suppose instead that investors derive a *nonpecuniary benefit from holding electricity firms*, and more so the lower their emission intensities. We show that if just 2% of wealth is held by investors who are willing to forgo less than 2% of wealth per year in equilibrium, they can achieve about one-quarter of the emission reductions of a \$100 carbon tax. The key property of equilibrium with nonpecuniary benefits is that incentives are provided where they are most effective. Traditional investors hold most firms and demand the same returns from these firms, providing no incentives. However, most firms are relatively clean, so the lack of incentives is irrelevant. Green investors end up subsidizing electricity firms, including some dirtier ones. Since the subsidy declines with emission intensity, incentives are provided in the right place.

We take away from these exercises that the marketing of green investment products should change. The current emphasis on negative screening, akin to a divestment campaign that punishes dirty firms with high costs of capital, is misguided. Instead, green investing should aim to incentivize emission reductions, especially in sectors that are currently dirty but can adjust. In our model, we quantify the effectiveness of different investment strategies by comparing alternative preferences of green investors. This approach is helpful because preferences for green products are still evolving and shaped in part by what products intermediaries offer. As green investing remains relatively new, there is both a strong demand for products that help reduce emissions, but no clear understanding of which approaches are most effective.²

The key to our equivalence results is that investors distinguish claims on firms not only by their payoffs, but also by the emissions the firms generate. Formally, we consider many asset markets indexed by future firm value and emissions. When households choose portfolios to maximize utility, they price in compensation for future carbon taxes or nonpecuniary costs. Their asset demand generates a price schedule that is independent of the production side of the economy. Equilibrium then matches households' asset demand with firms that maximize shareholder value. The equilibrium distributions of returns and emission intensities are jointly determined: technologies are funded via those markets where they achieve the highest value.³

²We note that a nonpecuniary benefit for electricity shares has features of a preference for *changes* in emissions, and hence some scoring rules developed in the financial sector. In particular, it rewards dirty firms to become cleaner. There is a subtle difference, however: if clean electricity is mostly produced by entrants that always have zero emissions, our scheme subsidizes them even though their intensities never change.

³The theme that asset characteristics and prices are jointly determined is familiar from other work. For example, Fostel and Geanakoplos (2015) study models with collateral constraints where markets clear to jointly determine interest rates and the default risk of the borrower. Rogerson, Shimer and Wright (2005) survey competitive search models where prices are determined jointly with market tightness, and thus the queue length of the seller. In our context, assets differ in endogenously determined convenience yields that reflect the issuer's future emissions.

To this end, firms' choice of production plans takes into account shareholders' required rates of return, which depend on their emission intensities.

To establish equivalence between a carbon tax and asset prices generated by investor preferences, we assume that utility is defined over consumption less a nonpecuniary cost that is proportional to emissions from the portfolio. An economy with such preferences but no carbon tax has the same equilibrium allocation as an economy with standard preferences and a carbon tax if the government spends all tax revenue on consumption. The functional form of utility implies that compensation for nonpecuniary cost generates the same replicating price schedule as compensation for carbon taxes collected from investors. Moreover, wasted carbon tax revenue has the same general equilibrium effects as the nonpecuniary cost.⁴

Incentives provided by replicating returns reward substitution to cleaner technologies both across firms and within firms—the return schedule does not simply penalize dirty firms. In fact, substitution within firms accounts for a substantial share of emission reductions in our quantitative exercises, especially in the electricity sector. The replication results require no assumptions about the boundaries of the firm. It does not matter, for example, whether greener production occurs in specific firms or in large conglomerates that also operate dirty production. As under a carbon tax, decarbonization may occur through the expansion of clean firms, which benefit from a lower cost of capital, or through the transformation of dirty firms that adopt cleaner production methods to access cheaper financing.

Our results require that financial markets are frictionless, and that the owners of a firm care about its emissions because of either taxation or preferences. We need no other assumptions on technology or goods market structure. While our baseline result is for a deterministic setting with perfect competition and no adjustment costs, we further show that replicating returns are available in the presence of risk, adjustment costs, or market power in goods markets. We can also accommodate features of technology that are nonstandard in growth models but often relevant for thinking about paths to decarbonization, such as vintage capital or the presence of exhaustible resources.

In particular, our framework allows for stranded assets. When investors perceive large nonpecuniary costs, allocations are the same as under a large carbon tax that induces firms to abandon dirty technologies and makes some capital goods worthless, or "stranded". When shareholders anticipate too high a cost from a dirty production plan, they choose not to undertake that plan. For some technologies, the best plan may then be to shut down. We note that it is not necessary to define emission intensities for off-equilibrium plans to evaluate them: in our framework, shareholders value plans based on the level of future emissions. Section 2.3 presents an explicit example where the shutdown of a dirty technology implies zero equilibrium prices of assets specific to dirty production, in particular vintage capital that cannot be repurposed

⁴We follow the literature in working with nonconsequentialist preferences, that is, investors dislike dirty assets for private reasons, not because of their aggregate real effects. Economies with green investors thus generate lower emissions because markets value those private costs, not because they directly address the climate externality. Nevertheless, results with green investors allow us to gauge how strong such taste has to be to replicate a serious carbon tax in the absence of government action.

and fuel reserves in the ground.

The production side of our quantitative model features an input-output structure with 72 sectors that makes carbon tax effects propagate through supply chains. The electricity sector consists of many plants that produce varieties of electricity with different fuels as well as at different emission intensities. We add this extra granularity to capture substitution potential in electricity more accurately, as well as to study how costs of capital change within and across individual electricity firms. The model shares key properties with many models in the literature. In particular, energy is complementary to capital, there is limited substitutability between energy sources, especially for transport, but relatively higher scope for substitution between electricity varieties.

We calibrate the model using sectoral data on cost and emissions from the Global Trade Analysis Project (GTAP) as well as at the firm level from TruCost and Compustat. Since there is considerable uncertainty about technology parameters such as long-run substitution elasticities, we benchmark our predictions for emissions, energy, and macro aggregates to a meta-study on carbon tax scenarios for the U.S. by the Energy Modeling Forum (EMF, see McFarland, Fawcett, Morris, Reilly and Wilcoxon 2018). We obtain results that are broadly similar to the typical model in the literature and also provide robustness analysis for key technology parameters that determine the intermittency of renewable electricity, the potential for ramping up nuclear power and technical progress in electricity.

All of our results assume that scope 1 emissions can be measured and are public information. Whether incentives come from a carbon tax collected from investors or investor preferences, emissions of a firm must be verifiable by the tax authority or investors, respectively. We view this as a relatively low burden. For large companies, measures of scope 1 emissions are already available from multiple data vendors, including the TruCost data we use for our calibration of the power sector. The SEC's March 2024 rule, currently inactive, requires that public companies disclose an estimate of scope 1 emissions with quarterly filings.

Our results also say that the scoring of firms for climate impact should focus on scope 1 emissions only, and do not require measurement or disclosure of more complicated concepts of emissions. The equivalence result relates a return schedule based on scope 1 intensity to a carbon tax on energy *use*. The return schedule should therefore not penalize energy producers for the dirty fuels they produce—part of their scope 3 emissions—but only for scope 1 emissions that arise in production. The argument is analogous to the principle that externalities can be taxed at the producer or user level, but taxing both does not make sense. Similarly, as long as electricity producers' cost of capital responds to their scope 1 emission intensity, it is not necessary to measure electricity users' scope 2 emissions.

Related literature. Our return schedule provides a benchmark for interpreting evidence on pollution premia — that is, cross-sectional differences in mean returns, and hence the costs of capital, between dirty and clean firms. Across various measurement approaches, estimated pollution premia tend to be modest—typically no more than a couple of percentage points—and thus far insufficient to provide strong incentives to the dirtiest firms. Giglio, Kelly and

Stroebel (2021) and Hong and Shore (2023) survey a large literature attempting to estimate these premia from asset return data. A key empirical challenge is that the ongoing transition toward greener investment makes it hard to estimate mean returns (Pástor, Stambaugh and Taylor 2021). Studies based on investor surveys reach similar conclusions on mean returns, while also highlighting considerable heterogeneity in preferences (for example, Riedl and Smeets 2017, Giglio, Maggiori, Stroebel, Tan, Utkus and Xu 2025, Aron-Dine, Beutel, Piazzesi and Schneider 2024). The most direct evidence on the channel we are interested in comes from Gormsen, Huber and Oh (2024), who infer managers' perceptions of costs of capital from corporate conference calls. Their findings highlight, in particular, that managers of utility companies perceive lower costs of capital for their cleaner divisions, consistent with the within-firm substitution effect in our model.

Our results on equilibria with investor tastes contribute to a growing literature on climate finance with heterogeneous investors with non-consequentialist preferences (see Pástor, Stambaugh and Taylor 2024 for an overview and Bonnefon, Landier, Sastry and Thesmar (2025) and Heeb, Kölbel, Paetzold and Zeisberger (2023) for evidence on non-consequentialist preferences.) Early work focused on *negative screening*, where some investors exclude some assets from their portfolios and accept higher risk as a result (Heinkel, Kraus and Zechner 2001; Geczy, Stambaugh and Levin 2021, Berk and van Binsbergen 2021). We follow instead a recent literature that models green investors as deriving *convenience yields* from their green assets, as in Pástor, Stambaugh and Taylor (2021), Pedersen, Fitzgibbons and Pomorski (2021), or Aron-Dine, Beutel, Piazzesi and Schneider (2024). While those papers study asset pricing and portfolio tilts, we focus on implications for real outcomes—specifically, emissions—in a production economy. We further establish that nonpecuniary *benefits* are particularly powerful because they reflect a willingness to subsidize green firms through lower costs of capital. By contrast, nonpecuniary costs imply higher returns for dirty assets, which are easier to erode when a small share of investors are indifferent to holding them.

Several other papers examine the relationship between carbon taxes and the cross-section of returns in financial markets. Iovino, Martin and Sauvagnat (2024) analyze how capital taxation affects an economy's carbon footprint through capital reallocation. Papoutsis, Piazzesi and Schneider (2024) study reallocation driven by the cross-sectional effects of unconventional monetary policy, focusing on the ECB's corporate bond purchase program. Pedersen (2025) studies how a carbon tax can be mimicked by "sustainable discount rates" that rise with emission intensities, which he defines as emissions at date $t + 1$ divided by firm value at date t . Since such intensities are not well-defined when the firm value is zero, he concludes that mimicking breaks down when there are stranded assets. Our general equilibrium framework clarifies that stranded assets present no special challenge for green finance: a carbon tax collected from investors or investor preferences make equity *prices* depend on the level of scope 1 emissions, so even a large carbon tax that strands assets is replicated.

The production side of our model builds on a long tradition of general equilibrium integrated

assessment models designed to assess carbon taxation and other policies.⁵ Studies that derive optimal carbon taxes from a tradeoff between environmental damage and adjustment cost typically feature a climate block with state variables such as temperature as well as a damage function that captures economic costs of warming (for example, Hassler and Krusell 2012, Golosov, Hassler, Krusell and Tsyvinski 2014, Barrage and Nordhaus 2024, Aghion, Boppart, Peters, Schwartzman and Zilibotti 2025). Since our objective is only to compare tax collection mechanisms and investor preferences for a given carbon tax rate, we keep the climate side of the model simple to minimally capture a climate externality. We share the focus on one country that engages in international trade with climate models of the U.S. economy such as Goulder and Hafstead (2017). We thus discuss carbon leakage only from the perspective of a small open economy (for carbon taxation in multicountry models, see, for example, Kortum, Weisbach, Wang and Yao 2023, Kortum and Weisbach 2024 or Ramondo, Garcia-Lembergman, Rodríguez-Clare and Shapiro 2025.)

The paper is structured as follows. Section 2 introduces a model with a representative agent. It shows how a carbon tax is reflected in asset returns when it is collected from investors as well as how the same returns arise without taxation from investor preferences. Section 3 presents our data and a first set of quantitative results: replicating returns when intensities do not adjust. Section 4 discusses the calibration of the model and derives equilibrium returns. Section 5 extends the model to allow for nonpecuniary benefits and costs as well as heterogeneous investors and presents quantitative results for this case.

2 Model

We consider a standard growth model with an input-output structure and climate externalities. We first lay out a physical environment and introduce a simple system of markets for which we show our equivalence result. For transparency, we abstract from a number of interesting frictions. We then discuss how the equivalence result holds for alternative decentralizations and extensions of the environment.

2.1 Physical environment

Preferences. A representative household lives forever, inelastically supplies one unit of labor and values a consumption good C_t as well as the state of the environment, summarized by the cumulative sum of emissions η_t . Preferences are represented by utility

$$\sum_{t=0}^{\infty} \beta^t (u(C_t) - v\eta_t), \quad (1)$$

where u is strictly increasing and $v > 0$. Adding damages separably in utility is a minimal way to introduce a climate externality. The precise nature of damages is not important for our

⁵See Nordhaus (2013) for a historical survey and IAMC (2025 website) for an overview of models currently used around the world.

exercise since we take the size of the carbon tax as given.

Technology. A consumption good C_t and an export good \tilde{Y}_t are produced using N domestic intermediate goods. Let y_t and \tilde{y}_t denote $N \times 1$ vectors of intermediate inputs in the production of the consumption and the export good, respectively. Each domestic intermediate good n is produced using labor l_t^n , capital k_t^n , domestic intermediates collected in an $N \times 1$ vector x_t^n , and \tilde{N} imported intermediates collected in a vector \tilde{x}_t^n . There are also N capital goods, each specific to one intermediate good. Capital good n is made from intermediates collected in a vector z_t^n . Capital in sector n can be first used one period after it has been installed and depreciates at rate $\delta_n \geq 0$.

Technology is described by production functions g , \tilde{g} , and f^n for the consumption good, the export good, and the n th domestic intermediate good that map nonnegative inputs into nonnegative outputs, as well as a production function h for the n th capital good that may also have negative inputs and outputs. For all goods, zero inputs imply zero output. Feasible allocations satisfy the resource constraints

$$\begin{aligned} C_t &= g(y_t), \\ \tilde{Y}_t &= \tilde{g}(\tilde{y}_t), \\ k_{t+1}^n &= (1 - \delta_n) k_t^n + h^n(z_t^n); \quad n = 1, \dots, N, \\ y_{t,n} + \tilde{y}_{t,n} + \sum_{m=1}^N (x_{t,n}^m + z_{t,n}^m) &= f^n(l_t^n, k_t^n, x_t^n, \tilde{x}_t^n); \quad n = 1, \dots, N. \end{aligned} \quad (2)$$

The last equation is the resource constraint for the n th domestic intermediate good, with gross output on the right-hand side and all uses of good n on the left-hand side. In particular, $x_{t,n}^m$ is the quantity of good n used to make good m , the n th element of the vector x_t^m that collects inputs into the production of good m . Since f^n describes gross output, good n may also be used as input into its own production. Similarly, $z_{t,n}^m$ is the quantity of good n used to produce new capital goods for sector m . Good n is further used as an input in the production of consumption goods $y_{t,n}$ and export goods $\tilde{y}_{t,n}$.

We have not made assumptions on the production function for capital goods h^n other than $h^n(0) = 0$. We thus allow for features of technology that are nonstandard in growth models, but sometimes arise in discussions of the green transition. Let $X^n = (l^n, k^n, x^n, \tilde{x}^n)$ denote a production plan for technology n . We can interpret the n th capital good as a fixed factor, such as land, and the n th intermediate good as services from that factor by setting both depreciation and h^n to zero, as well as $f^n(X^n) = k^n$. More generally, we can allow for irreversible investment by assuming h^n equals zero for all negative input bundles. Alternatively, we can interpret the n th capital good as an exhaustible resource and the n th intermediate good as extraction of that resource by setting $h^n(z) = -z_n$ and $f^n(X^n) = 0$. The cost of extraction can then be handled by defining an additional intermediate that is made from good n as well as capital, labor, and other inputs used in the extraction process.

Emissions. The production of intermediate goods may generate emissions. The emission func-

tion e^n indicates how many emissions are generated for every production plan for technology n . The evolution of the state of the environment can then be written as

$$\eta_t = \eta_{t-1} + \sum_{n=1}^N e^n (l_t^n, k_t^n, x_t^n, \tilde{x}_t^n) + \tilde{e}_t, \quad (3)$$

where \tilde{e}_t is an exogenous sequence that represents foreign emissions. The function e^n maps nonnegative inputs into the real line. Our notational convention that production of consumption, export, and capital goods does not generate emissions is not restrictive since we can always define additional intermediates.

International trade. The country takes as given prices of imported goods collected in a vector \tilde{p}_t . We abstract from foreign borrowing and assume that trade is balanced every period.⁶ We thus require that the value of imports of foreign goods purchased by all intermediate sectors equals the value of output of the export good:

$$\sum_{n=1}^N \sum_{m=1}^{\tilde{N}} \tilde{p}_{t,m} \tilde{x}_{t,m}^n = \tilde{P}_t \tilde{Y}_t. \quad (4)$$

Social planner problem. As a benchmark, consider a social planner who cares about local citizens' welfare. The social planner maximizes utility (1) subject to technology (2), the evolution of the state of the environment (3), the trade balance (4), and the resource constraint for labor

$$\sum_{n=1}^N l_t^n = 1. \quad (5)$$

The optimal allocation internalizes the negative impact of dirty energy use on households' utility due to higher cumulative emissions η_t . The planner as defined here does not take into account local citizens' substitution towards emission-intensive imported goods.

Let μ_t denote the planner's multiplier on the evolution of the state of the environment (3). The social cost of carbon, or the shadow cost of an additional unit of emissions in units of marginal utility, is given by $SCC_t = \mu_t / u'(C_t)$. The planner equates the marginal product of any input to its total marginal cost, including its contribution to emissions valued by the social cost of carbon. For example, Appendix A.1 derives the planner's FOC for $x_{t,m}^n$, the quantity of intermediate good m used for making good n as part of the production plan $X^n = (l_t^n, k_t^n, x_t^n, \tilde{x}_t^n)$ is

$$\frac{dg}{dy_m}(y_t) = \frac{dg}{dy_n}(y_t) \frac{df^n}{dx_m}(X^n) + SCC_t \frac{de^n}{dx_m}(X^n). \quad (6)$$

Standard arguments imply that a competitive equilibrium with a carbon tax rate equal to the

⁶This assumption is not essential for the results of this section—what matters there is that all investors who own domestic firms are subject to the same tax treatment or have the same preferences. For our quantitative work below, we consider steady states, so the assumption is natural.

social cost of carbon decentralizes the social optimum. Intuitively, the spot price of fuel works like the shadow cost and is augmented by the tax per unit of fuel, that is, the SCC times the emission factor.

2.2 Markets and equilibrium

Firms are owned by households, and their shares trade in frictionless financial markets. All goods trade in competitive markets. The consumption good serves as numéraire.

Carbon taxes. The government levies a proportional carbon tax τ_t on emissions. We distinguish two ways to collect this tax. A fraction ϕ of carbon taxes is collected directly from firms; a firm that emits e_t pays taxes $\phi\tau_t e_t$ at date t that are subtracted from capital income and hence lower enterprise value. A fraction $1 - \phi$ is instead collected from investors in proportion to their holdings of equity; an investor who owns shares s_t of the firm pays $(1 - \phi)\tau_t e_t s_t$. To implement the tax, the government has to measure emissions in period t . To collect from investors, it also has to know households' portfolios.

Firms and equity. Firms plan with an infinite horizon and may operate any number of technologies described by the production functions (2). To simplify notation in this section, we discuss only the problem of a firm that makes a single intermediate good n . The arguments readily extend to other types of firms; we discuss multiproduct firms in Section 2.3 below.

A production plan $X = (l, k, x, \tilde{x})$ for making good n at date $t + 1$ consists of capital installed at date t and variable inputs purchased at date $t + 1$. The pretax capital income generated by the firm at date $t + 1$ is

$$\pi_{t+1}^n(X) = p_{t+1,n} f^n(X) - \sum_{j=1}^N p_{t+1,j} x_j - \sum_{j=1}^{\tilde{N}} \tilde{p}_{t+1,j} \tilde{x}_j - w_{t+1} l. \quad (7)$$

It consists of revenue less the cost of variable inputs, namely expenditure on domestic and foreign intermediates at prices $p_{t,n}$ and $\tilde{p}_{t,n}$ as well as on labor at the wage w_t .

The cum-dividend value of the firm at date $t + 1$ is then

$$v_{t+1}^n(X) = \pi_{t+1}^n(X) - \phi\tau_{t+1} e^n(X) + V_{t+1}^n(k). \quad (8)$$

It consists of after-tax income, that is, pretax income less the carbon tax bill paid directly by the firm, plus a continuation value that depends on the capital stock k . If technology exhibits constant returns to scale, the continuation value $V_{t+1}^n(k) = p_{t+1}^k k(1 - \delta_n)$ is only the value of undepreciated capital. More generally, it can include rents earned by the firm, such as those resulting from decreasing returns.

Asset markets. If taxes are collected only from firms, stocks differ only by their after-tax value. When some taxes are collected from investors, in contrast, households distinguish firms also by their emissions. This is because firm ownership not only promises dividends but also comes with a carbon tax obligation for the investor. We thus consider a continuum of asset markets

indexed by the pair (v, e) . Let $q_t(v, e)$ denote the date t stock price of a firm that generates value v and emissions e at date $t + 1$. An investor who buys s_{t+1} shares at date t expects payoff $v s_{t+1}$ from the firm and a tax bill $(1 - \phi)\tau_{t+1}e s_{t+1}$ from the government, both at date $t + 1$.

Household problem. An individual household takes as given the evolution of the environment η_t and maximizes utility (1) by choosing consumption and equity holdings $s_{t+1}(v, e)$ subject to the budget constraint

$$C_t + \sum_{v,e} q_t(v, e) s_{t+1}(v, e) = \sum_{v,e} (v - (1 - \phi)\tau_{t+1}e) s_t(v, e) + w_t + T_t, \quad (9)$$

where w_t is the wage and T_t is a lump sum transfer from the government. For both the expenditure on equity on the left-hand side and the payoffs from equity on the right-hand side, we sum over all possible shares (v, e) .

Household optimality implies that stock prices satisfy the Euler equations

$$q_t(v, e) = \beta \frac{u'(C_{t+1})}{u'(C_t)} (v - (1 - \phi)\tau_{t+1}e) =: \frac{v - (1 - \phi)\tau_{t+1}e}{R_{t+1}(0)}. \quad (10)$$

For a clean firm with $e = 0$, the price only reflects the present value of dividends, with discounting by the *clean rate* $R_{t+1}(0)$, equal to the household's marginal rate of substitution (MRS). For dirty assets with $e > 0$, however, the anticipation of future taxes on the portfolio implies an additional additive price discount that compensates investors for the tax burden. This effect is absent only when all carbon tax collection is from firms, or $\phi = 1$.

Shareholder value. A firm that operates technology n chooses its production plan X at date t to maximize shareholder value

$$V_t(k_t^n) = \max_X \{-p_{t,n}^k (k - (1 - \delta)k_t^n) + q_t(v_{t+1}^n(X), e^n(X))\}. \quad (11)$$

The first term is capital investment, the only component of the plan X purchased already at date t . The second term is the present value to shareholders of owning the firm, or the date t ex-dividend value of its shares, where discounting is done with the function q_t .

The value of the firm depends on the market (v, e) where the firm raises funds. The two arguments of the price function q_t reflect the two components of the payoff to shareholders. The first is the future cum-dividend value, defined in (8). It contains, in particular, the carbon tax collected from firms. The second argument is emissions, which matter separately because of investors' concern with taxes on the portfolio. We note that substituting (8) into (11) delivers a Bellman equation for the firm. The only unusual feature is that discounting by q_t makes emissions matter beyond the direct tax bill.

Government. We assume that the government rebates all tax receipts lump sum to households. The government budget constraint is therefore

$$T_t = \sum_n \tau_t e^n (X^n). \quad (12)$$

The transfer is equal to the sum of carbon taxes paid on all technologies. We specify policy by the carbon tax rate τ_t as well as the share ϕ collected from firms. The transfer T_t then adjusts so the budget constraint holds in equilibrium.

Competitive equilibrium. Equity market clearing requires that the representative household owns all shares. Since every firm supplies one share, portfolio demand for equity of every type (v, e) must equal the number of firms that raise equity in market (v, e) :

$$s_{t+1}(v, e) = |\{n : (v, e) = (d_{t+1}^n(X^n), e(X^n))\}|. \quad (13)$$

With our convention of one firm per technology, at most N distinct equity claims are traded in equilibrium.

For given carbon tax rates τ_t and a share ϕ collected from firms, an equilibrium consists of an allocation together with vectors of intermediate goods prices, capital prices, an export good price, a wage, a government transfer, and a schedule of equity prices such that (i) households maximize utility (1) subject to the budget constraints (9), (ii) firms that buy capital maximize shareholder value (11), (iii) other firms maximize profits, (iv) households hold all equity of intermediate goods firms, (v) all domestic goods markets clear (2), (v) international trade is balanced (4), and the government budget constraint (12) holds.

We are now ready to state our main equivalence result.

Proposition 1. (i) *The equilibrium allocation does not depend on how carbon taxes are collected, that is, it is independent of the parameter ϕ .*

(ii) *When some carbon taxes are collected from investors ($\phi < 1$), then the equity price $q_t(v, e)$ is decreasing in emissions. Moreover, the rate of return on a firm with value v and emissions e is increasing in the emission intensity $\varepsilon = e/v$, and more so the higher the carbon tax rate:*

$$\frac{v}{q_t(v, e)} = \frac{R_{t+1}(0)}{1 - (1 - \phi)\tau_{t+1}\varepsilon} =: R_{t+1}(\varepsilon).$$

To establish part (i), we note first that the parameter ϕ appears only in household and firm optimization problems, but not in market-clearing conditions or the government budget constraint. We therefore only need to check that ϕ does not affect household and firm problems at equilibrium prices. Since equity prices satisfy household Euler equations (10), the ex-dividend

value of a firm that contemplates production plan X is independent of ϕ :

$$\begin{aligned} q_t(v_{t+1}^n(X), e^n(X)) &= \frac{v_{t+1}^n(X) - (1 - \phi)\tau_{t+1}e^n(X)}{R_{t+1}(0)} \\ &= \frac{\pi_{t+1}^n(X) - \tau_{t+1}e^n(X) + V_{t+1}^n(k)}{R_{t+1}(0)}, \end{aligned} \quad (14)$$

where the second equality follows from substituting for the cum dividend value v_{t+1}^n from (8).

The firm problem (11) is therefore independent of ϕ . Intuitively, when taxes are collected from investors, then asset prices adjust to compensate those investors. In frictionless markets, the adjustment lowers the value of a production plan by exactly the present value of the tax bill. Household optimality (10) requires that after-tax returns on all equity claims are equated in equilibrium. Household portfolio choice is therefore indeterminate, and the household faces the return on savings $R_{t+1}(0)$. Since ϕ does not affect production plans, it does not affect transfer income by the government budget constraint, and therefore does not affect the household problem either. This shows part (i).

While allocations are independent of ϕ , part (ii) says that asset prices are not. It follows directly from (10): holding fixed dividends, anticipation of taxes lowers the value of dirty firms as long as some carbon tax is collected from investors. When carbon taxes are collected from investors, the equity price schedule transmits incentives for cleaner production to firms. For firms that are active in equilibrium, in the sense that their values are nonzero, incentives are summarized by a schedule of returns, or costs of capital, that is increasing in the emission intensity. The higher the carbon tax, and the more taxes are collected from investors, the stronger the incentives must be from asset markets, and the steeper the return schedule slopes up with intensity.

Incentives from asset prices. For firm decision-making, it is not important whether prices and returns reflect emissions because of taxation or for other reasons, such as investor preferences. To isolate how asset markets replicate a carbon tax, we state:

Proposition 2. Holding fixed goods prices, a firm that faces no carbon tax and an equity price function

$$q_t(v, e) = \frac{v - \tau_{t+1}e}{R_{t+1}(0)} \quad (15)$$

chooses the same production plans as a firm that pays carbon tax at rate τ_{t+1} and faces a price function $\tilde{q}_t(v, e) = v/R_{t+1}(0)$ that discounts dividends independently of emissions.

The result follows directly from (14). A firm that faces no carbon tax but a price function that responds linearly to emissions behaves as if it has to pay a carbon tax at rate τ_{t+1} , the slope of the price function.

Carbon taxes collected from firms and stock prices that respond to emissions thus represent equivalent incentive mechanisms that encourage cleaner production. When firms pay carbon

taxes, they are valued for their after-tax payoff, discounted at $R_{t+1}(0)$ under the price function \tilde{q}_t . Shareholders who assess, say, a cleaner production plan, then trade off a lower tax burden against a possibly lower pretax payoff. Under the price function (15), in contrast, firms are valued for both pretax payoff and emissions. While a cleaner plan does not lower firms' tax burden, it makes the firm more valuable to shareholders. The firm can thus raise funds for the plan in a market (v, e) where its shares fetch a higher price for the same payoff v .

We note that our propositions do not require that the value of a firm is positive: the price function applies to any pair (v, e) . Even though an emission intensity and a rate of return can be defined only if $v \neq 0$ and $q_t(v, e) \neq 0$, markets guide investment also for other plans. It is not even important that the firm uses capital as an input: it is enough that it is owned by the representative household who cares about emissions. As a stark example, consider a technology that turns one unit of labor into one unit of a good and \bar{e} units of emissions. If the good is produced in equilibrium, it must be priced at marginal cost $p_{t,n} = w_t + \tau_t \bar{e}$. Owners of the firm earn income $p_{t,n} - w_t$ but value the firm at zero because of its emissions.

Incentives from investor preferences. As a concrete example for how asset price incentives can emerge without carbon taxation, suppose that investors incur nonpecuniary costs from holding dirty firms. Concretely, the representative agent has felicity

$$u(C_t - \gamma_t \sum_{v,e} e s_t(v, e)) - v \eta_t, \quad (16)$$

where the parameters $\gamma_t > 0$ capture the intensity of green preferences. This formulation, familiar from the climate finance literature, captures a form of nonconsequentialist preferences: investors not only suffer from climate change, but also directly dislike emissions connected to their own portfolio, without considering the effect of their decisions on aggregate emissions.

What is the equilibrium price function? Since the representative agent's portfolio is responsible for all the emissions in the economy, $E_t = \sum_n e^n (X_t^n)$, the price schedule (10) becomes

$$q_t(v, e) = \beta \frac{u'(C_{t+1} - \gamma_{t+1} E_{t+1})}{u'(C_t - \gamma_t E_t)} (v - \gamma_{t+1} e) =: \frac{v - \gamma_{t+1} e}{R_{t+1}(0)}. \quad (17)$$

The schedule takes the same form as in Proposition 2: it declines with emissions, and more so the more the household dislikes emissions, as described by the preference parameter γ_{t+1} . The clean rate $R_{t+1}(0)$ used to discount dividends of clean firms is given by a MRS that now reflects disutility from emissions.

An economy P, say, with green investors therefore generates the same price *differentials* across firms as an economy T with standard investors, a carbon tax rate $\tau_t = \gamma_t$, and so, by Proposition 2, the same firm choices for a given clean rate. However, whether the clean rate itself, and hence all general equilibrium effects, are also the same depends on what the government does with the tax revenue. Since green investors incur nonpecuniary costs, we cannot expect the same allocation as when the carbon tax is rebated lump sum. Nevertheless, there is a simple

characterization of equilibrium. The nonpecuniary costs $\gamma_t E_t$ look like the carbon tax revenues if the rate is γ_t . With our functional form for utility, we obtain the same allocation as when the government spends the tax revenues on consumption goods and throws them away.⁷ Formally, we have

Proposition 3. *Consider two economies. In economy P, the felicity of the representative investor includes nonpecuniary costs from dirty assets as in (16) with $\gamma_t = \tau_t$ and there is no carbon tax. In economy T, the felicity of the representative investor is standard as in (1), but the government collects carbon taxes at rate τ_t from either firms or investors and spends the tax revenues on the consumption good every period. Then*

(i) *Equilibria of the two economies share the same production plans and emissions. However, consumption in economy T is lower at $C_t^T = C_t^P - \tau_t E_t$.*

(ii) *Both economies share the same clean rate $R_{t+1}(0)$. Equity prices in economy P, given by (17) with $\gamma_t = \tau_t$, are decreasing in emissions. If carbon taxes are collected from investors, equity prices in economy T are the same as in economy P. If carbon taxes are collected from firms, equity prices are independent of emissions and given by $q_t(v, e) = v / R_{t+1}(0)$.*

The proof again compares equilibrium conditions. Start from an equilibrium of economy P. We show that the proposed allocation and clean rate satisfy the equilibrium conditions for economy T. With equal emissions and $C_t^T = C_t^P - \tau_t E_t$, the clean rate of economy P satisfies the household Euler equation in economy T. Proposition 2 says that the price function (17) leads to the same production plans as a carbon tax at rate τ_t collected from firms. Market clearing for all intermediates, capital, and labor continue to hold at optimal production plans, wages, and goods prices from economy P. In particular, since the government does not consume in P, output of the final good in economy T is C_t^P . Finally, the market for the final good clears since C_t^T is output C_t^P less government spending, given by carbon tax revenue $\tau_t E_t$.

Interpreting welfare. In economy P with felicity (16), the representative household incurs disutility from emissions in two ways: the climate externality term depends on aggregate emissions, whereas the private disutility depends on emissions connected to the portfolio. The equilibrium price schedule of economy P reflects the private disutility, and this lowers emissions as if there were a carbon tax. However, markets do not directly address the aggregate externality. This is different from economy T where the price schedule comes from a carbon tax that can implement a planner solution.⁸ This feature of nonconsequentialist preferences is not unique to our paper. Following the literature, we view such preferences as a useful device to generate return differentials observed in the data, and do not use them as a tool for studying optimal policy. Our point here is that green investing, as typically modeled, can reduce emissions just like a carbon tax by generating the right price schedule.

⁷As is common in models of fiscal policy, an alternative interpretation is that the government uses the revenue to produce a public good that enters utility separately so that its presence does not affect agents' choices.

⁸In economy P, a planner would choose the optimal carbon tax to reflect both sources of disutility: since agents suffer from higher temperature *and* their guilt of holding a portfolio that contributes to it, the tax rate would be higher than in economy T.

Replicating returns. Proposition 3 makes precise how asset returns can replicate the incentives of a carbon tax: investor preferences imply a price function, and hence a return schedule, that provides incentives to firms *as if* a carbon tax was in place. For quantifying the model, it is helpful to work with *net* returns $r_{t+1}(\varepsilon_{t+1}) = R_{t+1}(\varepsilon_{t+1}) - 1$. As long as returns and tax rates are small decimal numbers, their products are an order of magnitude smaller. We can then write the return schedule that replicates the carbon tax rate τ as

$$r_{t+1}(\varepsilon_{t+1}) - r_{t+1}(0) \approx \tau_{t+1}\varepsilon_{t+1}. \quad (18)$$

In other words, the premium on a dirty firm, defined as the spread between its net return and the clean net return, is given by the carbon tax rates.

Importantly, for any carbon tax *rate*, there is an entire replicating return *schedule*, and not only, say, one number for every firm. Of course, in any equilibrium, the cross section of returns consists of one number per firm, namely the return schedule evaluated at the equilibrium emission intensity. However, shareholder value maximization responds to incentives for *any* potential production plan. The price schedule values firms at off-equilibrium plans, and the return schedule records how the cost of capital at different plans varies with the emission intensity. The observed distribution of costs of capital then reflects not only the carbon tax rate being replicated, but also the features of the environment that shape the equilibrium distribution of intensities.

2.3 Discussion

In this section, we discuss the assumptions of our model and clarify that a number of features that our basic model abstracts from are not essential for asset markets to provide the same incentives as a carbon tax.

Boundaries of the firm. Our decentralization assumes that every technology is operated by a separate firm. However, the equivalence result also extends to *multigood* firms, as long as emission intensity is defined as total emissions divided by total firm value. Intuitively, with frictionless financial markets, substitution of technologies across firms operating a single technology is the same as substitution by an individual firm across multiple technologies it operates. To illustrate this, we introduce multigood firm types $m \in \mathcal{M}$, where a firm of type m produces all intermediate goods $n \in \mathcal{N}^m$, a subset of the set of all goods.

A multigood firm of type m chooses a production plan $X = (X^n)_{n \in \mathcal{N}^m}$ that specifies inputs for each technology n the firm operates. The firm's date $t + 1$ value for a given plan X is the sum of cum-dividend values (8) contributed by the individual technologies. The firm's overall value is thus

$$w_{t+1}^m(X) = \sum_{n \in \mathcal{N}^m} v_{t+1}^n(X^n).$$

The new margin is that the firm at date t makes a portfolio choice over how to allocate capital across the technologies n it operates.

On the asset demand side, nothing changes, so the household Euler equation for equity (10) continues to hold. For a given production plan X , the equity price of a multigood firm t depends on its future value $w_{t+1}(X)$ and its overall emissions $e(X) = \sum_{n \in \mathcal{N}^m} e^n(X^n)$. The firm thus solves

$$\max_X \left\{ - \sum_{n \in \mathcal{N}^m} p_{t,n}^k (k^n - (1 - \delta)k_t^n) + q_t(w_{t+1}^m(X), e(X)) \right\}. \quad (19)$$

This problem is directly analogous to the single-good firm problem (11). The only difference is that values and emissions are now sums over technologies.

Since the price function q_t is linear in both value and emissions, the problem of the multigood firm is separable: it chooses production plans in the same way as if there were a separate firm for each technology. It follows in particular that all of our propositions go through with arbitrary firms. In other words, incentives from asset markets induce substitution within multigood firms in the same way as substitution across firms that operate a single technology. Moreover, the cross section of rates of return depends on each firm's emission intensity, defined as total emissions over total value.

Adjustment costs. To capture that the choice of capital for technology n is subject to adjustment costs, we allow the investment function $h^n(z_t, k_t^n)$ to depend on previously installed capital k_t^n . While the cum-dividend value (8) looks as before, the expression for shareholder value takes into account that the firm chooses inputs for investment z in addition to the production plan X :

$$V_t(k_t^n) = \max_{(X,z):k=(1-\delta_n)k_t^n+h^n(z,k_t^n)} \left\{ - \sum_{j \in \mathcal{N}} p_{t,j} z_j + q_t(v_{t+1}^n(X), e^n(X)) \right\} \quad (20)$$

The firm must respect the evolution of the capital stock: it selects inputs z so as to put in place its new capital stock k , an element of the production plan X for the next date $t + 1$, given its previously installed capital stock k_t^n . Again, our propositions continue to hold, since the firm problem does not alter valuation by the linear function q_t .

Market power and other taxes. Our model assumes perfect competition in product markets. We also focus solely on carbon taxes and do not allow for any other taxes, such as corporate income taxes. Neither assumption is essential for the propositions. Introducing markups or other taxes would modify the intermediate goods firm problem (11) only through the definition of the firm value (8). As a result, the key step in our derivation that uses the linearity of the price function is unaffected.

Capital tax. An alternative tax scheme that delivers our replicating return schedule is a capital tax with a rate that increases linearly with the firm's emission intensity. With such a tax in place, investor optimality also equates after-tax returns as in (18). The return schedule, and hence incentives for all firms with positive value are therefore the same as under a carbon tax collected from investors. However, a capital tax of this type is only well-defined if the value of

a dirty production plan is nonzero. A carbon tax collected from investors, in contrast, depends on the future level of emissions and is always well-defined, leading to a very simple and general equivalence result. It also naturally connects to preferences used in the climate finance literature which depend on the level of emissions as well.

Risk. While our model so far is deterministic, the propositions extend to a setting with risk. What is important is only that financial markets are frictionless and shareholder value maximization is well defined. Suppose that households have standard expected utility preferences. We can allow for general exogenous shocks to technology: let ω_{t+1} denote the realization of shocks at date $t + 1$. To make the point with minimal extra notation, we assume that shocks take on a finite number of values and that households assign some probability over sequences of states ω_{t+1} that reflects their beliefs. All expectations below are computed using this probability.

In the presence of risk, production plans become partly state contingent because variable inputs—labor l , intermediates x , and foreign goods \tilde{x} —are all functions of the state ω_{t+1} . Capital remains a number k chosen at date t . We write the cum-dividend value of the firm at $t + 1$ as $v_{t+1}(X; \omega_{t+1})$, where the last argument captures the direct effect of exogenous technology shocks on profit. These shocks could, for example, influence productivity in the firm's production function. Since emissions are an output from production, we also allow them to depend on the state and write $e_{t+1}(X; \omega_{t+1})$. We note that carbon taxes collected directly from firms $\phi\tau_{t+1}e_{t+1}(X; \omega_{t+1})$ are then a possibly random component of cum-dividend value.

With risky payoffs and emissions, the cross-section of relevant equity markets must be indexed by *random variables* (v, e) , not simply numbers. At the same time, the household budget constraint (9) takes the same form as before, as long as variables are understood to be random variables. Household optimality (10) then becomes

$$q_t(v, e) = E_t \left[\frac{v - (1 - \phi)\tau_{t+1}e}{R_{t+1}(0)} \right]. \quad (21)$$

The random marginal rate of substitution $R_{t+1}(0) = \beta u'(C_t) / u'(C_{t+1})$ serves as a stochastic discount rate for assets $(v, 0)$ that do not generate emissions—it defines a random clean return. More generally, the value of an asset reflects the stochastic properties of both its payoff and its emission profile.

Shareholder value (11) looks as before, except that discounting by the function q_t now delivers the *expected* present value

$$V(k_t^n) = \max_X \left\{ -p_{t,n}^k (k - (1 - \delta)k_t^n) + q_t(v_{t+1}(X; \omega_{t+1}), e_{t+1}(X; \omega_{t+1})) \right\}.$$

It is immediate that all propositions carry over to the setting with risk. The key property that valuation is linear in value and emissions is preserved, except that it now holds state-by-state.

Border adjustment. Propositions 1-3 show how to replicate a carbon tax on domestically generated emissions. The response to such a tax naturally includes more imports of relatively

dirty goods that become more expensive to make domestically. This phenomenon—known as emissions leakage—can be addressed by also taxing imported goods based on their carbon content. The measurement problem here is nontrivial, as imports may be manufactured goods produced with dirty energy or through dirty industrial processes. If this problem can be overcome, however, then it is possible to replicate a carbon tax with border adjustment: we can modify the asset price schedule to add to emissions the carbon content of the intermediates \tilde{x}_t the firm uses. The firm's value then reflects not only the carbon tax bill but also the additional cost imposed by border adjustment.

Scope 2 emissions. So far, we have focused exclusively on scope 1 emissions—that is, taxation of emissions generated directly through the use of intermediate goods—and we have assumed that electricity use is untaxed. In contrast, scope 2 emissions attribute emissions from electricity generation to the *users* of electricity, rather than the *producers*. It is possible to derive an alternative replicating return schedule with this convention. Suppose we redefine the emission intensity for non-electricity firms to include both scope 1 and scope 2 emissions, computed as the total of these emissions divided by the firm's value. To avoid double-counting emissions, we set the emission intensity of electricity-producing firms to zero, thereby exempting them from carbon taxation. When taxes are collected from investors, we again obtain a return schedule that provides the same incentives as a tax collected from firms.

The distinction between returns based on scope 1 versus scope 2 emissions is the same as taxing electricity at the producer or consumer level, respectively. As long as emissions are not double-counted, the two approaches lead to the same allocations. In a scope 1 approach, dirty electricity firms face higher costs of capital, leading them to charge higher electricity prices, which in turn discourage electricity use. In contrast, in a scope 2 approach, electricity prices are lower, but users are discouraged by higher costs of capital that take their dirty electricity use into account. A practical difference between the two approaches is that measurement of scope 2 emissions requires accurately attributing emissions to electricity use. An advantage of the scope 1 approach we have emphasized is that such an attribution is not necessary.

Scope 3 emissions, fossil fuel producers, and reserves. Scope 3 emissions refer to a broad category of emissions associated not only with a firm's inputs but also with the use of its products. For example, both petroleum producers and car manufacturers are attributed scope 3 emissions from gas-fueled cars. We are interested in equivalence to a carbon tax, which taxes every emission exactly once. From this perspective, the broad concept of scope 3 emissions discussed in policy debates is not suitable, and indeed not necessary. Our approach instead says that for asset returns to provide incentives like a carbon tax, it is sufficient to develop simple scores based on the intensity of scope 1 emissions.

A simple variant of our equivalence result incorporates a portion of scope 3 emissions by increasing returns on fossil fuel producers. Suppose we eliminate taxes on fuel users and instead define the emission intensities of fossil fuel producers based on the downstream emissions generated when their fuels are burned—emissions that are categorized as scope 3 in emission accounting frameworks. We can then derive a set of replicating returns where only fossil fuel

producers pay positive return premia, whereas all fuel users pay the clean rate. However, this approach captures only the emissions associated with domestically produced fuels, including those burned abroad, and not imported fuels. An advantage of a return schedule based on scope 1 emissions is that it is easy to cover all domestic emissions, regardless of where fuels come from.

Returns based on scope 1 emissions are high for fossil fuel producers only to the extent they themselves burn fuel, say, during extraction. Consequently, in our baseline specification, replicating returns on fossil fuel producers are relatively small. This contrasts with much of the existing literature on green investing, which emphasizes divestment from fossil fuels as a key strategy. Moreover, Appendix A.2 presents an example where technology includes exhaustible resources as assets in the spirit of Hotelling. Fossil fuel reserves are then part of the value of fossil fuel companies. Fossil companies still do not pay higher returns (except to the extent that extraction is dirty), but their value falls in general equilibrium as taxes on fuel users lower cash flow.

Abandoning dirty technologies and stranded assets. Some polluting technologies may be phased out entirely in the green transition. For example, a technology may be profitable at low carbon taxes but not when the tax rate is sufficiently high. As a result, assets that are specific to the dirty technology—such as capital that has been irreversibly invested in the technology or assets used to produce inputs for the technology—become worthless. To clarify how asset stranding arises in our framework, consider a simple example: a clean good competes with a dirty good made from fixed capital. We lay out the minimal version of the example somewhat informally here, and present a more formal treatment of an extended version in the appendix.

Suppose the representative agent gets utility from a single consumption good that is made from two intermediate goods 1 and 2 using a linear production function with equal coefficients $C_t = y_t^1 + y_t^2$. Good 1 is made from labor according to the strictly increasing and concave production function $f(l_t^1)$. Good 2 is a Leontief aggregate of labor and a capital good that is in constant fixed supply \bar{k}^2 , so that $y_t^2 = \min\{l_t^2, \bar{k}^2\}$. Its production generates $\bar{\epsilon}$ emissions per unit of output. The only relevant decision in this economy is how much labor (out of the total supply of one unit) to employ in each technology—specifically, whether to allocate any labor to technology 2 or to shut it down entirely by setting $l_t^2 = 0$, in which case the capital good 2 is stranded.

Consider first a world without a carbon tax. To ensure that both intermediate goods are produced in equilibrium, we assume that $\bar{k}^2 < 1$. In other words, there is less capital than labor, so some labor is always allocated to technology 1. We also assume that $f'(1 - \bar{k}^2) < 1$: even when technology 2 gets \bar{k}^2 units of labor, its marginal product of one is still higher than under technology 1. It follows that output from technology 2 is maximized: the equilibrium allocation $l_t^2 = \bar{k}^2$ is determined by the scarcity of capital. The equilibrium wage is $w_t = f'(1 - \bar{k}^2) < 1$ and the aggregate value of production plans for technology 2 is $(1 - w_t + p_t^k)\bar{k}^2$, the sum of capital income and continuation value. Since consumption is constant, the interest rate is $R_t = \beta^{-1}$, so the price of capital is the discounted value of future capital income, $p_t^k = (1 - w_t)\beta / (1 - \beta) > 0$.

In contrast, technology 2 is shut down in equilibrium once there is a carbon tax with a sufficiently large tax rate $\tau > (1 - f'(1))/\bar{\epsilon}$. When $l_t^1 = 1$, prices must be $w_t = f'(1)$ and $p_t^k = 0$. To check that this indeed an equilibrium, consider a firm that contemplates making $y^2 > 0$ units of good 2 at these prices. Since it can get capital for free, it expects a pretax value $(1 - w_t)y^2 > 0$. However, it also anticipates a carbon tax bill $\tau\bar{\epsilon}y^2$. When the carbon tax is collected from investors, the value of the firm at t is

$$q_t((1 - w_t)y^2, \bar{\epsilon}y^2) = \beta (1 - f'(1) - \tau\bar{\epsilon}) y^2.$$

By our assumption on the tax rate, the value of the firm is negative for any plan with positive production and the optimal plan is indeed shutdown, or $y^2 = 0$.

The example clarifies that asset prices can replicate a carbon tax even in a situation where the tax pushes the equilibrium price of dirty capital to zero. The equilibrium value of firms that specialize in good 2 is also zero. Such firms simply hold worthless capital, the only production plan that generates nonnegative value. Firms choose this plan so the market for capital can clear, but they produce no output and no emissions. Since their value is zero in equilibrium, it is not possible to calculate an equilibrium emission intensity for such firms. This does not mean, however, that asset markets don't provide the right incentives: since every production plan firms contemplate is valued by the function q_t , off-equilibrium dirty plans are penalized as if a carbon tax was directly collected from firms.

In Appendix A.2, we extend the example to add fuel as a third good. We assume that the dirty good 2 is made from labor, capital, and fuel, and that fuel is extracted from a finite stock of resources. In addition to allocating labor, the economy must then decide how quickly the resource is extracted, which determines the path of output. With this extension, we can illustrate how asset markets bring about both types of stranded assets often discussed in the literature: a large enough tax not only strands capital that is exclusively usable in dirty production, but also strands fuel reserves that are left in the ground. Even if fuel production is not itself taxed, the price of fuel and hence the value of fuel reserves fall to zero when dirty production ceases.

3 Replicating return premia without adjustment

In this section, we begin our quantitative study of the distribution of return premia that replicate a given carbon tax. From Proposition 1, the key input to this calculation is the distribution of emission intensities. We proceed in two steps. In this section, we measure the *current* distribution of emission intensities and evaluate return schedules at those intensities. The resulting return distributions can be used to interpret results on pollution premia from the empirical literature. If financial markets were currently providing incentives that replicate some carbon tax, then we should see premia of the magnitude predicted by the replication result.

Our quantitative analysis assumes competitive markets and constant returns to scale, so the value of the firm is equal to the value of its capital stock. The observed intensity distribution is then relatively easy to measure: it requires only data on emissions and capital shares; it does

not require a detailed calibration of the model. At the same time, current emission intensities cannot tell us what premia will be earned in equilibrium if a substantial carbon tax were in place. This is because introducing a carbon tax will incentivize firms to substitute away from fossil fuels. Current intensities would only be relevant if all technologies are Leontief, so intensities are constant. Assessing the equilibrium return schedule requires taking a stand on substitution elasticities. Our second step in Section 4 below calibrates the model and provides results on equilibrium adjustment.

A back-of-the-envelope calculation. To get an initial idea of orders of magnitude, we start from the emission intensity per dollar of gross output $\hat{\varepsilon}^n$, a commonly reported number in emissions accounts. The emissions are measured as CO₂e, or CO₂ equivalents, which express the global warming impact of different greenhouse gases relative to carbon dioxide. Across 2-digit NACE sectors in 2017, $\hat{\varepsilon}^n$ ranged between zero and about 14kg CO₂e per dollar, with the highest numbers for water transport, air travel, and electricity. To obtain emissions per unit of capital ε^n , we divide by the capital share in output, and multiply by the user cost of capital ρ_t^n :

$$\varepsilon^n = \hat{\varepsilon}^n \frac{p_{t,n} y_{t,n}}{\rho_t^n k_t^n} \rho_t^n. \quad (22)$$

The user cost or rental rate ρ_t^n on capital is the interest rate plus the depreciation rate.

From Proposition 2, to replicate a carbon tax of \$10 per ton, we need 10 ε^n extra return on sector i . We have $\hat{\varepsilon}_i^n$ ranging from zero to .012 tons per dollar. The capital share in gross output is roughly 39%. With an interest rate of 1% and a depreciation rate of 10%, ρ_t^n is equal to 11%. The dirtiest firms thus require an extra return of

$$10(0.01) \frac{1}{0.39} 0.11 = 0.028.$$

In other words, to provide incentives equivalent to a \$10 higher carbon tax, the dirtiest firms have to worry that the spread can be 2.8% higher.

The example calculation shows that a modest carbon tax does not require large return differentials. However, since the required returns scale with the tax rate, sizable taxes like \$100 per ton call for substantial differentials. Moreover, moving from an intensity relative to output to an intensity relative to capital is not a simple rescaling that is the same for all sectors, say. Instead, for a given intensity relative to output, the translation depends on capital shares and user costs, both of which differ significantly across sectors. We now describe how we collect the relevant data and then present sector and firm-level results on replicating returns.

Sectoral data. We work with the GTAP 11 database from the Global Trade Analysis Project (Aguiar, Chepeliev, Corong and van der Mensbrugghe 2022) to measure costs and emissions at the sector level. This dataset includes a social accounting matrix, that is, a record of annual transactions between country-sector pairs within a unified and consistent framework. The GTAP-Power module contains details on the energy sector that go beyond usual input-output

matrices. In particular, the electricity sector is broken down by the type of fuel used to generate electricity. We also observe physical quantities for transactions of electricity, as well as other forms of energy. The GTAP-E module contains numbers on scope 1 emissions due to fossil fuel use. The primary sources for the data are national account systems, customs statistics, and energy balances. The reference year is 2017. We focus on transactions within the U.S. as well as those between the U.S. and the rest of the world, treated as a single aggregate.

We consider a total of 72 sectors listed in Table 6 in Appendix B that correspond largely to NACE 2-digit definitions. There are 15 primary sectors, 27 in manufacturing, and 18 in services. We distinguish seven subsectors for electricity generation based on fuel: coal, gas, oil, wind, solar, nuclear, and "other". We also have 5 fossil-fuel sectors: coal, natural gas, crude oil, refined petroleum & coke, and gas manufacturing & distribution. We define two additional sectors that describe production within households. A *household vehicles* sector consists of the share of automobiles recorded as household consumption in GTAP, plus household consumption of refined petroleum. To account for electric cars, we add the share of household electricity used for EVs⁹ as part of the energy mix used to power household vehicles. A *housing* sector comprises GTAP's dwellings sector and all other household energy consumption. This approach ensures that household capital is grouped with the energy used to run it, as is the case for firms.

Our quantitative analysis focuses on emissions generated by burning fossil fuels to produce CO_{2e}. We back out sectoral emission intensities relative to capital following (22). For every sector, GTAP provides monetary values for gross output, materials, and wages in 2017 dollars. Consistent with our model, we compute capital income for the sector as output less materials and wages. We assume that the user cost equals a common interest rate of 4% plus depreciation. To obtain sectoral depreciation rates, we merge GTAP to the *Fixed Asset Tables* of 2017 provided by the *Bureau of Economic Analysis*.

Sectoral results. Figure 1 presents sectoral distributions of capital and emissions together with replicating return premia by sector. Along the horizontal axis, we measure emission intensity in tons of CO_{2e} per million dollars. In both panels, every bar represents a sector or group of sectors. The height of the bars in the top panel reflects the sectors' share in the total capital of the economy, while the height in the bottom panel is the sectors' share in total emissions. To provide details for high-emission-intensity sectors, we group many low-emission-intensity sectors into a supersector, "Rest," which is shown to the left in the figure. To handle skewness, we also work with broken axes in the top panel.

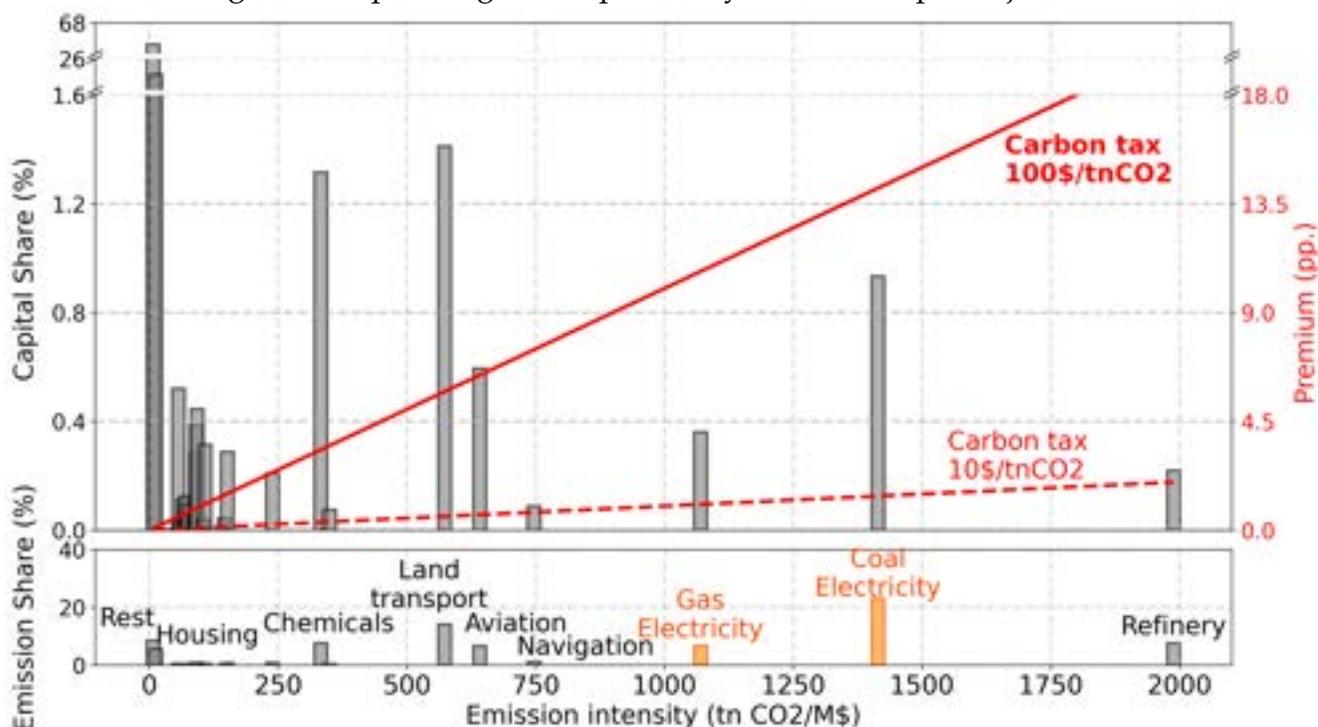
The figure makes two key points that are central to all our quantitative results. First, **the distribution of emission intensities is extremely right-skewed**. Only a handful of sectors, which account for a small share of total capital, sit at emission intensities above 100 tons per million dollars. They consist of refining, oil-fueled transport sectors (navigation, aviation, and land transport), coal- and gas-fueled electricity generation, and a few dirty manufacturing sectors, such as chemicals. Together, these sectors also account for a large share of total emissions.

⁹U.S. Energy Information Administration, *Electric Power Monthly*, December 2023, Table D.1. We use the first available data point, for 2018.

We note that households also contribute significantly to total emissions, especially through vehicle use. The majority of capital, in contrast, is held by relatively clean sectors, such as the large service sectors, which produce few emissions.

The second fact is that the **replicating return premium is small for most sectors but very large for a small number of really dirty sectors**. We plot two lines to indicate premia that replicate carbon taxes of \$10 and \$100 per ton of CO₂e, respectively. These replicating return premia (18) are (approximately) linear in intensity. For a more serious \$100 carbon tax, most capital commands a small premium of only a few basis points. The premium for chemicals is 3.5%, close to the upper bound of estimates on pollution premia in the literature. For the even dirtier transport and power sectors, premia are larger, topping out at 17% for the .2% of total capital that is invested in refineries. For dirty electricity subsectors using gas and coal, premia are 4% and 11%, respectively. Of course, the dirtiest sectors are served by multi-sector firms, which motivates a closer look at the firm level.

Figure 1: Replicating return premia by sector, no input adjustment



Note: The upper figure shows the capital share, the emission intensity ϵ^i , and the replicating premium of each sector. Each sector is represented by a bar, except for sectors with little intensity (Rest). The red solid line describes the premium schedule that replicates a \$100 carbon tax, while the dashed line describes the schedule that replicates a \$10 carbon tax. The lower figure measures the emission share of each sector along the left axis.

Firm- and division-level data for electricity generation. S&P Global Trucost provides firm-level information on electricity production and emissions by type of fuel for publicly traded

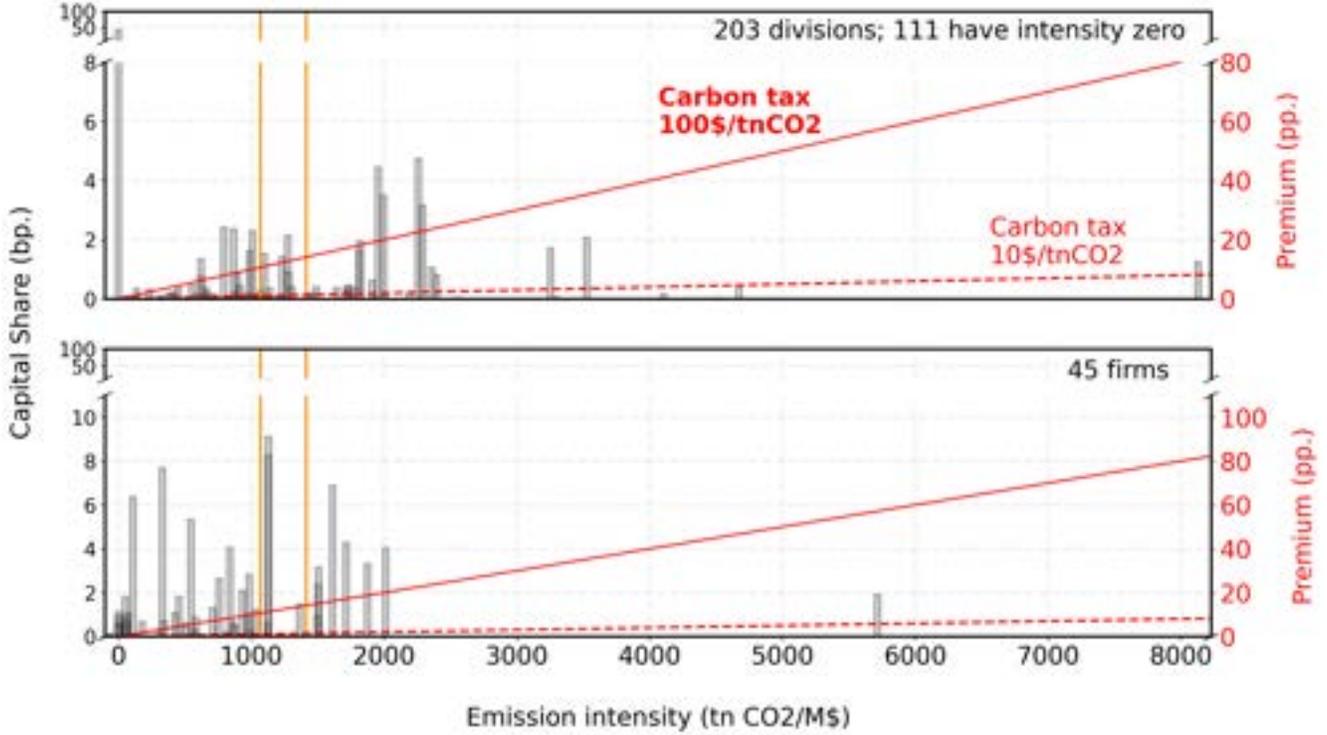
power companies. We consider active U.S. companies with positive electricity generation in 2017 and merge data on those firms to Compustat Fundamentals to add a firm-level measure of enterprise value, that is, the combined value of debt plus equity. Enterprise value is a model-consistent measure of total capital. At the firm level, we thus have a direct measure of the emission intensity relative to capital for 45 firms. In 2017, the 45 firms in our data produced 47% of the electricity in the U.S. as recorded by GTAP. Their enterprise value accounted for 44% of the capital we back out at the sectoral level.

We are also interested in the heterogeneity of power plants that use the same fuel. We define a *division* as the collection of all plants of the same firm that use the same fuel. Unfortunately, we do not have plant-level data on the value of capital. As a rough estimate, we divide the enterprise value of the firm among its divisions in proportion to the quantity of electricity produced. We obtain a sample of 203 divisions, with 111 clean divisions that do not produce any emissions and 92 dirty divisions. Only 3 firms contain only one division—the typical case is thus for one company to rely on multiple fuel types. In particular, 99% of total capital is held by firms that own both clean and dirty (gas, oil, or coal) plants. Those "mixed" firms own 99% of the total clean capital and virtually all of the dirty capital.

Firm- and division-level results Figure 2 presents the distribution of capital and replicating return premia for the electricity sector at the firm level (bottom panel) and division level (top panel). The heights of the bars indicate shares of total electricity sector capital. Most of the capital is employed in clean divisions: the leftmost bar at zero represents 43% of total electricity capital. We mark the sectoral averages for gas- and coal-fueled electricity from GTAP, already recorded in Figure 1, with orange vertical lines. Finally, we add two lines that allow us to read off annual replicating returns for carbon taxes of \$10 and \$100 per ton of CO₂e. We note that the range of the horizontal axis is much larger than at the sectoral level in Figure 1 since the dirtiest electricity firm produces more than four times as many emissions per unit of capital as the dirtiest sector.

We take away two main points from these more granular numbers. First, **within-sector heterogeneity in emission intensities and hence replicating return premia is substantial**. To illustrate, the division-level coefficients of variation in the emission intensity are 79% among gas-powered divisions and 72% among coal-powered divisions. If the dirtiest divisions were stand-alone firms, the replicating premia would thus have to be extremely large. Second, **co-ownership of clean and dirty plants compresses the (equally-weighted) distribution of replicating premia**. The firm-level histogram in the bottom panel shows hardly any capital at zero. While the capital-weighted premia distributions are equal by construction, the equally-weighted distribution of premia has a mean of 9% and an inter-quartile range of 7 pp among firms with positive emissions. In contrast, among divisions with positive emissions, the mean of equally-weighted premia is higher at 14%, and the inter-quartile range is wider at 12 pp.

Figure 2: Replicating return premia in electricity sector, no input adjustment



Note: The plot shows the capital share, the emission intensity ϵ^i , and the replicating premium for each division (upper plot) and each firm (lower plot) within electricity producers. Each division or firm is represented by a bar, except for the bar at zero intensity, which bundles many firms or divisions. For reference, we have added orange vertical lines, which represent sector-level references for gas-powered electricity and coal-powered electricity. The red solid line describes the premium schedule that replicates a \$100 carbon tax, while the dashed line describes the schedule that replicates a \$10 carbon tax.

4 Equilibrium replicating returns

We now compute equilibria of our model and determine equilibrium replicating returns for alternative carbon tax rates. This step requires taking a stand on more details of the production function, in particular, the elasticities of substitution between inputs.

4.1 Calibration and equilibrium adjustment

We assume that the production functions g , \tilde{g} , h^n , and f^n are nested CES functions. Each nest is parametrized by an elasticity of substitution and a set of nest weights that sum to one. Our choices for nests and elasticities follow the MIT U.S. regional energy policy (USREP) model, as described in Yuan, Rausch, Caron, Paltsev and Reilly (2019), another multisector model with additional energy detail. We present details in Appendix E, where we also demonstrate that our results are broadly similar to those of the EMF 32 study on carbon tax scenarios (McFarland,

Fawcett, Morris, Reilly and Wilcoxon 2018), a meta-study conducted by the Energy Modeling Forum using several leading models for U.S. environmental policy analysis. Here we provide a sketch that highlights key properties that drive the results.

Overview of energy use & emissions. To illustrate how energy use and substitution potential vary across the economy, we group sectors into six broad supersectors. Table 1 presents each supersector’s shares of capital and emissions—as in Figure 1—alongside their share of output, energy use, and energy expenditure as a proportion of total revenue. Two of these supersectors are energy producers. The *fossil* sector handles the extraction of primary energy sources, crude oil, natural gas, and coal, refines oil into secondary energy sources such as gasoline, and distributes gas to users. The electricity sector generates and distributes electricity. All forms of energy are measured in a common unit.

Table 1: Supersector overview

Supersector	Shares of aggregate (%)						Revenue shares (%)	
	Emissions	Capital	Coal	Oil	Gas	Electricity	Energy	Electricity
Fossil	9.1	2.3	4.5	52.2	19.0	3.2	49.3	1.9
Electricity	30.3	2.6	91.4	0.5	26.8	5.4	22.8	5.4
Transport	22.0	2.1	0.0	23.3	2.9	1.5	16.7	0.7
Dirty man.	9.5	2.1	2.2	5.3	13.9	6.3	7.5	2.7
Households	18.5	25.3	0.0	14.1	17.8	33.6	11.5	6.5
Other	10.7	65.7	1.8	4.5	19.6	50.0	1.0	0.8

Note: Summary of energy by user and producer. Shares of aggregate report the share of each supersector to the aggregate variable in each column. Share of sector revenue report input bills as a share of revenue of energy and electricity for each supersector.

At this broad level of aggregation, the pattern of pollution in the U.S. economy is relatively simple. Energy producers generate emissions by burning coal to generate electricity and by refining oil. Oil products are burned by households and firms for transportation. Gas is used throughout the economy, most intensively in electricity and "dirty manufacturing", defined as manufacturing sectors with high emission intensities, notably chemicals, iron and steel, non-ferrous metals, and other mineral products. Both households and the "other" supersector, which comprises mostly service-sector firms, use substantial amounts of gas for heating and also account for almost all electricity use. Energy expenditures represent only a small share of total revenue outside the energy sector itself.

Production functions in a typical energy-using sector allow only limited substitution between variable inputs. In particular, there is little scope for electrifying transport or manufacturing. As a result, energy users primarily adjust by reducing their overall energy consumption. The electricity sector is an exception: it features competition between varieties produced with very different emission intensities. The production function for final electricity takes into account the cost of the power grid, constraints on hydro and nuclear power, and the intermittency of renewables. Still, it allows for substantial "cleaning" of electricity production. Standard trade

elasticities enable relatively easy substitution between domestic and imported goods. Finally, because energy represents a small share of total revenue, energy-price changes have only a minor effect on output in most sectors.

Equilibrium allocations. We conduct comparative statics across steady states that differ only in the level of the carbon tax. The top panel of Table 2 reports aggregate statistics for carbon taxes of \$10, \$50, and \$100 per ton. In our benchmark case of a \$100 carbon tax, emissions fall by nearly 50%, accompanied by a modest 2% decline in output. We break down this emission reduction into three components: percentage changes in capital K , energy intensity (final energy use F relative to capital), and emissions E per unit of final energy.¹⁰ The aggregate economy lowers emissions through a slight contraction combined with substantial energy savings—a 12.5% drop in the aggregate emission intensity—and most importantly a shift towards cleaner energy. Aggregate electricity declines by 9%.

Table 2: Aggregate effects under different carbon taxes

Panel A: Aggregate effects under different carbon taxes							
Carbon Tax (\$)	Percent Changes				Components of Emission Reduction		
	ΔC	ΔY	Electricity	E	ΔK	$\Delta F/K$	$\Delta E/F$
10	-0.1	-0.2	-1.8	-9.6	-0.4	-3.4	-6.0
50	-0.7	-1.1	-6.4	-33.3	-1.7	-10.0	-24.6
100	-1.4	-2.0	-8.8	-48.4	-2.7	-12.5	-39.3

Panel B: Decomposition by supersectors at \$100 carbon tax							
Supersector	Share	Emissions		ΔK	Components		Output ΔY
		$\Delta E/\text{total } E$	$\Delta E/E$		$\Delta F/K$	$\Delta E/F$	
Fossil	9.1	-5.1	-55.9	-29.7	12.3	-44.2	-32.9
Electricity	30.3	-26.2	-86.3	38.1	-20.3	-87.6	-8.8
Transport	22.0	-5.0	-22.6	-4.7	-18.7	-0.1	-8.2
Dirty manuf.	9.5	-3.1	-32.3	-5.0	-26.0	-3.7	-7.4
Households	18.5	-5.5	-29.6	-2.8	-22.3	-6.8	-6.9
Other	10.7	-3.6	-33.8	-3.3	-19.4	-15.1	-1.6

Note: Panel A reports effects of carbon taxes on macroeconomic variables and decomposition of emission reductions in changes in capital, changes in energy intensity, and changes in emission intensity as percentage point contribution to the overall emission reduction. Panel B decomposes the emission reduction by supersectors. The first part reports emissions reduction as a share of total economy emissions and as a share of own emissions. The second part reports changes in capital, changes in fuel use per capital, and changes in emissions per fuel use. The third part shows output reduction in each supersector.

The bottom panel of Table 2 shows how different sectors contribute to the overall emission reduction. The electricity sector alone accounts for roughly half of the total decline. It achieves this mostly by substituting towards cleaner fuels, rather than saving energy or shrinking—in fact, its capital stock grows substantially. In contrast, the fossil sector contracts sharply.

¹⁰We report percentage changes rather than log changes, so there is a residual not shown in the decomposition. This residual can be significant when changes are large. Final energy represents all energy used, including from secondary sources such as gasoline or electricity.

Non-energy sectors maintain roughly the same scale but lower emissions, mainly by reducing energy consumption. The "other" and household sectors partially substitute gas with electricity. However, households and the transport sector continue to rely on oil and thus substitute little.

Smaller carbon taxes lead to correspondingly smaller reductions in emissions and economic activity. When we lower the tax rate from \$100 to \$50 (or \$10), the resulting losses in consumption and output scale roughly in proportion—declining by about one-half (or one-tenth). Emissions, however, fall less proportionately, decreasing by only one-third (or one-fifth). This nonlinearity reflects the skewed distribution of emissions and emission intensities across sectors, which makes the marginal emission benefit of targeting the dirtiest sectors disproportionately larger than that of other sectors. Moreover, the decomposition of emission-intensity changes reveals that energy savings (a lower energy intensity relative to capital) contributed more prominently to emission reductions at lower intensities. In contrast, fuel substitution becomes more important at higher intensities.

Our model includes no adjustment costs beyond a one-period time-to-build constraint. Along a transition path, gradual adjustment would therefore only come from consumption smoothing. Since aggregate capital moves little even in response to a large carbon tax, the transition is likely to be fast. From a macro perspective, the introduction of the tax is more like a short recession than a structural change requiring a long adjustment period. The adjustment that generates the large reduction in emissions is sectoral reallocation, particularly through the substitution of variable inputs such as primary energy. This helps explain why even models with richer adjustment costs, as featured in the EMF 32 study, tend to exhibit relatively fast transitions, with most environmental gains realized after a few years (see Appendix E for references and further details). We thus interpret our steady state comparison not as characterizing long-run outcomes, but rather as capturing the economy's response over a 5-10 year horizon.

4.2 Equilibrium returns

Equilibrium replicating returns are significantly smaller than replicating returns without endogenous adjustment of emission intensities. The reason is that a carbon tax creates strong incentives to substitute away from dirty fuels. Table 3 shows equilibrium replicating returns under a \$100 carbon tax and compares them to returns without adjustment. The leftward shift of the return distribution is clear, and the reduction in premia is especially pronounced for dirty firms in the upper tail. These firms adjust relatively more: for example, the median premium drops from 6bp to 4bp, whereas the 95th percentile drops from 333bp to 61bp. The bottom rows of the table report premia for the electricity sector, where the adjustment is particularly dramatic: premia in the right tail decline by more than 10 percentage points.

To get intuition for what sectors see larger differences between equilibrium returns and those without adjustment, consider the example of a sector with a CES production function with equal weight on only two inputs, capital and some fossil fuel j that trades at the price p_j , say. The premium for a firm producing at intensity ε is $\tau \varepsilon$. Let ψ_j be the derivative of the firm's emission function with respect to the fossil fuel j . Differentiating the first-order condition of such a firm,

Table 3: Premia with and without equilibrium adjustment

Quantiles	50th	75th	90th	95th	99th
No Adjustment, Sectors	6	13	45	333	642
Equilibrium, Sectors	4	7	36	61	465
No Adjustment, Electricity Firms	975	1501	1613	1876	2013
Equilibrium, Electricity Firms	21	74	109	165	368

Note: Premia in basis points. All premia calculated from a \$100 tax per ton of CO₂e.

a small increase in the carbon tax leads to a change in its return premium by a factor

$$\frac{d(R_t(\varepsilon) - R_t(0))}{d\tau} = \varepsilon \left(1 - \sigma \frac{\tau\psi_j}{p_j + \tau\psi_j} \right). \quad (23)$$

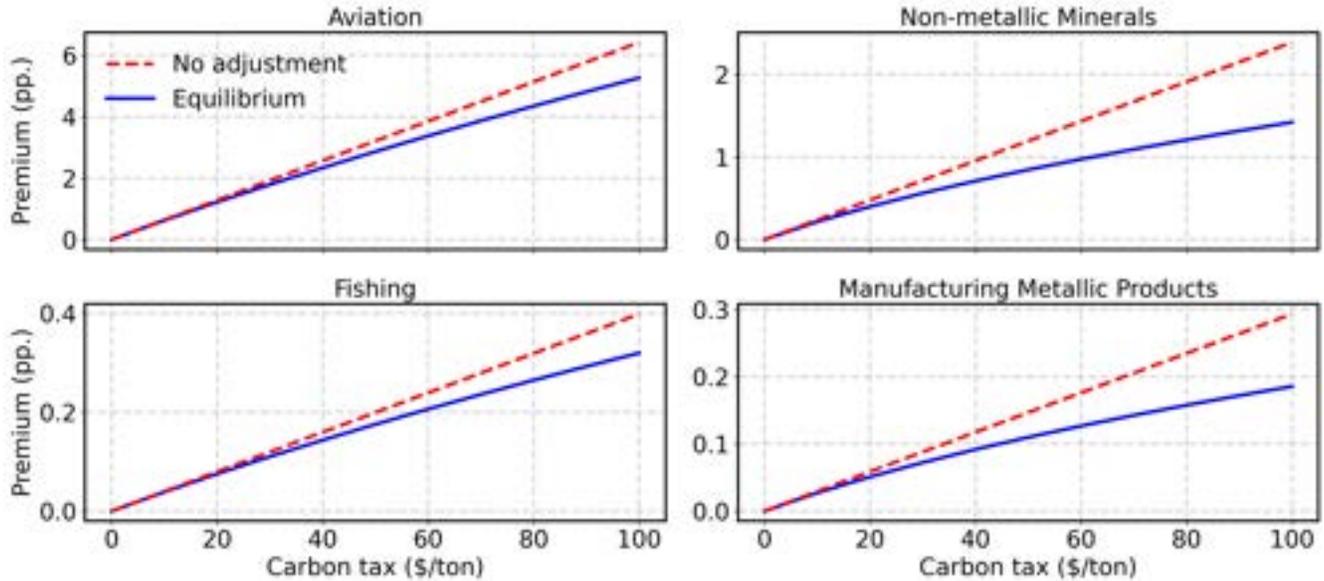
A direct effect increases the spread by ε . In addition, equilibrium adjustment lowers the equilibrium intensity, adding a second, negative effect.

The size of equilibrium adjustment increases with two features of technology. The first is substitutability of inputs: a higher σ means that when fuel becomes more expensive, firms substitute more easily towards cheaper inputs and thereby lower their emission intensity more. As a result, the extra spread required to replicate a higher carbon tax is smaller. The second feature is the share of carbon taxes in total fuel cost. If the carbon tax represents a larger share of the total fuel cost, then a change in the tax rate translates into a larger percentage change in the fuel cost and, hence, a larger reduction in intensity. Holding fixed the production function, a firm using dirtier fuel with a lower spot price thus pays a lower replicating spread in equilibrium.

Figures 3 and 4 illustrate the shape the cross section of replicating premia for firms that use many more inputs than the simple firm characterized by (23) and that are faced with larger changes in the carbon tax. Figure 3 considers several non-energy sectors, while Figure 4 looks at electricity. In each panel, the solid blue line shows the sector-specific premium as a function of the carbon tax rate, measured along the horizontal axis. The slope and hence the shape of the curve can be understood from (23). At zero initial tax, a small tax increase has a negligible effect on the fuel cost, so the response is the same as without any adjustment ($\sigma = 0$). The slope of the curve at zero, therefore, isolates the direct effect, which is linear in the tax rate, and is drawn as a dashed red line. At low taxes, a linear approximation is fairly accurate. At higher taxes, however, adjustment begins to matter, and the blue curve diverges from the red line.

The figures clarify the importance of fuel cost for adjustment. Coal stands out because its emission factor is high relative to its spot price—burning coal generates emissions more cheaply than burning natural gas, and especially more cheaply than burning oil. As a result, coal-using sectors exhibit a particularly large equilibrium adjustment in returns. To replicate a \$100 carbon tax on coal-based electricity, for example, requires an 11% premium without adjustment, but less than 5% in equilibrium. This reduction is much larger than for gas-fueled electricity, which has similar substitution elasticities between fuel and other inputs. Oil-using

Figure 3: Replicating return premia as a function of the carbon tax rate



Note: Replicating premia in percent per year for different carbon tax rates (horizontal axis) and selected sectors. Dashed red lines show premia without adjustment, solid blue lines show premia at equilibrium intensities.

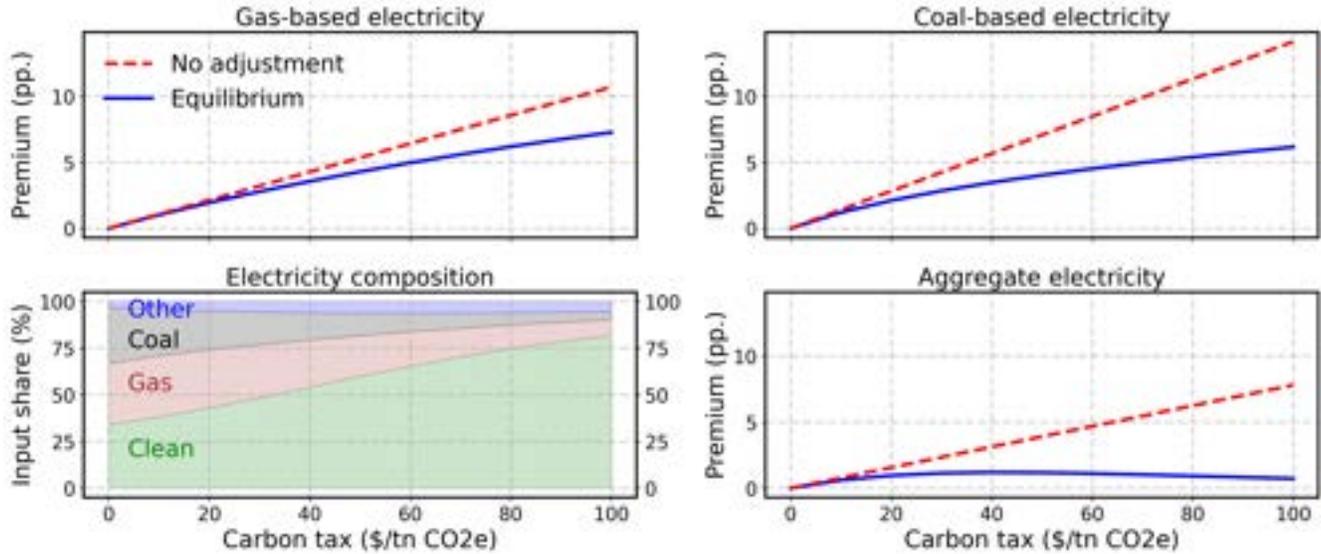
sectors such as aviation and fishing in Figure 3 show minimal adjustment, and equilibrium premia resemble those without adjustment.

The bottom left panel of Figure 4 summarizes substitution of the electricity sector across fuel types. As expected in models like ours, a lot of substitution is possible by reducing dirty power generation and increasing clean power, especially wind and solar. The bottom right panel shows replicating premia for the aggregate electricity sector. They capture both the adjustment at the fuel-type level and the shift to clean fuels. Overall, we obtain a very large adjustment effect: even high carbon taxes require only small replicating premia. In fact, the equilibrium premium schedule is decreasing beyond a \$50 carbon tax. Intuitively, when the potential for substitution is high, endogenous cleaning of the electricity sector can generate lower equilibrium premia.

Table 4 shows how different margins of adjustment in the electricity sector contribute to the overall reduction in emissions, or the case of a \$100 carbon tax. For each fuel type, we decompose the overall change in emissions into three components. We first ask how much emissions would change from only capital reallocation across firms, holding fixed firms' capital allocation to different technologies and their emission intensities. This component is positive for all fuel types; that is, initially, dirtier firms end up with more capital as a result of the carbon tax. The numbers are also small, indicating that reallocation of capital across firms of different average intensity is not an important effect.

The second component is reallocation of capital across divisions within a firm. This is clearly where the key emission reductions come from: firms substitute away from coal and gas by shutting down dirty plants, shrinking those divisions, and ramping up investment in

Figure 4: Replicating return premia as a function of the carbon tax rate



Note: Replicating premia and composition of production in electricity sector. Top line shows replicating premia in percent per year as a function of the carbon tax rate for gas- and coal-powered electricity sectors, bottom right panel shows replicating premia for the aggregate electricity sector. Dashed red lines show premia without adjustment, solid blue lines premia at equilibrium intensities. Bottom left panel shows shares of total electricity output contributed by subsectors as a function of the carbon tax rate.

renewables. By construction, growth in renewables here reflects wind and solar power, as we have constrained the capital stock of nuclear and hydro power. An experiment that lifts these constraints implies a large shift into nuclear and hydro instead, resulting in slightly lower electricity prices, but otherwise relatively similar results. The final component is the change in intensity, which contributes little, as one would expect for subsectors that specialize in particular fuels.

Table 4: Sectoral Contributions to Emissions Changes

Electricity Type	Firm capital	Division Share	Intensity	Total	Emission share (baseline)
Coal	7	-72	-5	-70	76
Gas	9	-25	-2	-18	22
Other	1	0	0	1	2
Total	17	-97	-7	-86	100

4.3 The role of technology, trade and taxes on households

In Table 5 , we present results on aggregates and equilibrium under a number of alternative scenarios to highlight the role of some of our assumptions.

Electricity. Since substitution within and towards electricity is a key force in emission reduction, Table 5 reports results for alternative electricity production functions. First, we make it easier for energy users to substitute electricity for other fuels. The *high electrification elasticity* scenario raises the substitution elasticity to one, more than double the baseline value of 0.4. Electricity production declines by less than in the baseline, and emissions decline by an extra 3pp, indicating a relatively modest sensitivity to this parameter. Second, the *low intermittent electricity* scenario reduces the elasticity of substitution between intermittent (wind and solar) and baseload electricity to zero. As cleaning of electricity becomes more expensive, electricity falls by more and emissions increase by 10pp less than in the baseline. The contribution of the electricity sector to emissions reduction is halved, although the overall reduction still amounts to a substantial 38.9pp.

Table 5: Scenario Analysis Results

Scenario	Output (%)	Emissions (%)	Premia (bp)		Elec. firms (bp)		Extra Metrics
			p50	p95	p50	p95	
Carbon Tax 100	-2.0	-48.4	4	70	21	165	Imported Carbon 4% Electricity -9% Nuclear Energy 0% Solar Energy 684%
High Electrification Elasticity	-2.0	-51.1	3	72	22	200	Electricity -4%
Low Intermittent Elasticity	-2.3	-38.9	4	143	334	596	Electricity -13%
Flexible Nuclear	-2.1	-50.4	4	57	19	73	Nuclear Energy 280%
Solar Innovation	-0.3	-52.2	3	56	0	0	Solar Energy 6997%
High Trade Elasticity	-1.5	-52.7	4	70	21	165	Imported Carbon 5%
Low Trade Elasticity	-2.1	-47.4	4	70	21	165	Imported Carbon 4%
No Household Taxed	-1.3	-43.1	4	70	22	165	-
Only Electricity Taxed	-0.5	-26.4	-	-	23	164	Electricity -8%

Note: Output and emissions changes are in percentage change relative to no carbon tax. Premium values are in basis points (bp.). p50 and p95 refer to the 50th and 95th percentiles of the premium distribution, respectively. Imported carbon is total emissions generated in foreign countries from the production of U.S. imports, relative to domestic emissions without carbon tax. Electricity is the percentage change in electricity consumption, relative to no carbon tax. Nuclear and solar energy are the percentage change in nuclear energy consumption, relative to no carbon tax.

Technical progress in electricity. We also explore two scenarios for technical progress that lower the cost of expanding clean electricity. The *flexible nuclear* scenario removes a constraint on the capital stocks of nuclear power that we maintain in the baseline, along with a constraint on hydro, motivated by the recent stagnation in these technologies. If nuclear power is instead unconstrained, it emerges as a competitive clean energy source, leading to a 2pp extra reduction in emissions. Nuclear generation increases substantially, partially replacing wind and solar, which play a larger role in the baseline case. The second *solar innovation* scenario assumes a technological improvement in solar power, modeled as a 10% increase in solar TFP. This results in a 4pp additional reduction in emissions, highlighting the potential of continued innovation in renewables to support decarbonization.

While all specifications in Table 5 study the same \$100 carbon tax, alternative assumptions on technology matter for equilibrium replicating returns. Interestingly, higher emissions need

not go along with higher premia in the right tail. For example, high electrification makes dirty electricity more profitable, so more dirty capital is retained and premia in the right tail are higher than in the baseline, although emissions overall are lower. Substitution to cleaner power shows up as a 1bp reduction in premia at the median, a small reduction spread over many firms. In contrast, when clean electricity is more expensive because of its intermittency, then dirty electricity firms produce more and their premia rise along with emissions. Similarly, when technical progress crowds out dirty electricity firms, the right tail of premia flattens out.

International trade. Since our exercise introduces a carbon tax in the U.S. only, while holding the prices of foreign goods fixed, it makes sense to import dirty goods that are now more expensive to produce domestically. To roughly assess the resulting leakage of emissions, we define a notion of "imported carbon" that can be computed with the data we have: we sum up all extra imports after the tax is imposed and compute emissions that would have occurred if those goods had been produced in the U.S. at the old pretax emission intensities. For the baseline exercise, imported carbon amounts to 3.9% of pretax emissions. When we increase the Armington trade elasticity to 8 from its baseline value of 2.5 in the *high trade elasticity* scenario, domestic emissions decline by 4pp via substitution to foreign goods, while imported carbon increases by 1pp. We take away that leakage is present but relatively small. Moreover, it does not affect replicating returns, which remain unchanged across scenarios. The *low trade elasticity* scenario reduces the Armington trade elasticity to 1, the Cobb-Douglas case. Here, the domestic emissions increase by 1pp, while leaving imported carbon unaffected.

Selective taxes. A final set of robustness checks removes carbon taxes on select sectors. When we remove the tax on households in the *no household taxed* scenario, the emission reduction declines by about 5pp. When we remove taxes on all sectors but electricity in the *only electricity taxed* scenario, we retain an emission reduction of 26pp. In both cases, the numbers are close to the sectors' contribution to the baseline result. Moreover, replicating returns change little. We conclude that decompositions of the effects into households, industry, and power sector are close to linear, and not subject to strong interaction effects. These results show that if a return schedule can reach only firms, and not households, it nevertheless generates substantial emission reductions. Moreover, a tax that focuses on electricity can reduce emissions without large spillovers on other sectors—we return to this point when we turn to heterogeneous agent economies where some agents obtain nonpecuniary benefits from owning electricity capital.

5 Heterogeneous tastes for clean assets

We now extend our model to heterogeneous investors. We distinguish between two types of preferences. Agents with *nonpecuniary costs* experience discomfort when their portfolio contains equity in dirty companies. In contrast, agents with *nonpecuniary benefits* experience pride from holding particularly clean portfolios. We further allow for *sector-specific preferences*, where agents may derive utility from holding clean assets in specific sectors, such as clean electricity firms. As a result, equity markets in this section distinguish firms not only by their values and emissions, but also by sector. A portfolio is thus represented as a collection of nonnegative equity holdings,

where $s_{t,n}(v, e)$ denotes the equity held at time t in firms in sector n with value v and emissions e .

5.1 Extending the model

Our main interest is in the effects of cross-sectional heterogeneity in investor preferences on aggregate emissions. To isolate this mechanism, we shut down any dynamics of the wealth distribution that would result from differential portfolio performance. Specifically, we adopt the common modeling device of a large family that pools resources. We now describe the optimization problem of the family, which delivers asset pricing conditions and portfolio allocations for family members with different tastes for clean investment. The supply side of the model remains the same as above, so the definition of equilibrium changes only minimally. For simplicity, we assume throughout this section that all capital goods have the same price as the final consumption good.

Household sector: a large family of investors. More formally, a continuum of households of mass one constitutes a large family. Individual member households differ in their concern for the environment. A share κ^g of *green* households dislike equity more if the firm generates more emissions and also may have a preference for particular sectors. The felicity of a green household who consumes c_t^g and holds a portfolio with shares $s_{t,n}^g(v, e)$ of firms in sector n with values v and emissions e is given by

$$u \left(c_t^g - \gamma \sum_{n,v,e} es_{t,n}^g(v, e) + \gamma \sum_n \varepsilon_n^* \sum_{v,e} vs_{t,n}^g(v, e) \right) - v\eta_t, \quad (24)$$

where the scalar $\gamma > 0$ governs the strength of the green investment motive and $\varepsilon_n^* > 0$ captures preference for sector n .

In addition to pecuniary payoffs, ownership of firms affects green investors in two ways, represented by the second and third terms that enter u . The second term captures *nonpecuniary costs*, or disutility, from emissions, as introduced in the felicity (16) of the representative agent in Proposition 3. The third term is a preference for holding claims on firms in certain sectors n . Here, the weight ε_n^* on sector n multiplies the total value of shares invested in sector n . This specification reflects investors' ambition to "clean up" particularly dirty sectors. In our quantitative example below, this sector is electricity. With sector-specific preferences, a green investor may prefer a dirtier firm within their favored sector n to a cleaner firm in another sector. We assume that there are always some sectors n such that $\varepsilon_n^* = 0$. The functional form treats consumption and the nonpecuniary benefits of green investing as *perfect substitutes*, which rules out wealth effects and isolates the role of taste heterogeneity.

The remaining share of households, $\kappa^t = 1 - \kappa^g$, are *traditional* households who derive utility solely from consumption, with felicity $u(c_t^t)$. The family pools resources and chooses the aggregate capital allocation s_t , but individual members are responsible for investing the same per capita amount of that total. Specifically, each member must invest the same per capita

amount s_t . Type i households together invest $\kappa^i s_t$ where i indexes green and traditional taste. Formally, equity holdings for each type must satisfy:

$$\kappa^i \sum_{n,v,e} q_{t,n}(v,e) s_{t,n}^i(v,e) = \kappa^i s_t \quad , i = g, t. \quad (25)$$

Prices now condition on the sector as well as value and emissions. Family members' nonpecuniary costs or benefits in (24) depend only on investments that they themselves are responsible for. We also impose short sale constraints $s_{t,n}^i(\varepsilon) \geq 0$ on all family members and types of capital.

The family's objective is to maximize the equally weighted sum of utilities across its members. At date t , the family must respect the aggregate budget constraint:

$$\kappa^g c_t^g + \kappa^t c_t^t + s_t = \sum_{i=g,t} \kappa^i \sum_{n,v,e} v s_{t,n}^i(v,e) + w_t + T_t, \quad (26)$$

where the left-hand side is the overall consumption (by both green and traditional members) plus total savings, and the right-hand side represents the total available resources—payoffs from all portfolios chosen by members in the previous period, wage income, and government transfers. The family maximizes its objective by choosing capital and consumption subject to the sequence of members' equity constraints (25) and budget constraints (26).

Portfolio choice and savings. Let λ_t denote the Lagrange multiplier on the family's date t budget constraint (26). Since the family maximizes an equally weighted sum of member utilities and treats all members of type i symmetrically, λ_t also represents the common marginal utility of consumption across members. Let μ_t^i denote the multiplier on type i 's date t equity constraint (25). As is standard in models with short-sale constraints, the optimality conditions for asset holdings take the form of Euler *inequalities*. In particular, the first-order conditions for equity in a firm in sector n with future value v and emissions e imply

$$\mu_t^t q_{t,n}(v,e) \geq \beta \lambda_{t+1} v, \quad \mu_t^g q_{t,n}(v,e) \geq \beta \lambda_{t+1} (1 + \gamma (\varepsilon_n^* v - e)). \quad (27)$$

The left-hand side of each inequality is the marginal cost of buying a share in the firm. It must be greater or equal to the marginal benefit on the right-hand side. The conditions hold with equality if family member i invests a positive amount in the firm.

It is convenient to characterize equilibrium in terms of emission intensities and returns, as opposed to levels of emissions, payoffs, and prices. Every firm that operates in equilibrium is owned by some investor type, so one of the two conditions holds with equality. It follows that the price function is linear in (v, e) for all operating firms. The return $v/q_{t,n}(v, e)$ therefore depends only on the emission intensity $\varepsilon = v/e$. We can write the conditions for equity in a firm in sector n operating at emission intensity ε as

$$\mu_t^t \geq \beta \lambda_{t+1} R_{t+1,n}(\varepsilon), \quad \mu_t^g \geq \beta \lambda_{t+1} R_{t+1,n}(\varepsilon) (1 + \gamma (\varepsilon_n^* - \varepsilon)). \quad (28)$$

For traditional investors, the benefits from investment are purely pecuniary and hence given by the financial return multiplied by the present value of resources next period $\beta\lambda_{t+1}$. For green investors, by contrast, there are additional nonpecuniary benefits or costs that depend on the emission intensity and the sector of the firm. A convenient feature of our functional form for utility is that the nonpecuniary component is also proportional to the shadow value λ_{t+1} . Pooling of resources by the family implies that the effect of taste heterogeneity is purely myopic: green family members are compensated for the disutility from dirty capital only for one period, or pay for the privilege of investing in electricity only for one period.

The family's optimal savings and consumption decision depends on the average return on total household wealth R_{t+1}^w implied by individual households' decisions, that is, the value-weighted average returns on all firms. Let $\bar{\varepsilon}_{t,g}$ denote the average emission intensity when equity is weighted by green investors' portfolio weights and let $\bar{\varepsilon}_{t,g}^*$ denote the average of the preference parameters ε_n^* , weighted by green investors' portfolio weights on the different sectors n . Combining first-order conditions, we then have

$$\lambda_t = \beta\lambda_{t+1} \left(R_{t+1}^w + \gamma\kappa^g \left(\bar{\varepsilon}_{t,g}^* - \bar{\varepsilon}_{t,g} \right) \right). \quad (29)$$

The average effective return (including the average nonpecuniary component) must equal the aggregate marginal rate of substitution $\lambda_t/\beta\lambda_{t+1}$. In steady state, this MRS equals the discount rate β^{-1} . The pecuniary return R_{t+1}^w is higher (lower) than the discount rate if green investors' nonpecuniary costs (benefits) dominate *for the average firm held by green investors*. Nonpecuniary costs require returns higher than the discount rate to compensate green investors for the disutility incurred from investing in dirty capital, meaning that green preferences work *as if* a carbon tax were collected from investors. Green investors who obtain nonpecuniary benefits, in contrast, require less compensation. In this case, green taste acts like a capital subsidy, reducing the required financial returns on clean investments.

Green representative agent. When all investors have green preferences, $\kappa^g = 1$, they hold all equity. In this case, the green asset pricing equation—the second inequality in (28)—holds with equality for all firms. As a result, all effective returns, including the nonpecuniary component, are equal to the clean rate on a firm with zero emissions

$$R_{t,n}(0) = R_{t,n}(\varepsilon) (1 + \gamma(\varepsilon_n^* - \varepsilon)).$$

The key implication is that financial returns must be higher on dirtier firms. The parameter ε_n^* shifts the entire return schedule of sector n . A higher ε_n^* lowers the return in that sector, effectively subsidizing investment. With $\varepsilon_n^* = 0$, we obtain the special case of Proposition 3, a green representative agent who generates the same return schedule as a carbon tax. Since the preference parameter γ plays the same role as the carbon tax rate, we measure it in what follows in dollars per ton of CO₂e.

Heterogeneous agent portfolios. More generally, the investor optimality conditions (28) provide a simple characterization of equilibrium portfolios. First, all firms held by traditional investors must deliver the same return, and this return must exceed the return on any firm held exclusively by green investors. This follows from the first equation in (28), which says that the ratio $\mu_t^t / \beta\lambda_{t+1}$ sets an upper bound for all returns. Since traditional households hold some capital in equilibrium, the upper bound must bind for all the firms they fund. Second, by a similar argument, all firms held by green investors must earn the same *effective* return. Because the nonpecuniary component of green investors' effective return declines with emission intensity, firms with higher emission intensities must deliver higher financial returns to green investors.

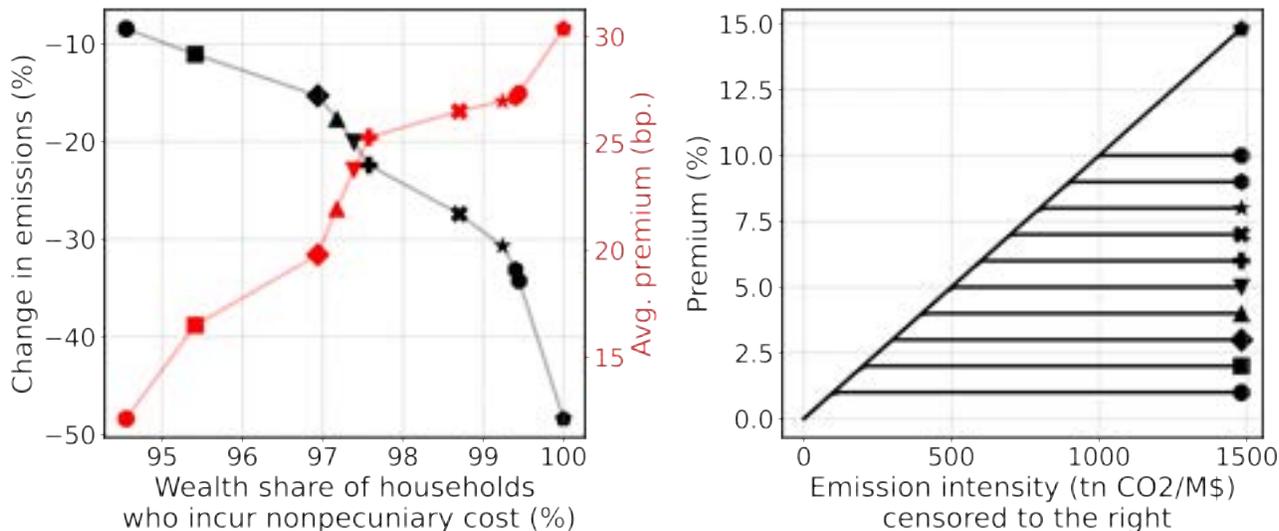
Together, these two properties imply a clear segmentation of firm ownership within each sector. Traditional investors only hold the dirtier firms, while green investors only hold the cleaner ones. As a result, the equilibrium return schedule within a sector is piecewise linear, with a kink at some threshold intensity, denoted $\bar{\varepsilon}_n$. Below this threshold, returns rise with emission intensity at a slope of γ ; above the threshold, returns are flat, as only traditional investors hold firms at those emission intensities. In equilibrium, green (traditional) investors only hold firms with emission intensities below (above) the threshold. This structure reflects the fact that green investors find higher-emission firms less desirable and demand higher financial returns to hold them. The return schedule, therefore, adjusts to reflect their preferences over the range of firms they fund. Importantly, this effect stems from relative valuation by green investors and holds regardless of whether their preferences involve nonpecuniary costs or benefits.

The location of the thresholds $\bar{\varepsilon}_n$ depends on the relative wealth of green and traditional investors, which is captured in our setting by the parameter κ^g . Because green investors must hold all capital below the threshold, a larger κ^g shifts the threshold up: when green investors are relatively wealthier, the segment of the schedule that is responsive to emission intensities is longer, and a smaller share of capital in the economy, particularly the dirtiest firms, earns the same flat return. A key difference from the representative agent case is that the demand for equity is no longer perfectly elastic at the returns provided by the schedule. Instead, equity demand for clean or dirty firms is shaped by the wealth of their respective investor base. In equilibrium, this demand must match the firm-side supply of equity, which itself responds endogenously to the return schedule.

Nonpecuniary costs versus benefits. Whether green investors' preferences reflect nonpecuniary costs or benefits determines the *intercept* of the return schedule. Consider first the case of nonpecuniary costs in all sectors, that is $\varepsilon_n^* = 0$ for all n . In this case, there exists an economy-wide clean rate $R_t(0)$, which is the return earned by the cleanest firms with zero emissions. All other firms must deliver higher returns to compensate green investors for the disutility of their emissions. As shown in the right panel of Figure 5, the return schedule rises linearly with emission intensity up to the threshold and remains flat beyond it. Notably, firms with emission intensities above the threshold—owned solely by traditional investors who are indifferent to emissions—also pay returns strictly above the clean rate. While these investors would be willing to hold dirty firms at lower returns, high equilibrium returns on the dirtiest firms are

necessary for markets to clear. The dirtiest firms must offer sufficiently high returns to prevent traditional investors from reallocating their capital to slightly cleaner firms held by green investors, which offer a premium over the clean rate.

Figure 5: Equilibria with nonpecuniary costs ($\gamma = \$100/\text{tn CO}_2\text{e}$)



Note: Each symbol represents the equilibrium of an economy with a different share of green investors κ^g , with γ chosen s.t. a representative green investor generates a \$100 carbon tax. Left panel shows change in emissions relative to the baseline with no green investing (black line, left axis) and the average premium over the clean rate paid by a firm that generates zero emissions (red line, right axis). Right panel shows equilibrium return schedules for each economy by plotting premia relative to the clean rate. Each schedule consists of a sloped segment starting at its symbol and a flat segment at high intensities.

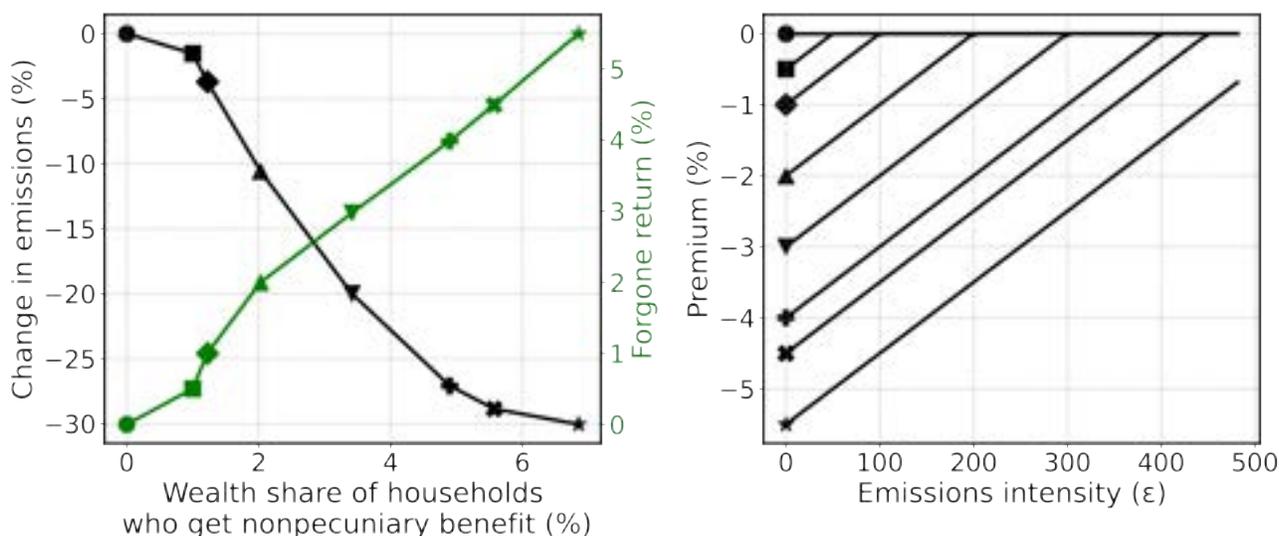
Now suppose instead that green investors derive nonpecuniary benefits $\varepsilon_n^* > 0$ from investing in a single preferred sector n . In this case, the cleanest firms in that sector earn returns that are not only lower than those of dirtier firms within the same sector, but also lower than the common return earned across all other sectors. In the preferred sector, we again have a threshold intensity beyond which the returns flatten. A new feature with sector-specific preferences is that the location of the threshold now depends not only on the wealth of green investors but also on the relative size and emissions distribution of the preferred sector compared to other sectors in the economy. In particular, if there are very few green investors relative to the supply of clean capital in the preferred sector, a jump in the return schedule may occur at the threshold emission intensity. Green investors buy the cleanest firms, with emission intensities up to the threshold, and obtain nonpecuniary benefits that reduce the financial returns. Traditional investors hold dirtier firms beyond the threshold, which offer flat returns that are discretely larger.

Our goal in this section is to clarify the potential of a small group of green investors to reduce emissions, rather than providing an exhaustive characterization. We thus focus on economies where the return schedule is continuous at the threshold: we choose wealth shares

such that green investors hold only electricity firms and generate a continuous equilibrium return schedule. The right panel of Figure 6 illustrates the class of return schedules for the electricity sector we consider. The benchmark rate for measuring premia here is the rate earned by traditional investors in equilibrium, which is also the clean rate for all sectors other than electricity.

Different schedules start at different intercepts, but all share an increasing piece with the same slope and flatten out once they reach zero. The variation in intercepts reflects the magnitude of the nonpecuniary benefits of green investors from holding the cleanest electricity firms. Compared to the patterns in the right panel of Figure 5, the key new feature is that relatively (but not perfectly) clean electricity firms pay lower returns than firms in all other sectors (clean or not). As a result, the return schedule is responsive to emission intensity even for electricity firms with relatively high emissions.

Figure 6: Equilibria with nonpecuniary benefits from electricity ($\gamma = \$100/\text{tn CO}_2\text{e}$)



Note: Each symbol represents the equilibrium of an economy with a different share of green investors κ^g , all with $\gamma = \$100/\text{tn CO}_2\text{e}$, but different ϵ_n^* for the electricity sector. Left panel shows change in emissions relative to the baseline with no green investing (black line, left axis) and difference in return on wealth between traditional and green investors (green line, right axis). Right panel shows equilibrium return schedules for the electricity sector of each economy by plotting premia relative to the return earned by traditional investors. Each schedule consists of a sloped segment starting at its symbol and a flat segment at high intensities.

Firm behavior. The firm problem in the economy with green investor preferences is the same as before. In particular, firms face return schedules that depend on their emission intensity and take these schedules into account when choosing their production plans. A technical difference is that, when the premium schedule is concave, as in Figures 5 or 6, the firm problem is no longer convex. First-order conditions are therefore no longer sufficient to characterize the solution. However, a solution still exists and is easy to characterize, since the schedule faced

by any firm depends on a finite number of segments. For any given schedule, we solve the firm problem by taking cases: we check the optimal solution within each segment and compare solutions across cases.

Another interesting consequence of a concave return schedule is that the problem of a multi-product firm is no longer equivalent to the weighted problem of its single-product subsidiaries. For example, an electricity firm that owns both a zero-emission solar installation and a gas-fired power plant may face worse terms in funding markets than two stand-alone clean and dirty subsidiaries. In other words, it can make sense for a "mixed" firm with clean and dirty operations to spin off one of the two subsidiaries. We thus work in this section with single-technology firms. In particular, all electricity firms employ only one type of fuel. We view this as a conservative assumption that makes sure that no results are driven by arbitrary restrictions on capital structure.

Equilibrium. An equilibrium consists of an allocation and prices, including a return schedule, so that all households and firms optimize and markets clear. The return schedule is determined jointly with the distribution of capital. On the household side, the exogenous distribution of member types determines how much equity is held by green versus traditional investors. Within the holdings of each type, the demand for equity by intensity is perfectly elastic, provided that the Euler equations (27) hold. On the firm side, a return schedule gives rise to a collection of production plans and hence a distribution of equity by emission intensity. In equilibrium, the distributions and returns are consistent.

5.2 Quantifying the impact of heterogeneous preferences

Our goal is to assess how small amounts of heterogeneity matter for both the return distribution and equilibrium emissions. We focus on green investors with a strong taste for clean capital. We set the preference parameter γ so that if all investors were green ($\kappa^g = 1$), their behavior would work like a \$100 carbon tax in the case of nonpecuniary costs or benefits. We study two sets of counterfactuals. The first explores the robustness of an equilibrium in which most investors incur nonpecuniary costs (all intercepts are zero, $\varepsilon_n^* = 0$). The second set of exercises assumes that there is a positive intercept in all subsectors of electricity generation. As before, we consider comparative statics of steady states.

Nonpecuniary costs. Figure 5 shows how equilibrium emissions vary with the share of green investor households, κ^g , on the horizontal axis. When all investors are green ($\kappa^g = 100\%$), we get the green representative-agent benchmark, achieving an emission reduction of nearly 50%—equivalent to a \$100 carbon tax. Moving left, the solid black line traces steady-state emission reductions for lower values of κ^g . The corresponding return schedules are marked in the right panel with the same symbol. The red symbols further show the average return premium, in basis points, measured along the right vertical axis. The main result from the figure is that **emissions are extremely sensitive to a small drop in the share of green investors below 100%**. A drop to $\kappa^g = 99\%$ already lowers the emission reduction to 28pp, close to one half of the representative agent benchmark. A five percent lower share means the emission

reduction is 10pp or one fifth of the benchmark.

This result is driven by the extreme skewness of the distributions of both emissions and emission intensities. A carbon tax achieves large emission reductions precisely because it creates strong incentives for the dirtiest firms to lower their emission intensities. Replicating returns can deliver the same effect, but only if they continue to rise steeply in the right tail of the distribution of emission intensities. The presence of even a small number of indifferent investors removes this effect. As we have seen, those investors must earn the highest equilibrium returns, and are not responsive to intensity, so the return schedule flattens out. As a result, the strongest incentives targeting the dirtiest firms disappear. Since emissions themselves are also heavily skewed, there is a large cost in terms of total emissions.¹¹

Nonpecuniary benefits from electricity. Figure 6 analyzes economies with a small share of green investors, again measured along the horizontal axis. The benchmark case is the leftmost point, where $\kappa^g = 0$, corresponding to a representative agent without a carbon tax. Unlike the previous exercise, green investors here derive nonpecuniary benefits from holding equity in electricity firms. This setup requires specifying two preference parameters: the slope γ and the intercept ε_n^* . We study a set of economies with continuous return schedules by jointly varying κ^g and ε_n^* , so that the wealth of green investors is just enough to purchase all electricity capital below the threshold. These parameter combinations imply the return schedules shown in the right panel of Figure 6. The left panel displays, as green symbols, the foregone financial return per unit of wealth for the green investors, measured in percentage points along the right vertical axis.

The main takeaway from the economies with nonpecuniary benefits is that **a small wealth share of green investors can lead to substantial emission reductions—provided their preferences are focused on a dirty sector with significant substitution potential**, such as electricity. Remarkably, with just a 2% wealth share, emissions fall by more than 10% relative to baseline. When green investors hold 6% of total wealth, emissions decline by over half of what a \$100 carbon tax would achieve. The key effect is that electricity firms have strong incentives to expand clean production. The return schedule subsidizes clean electricity generation relative to all other sectors. The green line in Figure 6 records the amount of subsidy along the right vertical axis, the cost to green investors for holding a unit of capital. Of course, since the incentives, and hence equilibrium subsidies, are now provided by a few green investors, the foregone return per unit of wealth for those investors is a lot larger than the average premium in the nonpecuniary cost case above. Still, losses remain moderate, on the order of a few percentage points.

5.3 The role of short sale constraints

Short-sale constraints are important for the results in this section. In our model, they generate return schedules with different segments, a flat segment where capital is held by traditional investors, and a sloped segment where capital is held by green investors. Only the sloped

¹¹Put differently, the return premium can be interpreted as a carbon tax rate times an emission intensity. When the premium flattens, the implied carbon tax rate effectively declines for higher emission intensities.

segment provides incentives to reduce emission intensities, so green investing is effective when many firms either lie on this segment or can adjust to reach it. We briefly discuss the robustness of our results to alternative assumptions.

Shorting fees. The subsidies required to generate the largest emission reductions in Figure 6 are large relative to typical estimates of shorting fees. Estimates of fees have increased in recent decades. The conventional wisdom at the turn of the millennium was that shorting fees are small: for example, D’Avolio (2002) shows that fees on most stocks are less than 1%, and that only about 10% of stocks are hard to borrow and involve much higher fees. More recently, Muravyev, Pearson and Pollet (2025) report an average shorting fee of roughly 2% and double-digit shorting fees in the tail of the distribution. Engelberg, Reed and Ringgenberg (2018) document substantial time variation in shorting fees and high persistence. As a consequence, there is substantial risk in shorting stocks over longer horizons. A conservative approach therefore restricts attention to return schedules with return differentials below 2 percentage points. Emission reductions are still sizeable: the schedule with the triangle symbol, for example, achieves a 10% emission reduction, the headline number quoted in the introduction.¹²

Risk aversion. In our deterministic model, a stark division of the return schedule into segments arises because assets are perfect substitutes, so any difference in tastes for green assets implies that only one investor type holds the asset whereas the other hits the short-sale constraint. An extension of the setup to risk-averse agents can give rise to equilibria where short-sale constraints do not bind, or bind for fewer assets. For example, suppose there are no short-sale constraints, there are shocks to firms’ TFP, and investors are equally risk averse, but green investors incur nonpecuniary costs from dirty portfolios, a case often considered in the literature. In equilibrium, green (traditional) investors still hold more clean (dirty) firms, as in our model. The difference is that they also short dirty (clean) firms.

In a setup with risk aversion and no shorting, equilibrium prices typically reflect wealth-weighted averages of individual investors’ valuations. Returns would thus in particular reflect a wealth-weighted average of different investors’ nonpecuniary costs. One prominent force of this section, that investors with a small wealth share matter disproportionately, is then muted. A large share of green (traditional) investors would tend to push returns close to the case of a representative green (traditional) investor. In addition, returns could reflect compensation for risk—details here would depend on the correlation of returns across sectors and firms, and on how risk is shared across investor types.

Although these caveats might affect the quantitative results, the main takeaways from this section are likely to carry over to a setting with risk. Consider our first result that when preferences reflect only nonpecuniary costs, green investing does not effectively reduce emissions. As we have seen, preferences that replicate a \$100 carbon tax require the willingness to sacrifice

¹²Alternative parametrizations that also keep return differentials below 2% weaken green investors’ intensity γ . The return schedule is then flatter for electricity firms, but premia may be below zero for a larger range of intensities than under the triangle schedule in Figure 6. We have found that as long as return differentials respect the same bound, the overall emission reduction is very similar to the triangle schedule.

double-digit annual returns on the dirtiest firms. A large enough coalition of such investors may not assemble even if short selling is free. Moreover, when return differentials become large, short-sale constraints also bind in models with risk, so the sorting mechanism reemerges. For example, Aron-Dine, Beutel, Piazzesi and Schneider (2024) study a quantitative model with heterogeneous risk-averse agents and nonpecuniary costs where sorting occurs because investors disagree about the expected returns of green assets.

Our second result is that a small coalition of green investors can have impact if they favor clean electricity firms. In a model with risk and short-sale constraints, the same mechanism is relevant: if green investors bid up prices of electricity firms so that average returns are lower, traditional investors will find those firms less attractive to hold. Since traditional investors can be well-diversified even if they do own electricity firms, a tiny share of capital, their incentive to hold such firms is almost as small as in our model. A difference is that the cost to green investors of focusing on electricity may be larger if investors also forego diversification benefits. While this requires recomputing what willingness to pay is required, it leaves in place the principle that a small share of green investors can matter.

6 Conclusion

To conclude, we summarize the main takeaways for the impact of climate finance and point to some areas for future research.

Our replication results show that, in principle, carbon taxation through asset returns can work, that is, asset markets can provide optimal incentives to address a climate externality. In particular, replication is consistent with exactly the type of nonconsequentialist preferences that are widely studied in the literature and supported by experimental evidence—nonpecuniary costs from dirty asset holdings—provided that a representative agent exhibits such preferences. What is important, moreover, is that investors’ concern for emissions reflects a score that counts all emissions once, such as scope 1 emissions.

Our result provides a recipe for computing replicating returns for any given carbon tax rate that can shed light on the role of market incentives in other contexts. In particular, our formula allows two types of exercises. On the one hand, to assess the current incentives provided by financial markets in a particular country or sector, we can evaluate the replicating return schedule at current emissions intensities and compare the distribution to (appropriately risk-adjusted) returns in the data.

On the other hand, we can evaluate a replicating return schedule at model-implied emission intensities from any integrated assessment model that makes predictions about emissions and firm value, and under any scenario for carbon taxes and the trajectory of technical progress. This type of exercise restates a path of carbon taxes as a path of return distributions, and hence asset price paths that show what it takes for climate finance to provide optimal incentives. We have emphasized that since intensities are endogenous, the answer generally varies with features of the environment.

Our quantitative work identifies two features of modern economies that we expect to be widely relevant for any such exercises, as well as more generally for thinking about the potential impact of climate finance solutions. First, scope 1 emission intensities relative to enterprise value are extremely right-skewed, with the highest intensities in sectors that produce the most emissions. This means that the success of climate finance turns on whether it can provide strong incentives in the right tail of the intensity distribution.

Second, some sectors, such as electricity, produce a lot of emissions at high intensities, but are also relatively responsive to a carbon tax, due to high elasticities of substitution in technology. The presence of any such sectors means that a lot of emission reduction is lost if incentives cannot be provided to all firms, but only to all but the dirtiest firms, for example, because there is a small share of neutral investors. This leads to our conclusion that carbon taxation through asset returns is ultimately not effective.

These two features also lead to our final takeaway: climate finance can be effective even with a few green investors if funds are marketed to incentivize dirty firms that can substitute. This alternative approach moves away from the replication of carbon taxation, which simply penalizes dirty firms. We have formally studied preferences that reflect nonpecuniary benefits from investing in electricity. We view this approach as describing an alternative trading strategy that delivers results for investors willing to sacrifice some returns to advance the green transition.

Importantly, targeting sectors that can substitute still requires a score that investors respond to, as do replicating returns. Our recommendation is that scoring should focus exclusively on the level of scope 1 emissions. In addition to avoiding double-counting, scope 1 emissions are relatively easy to measure. Scope 2 emissions, in contrast, require attributing electricity to how dirty the fuel is, while measuring scope 3 emissions involves complex calculations of the upstream and downstream emissions associated with producing a good. Moreover, any blending of multiple scopes, such as averaging, should be avoided to prevent double taxation.

We emphasize two more properties of our proposed mechanism. The first is that it creates incentives through a cost-of-capital effect, just like replicating returns do. It is, therefore, conceptually different from activist investing, where a few investors effectively take an industry private and change its technological profile. The advantage of a cost-of-capital mechanism is that it does not have to work through equity markets, unlike activist investing, which requires control rights. Instead, return differentials that incentivize substitution can equally work in fixed-income markets. This broadens the investor base from which green investors can be drawn.

A second key property is that, in principle, green investors can come from other countries. If a dirty sector that can substitute relatively easily can fund itself in international capital markets, then the taste for green investing can thus be "imported" in the form of lower costs of capital. This observation helps address a common criticism of climate finance: green investing is currently growing, especially in areas such as Europe, where other measures to combat climate change are also being implemented. Through international financial markets, it is possible to direct the willingness to pay for climate to other countries where the local population has relatively more

neutral preferences, thus creating gains from trade.

References

- Aghion, Philippe, Timo Boppart, Michael Peters, Matthew Schwartzman, and Fabrizio Zilibotti (2025) "A Theory of Endogenous Degrowth and Environmental Sustainability," Working Paper w33634, National Bureau of Economic Research.
- Aguiar, Angel, Maksym Chepeliev, Erwin Corong, and Dominique van der Mensbrugge (2022) "The global trade analysis project (GTAP) data base: Version 11," *Journal of Global Economic Analysis*, 7 (2).
- Aron-Dine, Shifrah, Johannes Beutel, Monika Piazzesi, and Martin Schneider (2024) "Household Climate Finance: Theory and Survey Data on Safe and Risky Green Assets," Working paper, Stanford.
- Barrage, Lint and William D. Nordhaus (2024) "Policies, projections, and the social cost of carbon: Results from the DICE-2023 model," *Proceedings of the National Academy of Sciences*, 121 (13), 2655–2689.
- Berk, Jonathan and Jules H van Binsbergen (2021) "The Impact of Impact Investing," Working Paper, Stanford GSB.
- Bonnefon, Jean-François, Augustin Landier, Parinitha Sastry, and David Thesmar (2025) "The moral preferences of investors: Experimental evidence," *Journal of Financial Economics*, 163.
- D'Avolio, Gene (2002) "The market for borrowing stock," *Journal of Financial Economics*, 66, 271–306.
- Engelberg, Joseph E., Adam V. Reed, and Matthew C. Ringgenberg (2018) "Short-Selling Risk," *Journal of Finance*, 73, 755–786.
- Fostel, Ana and John Geanakoplos (2015) "Leverage and Default in Binomial Economies: A Complete Characterization," *Econometrica*, 83, 2191–2229.
- Geczy, Christopher, Robert F Stambaugh, and David Levin (2021) "Investing in Socially Responsible Mutual Funds," *Review of Asset Pricing Studies*, 11 (2), 309–351.
- Giglio, Stefano, Bryan T. Kelly, and Johannes Stroebel (2021) "Climate Finance," *Annual Review of Financial Economics*, 13 (1), 15–36.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, Zhenhao Tan, Stephen Utkus, and Xiao Xu (2025) "Four facts about ESG beliefs and investor portfolios," *Journal of Financial Economics*, 164.
- Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2014) "Optimal Taxes on Fossil Fuel in General Equilibrium," *Econometrica*, 82 (1), 41–88.

- Gormsen, Niels Joachim, Kilian Huber, and Sangmin Oh (2024) "Climate Capitalists," Working Paper, Chicago Booth.
- Goulder, Lawrence H. and Marc A. C. Hafstead (2017) *Confronting the Climate Challenge: U.S. Policy Options*, New York: Columbia University Press.
- Hassler, John and Per Krusell (2012) "Economics and Climate Change: Integrated Assessment in a Multi-Region World," *Journal of the European Economic Association*, 10 (5), 974–1000.
- Heeb, F., J. F. Kölbel, F. Paetzold, and S. Zeisberger (2023) "Do investors care about impact?" *Review of Financial Studies*, 36, 1737–87.
- Heinkel, Robert, Alan Kraus, and Josef Zechner (2001) "The Effect of Green Investment on Corporate Behavior," *Journal of Financial and Quantitative Analysis*, 36 (4), 431–449.
- Hong, H. and E. Shore (2023) "Corporate social responsibility," *Annu. Rev. Financ. Econ*, 15, 327–50.
- Iovino, Luigi, Thorsten Martin, and Julien Sauvagnat (2024) "The Environmental Bias of Corporate Income Taxation," Working paper, Bocconi.
- Kortum, Samuel and David Weisbach (2024) "Optimal Unilateral Carbon Policy," Working Paper, Yale.
- Kortum, Samuel, David Weisbach, Michael Wang, and Yujia Yao (2023) "Trade, Leakage, and the Design of a Carbon Tax," *Environmental and Energy Policy and the Economy*, 04, 43–90.
- Macaluso, Nick, Sugandha Tuladhar, Jared Woollacott, James R. McFarland, Jared Creason, and Jefferson Cole (2018) "The impact of carbon taxation and revenue recycling on U.S. industries," *Climate Change Economics*, 9 (1).
- McFarland, James R., Allen A. Fawcett, Adele C. Morris, John M. Reilly, and Peter J. Wilcoxon (2018) "Overview of the EMF 32 study on U.S. carbon tax scenarios," *Climate Change Economics*, 9 (1).
- Muravyev, Dmitriy, Neil D. Pearson, and Joshua M. Pollet (2025) "Anomalies and their short-sale costs," *Forthcoming Journal of Finance*.
- Nordhaus, William (2013) "Integrated Economic and Climate Modeling," *Handbook of Computable General Equilibrium Modeling*, Elsevier, 1, 1069–1131.
- Papoutsis, Melina, Monika Piazzesi, and Martin Schneider (2024) "How unconventional is green monetary policy?", Working paper, Stanford.
- Pástor, Luboš, Robert F Stambaugh, and Lucian A Taylor (2021) "Sustainable investing in equilibrium," *Journal of Financial Economics*, 142 (2).

- Pástor, Luboš, Robert F. Stambaugh, and Lucian A. Taylor (2024) “Sustainable investing,” Working paper, Chicago Booth.
- Pedersen, Lasse (2025) “Carbon Pricing versus Green Finance,” Working Paper, Copenhagen Business School.
- Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski (2021) “Responsible investing: The ESG-efficient frontier,” *Journal of Financial Economics*, 142, 572–97.
- Ramondo, Natalia, Garcia-Lembergman, Andrés Rodríguez-Clare, and Joseph Shapiro (2025) “Carbon Emissions in the Global Economy,” Working paper, Boston University.
- Riedl, Arno and Paul Smeets (2017) “Why Do Investors Hold Socially Responsible Mutual Funds?” *The Journal of Finance*, 72 (6), 2505–2550.
- Rogerson, Richard, Robert Shimer, and Randy Wright (2005) “Search-theoretic models of the labor market: A survey,” *Journal of economic literature*, 43, 959–988.
- Yuan, Mei, Sebastian Rausch, Justin Caron, Sergey Paltsev, and John Reilly (2019) “The MIT US Regional Energy Policy (USREP) model: the base model and revisions,” *Joint Program Technical Note TN*, 18.

APPENDIX

A Derivations and proofs

In this section we collect formal arguments not contained in detail in the text.

A.1 Social planner problem

The social planner maximizes utility (1) subject to technology (2), the evolution of the state of the environment (3), the trade balance (4), and the resource constraint for labor (5). Let λ_t be the Lagrange multiplier on the first equation in (2), while $\lambda_{t,n}$ is the multiplier on the resource constraint for good n , the last equation in (2), and μ_t is the multiplier on the evolution of the state of the environment (2). The first-order conditions for C_t , $y_{t,n}$ and $x_{t,m}^n$ are

$$\begin{aligned} u'(C_t) &= \lambda_t, \\ \lambda_t \frac{dg}{dy_{t,n}}(y_t) &= \lambda_{t,n}, \\ \lambda_{t,m} &= \lambda_{t,n} \frac{df^n}{dx_m}(l_t^n, k_t^n, x_t^n, \tilde{x}_t^n) + \mu_t \frac{de^n}{dx_m}(l_t^n, k_t^n, x_t^n, \tilde{x}_t^n). \end{aligned}$$

Defining the social cost of carbon as $SSC_t = \mu_t / u'(C_t)$, we combine the three equations to obtain equation (6).

A.2 Abandoned technologies and stranded assets

In this appendix we characterize an extended version of the example in Section 2.3 that shows how technologies can be abandoned and assets can be stranded. We consider an economy with three sectors: a clean good 1 made from labor and a dirty good 2 made from labor, capital, and a fuel good 3. Capital is in fixed supply and does not depreciate. The fuel good 3 is extracted from a stock of reserves. The example thus also illustrates how our framework deals with fixed factors and exhaustible resources. Technology and emissions are described by

$$\begin{aligned} C_t &= y_t^1 + y_t^2 \\ y_t^1 &= f(l_t) \\ y_t^2 &= \min \left\{ l_t^2, a(k_t^2, y_t^3) \right\} \\ k_t^2 &= \bar{k}^2 \\ k_{t+1}^3 &= k_t^3 + z_t; \quad z_t \leq 0 \\ z_t + y_t^3 &= 0 \\ e^2(k_t^2, l_t^2, y_t^3) &= \bar{\varepsilon} \min \left\{ l_t^2, a(k_t^2, y_t^3) \right\} \end{aligned}$$

Here the first three equations are production functions for the consumption good and goods 1 and 2, respectively. The fourth equation fixes the quantity of capital for technology 2. The next two equations describe the depletion of reserves. The quantity k_t^3 of capital good 3 represents fuel reserves; here the gross investment function h is chosen so reserves can only be reduced. The next to last equation describes a dummy good with production function $f^n(X) = 0$ that serves to translate the resource extraction $-z_t$, modeled as disinvestment of capital good 3, into intermediate output y_t^3 . Finally, emissions from technology 2 are proportional to output—the emission intensity relative to gross output \bar{e} is independent of the scale of production.

The example reduces to the one discussed in Section 2.3 when the fuel good 3 is omitted. We again make assumptions so both goods are produced and capital earns rents: suppose f is strictly increasing and strictly concave with $f'(0) = \infty$ and $f'(1) < 1$ and that a is homogeneous of degree 1, strictly increasing, quasi-concave and smooth with $a(k, 0) = 0$, $\lim_{y^3 \rightarrow 0} \frac{da}{dy^3}(k, y^3) \rightarrow \infty$ for all $k > 0$ and $\lim_{k \rightarrow 0} \frac{da}{dk}(k, y^3) \rightarrow \infty$ for all $y^3 > 0$.

Consider an equilibrium without a carbon tax. Given our assumptions on production functions, both goods 1 and 2 are produced in equilibrium. Indeed, the marginal product of labor for technology 1 becomes very high when labor is small, and falls below one if fuel extraction is very small.

Since goods 1 and 2 are perfect substitutes in making the numeraire consumption good, their equilibrium prices are equal to one. Firms that operate technology 1 choose labor l to maximize

$$f(l) - wl,$$

and firms that operate technology 2 choose capital and fuel to maximize

$$-p_{t-1,2}^k k^2 + R_t^{-1} \left(a(k^2, y^3) (1 - w) + p_{t,2}^k k^2 - p_{t,3} y^3 \right).$$

Firms that operate technology 3 choose $k^3 \geq 0$ and $z \leq 0$ to maximize

$$-p_{t,3} z^3 - p_{t,3}^k k^3 + R_{t+1}^{-1} p_{t+1,3}^k k_{t+1}^3 = -p_{t,3} z^3 - p_{t,3}^k k^3 + R_{t+1}^{-1} p_{t+1,3}^k (k^3 + z).$$

When the firm extracts reserves, the first-order conditions are $p_{t+1,3}^k = R_{t+1} p_{t,3}^k$ and $p_{t,3} = R_{t+1}^{-1} p_{t+1,3}^k = p_{3,t}^k$. In particular, we recover the familiar Hotelling rule $p_{t+1,3} = R_{t+1} p_{t,3}$.

Putting together first-order conditions of firms and households, an equilibrium is a solution

$(w_t, p_{t,2}^k, p_{t,3}, y_t^3, C_t, R_t)$ to the difference equation

$$\begin{aligned}
w_t &= f' \left(1 - a \left(\bar{k}^2, y_t^3 \right) \right) \\
\frac{da}{dk^2} \left(\bar{k}^2, y_t^3 \right) &= R_t p_{t-1,2}^k - p_{t,2}^k \\
\frac{da}{dy^3} \left(\bar{k}^2, y_t^3 \right) &= p_t^3 \\
p_{t,3} &= R_t p_{t-1,3} \\
u' (C_t) &= \beta R_{t+1} u' (C_{t+1}) \\
C_t &= f \left(1 - a \left(\bar{k}^2, y_t^3 \right) \right) + a \left(\bar{k}^2, y_t^3 \right)
\end{aligned} \tag{30}$$

with initial condition k_0^3 given and a second boundary condition from the household's lifetime budget constraint.

We show that there is a unique equilibrium. For any initial fuel price $p_{0,3}$, the fourth equation delivers a path of resource prices. The third equation then delivers a path of extraction, and the first a path of wages. Consumption is also determined once extraction is known, and we can derive a sequence of interest rates. The price of capital is found by solving the second equation forward to obtain the present value of all marginal products of capital $\frac{da}{dk^2} (\bar{k}^2, y_t^3)$.

Finally, the household's lifetime budget constraint must hold, which pins down the initial price of the natural resource:

$$\sum_{t=0}^{\infty} \prod_{\tau=0}^{t-1} C_t = p_{0,2}^k k_0^2 + p_{0,3} k_0^3 + \sum_{t=0}^{\infty} \prod_{\tau=0}^{t-1} w_t.$$

Now consider an equilibrium with a large carbon tax, defined as $\tau > \frac{1-f'(1)}{\bar{\epsilon}} > 0$. We show that there is an equilibrium where production of the dirty sector 2 shuts down and both capital and the resource become worthless. In any such equilibrium, technology 1 is used so its first-order condition for labor holds. Equilibrium is characterized by

$$\begin{aligned}
w_t &= f' (1) \\
y_t^2 &= y_t^3 = 0 \\
p_t^3 &= p_{t,2}^k = 0 \\
C_t &= f (1)
\end{aligned}$$

and interest rates are constant at $R_t = \beta^{-1}$.

We need to verify that no firm wants to enter and make good 2 at these prices. A production

plan that employs technology 2 yields after-tax value

$$R_t^{-1} a(k_t^2, y_t^3) (1 - w_t - \tau \bar{\varepsilon}) = R_t^{-1} a(k_t^2, y_t^3) (1 - f'(1) - \tau \bar{\varepsilon}) < 0$$

and is therefore not worth undertaking, even if capital can be purchased at zero price. We note that the wage is lower in the shutdown equilibrium: since taxation penalizes the use of fuel and capital good 2 that are complementary to labor, the demand for labor declines and its price falls.

Suppose the carbon tax is collected from investors. A firm that enters at date $t - 1$ and employs technology 2 expects a pretax value at date t of

$$v_{t,2} = a(k_t^2, y_t^3) (1 - w_t) > 0$$

The firm's emission intensity relative to value is therefore

$$\varepsilon_t = \frac{\bar{\varepsilon}}{1 - w_t}$$

which is independent of scale but does depend on the wage. Owners of the firm therefore expect an after-tax payoff

$$v_{t,2} (1 - \tau \varepsilon_t) = v_{t,2} \left(1 - \frac{\tau \bar{\varepsilon}}{1 - w_t}\right) = v_{t,2} \left(1 - \frac{\tau \bar{\varepsilon}}{1 - f'(1)}\right) < 0$$

Running the firm is not profitable under any discount rate. It can also be derived as the pretax value discounted at the rate $R_t (1 - \tau \varepsilon_t)^{-1}$.

B Input-Output structure

Here we detail the input-output structure of the quantitative model used in Section 30. Much of the production structure follows the multisector model developed by Yuan, Rausch, Caron, Paltsev and Reilly (2019). We depart from their setup to accommodate firm-level heterogeneity in the electricity sector and to match the slightly different breakdown of our sectoral data. Table 6 lists all sectors in our model.

Figure 7 describes a generic non-electricity sector. Each node represents an intermediate input and each arrow an input-output linkage of a sector-specific production function. Each production stage (illustrated by a set of arrows that point to a single node) is described by a CES production function. For easier reference, we include the label of the CES elasticity of substitution σ at each stage of production. For example, in the lower left corner of the diagram, dirty fuels are combined into a dirty fuel bundle in a CES nest with substitution elasticity σ_{Dirty} , which in turn is combined with electricity into an energy bundle in a CES nest with elasticity σ_{Energy} . Elasticity numbers are collected in Table 7 below.

Beyond energy, the nested CES structure for a generic sector has several additional layers.

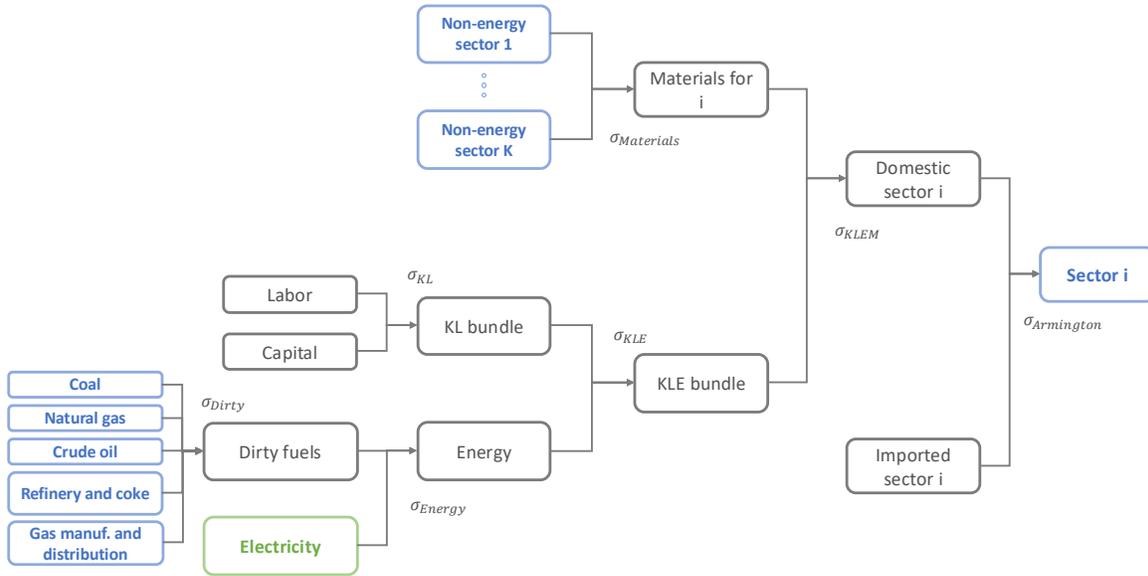


Figure 7: Input-output structure of a generic sector

Labor and capital are combined to make a “KL” bundle, which is then combined with the energy bundle to produce “KLE”. All non-energy related goods are combined into a materials bundle, which is combined with the KLE bundle to produce the final output of the domestic sector. Finally, domestic output is combined with imported goods to produce the total supply of the good made by the sector. This generic structure applies to the 60 non-energy sectors in the model. It also applies to the electricity transmission and distribution (ETD) sector, 4 non-carbon electricity sectors, and all 5 dirty fuel sectors. Nodes with a blue border highlight which products have a generic production structure.

Figure 8 describes the input-output structure of the electricity sector. Electricity production starts from 7 electricity varieties: coal-based, gas-based, hydro, nuclear, solar, wind, and other¹³. Each electricity type is aggregated in two groups based on technological constraints on storage and readiness-to-dispatch upon changes in demand. We bundle solar and wind as intermittent sources, and the rest as baseline sources. Intermittent and baseline sources are combined into electricity generation, which in turn is combined with the output of the ETD sector to produce the total supply of electricity in the economy. Green-bordered nodes highlight the final electricity sector.

Carbon-based electricity varieties have a production function slightly different from the generic sector structure. Leveraging firm- and division-level micro-data, as described in D below, we include an extra production layer in each electric variety to account for firm level

¹³Other is a residual that will absorb electricity production based on oil, biomass, geothermal and any other source not considered in our model.

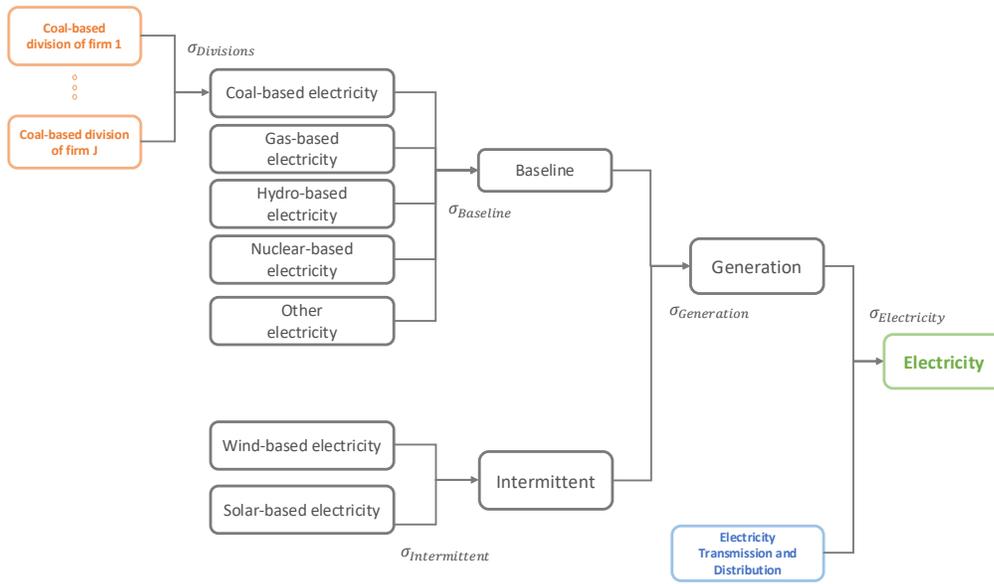


Figure 8: Input-output structure of electricity

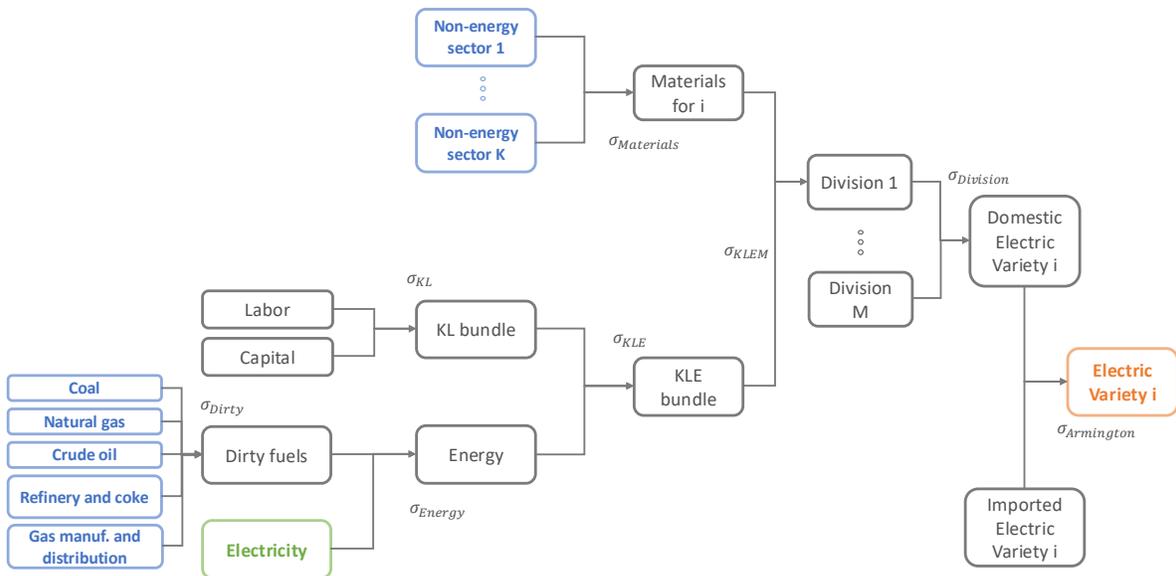


Figure 9: Input-output structure of a generic electric variety

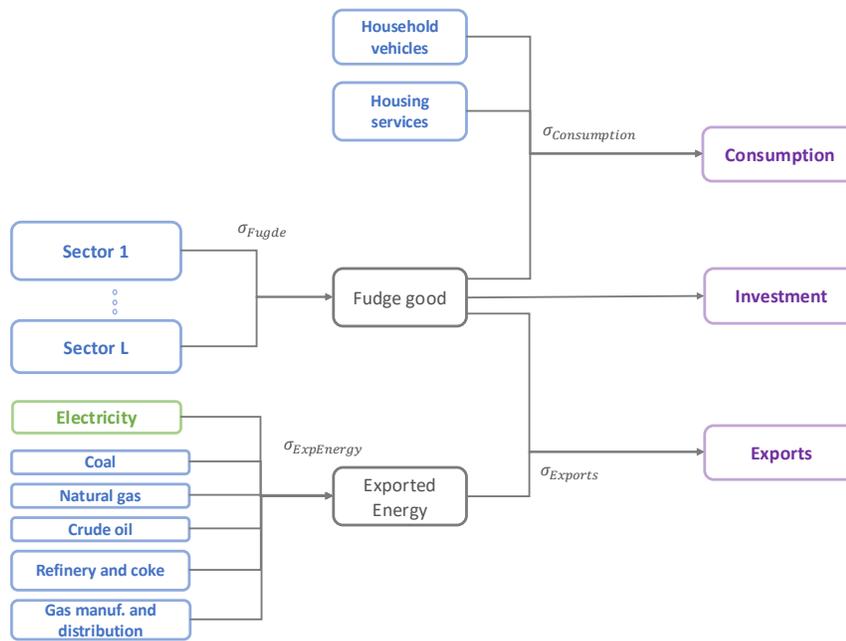


Figure 10: Input-output structure of final demand

heterogeneity. Figure 9 describes the structure of a single electric variety (for example, coal-based electricity). Each ‘division’ represents a firm-variety pair. The output of each division is bundled into a domestic electric variety. The rest of the input-output structure replicates the generic sector. Orange-bordered nodes denote an electric variety.

Finally, Figure 10 illustrates the structure of the final demand. A “fudge” good bundles output of all non-energy sectors, except for housing services and household vehicles. This good is used in the production of the final consumption good, and it is directly used as capital and exported. In the case of final consumption, the “fudge” good is combined with two additional goods: housing services and household vehicles. In the case of exports, “fudge” is combined with exported energy goods.

Table 6: Sectors and Aggregators

Primary	Sectors				Aggregators
	Carbon Fuels	Manufacturing	Services	Electric Varieties	
Paddy rice	Coal	Bovine meat products	Water	Coal-BL	Electricity
Wheat	Oil	Other Meat products	Construction	Gas-BL	Final Demand
Other cereal grains	Gas	Vegetable oils and fats	Trade	Other-BL	
Vegetables, fruit, nuts	Gas manufacturer	Dairy products	Accommodation	Hydro-BL	
Oil seeds	Refinery	Processed rice	Commercial land transport	Wind-BL	
Sugar cane, sugar beet		Sugar	Water transport	Solar-BL	
Plant-based fibers		Other Food products	Air transport	Nuclear-BL	
Other Crops		Beverages and tobacco products	Warehousing		
Cattle		Textiles	Communication		
Other Animal products		Wearing apparel	Financial services nec		
Raw milk		Leather products	Insurance		
Wool, silk-worm cocoons		Wood products	Real estate activities		
Forestry		Paper products, publishing	Other Business services		
Fishing		Chemical products	Recreational		
Other Mining		Basic pharmaceutical products	Public Administration		
		Rubber and plastic products	Education		
		Other Mineral products	Human health and social work		
		Ferrous metals	Housing services		
		Other Metals	Household vehicles		
		Metal products	Electricity Transmission and Distribution		
		Computer, electronic and optics			
		Electrical equipment			
		Other Machinery and equipment			
		Motor vehicles and parts for firms			
		Other Transport equipment			
		Other Manufactures			

C Data

In this section, we describe our data sources, and how we clean the data and map them to the model. Our main sources of data are GTAP 11 database from the *Global Trade Analysis Project*, and Trucost and Compustat databases from *S&P Global*. We also describe some complementary data sources, used to pin down specific parameters of the model.

GTAP 11 database GTAP 11 is a dataset developed by the *Global Trade Analysis Project*. GTAP is a global network of researchers conducting quantitative analysis of international trade. It is coordinated by the Center for Global Trade Analysis in Purdue University's Department of Agricultural Economics. For further reference, see Aguiar, Chepeliev, Corong and van der Mensbrugge 2022. It collects in an unified and consistent framework annual transactions across countries and economic sectors, based on official national accounting and trade statistics. It also includes physical transactions of energy-related goods, and their related CO₂ emissions, based on official trade balances and consistent with 2006 IPCC guidelines Tier 1. Energy and emission related information are included in two satellite datasets called GTAP-Power and GTAP-E, respectively.

The most disaggregated version of GTAP 11 reports transactions of 141 countries and 76 sectors. There are 17 energy-based goods. The 12 electricity-producing energy sectors comprise 7 *base load* (BL) technologies (NuclearBL, CoalBL, GasBL, HydroBL, OilBL, WindBL, and OtherBL), 4 *peak load* (P) technologies (GasP, OilP, HydroP, and SolarP) as well as a transmission and distribution sector (TnD). The 5 carbon-based energy sectors are Coal, Gas, Crude Oil, Petroleum and Coal Products, as well as Gas Manufacture and Distribution. Physical quantities are reported for energy-related goods, except for transport and distribution of electricity. Emissions are reported for carbon-based energy goods only. The year of reference is 2017. All monetary values are reported in million 2017 US dollars, energy physical units are in thousands of tons of oil equivalent, and emissions are in million of tons of CO₂.

Data extraction. We extract the following data for the United States from the GTAP 11 database. We obtain the monetary value of (i) intra-industry and final purchases (consumption, investment, government and exports) of domestic and foreign goods, (ii) factor purchases (skilled labor, unskilled labor, capital, land and natural resources), as well as taxes and subsidies of domestic sectors and finally (iii) residuals at sector level that accounts for statistical differences between total sales and costs. By construction, total sales and costs of each domestic sector are equal, up to the residual term. We further extract physical quantities of domestically produced and imported energy-related goods, purchased by domestic sectors and the final demand as well as emissions of CO₂ related to the use of energy-related goods.

Mapping data to model: domestic intermediates. In order to calibrate the model, we need observable counterparts for factor shares and input shares of all goods. We measure the value added of a domestic sector as the sum of all factor costs originally reported (skilled labor, unskilled labor, capital, land, and natural resources), plus taxes and subsidies paid by the

sector, and statistical differences between total sales and costs. We then split value added into labor and capital, in proportion to the sum of skilled and unskilled labor costs relative to the sum of capital, land and natural resources.

International trade. As part of the input-output matrix, GTAP records the value of imports of every good by every domestic sector. We aggregate across buyer sectors to obtain the total domestic sales of the imported good. On the demand side, we aggregate purchases across origins, so all buyers – whether households or firms of different sectors – buy a single Armington bundle of any given good, instead of choosing buyer-specific combinations of local and imported goods. In addition, our model assumes a single price vector for energy-related goods (dirty fuels and electric varieties) for which we also observe quantities. We compute prices by dividing the value of total sales of a good by the total quantity good sold. Based on this normalized price, we then redefine quantities at the buyer-seller level to match values. This step removes small discrepancies in reporting of international trade.

Electricity. We simplify the electricity sector by reducing the originally reported 12 electricity-producing energy sectors to 7 electricity varieties. In particular, we merge GasP and GasBL into Gas, OilP, OilBL and OtherBL into Other and HydroP and HydroBL into Hydro. We further add an aggregator (CES nest) called “electricity” that collects the physical output of the 7 electric varieties and combines revenues of these varieties plus the transmission and distribution sector. We replace each purchase of an electric variety in the original dataset with a purchase of this newly created aggregate electricity sector. We normalize the price paid across sectors to match the ratio of total sales to total quantity of electricity sold.

Vehicles and housing. We define two automobile sectors: one that produces vehicle services consumed by households, and another that produces cars as an intermediate input for firms. The latter sector is derived by first computing the share of automobiles purchases as an input by domestic sectors other than households, and then multiplying GTAP automobile sector input and factor costs by that share. Output of the vehicle services sector consists of the remaining share of auto sector output, plus GTAP final consumption of petroleum and coke. ELEC We further create a housing services sector by adding to GTAPS’s dwelling sector all consumption of energy by the final demand, except for petroleum & coke, which is assumed to power cars.

Consumption, exports and capital. As displayed in Figure 10, we bundle all non-energy, non-car, non-housing goods purchased by GTAP’s final demand sector into a single *fudge* good. The fudge good is the numeraire for the quantitative exercises and serves three purposes. First, it is combined with housing services and vehicles services to produce final consumption. In addition, it is directly installed as capital and exported. The quantitative exercises thus abstract from difference in goods baskets used to make sector-specific capital goods or exports. Robustness analysis suggests that such differences do not have an important effect on the results.

BEA capital stock and depreciation. To complement GTAP value and input numbers, We extract depreciation and capital stock information from *Fixed Asset Tables* from *Bureau of Economic*

Analysis (BEA) to construct sector-specific depreciation rates. For consistency with the rest of the data sources, we use current-cost measures of 2017. We extract capital stocks from *Table 3.1ESI* at NAICS level. We add government capital stocks from *Table 1.1*. Annual depreciation is from *Table 4.4ESI* at NAICS level and depreciation of government assets is from *Table 1.3*. We match BEA sectors to our model sectors and compute the depreciation rate of each sector as the ratio of depreciation to the capital stock.

D Firm- and division-level data for the electricity sector

Firms and divisions. Trucost is a subsidiary of *S&P Global* that provides data on environmental impact of companies. Trucost breaks down electricity generation of every firm by the type of fuel used, in particular, coal, gas, solar, geothermal, hydro, oil, wind, nuclear, landfill, biomass, LNG, tidal, and other. We also observe emissions per type of fossil fuel, that is, coal, gas, oil, LNG, and LPG. We can therefore divide up firms into *divisions* that aggregate all generation of the same firm that use the same fuel. To map fuel types to our model, we aggregate Trucost’s geothermal, landfill, biomass, tidal, LNG and LPG and other energy types into the *other* category of the model.

We merge TruCost data to Compustat using the internal *S&P Global* firm identifier *gvkey*. We focus on US-based active companies with positive electricity generation in 2017. We measure enterprise value as the sum of market capitalization (“mkval”), long-term debt (“dltt”), short-term debt (“dlc”) and preferred stock (“pstkrv”) less cash (“che”). We end up with a dataset of 45 energy-producing firms and 203 divisions. For each division, we observe enterprise value at the firm level, as well as generation by type of fuel and emissions by type of fuel at the division level. We impute division-level enterprise value in proportion to the division-level generation of each firm.

Table 7: Elasticities of substitution

Generic sector (Figure 7)		Final demand (Figure 10)	
Parameter	Value	Parameter	Value
$\sigma_{\text{Armington}}$	2.5	$\sigma_{\text{Consumption}}$	1
σ_{KLEM}	0	σ_{Fudge}	1
σ_{KLE}	0.4	$\sigma_{\text{ExpEnergy}}$	0
$\sigma_{\text{Materials}}$	0	σ_{Exports}	0
σ_{KL}	1		
σ_{Energy}	0.5		
σ_{Dirty}	1		

E Model quantification

Our quantification of the model follows standard practice in the literature. We assume CES production functions and calibrate substitution elasticities based on values in existing models. The weights in the production functions are determined from a "social accounting matrix" that captures the flow of inputs and outputs across sectors and firms. For non-energy goods, we normalize units by setting their prices equal to one in an initial equilibrium. This choice is without loss of generality as long as counterfactuals focus on percentage changes in quantities and prices, rather than physical units, as we do for non-energy goods in our model.

For energy goods, where emissions depend on physical units, we rely on data for both prices and quantities. The GTAP database reports quantities of emission-generating energy goods in units of *coal equivalents*, allowing us to compare energy content across different fuels. This standardization underlies the construction of Table 1. For nuclear, solar, and wind electricity, we only observe generation data, not the energy content of their inputs. To make the table, we measure energy content using the *substitution method*, that is, we divide generation by a factor of approximately .4—treating these technologies as if they were as inefficient as fossil fuels. This convention is not relevant for quantifying production functions, since those include only inputs with positive prices.

We use observable prices for energy goods, labor, and capital. For capital, we measure user costs as the sum of the interest rate and the depreciation rate, as described above. Given the CES nesting structure in production, we can back out nest weights and TFP for each sector using sectoral data on gross output, material inputs, and labor compensation by sector. In particular, expenditure shares identify the nest weights, while TFP is pinned down by equating the unit cost of a sector to its price.

Although our model assumes balanced trade, U.S. data in 2017 show a trade deficit of about 2.8% of GDP. To reconcile this with aggregates, we scale down the final good input to consumption and increase exports. This adjustment leads to a slight understatement of consumption in the model. However, because our analysis focuses on comparative steady states, this convention is not essential for our results.

Technology parameters. Key elasticity values are listed in Table 7. For a typical non-electricity sector, the location of each elasticity in the nesting structure is displayed in Figure 7. Domestic and imported varieties are substitutes, with an Armington elasticity of 2.5. The domestic variety is produced using materials and a *KLE bundle* that combines capital (K), labor (L), and energy (E). Materials and the KLE bundle must be combined in fixed proportions. Within the KLE bundle, substitution elasticities are relatively low between the capital-labor composite and energy, as well as between electricity and a bundle of dirty fuels. We assume an elasticity of one between capital and labor, as well as among the different dirty fuels.

The production technology for the electricity sector differs from that of other sectors, as it is designed to accommodate detailed division-level data. For each fuel type, we specify a nest over multiple varieties, where each variety is produced using materials and a KLE bundle. The

number of varieties per fuel type equals the number of divisions observed in the TruCost dataset, plus one residual variety used to reconcile the division-level data with sectoral aggregates from GTAP. We further introduce an aggregate electricity sector that combines electricity generated from different fuels. This aggregation accounts for the imperfect substitutability across fuel types, and incorporates the additional resources required to operate and maintain the electricity grid.

At the division level, we observe electricity output and fuel inputs, but not other inputs such as materials and labor. We describe each division's production technology using nested CES production functions with common elasticities of substitution and common relative weights on capital, materials, and labor. However, we allow the weight on fuel relative to the other inputs to vary across divisions. We calibrate the relative weights of capital, labor, and materials to the aggregate data for all plants with the same fuel from GTAP. We then adjust the weight on fuel at the division level to match observed emission intensities relative to output, thereby capturing heterogeneity in fuel use across divisions.

The elasticities for electricity varieties, illustrated in Figure 8, are largely consistent with those of a generic sector, with a few exceptions. We use a lower Armington trade elasticity of .3 for electricity. In our baseline, we also assume that nuclear and hydro power operate with fixed capital stocks. For these technologies, we set substitution elasticities for the KL and the KLE bundle to zero. We thus ensure that electricity from nuclear and hydro power remains constant. For the aggregation of electricity, we assume an elasticity of 6 for substitution among divisions using the same fuel, and an elasticity of 10 for substitution across different fuels and between generated electricity and distribution services. Our robustness exercises in Section 4.3 vary these numbers.

The right-hand side of Table 7 lists elasticities relevant for the aggregation of goods, corresponding to the nesting structure shown in Figure 10. We define the *fudge good* as a Cobb-Douglas bundle of intermediates that represents all capital goods and the export good. Accordingly, the production functions \tilde{g} and h^n are identical. The fudge good serves as numeraire in our calibration. The consumption good is a Cobb-Douglas aggregate of the fudge good, vehicle services consumed by households, and housing services.

Benchmarking to EMF 32 meta-study. The Stanford Energy Modeling Forum (EMF) conducted a coordinated analysis of various carbon tax scenarios using 11 dynamic computable general equilibrium models specifically designed for evaluating environmental policy. Results were published in a special issue of *Climate Change Economics*, with headline numbers summarized in McFarland et al. (2018), henceforth McF, with additional detail in Macaluso et al. (2018).

The scenarios analyzed in the meta-study are broadly comparable to our exercise: the carbon tax is applied to U.S. CO₂ emissions from fossil fuels, with no concurrent policy changes abroad, and tax revenues are rebated to households. However, there are important differences. The EMF models prescribe a baseline with economic growth and a dynamic path for the carbon tax. The models also have richer dynamics than ours, in particular adjustment costs or explicit

tracking of vintages of capital goods.

We focus on the scenario of a \$50 per ton carbon tax that continues to grow at one percent after the initial introduction. Emissions trajectories in McF Figure 3 show that most of the EMF models feature fast adjustment to a state where emissions are close to constant. In many cases, it takes only one or a few periods for the model to reach a flat path. Similarly, the mix of primary energy is close to constant after the tax is imposed. These patterns suggest that transition dynamics are not a central feature in the leading models. Therefore, it is reasonable to compare a version of our exercise—differences of steady states with and without a \$50 carbon tax—to the models' predictions for the state after their initial short adjustment. This supports the validity of using comparative statics across steady states for evaluating medium-run policy impacts.

Our model delivers results that are quantitatively in line with the typical outcomes reported in the meta-study along several key dimensions: (i) emission reductions overall and by broad sector (ii) primary energy and electricity consumption and (iii) output loss and carbon tax bill.

Specifically, McF Figure 4 reports percentage reductions of emissions relative to a no-tax baseline, both total and broken down by electricity, transport as well as a combined category of households & industry. Across the 11 EMF models, economy-wide emission reductions range from 20% to 40%. Sectoral reductions fall between 50% and 70% for electricity, remain below 15% for transport, and lie between 15% and 30% for households & industry. In comparison, our analysis produces a 33% reduction in aggregate emissions, with sectoral contributions of 62% from electricity, 14% from transport, and 24% from the combined households and industry sectors. These results place us squarely within the range of the EMF models.

Regarding primary energy consumption, all EMF models project strong reductions in coal use—ranging from 30% to 90%—and modest reductions in oil use, below 10% (Macaluso 2018, Figure 2). Our results fall well within these ranges, with a 70% reduction in coal consumption and a 7% reduction in oil use. In contrast, EMF model predictions for natural gas vary widely: some show a 20% increase, while others predict a decline of up to 50%. Our results show a 35% reduction in gas use. This is partly driven by substitution toward renewable energy. However, even in the robustness scenario that effectively shuts down such substitution, gas use still falls by 20%. A key factor behind this result is the heterogeneity of power plants in our model. By allowing substitution away from the least efficient gas-fired electricity varieties, we capture an important channel for reducing gas use—even when renewables are constrained.

Across EMF models, aggregate energy and electricity consumption decline by 4–15%, as shown in Macaluso et al. (2018), Figure 5. In our model, agents consume 11% less energy and 6% less electricity—well within this reported range. The EMF models further report a few basis points lower growth per year on average over 25 years relative to the reference scenario. This number aligns closely with our result of a .7% drop in steady-state consumption, which translates to a similar cumulative effect over the medium run. Finally, the carbon tax revenue reported in the EMF models stabilizes after the initial adjustment phase and averages \$200–\$250 billion per year in 2010 dollars, or 1.1%–1.7% of initial GDP. In our model, the carbon tax

generates total revenue of about 1.2% of initial GDP, within this range.