Asset returns as carbon taxes*

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

April 2025

Abstract

In frictionless financial markets, a carbon tax on energy users provides the same incentives as a "replicating" cost-of-capital schedule that depends on firms' emission intensities, defined as scope 1 emissions relative to enterprise value. At currently observed emission intensities for the US economy, the fat right tail of the replicating return distribution is far beyond empirical estimates of pollution premia. Nevertheless, endogenous adjustment implies high carbon taxes are consistent with modest equilibrium return differentials. With heterogeneous preferences over dirty portfolios, emissions fall by more if a few green investors obtain nonpecuniary benefits from clean investments than when a majority perceives nonpecuniary costs.

*We thank Hunt Allcott, Niels Gormsen, Lars Hansen, Ben Hebert, Kilian Huber, Kjetil Storesletten, Annette Vissing-Jorgensen as well as conference participants at Chicago, MIT, Normac, 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 through costs of capital. If investors require, say, pollution premia on dirtier firms, it becomes more expensive for such firms to raise capital, and they should grow more slowly. A recent empirical literature measures return differentials between clean and dirty 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 sizeable fraction of households is willing to give up some return to benefit the environment. What is not clear, however, is how large such taste for green assets has to be for this to make 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 a large class of models with frictionless firm financing, a carbon tax on energy users provides the same incentives as a *replicating* cost-of-capital schedule that depends linearly on a firm's emission intensity, defined as scope 1 emissions relative to enterprise value. We use this equivalence result to calculate replicating costs of capital for the US economy. While premia over the cleanest sector are small for most sectors, the extreme skewness of emission intensities requires large premia for a few dirty sectors, such as power, transportation, and parts of manufacturing.

Formally, a carbon tax is equivalent to taxing capital at a rate that rises linearly with the emission intensity. For example, a carbon tax of \$100 per ton of CO_2 equivalent requires that the cost of capital rises by \$100 per unit of emission intensity. While the aggregate intensity for the US economy is currently 49 tons per million dollars or .000049 tons per dollar enterprise value, currently observed intensities at the 2-digit sector level range between zero and .0017tn/\$. The average replicating return premium is therefore 49bp and premia rise across sectors from a median of 5bp to 43bp at the 90th percentile and to 5% for the aviation sector at the 99th percentile. For publicly traded electricity companies, also located in the right tail of intensities, replicating premia rise from a median of 10% to 16% at the 90th percentile. This calculation with observed emission intensities provides an upper bound on required returns.

Once we take the equilibrium adjustment of emission intensities into account, replicating premia become significantly smaller, especially in sectors that offer potential for substitution to cleaner technologies. To predict adjustment, we use a multisector growth model with a detailed energy sector calibrated to sector-level data as well as firm-level data from the electricity sector. While the median sector requires a 4bp premium, the 90th and 99th percentiles drop to 31bp and 4%, respectively. Within electricity, adjustment is particularly strong as firms abandon coal and gas in favor of renewables to lower their costs of capital: the median premium drops to 23bp and the 99th percentile to 1.4%. Most of this adjustment occurs within firms—the replicating return schedule creates incentives for dirty firms to clean up their production.

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

After quantifying the effect of a replicating return schedule on the supply side of the economy, we ask when such a schedule emerges in equilibrium due to investor tastes. We first point out that when *all* households have preferences that incorporate a linear nonpecuniary cost from holding dirty assets, equilibrium returns are exactly replicating, that is, linear in emission intensity. In other words, we describe an economy where investor preferences imply a schedule of returns that provides the same incentives to firms *as if* a carbon tax were in place. Allocations are thus the same as with the carbon tax, although welfare is different since agents perceive nonpecuniary costs from dirty capital.

We then extend the model's asset demand side to consider heterogeneous investors who face short sale constraints. One important feature we incorporate is that not all investors share a concern about dirty capital. Moreover, we do not equate green investing motives with nonpecuniary costs of dirty capital, as is common in the literature, but instead consider a broader class of preferences: we allow some agents to obtain a nonpecuniary *benefit* from holding capital in particular sectors. The idea here is to compare different ways in which climate finance products could be marketed. The key common property of our heterogeneous agent economies is that any return schedule is capped: the presence of unconcerned investors places an upper bound on returns that dirty firms have to pay.

Quantitative analysis of our model with heterogeneous investors shows that a small number of investors can have a big impact. This is because the distribution of emissions is skewed, so the effectiveness of green investing turns on changing the behavior of a small share of dirty firms (weighted by capital). We show that nonpecuniary costs cannot generate a high-impact return schedule: a small share of unconcerned investors mutes incentives since the cap implies weak incentives for the dirtiest firms that produce most emissions. Quantitatively, the presence of just 2% unconcerned investors takes away one half of the emission reduction that a representative agent with taste for a $100-per-ton carbon tax would implement. We conclude that popular marketing of green investment products based on negative screening, akin to a divestment campaign, is misguided.

Instead, nonpecuniary benefits from investing in clean electricity can have substantial effects even if green investors account for only a small share of total wealth. While nonpecuniary costs work like a capped tax on capital with a constant rate for the dirtiest firms, nonpecuniary benefits work like a targeted subsidy that dirty firms like to take advantage of. Quantitatively, 3% of investors who forego 3% of wealth per year in equilibrium can achieve about one half the reductions of a $100-per-ton carbon tax. What is crucial here is that green investors obtain nonpecuniary benefits also from holding dirty electricity companies, and as a result subsidize those firms at the expense of cleaner non-electricity firms. Favoring electricity this way allows an incentive scheme for the dirtiest firms that is not diluted by the majority of unconcerned investors.

To see how our equivalence result works formally, compare a carbon tax on fossil fuel use to a tax on equity that rises with emission intensity. With a tax on equity in place, investors and firms distinguish stocks by the emission intensities of the firms that issue them. We thus consider a continuum of markets for stocks indexed by emission intensity. When households choose portfolios to maximize utility, they price in return differentials that compensate them for the capital tax. Investor optimality thus implies a return *schedule* for a cross-section of equity markets that differ by convenience yields, here due to tax treatment. This replicating return schedule is independent of the production side of the economy.

It is crucial for the result that emission intensities are endogenous: firms choose intensities as part of production plans that include fossil fuel inputs. Equilibrium matches households' asset demand, that is, the return schedule, with firms that maximize shareholder value. The distributions of returns and emission intensities are then jointly determined: technologies are funded via those markets where they achieve the highest value. To this end, firms' choice of production plans takes into account shareholders' required after-tax rate of returns, which depend on the emission intensity. In particular, when the tax rate per dollar of equity rises linearly with intensity, then the tax bill—equity multiplied by the tax rate—becomes independent of the value of equity and enters the firm's problem exactly like the carbon tax bill.

The replicating schedule we characterize need not literally come from capital taxes. More generally, it provides a ballpark estimate of what return differentials are required to make climate finance have an impact, regardless of where those differentials ultimately come from. On the firm side, the equivalence result requires shareholder value maximization but not specific assumptions on technology or market structure. It is thus robust to many extensions, including heterogeneous capital goods, risk from productivity or preference shocks, or adjustment costs to capital. It also holds in the presence of rents from product market power or distortionary taxes. Our approach can therefore be applied to compute replicating returns implied by most existing computable general equilibrium models used to study environmental policy.

Since asset demand in our model defines a return schedule, it provides an incentive scheme, and not simply a penalty on currently dirty firms. The incentive scheme rewards substitution within firms just like substitution across firms. In fact, the former drives a large share of emission reductions in our quantitative exercises, especially in the electricity sector. The equivalence result requires no assumptions on the boundaries of the firm: it does not matter, for example, whether greener production occurs in specific firms or whether it is done in large conglomerates that also engage in a lot of dirty production. As with a carbon tax, decarbonization may occur because clean firms grow at the expense of dirty firms since their cost of capital is lower or because initially dirty firms transform themselves by cleaning up their production in order to enjoy a lower cost of capital.

All of our results assume that scope 1 emissions can be measured and are public information. Whether incentives come from a capital tax or investor tastes, the emission intensity relative to the value of the 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. We also note that our results rely only on emissions generated by fossil fuel use, which are arguably even easier to measure since they are directly related to purchases of inputs.

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 equates 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.

In our heterogeneous agent economies, the return schedule reflects investor sorting. Preferences imply that green investors prefer to hold cleaner to dirtier capital, holding fixed pecuniary returns: green investors must be compensated for holding the dirtier capital. In equilibrium, they hold all capital employed at emission intensities below some threshold, while neutral investors hold capital employed at higher intensities. While neutral investors perceive clean capital as too expensive, short-sale constraints prevent them from equating returns. This is what allows green investors who obtain nonpecuniary benefits from clean electricity to provide a powerful incentive scheme. If green investors are willing to hold both clean and dirty electricity firms, they can push down returns so they can be responsive to emission intensity even in the presence of many neutral investors.

Related literature. Our modeling framework follows the long tradition of general equilibrium integrated assessment models designed to assess policies such as carbon taxes (see Nordhaus 2013 for a historical survey and IAMC (2025 website) for an overview of models currently used around the world). Since our interest is in comparing tax schemes and investor preferences rather than, say, computing optimal tax rates, we keep the climate side of the model simple and only minimally capture a climate externality. We share the focus on one country that engages in international trade with climate models of the US economy, for example Goulder and Hafstead (2017) or Yuan et al. (2019). In particular, we also specify an input-output structure that allows tax effects to propagate through supply chains. For the electricity sector, we further bring in firm-level information to model substitution, in the spirit of detailed "bottom-up" models used in energy economics.

A number of other papers have discussed the relationship between carbon taxes and the cross-section of returns in financial markets. Iovino et al. (2024) study the effects of capital taxation on an economy's carbon footprint via capital reallocation. Papoutsi et al. (2024) consider reallocation due to cross-sectional effects of unconventional monetary policy, in particular the ECB's corporate bond purchase program. Pedersen (2024) compares carbon taxes to return differentials under various assumptions on firm objectives. A key difference between our setup and these papers is the concept of equilibrium: we emphasize that when investors distinguish stocks by emission intensity, capital market equilibrium involves a cross-section of markets that clear to jointly determine prices and properties of traded assets, here emission intensities.[2] This is what leads to a very general equivalence result.

Our return schedule provides a benchmark to interpret evidence on pollution premia, that is, cross-sectional differences in returns, and hence costs of capital, between dirty and clean firms. Evidence from distinct measurement approaches points to pollution premia of at most a couple of percentage points—as yet too little to provide strong incentives to the dirtiest firms. Hong and Shore (2023) survey a large literature that tries to measure the pollution premium in return data. A major difficulty is that the transition towards greener investing makes it hard to estimate means (Pástor, Stambaugh and Taylor 2021). Studies based on investor surveys reach similar conclusions on averages, and further point to substantial heterogeneity in tastes (for example, Riedl and Smeets 2017, Giglio et al. 2025, Aron-Dine et al. 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. They emphasize 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 (for an overview, see Pastor et al. 2024).[3] Early work studied the effects of negative screening: some investors prefer not to hold some assets and end up with riskier portfolios (Heinkel et al. 2001; Geczy et al. 2021). Berk and van Binsbergen (2021) show that if only few investors apply negative screens, the aggregate impact of climate finance is quantitatively small. Our model assumes instead that green investors earn a convenience yield from holding green assets, as in Pástor, Stambaugh and Taylor (2021), Pedersen, Fitzgibbons and Pomorski (2021), or Aron-Dine et al. (2024). While those papers study asset pricing and investors' green portfolio "tilts", we focus on emission reductions in a production economy. We further emphasize the distinction between nonpecuniary costs and benefits: in the latter case, even a few investors can have a large impact.

The paper is structured as follows. Section 2 introduces a model with a representative agent and standard preferences and shows the equivalence between a carbon tax and a capital tax with a rate that responds to emission intensity. 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]This theme is familiar from other work. For example, Fostel and Geanakoplos (2015) describe models with collateral constraints where markets clear to jointly determine interest rates and the default risk of the borrower. Rogerson et al. (2005) survey models of competitive search markets 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 emission intensity of the issuer.

[3]We follow most of this work in assuming *non-consequentialist* preferences, that is, investors like clean assets for private reasons, not because of their aggregate real effects, consistent with the evidence in Bonnefon et al. (2025) and Heeb et al. (2023).

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 as well as 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

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

where $v > 0$. This is a minimal way to introduce a climate externality. The precise nature of damages is not important for our exercise since we take the size of the carbon tax as given.

Technology. Production of the consumption good C_t uses a vector of N intermediate goods y_t where, $n \in \mathcal{N} := \{1, \ldots, N\}$. Intermediate good n is made from labor l_t^n, capital k_t^n, a collection of intermediate goods x_t^n and $\tilde{N}$ foreign goods collected in a vector $\tilde{x}_t^n$. The economy exports a good $\tilde{Y}_t$ made from intermediates. There are also N capital goods, one specific to each intermediate good, with capital good n made from intermediates collected in a vector z_t^n. Capital can be first used one period after it has been installed and depreciates at rate δ_n in sector n. Production functions for all goods are homogeneous of degree one in inputs.

Technology is described by the production functions g, $\tilde{g}$, h^n, and f^n for the consumption good, the export good, the nth capital good, and the nth intermediate good, respectively. The resource constraints are

$$\begin{aligned} C_t &= g \left(y_t \right), \\ \tilde{Y}_t &= \tilde{g} \left(\tilde{y}_t \right), \\ k_{t+1}^n &= \left(1 - \delta_n \right) k_t^n + h^n \left(z_t^n \right), \qquad n \in \mathcal{N} \\ y_{t,n} + \tilde{y}_{t,n} + \sum_{j \in \mathcal{N}} \left(x_{t,n}^j + z_{t,n}^j \right) &= f^n \left(l_t^n, k_t^n, x_t^n, \tilde{x}_t^n \right), \qquad n \in \mathcal{N}. \end{aligned} \tag{2}$$

The last equation is the resource constraint for intermediate good n. The production of good n is used for consumption $y_{t,n}$, export goods $\tilde{y}_{t,n}$, other intermediates $x_{t,n}^j$ and investment $z_{t,n}^j$.

Fossil fuels and emissions. A subset of goods $\mathcal{F} \subset \mathcal{N}$ are fossil fuels. They are used in the production of intermediates, but not in making the final good or capital. This convention is not restrictive since we can always define additional intermediates. Using fossil fuel $f \in F$ in production produces ϕ_f emissions per unit of fuel. The emissions e_t^n of sector n are the sum of emissions from each of its fossil fuel inputs. The evolution of the state of the environment can then be written as

$$\eta_t = \eta_{t-1} + \sum_{n \in \mathcal{N}} e_t^n + \tilde{e}_t = \eta_{t-1} + \sum_{n \in \mathcal{N}} \sum_{f \in \mathcal{F}} \phi_f x_{t,f}^n + \tilde{e}_t, \tag{3}$$

where $\tilde{e}_t$ is an exogenous sequence that represents foreign emissions. To keep notation simple, fossil fuels and other goods (such as foreign goods $\tilde{x}$) are disjoint sets of goods. This is not restrictive since we can define a fossil fuel that is produced 1-1 from a foreign good.[4]

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. 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:

$$\sum_{n \in \mathcal{N}} \sum_{m \in \tilde{\mathcal{N}}} \tilde{p}_{t,m} \tilde{x}_{t,m}^n = \tilde{P}_t \tilde{Y}_t. \tag{4}$$

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

$$\sum_{n \in \mathcal{N}} l_t^n = 1. \tag{5}$$

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 (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 fossil fuel j in sector n, say, to its total marginal cost, which consists not only of the shadow cost of the fuel, that is, the multiplier $\lambda_{t,j}$ on the jth resource constraint in the last equation of (5), but also the social cost of carbon. Writing h_i for the ith partial derivative of a function h, we have

$$g_n(y_t) f_j^n \left(l_t^n, k_t^n, x_t^n, \tilde{x}_t^n \right) = \lambda_{t,j} + \phi_j SCC_t.$$

Standard arguments imply that a competitive equilibrium with a carbon tax rate equal to the social cost of carbon decentralizes the social optimum. Intuitively, the spot price of fuel works

[4]The benefit of this notation is that the elements of $\tilde{x}$ do not cause emissions in equation (3). For example, $\tilde{x}$ may contain "gas from Canada". Before this gas can be burned domestically, it has to be transformed 1-1 into a domestic intermediate variety "gas from Canada ready for use". Our convention implies that the burning of this domestic variety causes the emissions.

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

We consider competitive firms owned by households and traded in frictionless financial markets. All goods are traded in competitive markets. The consumption good serves as numéraire.

Firms and equity. We assume one firm type for each production function in (2). Since firms that make the same good are all identical, we describe one representative firm per good. Producers of the consumption goods and the N capital goods exist for one period. They buy intermediate goods at prices $p_{t,n}$. Producers of the nth capital good sell at the price $p^k_{t,n}$. Intermediate goods firms exist for two periods. Firms enter at date t, raise funds by issuing equity and use those funds to purchase capital goods, both at date t . At date $t+1$, they purchase intermediates, hire labor, produce, and sell undepreciated capital. We denote by $v_{t,n}$ the cum-dividend value of a firm in sector n to shareholders at date t.

Capital and carbon taxes. The government levies a proportional carbon tax τ^c_t on firms per unit of emissions. In addition, there are taxes on households' holdings of equity. The tax rate per unit of equity depends on the emission intensity of the firm that has issued the equity. The emission intensity of a firm that is worth v_t and uses fossil fuels that generate emissions e_t is

$$\varepsilon_t = \frac{e_t}{v_t}$$

The capital tax schedule applied at date t is given by a weakly increasing function $\tau^k_t(\varepsilon_t)$. In order to implement an actual tax with this schedule, the government would have to measure emissions over some period of time as well as the value of the firm at the end of that period. The exact timing convention is not crucial but simplifies arguments below.

Equity markets. Households save by investing in equity. Since capital taxes depend on the intensity of production at date $t+1$, firms and households at date t distinguish equity claims by their intensity ε_{t+1}. We therefore consider a continuum of equity markets indexed by ε_{t+1}. We quote prices in those markets in terms of returns: let $R_{t+1}(\varepsilon)$ denote the gross rate of return on equity of intensity ε earned by investing one unit of numéraire from date t to 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 nonnegative equity holdings $s_{t+1}(\varepsilon)$ subject to the budget constraint

$$C_t + \int s_{t+1}(\varepsilon)\, d\varepsilon = \int R_t(\varepsilon)\left(1-\tau^k_t(\varepsilon)\right) s_t(\varepsilon)\, d\varepsilon + w_t + T_t, \tag{6}$$

where w_t is the wage and T_t is a lump sum transfer from the government. Here the integrals sum up over all equity markets the household can invest in. Households care about after-tax returns.

Since the problem is deterministic, after-tax returns on all equity claims must be equal. Asset prices thus satisfy the Euler equations

$$R_{t+1}\left(0\right) = \beta u'\left(C_t\right)/u'\left(C_{t+1}\right), \tag{7}$$

$$R_{t+1}\left(\varepsilon_{t+1}\right)\left(1-\tau^k_{t+1}\left(\varepsilon_{t+1}\right)\right) = R_{t+1}\left(0\right). \tag{8}$$

We label $R_t\left(0\right)$ the "clean rate", that is, the rate of return on equity of a firm that does not use any fossil fuels. Since that equity claim is not taxed, its rate of return is equal to the household's marginal rate of substitution between consumption at date t and $t+1$. Returns on other equity claims are typically higher than the clean rate to compensate investors for taxes, and more so the dirtier the firm is, in the sense of higher ε.

We note that the shape of the capital tax schedule is inherited to a close approximation by the schedule of *net* returns $r_{t+1}\left(\varepsilon_{t+1}\right) = R_{t+1}\left(\varepsilon_{t+1}\right) - 1$. As long as returns and taxes are small decimal numbers, their products are an order of magnitude smaller. We can then write (8) as

$$r_{t+1}\left(\varepsilon_{t+1}\right) - r_{t+1}\left(0\right) \approx \tau^k_{t+1}\left(\varepsilon_{t+1}\right). \tag{9}$$

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 its capital tax rate. In particular, the case of linear capital tax rates that is central to our analysis below then implies an approximately linear return schedule. When we use language like "linear return schedule" in what follows, it is understood that this approximation holds.

Intermediate goods firm problem. Consider a firm that makes good n with a production plan $X = (l, k, x, \tilde{x})$, that is, a list of labor, capital, intermediates and foreign goods. The *pre-carbon-tax* cum-dividend value of the firm at date $t+1$ generated by the plan is

$$V_{t+1,n}\left(X\right) = p_{t+1,n} f^n\left(l,k,x,\tilde{x}\right) - \sum_{j\in\mathcal{N}} p_{t+1,j}x_j - \tilde{p}_{t+1,j}\tilde{x}_j - w_{t+1}l + p^k_{t+1,n}k\left(1-\delta_n\right). \tag{10}$$

Firm value includes pretax capital income, that is, revenue less expenditure on domestically produced intermediates and foreign goods, as well as the wage bill, plus the value of undepreciated capital. The firm further pays a proportional carbon tax τ^c_{t+1} on its emissions.

At time t, firms maximize shareholder value by choosing a production plan. They understand that every production plan implies a particular emission intensity ε, which determines the equity market in which they raise funds. Firms take as given the entire return schedule $R_{t+1}\left(\varepsilon\right)$ that describes how returns respond to intensity. Firms thus compute shareholder value using the discount factor of their particular shareholders. Firms solve

$$\max_X \left\{ -p^k_{t,n} k + R_{t+1}(\varepsilon)^{-1} v_{t+1,n} \right\}$$
$$\text{s.t. } v_{t+1,n} = V_{t+1,n}(X) - \tau^c_{t+1} \sum_{f \in \mathcal{F}} \phi_f x_f, \qquad \varepsilon = \frac{\sum_{f \in \mathcal{F}} \phi_f x_f}{v_{t+1,n}}. \tag{11}$$

Here, the first constraint defines the *after-carbon-tax tax* date $t+1$ value as the pre-tax value (10) less the carbon tax bill. The second constraint is the definition of emission intensity relative to (after tax) value. The firm's emissions in the numerator and its value in the denominator both depend on the production plan.

Since households' Euler equations (8) hold in any equilibrium, we can substitute into (11) and rewrite shareholder value as

$$-p^k_{t,n} k + R_{t+1}(0)^{-1} \left(1 - \tau^k_{t+1}(\varepsilon)\right) v_{t+1,n}.$$

A firm that discounts using the cost of capital $R_{t+1}(\varepsilon)$ thus behaves like a firm that discounts using the clean rate $R_{t+1}(0)$, but pays a tax on firm value that depends on the emission intensity ε. When the firm contemplates a production plan with more fossil fuels x_f, it therefore not only takes into account the standard effect on profits and the carbon tax bill, but also the fact that its cost of capital rises with the emission intensity ε.

The case of linear capital tax rates. The effect of capital taxes is particularly simple when the capital tax rate schedule is linear in the emission intensity, that is, $\tau^k_{t+1}(\varepsilon) = \bar{\tau}^k_{t+1}\varepsilon$ for some slope $\bar{\tau}^k_{t+1}$, say. In this case, we can substitute out ε and $v_{t+1,n}$ using the constraints of problem (11) to obtain

$$-p^k_{t,n} k + R_{t+1}(0)^{-1} \left(V_{t+1,n}(X) - \left(\tau^c_{t+1} + \bar{\tau}^k_{t+1}\right) \sum_{f \in \mathcal{F}} \phi_f x_f \right). \tag{12}$$

Since emission intensity is defined relative to value, a linear tax implies that value cancels. As a result, the slope of the capital tax schedule now enters as a proportional tax on emissions, exactly like the carbon tax. This observation is the basis of our main equivalence result.

With a linear capital tax rate schedule, the firm's first-order conditions for capital and fossil fuel f are given by

$$p^k_{t,n} = R_{t+1}(0)^{-1} \left(p_{t+1,n} \frac{df^n}{dk}(l,k,x,\tilde{x}) + p^k_{t+1,n}(1-\delta_n) \right),$$
$$p_{t+1,n} \frac{df^n}{dx_f}(l,k,x,\tilde{x}) = p_{t+1,f} + \phi_f \left(\tau^c_{t+1} + \bar{\tau}^k_{t+1}\right), \tag{13}$$

respectively. The first-order condition for capital is standard: it equates the user cost of capital to its marginal product df^n/dk. This is a result of the linear rate schedule: in general, investment in a sector incurs a capital tax and moreover lowers the emission intensity for that sector. With a linear tax rate, the two effects cancel so the capital tax schedule does not distort the firm's investment decision. We return to this issue when we discuss concave schedules in Section 5 below. We also note that the firm's other first-order conditions for variable inputs are the same as without taxes, and we do not write them out explicitly.

The optimal choice of a fossil fuel equates its marginal product to its market price plus the marginal tax rate. The carbon tax appears here because it is proportional to emissions. With a capital tax rate schedule that depends on the emission intensity, capital taxes also matter: higher emissions increase the emission intensity and thereby the cost of capital. The special feature of a linear capital tax schedule is that the marginal effect on the cost of capital is independent of the emission intensity, it only depends on the slope. As a result, the slope of the capital tax schedule plays exactly the same role as the carbon tax rate.

Other firms and government. Firms that make consumption, export and capital goods choose production plans to equate marginal products of intermediates to intermediates' prices. To write their first-order conditions, we again denote the nth partial derivative of a function g by g_n. We then have

$$p_{t,n} = g_n\left(y_t\right) = \tilde{P}_t \tilde{g}_n\left(\tilde{y}_t\right) = p^k_{t,m} h^m_n\left(z^m_t\right),\ m \in \mathcal{N}. \tag{14}$$

The first two equalities are FOCs for consumption and export goods producers, respectively. The last equality says that the marginal product of good n for a capital goods producer that sells capital to sector m firms at the price $p^k_{t,m}$ is also equal to the price of intermediate good n.

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

$$T_t = \sum_n \left(\tau^k_t\left(\varepsilon^n_t\right) v_{t,n} + \tau^c_t \sum_{f \in \mathcal{F}} \phi_f x^n_{t,f} \right). \tag{15}$$

The transfer is equal to the sum of capital tax income and carbon tax income. We specify policy by a capital tax schedule τ^k_t, possibly nonlinear in emission intensity, and a carbon tax rate τ^c_t. The transfer T_t then adjusts so the budget constraint holds in equilibrium.

Competitive equilibrium. Equity market clearing requires that households' portfolio demand for equity of emission intensity ε is the same as the capital demand of all firms that choose production plans with emission intensity ε. Since we have a finite number of sectors, trade will occur only on a finite number of equity markets. Denoting by ε^n_{t+1} the emission intensity in sector n, we write market clearing in any market with positive trade as

$$s_{t+1}\left(\varepsilon\right) = \sum_{n \in \mathcal{N}:\ \varepsilon^n_{t+1} = \varepsilon} p^k_{t,n} k^n_{t+1}. \tag{16}$$

Here the right-hand side is the supply of equity by all sectors that choose emission intensity ε.

For a given capital tax schedule and carbon tax rate, an equilibrium consists of an allocation together with vectors of intermediate goods prices, capital prices, a wage, a government transfer and a return schedule such that (i) households maximize utility (1) subject to the budget constraints (6), (ii) intermediate goods firms maximize shareholder value (11), (iii) other firms optimize so first-order conditions (14) holds, (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 (15) holds.

We are now ready to state our equivalence result.

Proposition 1. *For any sequence of carbon tax rates τ_t^c, an equilibrium with carbon tax rate τ_t^c and no capital tax leads to the same allocation as an equilibrium with carbon tax rates $\hat{\tau}_t^c < \tau_t^c$ and a capital tax schedule*

$$\tau_t^k(\varepsilon) = (\tau_t^c - \hat{\tau}_t^c)\,\varepsilon. \tag{17}$$

In particular, if τ_t^c equals the social cost of carbon, then the equilibrium with capital tax schedule $\tau_t^k(\varepsilon) = \tau_t^c \varepsilon$ achieves the solution to the planner problem.

The proof follows from comparing equilibrium conditions under the two tax schemes. Only households and intermediate goods firms face different problems: households pay capital tax, whereas intermediate goods firms pay carbon tax. When the capital tax rate schedule is linear, decisions on investment and fossil fuels are determined by (7) and (13). In these conditions, the slope of the capital tax rate schedule and the carbon tax enter as a sum. All other firm first-order conditions are identical under the two tax schemes. The only other equilibrium condition affected by taxes is the government budget constraint.(15). With a linear capital tax rate schedule $\tau_t^k(\varepsilon) = \bar{\tau}_t^k\,\varepsilon$, we can cancel the firm values $v_{t,n}$ so tax revenue on the right-hand side can be written as

$$\left(\bar{\tau}_t^k + \tau_t^c\right) \sum_{f \in \mathcal{F}} \phi_f x_f^n$$

and the two parameters again enter only as a sum.

Proposition 1 compares the equilibrium effects of alternative tax schemes and hence the interaction of firms, households, and the government. To interpret the result, it is helpful to isolate the contribution of firms. The role of households and the government in the model is to provide a return schedule $R_{t+1}(\varepsilon)$ that incentivizes firms to save on emissions. The effects of this schedule on production do not depend on how exactly it comes about. More formally, we have

Proposition 2. *For any return schedule $R_t(\varepsilon) = R_t(0)\,(1 - \tau_t^c\,\varepsilon)^{-1}$, firms' optimization together with balanced trade as well as market clearing for intermediates and labor imply the same prices and production plans and hence emissions as under the carbon tax τ_t^c.*

Per our discussion above, the objective in intermediate goods firms' problem (11) reduces to (12) when the capital tax schedule is linear, so the slope $\bar{\tau}_t^k$ and the carbon tax rate τ_t^c enter as a sum. All other firm problems and the market clearing conditions for intermediates and labor are unchanged.

The propositions underscore that incentives provided by a carbon tax and a capital tax schedule are perfect substitutes. Proposition 1 says that if we have a desired target carbon tax τ_t^c then, if an actual carbon tax $\hat{\tau}_t^c$ can be implemented, the slope of the capital tax schedule needed to get to the target should be $\tau_t^c - \hat{\tau}_t^c$. Proposition 2 provides an interpretation in terms of climate finance: if we have a desired carbon tax rate τ_t^c, but can impose only a carbon tax $\hat{\tau}_t^c$, then financial markets have to deliver a return schedule $R_t(\varepsilon) = R_t(0)\left(1 - \left(\tau_t^c - \hat{\tau}_t^c\right)\varepsilon\right)^{-1}$ in order to achieve the target carbon tax. The cross section of returns thus has to respond less to the emission intensity the higher is the actual carbon tax $\hat{\tau}_t^c$.

It is important that there is a capital tax rate *schedule*, and not a flat tax rate. Capital owners and firms internalize the externality because the former worry that their after-tax returns are too low, and the latter worry that that their cost of capital is too high when their emissions intensity is too high. To understand what incentives the tax schedule has to provide, we thus not only want to know what capital taxes firms would pay *in equilibrium* with an efficient capital tax system, but also what off-equilibrium threats such a system must provide. It is natural to focus on two numbers per firm: its equilibrium tax rate, and the rate it would have to pay if it did not adjust its emission intensity away from the initial equilibrium. Using a calibrated model, we can compute both numbers for every tax rate.

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 the equivalence between incentives provided by the two tax schemes.

Objective of the firm and equilibrium concept. We have assumed so far that firms live for two periods: they raise funds to buy capital in the first period, then produce, and pay out the resulting capital income in the second period. This setup is a minimal way to talk directly about equity and have return premia depend on the emission intensity relative to market value. However, the Modigliani-Miller theorem holds in our model, so many equivalent alternative structures for the firm are possible. For an alternative example that is helpful to clarify the equilibrium concept, suppose firms live for only one period and rent capital that is owned by households. In this case, capital taxes are levied on households' capital holdings rather than their equity holdings.

With this alternative market structure, households care about emission intensities when they rent out their capital, not when they invest. Instead of many equity markets indexed by the emission intensity ε, we thus have many rental markets for capital. Since emission intensity is observable, rental rates depend on it, and household optimality determines a *schedule* of rental rates. Rental rates depend on emission intensity since households pass on capital taxes to firms by charging higher rates. The decentralization is familiar from models with variable capacity utilization, where depreciation, and hence the rental rate of capital, also depends on features of the production plan, namely utilization. Rental contracts for cars that depend on mileage also work this way.

The common denominator of the rental and equity market decentralizations is that competitive markets can coordinate sorting of resources by emission intensity when the latter is public information. In the rental market case, the rent schedule provides price signals that guide tradeoffs between dirty use of capital and tax payments. In the equity market case, this job is performed by the return schedule. The market transactions that take place to implement coordination are of course different. In the rental market case, we have within-period trades between firms and households. In the equity market case, coordination involves firm decisions made by shareholders who acquire the firms. But the resulting allocation reflects the same optimal tradeoff.

Rents and other taxes. Our model assumes perfect competition in product markets and constant returns to scale, so firms do not earn any rents. We also focus exclusively on capital taxes and carbon taxes and do not allow for any other taxes such as corporate income tax. Neither assumption is important for the equivalence result. Indeed, decreasing returns, markups, or other taxes would affect the intermediate goods firm problem (11) only via the definition of the pre-carbon-tax firm value (10). As a result, the key step in our derivation that makes the slope of a linear capital tax schedule enter like a carbon tax rate is unaffected.

Boundaries of the firm. Our decentralization assumes that every technology is operated by a separate firm. However, the equivalence result holds also with multigood firms, with emission intensity 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, we introduce multigood firm types $m \in \mathcal{M}$, where a firm of type m makes all intermediate goods $n \in \mathcal{N}^m$, a subset of the set of all goods. Multigood firms can coexist in any sector n with firms that operate only one technology.

A multigood firm of type m chooses a production plan $X = (X_n)_{n \in \mathcal{N}^m}$ that specifies inputs for all the technologies n the firm operates. Its date $t+1$ value for a given plan X is the sum of values contributed by the individual technologies. Each component is given by profits plus undepreciated capital, as in (10). The overall pretax firm value is the sum

$$W_{t+1,m}(X) = \sum_{n \in \mathcal{N}^m} V_{t+1,n}(X_n)$$

The key new margin is that the firm at date t makes a portfolio choice on how much capital to install to operate each of the technologies n. As before, however, capital taxes are levied at the firm level, and therefore depend on the firm's emission intensity: the ratio of its aggregate emissions relative to firm value $W_{t+1,m}$.

A multigood firm chooses an equity market in which to raise equity, given the return schedule that depends on firm-level emission intensities ε. Household Euler equations for equity (7)-(8)

hold as before, since the asset demand side has not changed. The multigood firm thus solves

$$\max_{X}\left\{-\sum_{n\in\mathcal{N}^m} p^k_{t,n}k^n + R_{t+1}(\varepsilon)^{-1} w_{t+1,m}\right\}$$
$$\text{s.t. } w_{t+1,m} = W_{t+1,m}(X) - \tau^c_{t+1}\sum_{n\in\mathcal{N}^m}\sum_{f\in\mathcal{F}}\phi_f x^n_f, \qquad \varepsilon = \frac{\sum_{n\in\mathcal{N}^m}\sum_{f\in\mathcal{F}}\phi_f x^n_f}{w_{t+1,m}}, \quad k=\sum_{n\in\mathcal{N}^m}k^n. \tag{18}$$

The multigood firm problem is analogous to the single-good firm problem (11). The only difference is that aggregate emissions are now a sum over emissions contributed by the different technologies n. The equivalence result thus follows by the same argument.

Risk. While our model so far is deterministic, the equivalence result extends 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 do not introduce probability distributions for the shocks. We instead assume that shocks take on a finite number of values, and that there is some probability over sequences of states ω_{t+1} that reflects household beliefs. All expectations formed below are computed using this probability.

Introducing risk only matters for households and intermediate goods firms who make intertemporal decisions; within-period decisions by other firms work as before. For an intermediate goods firm, the production plan chosen at date t is now partly state-contingent because its 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 $V_{t+1,n}(X;\omega_{t+1})$. Here the last argument captures the direct effect of exogenous technology shocks on profit. It could reflect, for example, the effect of productivity shocks on the production function.

Since firms' choice of fossil fuel inputs and their value both respond to shocks, so does their emission intensity: ε is also now understood to be a function of ω_{t+1}. When capital taxes are in place, the tax rate is therefore stochastic from the perspective of date t. The cross section of relevant equity markets is then identified by all possible *random variables* ε, not simply numbers. In the presence of taxes, an equity claim is different from another as long as it is taxed at a different rate with positive probability. Households and firms thus assess equity claims with different stochastic emission intensity profiles.

The household budget constraint (6) takes the same form as before, as long as variables are understood to be random variables. Similarly, the Euler equations (7)-(8) hold in the same form. With a representative agent who prices all assets, 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 all contingent claims in the absence of capital taxation—it defines a random clean return. Returns on equity differ from this benchmark state by state depending on the emission intensity profile of the firm.

Shareholder value is now the *expected* present value

$$-p_{t,n}^{k}k^{n}+E_{t}\left[R_{t+1}\left(0\right)^{-1}\left(1-\tau_{t+1}^{k}\left(\varepsilon\right)\right)V_{t+1,n}\left(X;\omega_{t+1}\right)\right]$$

where we have already used the Euler equation for the discount rate $R_{t+1}(\varepsilon)$. It is immediate that the problem with a linear capital tax schedule reduces again to a problem with a proportional carbon tax.

Scope 2 emissions So far, we have focused exclusively on scope 1 emissions, that is, taxation of fossil fuels happens when they are used, and not, say, when they are produced. Moreover, there is no taxation of electricity use. Scope 2 emissions attribute emissions generated in the production of electricity to the user of electricity, and not the producer as in our setting. It is possible to derive an alternative equivalence results with this convention. Suppose that we define the emission intensity for non-electricity firms as the sum of scope 1 and scope 2 emissions over enterprise value for all firms. We then introduce a linear capital tax rate as before. In order not to double-count emissions, we further set emission intensities for electricity firms to zero. We would then again obtain equivalence to a carbon tax.

The difference between taxes on capital based on scope 1 versus scope 2 emissions is the same as taxing electricity at the producer or consumer level, respectively. As long as we avoid double-counting, there is no difference in allocations. When intensities are defined based on only scope 1 emissions, dirty electricity firms pay higher costs of capital and hence charge higher prices for electricity, which discourages electricity use. When intensities are based on scope 2 emissions as well, then dirty electricity prices are lower, but users are discouraged by capital taxes that take electricity use into account. A practical difference between the two cases 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 comprise a large number of additional measurement approaches, including dirtiness of inputs as well as how many emissions are generated when using a firm's product. For example, both petroleum producers and car manufacturers are attributed scope 3 emissions generated by gas-fueled cars. We are interested in equivalence to a carbon tax, which taxes every emission exactly once. As a result, 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.

There is one simple variant of our equivalence result that uses part of scope 3 emissions: a tax on capital of fossil fuel producers. Suppose we remove any taxes on users, and define emission intensities for fossil fuel producers based on emissions generated by their downstream use. In emission accounting frameworks, such downstream emissions are counted as part of scope 3 emissions. We can then derive a set of replicating returns where only fossil fuel producers pay positive premia, whereas all fuel users pay the clean rate. We note however, that such a tax on its own covers only emissions due to domestically produced fuels, including those burnt abroad, and not imported fuels. An advantage of a capital tax based on scope 1 emissions is that it is easy to cover all domestic emissions, regardless of where fuels come from.

Capital taxes based on scope 1 emissions apply to fossil fuel producers only to the extent they themselves burn fuel, say, during extraction. As a result, replicating returns on fossil fuel firms in our baseline specification are relatively small. This is in contrast to a lot of existing literature on green investing that focuses on divestment from fossil fuels as a key strategy. We also note that it is straightforward to extend our model to add exhaustible resources as assets a la Hotelling. Fossil fuel reserves would then comprise part of the value of fossil fuel companies, and some of those assets would become stranded when a carbon tax is implemented. This effect does not change our results: as with a carbon tax, the drop in the value of reserves would come from future low cash flows, as opposed to a return differential due to capital taxes.

Adjustment costs. Suppose adding new capital goods i_t^n to an existing capital stock k_t^n installed for making intermediate goods in sector n incurs adjustment costs $c^n\left(i_t^n/k_t^n\right)k_t^n$ in units of the final good, where c^n is convex with a minimum at δ_n such that $c^n\left(\delta_n\right) = 0$. To illustrate the dynamics, we assume an infinitely lived intermediate goods firm. We formulate the firm problem recursively, with capital as the only endogenous state variable. Since technology exhibits constant returns to scale, the value function of the firm is linear in capital; we write it as $q_{t,n}k_t^n$, where $q_{t,n}$ is the shadow value of capital installed in sector n.

At date t, a firm comes in with capital stock k_t^n and then chooses a production plan $X = (l, k, x, \tilde{x})$ for the following period. This plan matters for the equity market in which the firm is traded at date t. The rate of return required by shareholders in that market determines the appropriate discount rate the firm uses to evaluate production plans. As before, it is convenient to define the firm's pre-carbon-tax cum-dividend value at date $t+1$

$$V_{t+1,n}\left(X\right) = p_{t+1,n}f^n\left(l, k, x, \tilde{x}\right) - \sum_{j\in\mathcal{N}} p_{t+1,j}x_j - \tilde{p}_{t+1,j}\tilde{x}_j - w_{t+1}l + q_{t+1,n}\,k. \tag{19}$$

Firm value takes the same form as without adjustment costs—the only difference is that it incorporates the continuation value of the firm $q_{t+1,n}\,k$. The continuation value, in particular, captures the effect of future discount rates that respond to future production plans on the present value. Since capital is the only endogenous state variable, only the marginal value of installed capital $q_{t+1,n}$ is needed to record this effect.

We can now rewrite the shareholder value maximization problem at date t as

$$q_{t,n}\, k_t^n = \max_{i,X,\varepsilon} \left\{ -p_{t,n}^k i - c^n \left(i/k_t^n\right) k_t^n - R_{t+1}\left(\varepsilon\right)^{-1} v_{t+1,n} \right\}$$
$$\text{s.t. } v_{t+1,n} = V_{t+1,n}\left(X\right) - \tau_{t+1}^c \sum_{f\in\mathcal{F}} \phi_f x_f, \quad k = \left(1-\delta\right) k_t^n + i, \qquad \varepsilon = \frac{\sum_{f\in\mathcal{F}} \phi_f x_f}{v_{t+1,n}} \tag{20}$$

The new feature here is that the firm pays adjustment cost, in addition to buying new capital. Moreover, the capital stock k in the production plan X for next period is the sum of new investment and undepreciated old capital. We note that we have not assumed a particular dividend policy: it does not matter whether the firm pays out profits and raises new money for investment, or whether the investment is funded with retained earnings.

Problem (20) shows that the equivalence between carbon tax and capital taxes extends to the setting with adjustment costs. Indeed, nothing has changed on the household side, so Euler equations $R_{t+1}\left(\varepsilon\right)^{-1} = R_{t+1}\left(0\right)^{-1}\left(1-\bar{\tau}_{t+1}^k \varepsilon\right)$ continue to hold. Substituting into the problem and using the definition of emission intensity again cancels the value $v_{t+1,n}$. The difference to our baseline case is that the relevant intensity depends not only on the market price of capital (that is, the replacement cost), but the value of installed capital to the firm. At the same time, we also have the standard result that the value of capital equals the market value of the firm.

Intuitively, adjustment costs do not alter the key mechanism we exploit: emissions arise when capital is used together with fossil fuels. A premium can then provide the same incentives as a carbon tax when it reflects the current use of fuel. It is not important that the funds raised might be used to purchase capital that produces future emissions. All effects of future emissions on the value of the firm are reflected in the value of installed capital.

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 2, the key input to this calculation is the distribution of emission intensities relative to capital. We proceed in two steps. In this section, we measure the *current* distribution of emission intensities and hence returns if intensities remained constant—replicating returns without adjustment. This distribution is relevant because firms' choices depend on the entire schedule they face. It is also 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 by dirtier firms *in equilibrium*. This is because introducing a carbon tax will incentivize firms' substitution 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 $\hat{\varepsilon}^i$, a commonly reported number in emissions accounts. The emissions are measured as CO_2e, or CO_2 equivalents, which express the global warming impact of different greenhouse gases relative to carbon dioxide. Across 2 digit NACE sectors in 2017, $\hat{\varepsilon}^i$ ranged between zero and about 12kg CO_2e per dollar, with the highest numbers for water transport, air travel, and electricity. To obtain emissions per unit of capital ε^i, we divide by the capital share in output, and multiply by the user cost of capital ρ_t^i :

$$\varepsilon^i = \hat{\varepsilon}^i \frac{p_{t,i} y_{t,i}}{\rho_t^i k_t^i} \rho_t^i. \tag{21}$$

In the absence of capital taxes, the user cost or rental rate ρ_t^i on capital is the interest rate plus the depreciation rate.

From Proposition 2, to replicate a carbon tax of 10\$ per ton, we need $10\,\varepsilon^i$ extra return on sector i. We have $\hat{\varepsilon}_t^i$ 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^i 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 2023) 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 detail on the energy sector that goes beyond usual input-output matrices. In particular, the electricity sector is broken down by the type of fuel used to make electricity. We also observe physical quantities for transactions of electricity as well as other 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 US as well as between the US and the rest of the world, treated as an aggregate.

We consider a total of 72 sectors listed in Table 4 in Appendix A.3 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 EVs[5] 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.

We back out sectoral emission intensities relative to capital following (21). 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_2 equivalents 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. In order to provide detail for high emission-intensity sectors, we group many low-intensity sectors into a supersector, "Rest," which is shown towards 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 that account for a small share of total capital sit at emission intensities above 100 tn/$mn. 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, in particular chemicals. Together these sectors also account for a large share of total emissions. 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 dollars per ton of $CO2e$, respectively. It follows from (9) that replicating premia are (approximately) linear in intensity. For the serious carbon tax of

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

100$/ton, most capital should command 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 even 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

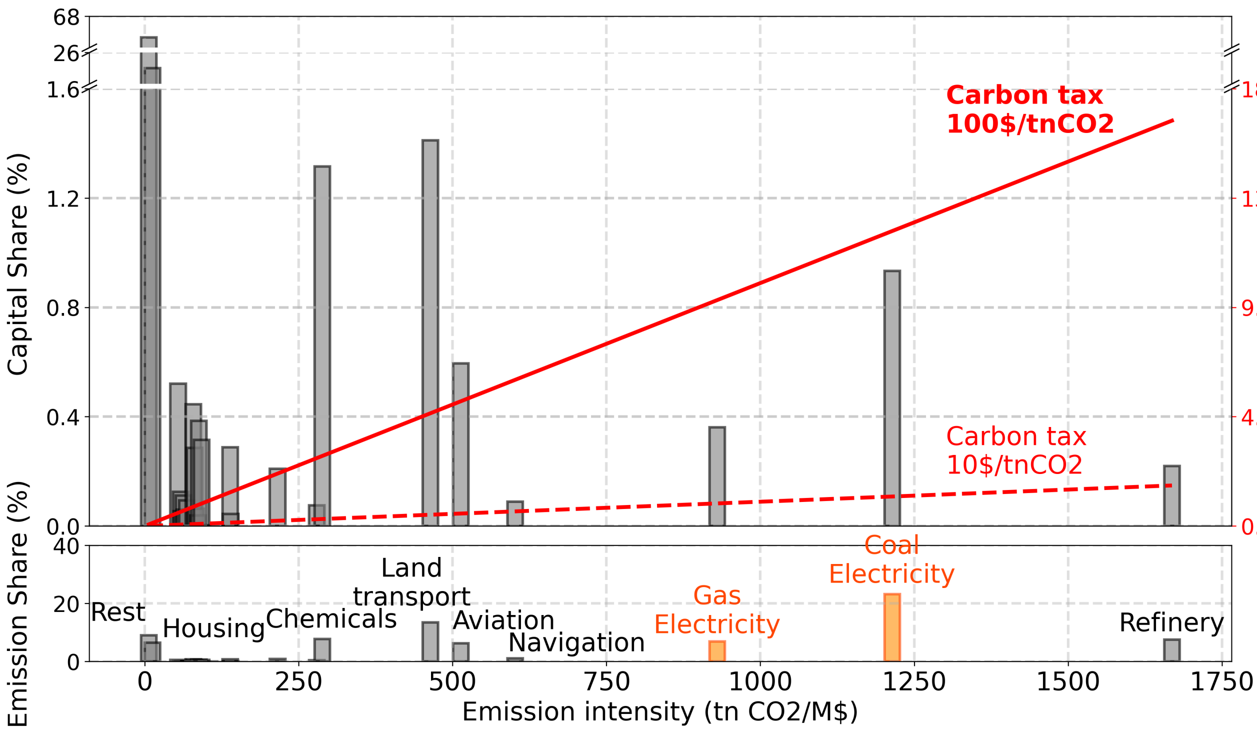

Note: The upper figure shows the capital share, the emission intensity ε^i, and the replicating premium of each sector. Each sector is represented by a bar, except for sectors with little intensity (Rest). The red solid line describes the premium schedule that replicates a carbon tax of \$100, while 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 US 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 US as recorded by GTAP. Their enterprise value accounted for 42% of the capital we back out at the sectoral level.

We are also interested in 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 45% of total electricity capital. We mark by yellow vertical lines the sectoral averages for gas- and coal-fueled electricity from GTAP already recorded in Figure 1. 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 than 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

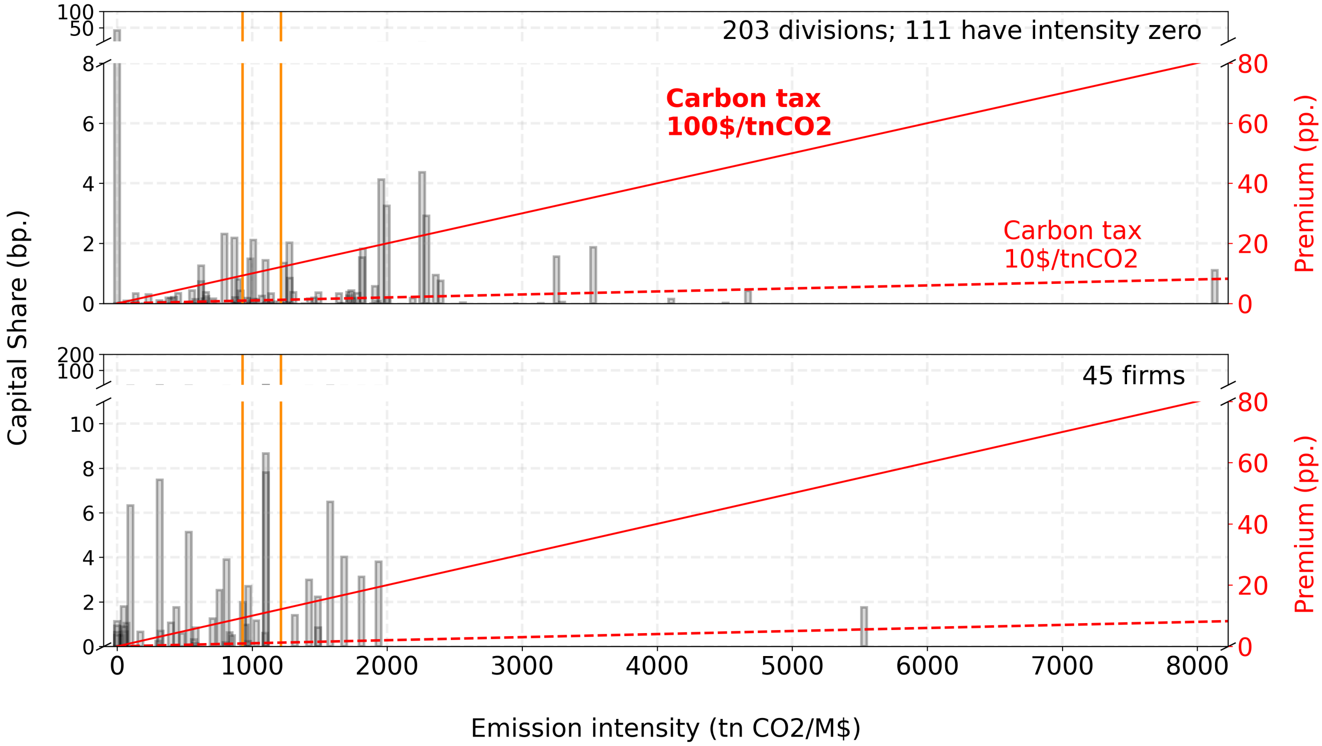

Note: The plot shows the capital share, the emission intensity ε^i, and the replicating premium for each division (upper plot) and each firm (lower plot) within electricity producers. Each division or firm is represented by a bar, except for the bar at zero intensity, which bundles many firms or divisions.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

4 Equilibrium replicating return premia

We now compute equilibria of our model and determine equilibrium replicating returns for a carbon tax of $100 per ton of *CO2e*. This step requires taking a stand on more details of the production function, in particular elasticities of substitution between inputs.

Calibration. 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. We allow for sector-specific total factor productivity (TFP) in the outmost nest for each sector. Our choices for the nesting structure and many elasticities follow the MIT US regional energy policy (USREP) model described in Yuan, Rausch, Caron, Paltsev and Reilly (2019), another multisector model with extra energy detail. We present details on nests and all numbers in Appendix A.1 – here we provide a sketch that highlights key properties that drive the results.

Carbon or capital taxes affect the relative prices of inputs, such as capital and energy goods. In response to such price changes, producers of non-energy goods substitute cheaper inputs for expensive inputs to reduce costs. For a generic non-energy good, this input substitution compares domestically sourced goods with the option to import the goods. For internationally tradable goods, a domestic and an imported variety are substitutes with an Armington elasticity of 2.5. The domestic variety is produced using materials, as well as a *KLE bundle* that combines capital (K), labor (L), and energy (E). The KLE bundle and materials must be combined in fixed proportions. Substitution elasticities are relatively low between a capital-labor bundle and energy, as well as between a dirty fuel bundle and electricity. They are equal to one between capital and labor as well as among alternative dirty fuels.

We follow standard practice in the literature and choose units for all non-energy goods by assuming that their prices in an initial equilibrium are equal to one.[6] We use observable prices only for energy goods, labor, and capital. For capital, we measure user costs as the interest rate plus the depreciation rate, as described above. Given a CES nesting structure, we can then back out nest weights and TFP for each sector from sectoral data on the value of gross output, materials, and labor. In particular, expenditure shares identify the nest weights. Moreover, TFP is pinned down by setting the unit cost of a sector equal to one.

The technology for the electricity sector is different since we allow for many varieties to accommodate our division-level data. For every fuel type, we specify a nest of varieties, each of which is made from materials and a KLE bundle. The number of varieties per fuel type is one plus the number of divisions we observe in TruCost, with one residual variety to reconcile the division-level data with the GTAP sectoral aggregates. We further introduce an aggregate electricity sector that combines electricity generated from different fuels. Here, we incorporate imperfect substitutability due to the intermittency of renewable technologies such as solar and wind.

At the division level, we observe the output of electricity and fuel inputs, but not other inputs such as materials and labor. We describe divisions' technologies using nested CES production functions with the same elasticities of substitution and the same relative weights on capital, materials, and labor, but different weights on fuel versus the other inputs. We calibrate relative weights of capital, labor, and materials to the aggregate data for all plants with the same fuel from GTAP. We then modify the weight on fuel to match emission intensities relative to output at the level of the individual division.

We work with a common *final good*, a Cobb-Douglas bundle of intermediates that represents all capital goods and the export good, that is, the production functions $\tilde{g}$ and h^n are identical. The aggregate good serves as the numeraire in our calibration. The consumption good is a Cobb-Douglas aggregate of the final good, vehicle services consumed by households, and housing services. While our model assumes balanced trade, the US data in 2017 show a trade deficit of about 2.8% of GDP. To reconcile aggregates, we scale down the final good input to consumption

[6]This approach is not restrictive as long as counterfactuals focus on percentage changes in quantities and prices, and not physical units. This is the case for intermediates. For energy, we are interested in emissions and therefore physical units, so we work with data on both prices and quantities.

and increase exports. The model therefore slightly understates consumption. Since our results are about changes across steady states, they are not sensitive to this convention.

Equilibrium allocations We perform comparative statics across steady states that differ only by a carbon tax. In addition to the fixed factor labor, we assume that capital stocks for hydro and nuclear power are fixed at their initial equilibrium values. A version of the model without these restrictions would imply a large shift of electricity production towards hydro and nuclear, which are clean but not intermittent sources of power. However, given the current regulatory and technological environment, those fuel types are unlikely to strongly expand. Since our model otherwise has few adjustment costs and constant returns, the allocation of capital and labor across sectors is likely to be relatively similar along a transition path, as long as the transition is slow and the interest rate does not deviate much from the steady state value.

At the macro level, our model produces responses to a carbon tax that are similar to the existing literature. We present key statistics in Table 1 for carbon taxes of \$10 and \$100 per ton. Our benchmark case of a \$100 per ton carbon tax delivers close to 50% fewer emissions together with a 1.5% output loss. About half of this emission reduction is contributed by the electricity sector. While the overall electricity production rises only modestly, there is a large shift in how it is produced, with firms shifting to much more renewables. When the tax rate is reduced tenfold, the consumption and output loss and the numbers for electricity substitution scale approximately by a factor of close to 10 as well. The emission reduction, however, is still almost 10%, less than a factor of five lower. This result is due to the skewness of emission intensities and emissions themselves: the marginal benefit of affecting the dirtiest sectors is large.

Carbon Tax	Consumption	Output	Capital	Emissions	Electricity	Clean Electricity	Dirty Electricity
10	-0.04%	-0.2%	-0.36%	-8.99%	0.48%	10.28%	-4.59%
100	-0.49%	-1.53%	-2.28%	-45.82%	1.93%	127.58%	-63.19%

Table 1: Aggregate economic outcomes under different carbon tax levels.

The other margins to avoid fossil fuels, and hence reduce emissions, in our model are the reallocation of inputs towards cleaner sectors, electrification, and reduction of emission intensities at the sector level. Outside electricity, the main contributions to emissions reductions come from households (5pp, about equal parts from dwellings and motor vehicles), transport (5pp, with the largest chunk from aviation), and chemicals (2.5%). Another 10pp is accounted for by cumulating small adjustments over the many remaining sectors. Capital reallocation, that is, changes in sectoral shares of overall capital, takes place mostly within electricity, away from gas and coal and towards wind and solar; it is small elsewhere. Changes in revenue shares, in contrast, are larger. Here, limited substitutability implies that many dirty goods, such as transport, that become more expensive with the carbon tax see increased revenue shares. For example, both air and water transport see price increases of about 8% which increases their revenue shares.

We perform a number of robustness exercises to investigate the role of alternative elasticity specifications. First, we vary elasticities related to the use of energy both in electricity and in other sectors. Naturally, more (less) substitution potential entails more (less) emission reductions. Quantitatively, changes of elasticities by a factor of 1.5 or .5 in either direction do not have a strong effect; they change the baseline emission reduction by within 10pp. We take away that the order of magnitude of emission reductions is relatively robust to changes in these parameters. Second, we remove carbon taxes on select sectors. When we remove the tax on households, emission reduction declines by about 5pp. When we remove taxes on all sectors but electricity, we retain an emission reduction of 25pp. In both cases, the numbers are close to the sectors' contribution to the baseline result. We conclude that decompositions of the effects into households, industry, and power sector are close to linear, and not subject to strong interaction effects.

Equilibrium returns. Equilibrium replicating returns are significantly smaller than replicating returns without endogenous adjustment of emission intensities. This is because the carbon tax provides strong incentives for substitution away from dirty fuels. Table 2 presents equilibrium replicating returns for carbon taxes of \$100 per ton and compares them to returns without adjustment. The leftward shift of the distribution is readily apparent. In particular, the basis point reduction in return premia is larger for dirty firms in the upper tail. These firms adjust relatively more: for example, the median premium drops from 5bp to 4bp, whereas the 95th percentile drops from 288bp to 68bp. The bottom two lines in the table display premia for the electricity sector. Here the difference in premia is extremely strong with premia declining in the right tail by more than 10pp.

Premia with and without Equilibrium Adjustment

Quantiles	50th	75th	90th	95th	99th
No Adjustment, Sectors	5	12	43	288	513
Equilibrium, Sectors	4	8	31	68	385
No Adjustment, Electricity Firms	970	1316	1579	1810	1936
Equilibrium, Electricity Firms	23	97	139	231	427

Table 2: Premia in basis points. All premia calculated from a \$100 tax per ton of CO_2.

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 $\tau_c\,\varepsilon$. Differentiating the first-order condition of such a firm, a small increase in the carbon tax leads to a change in its return premium by a factor

$$\frac{d(R_t(\varepsilon) - R_t(0))}{d\tau^c} = \varepsilon \left(1 - \sigma \frac{\tau^c \phi_j}{p_j + \tau^c \phi_j}\right). \tag{22}$$

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: 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.

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

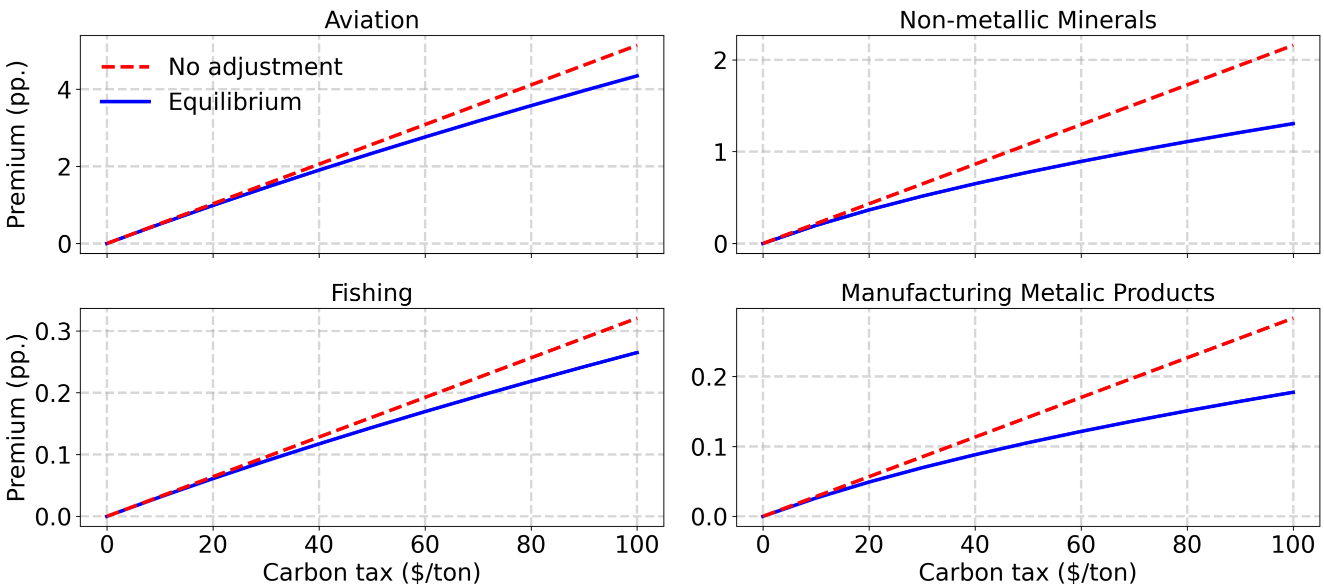

Note: Replicating premia in percent per year for different carbon tax rates (horizontal axis) and selected sectors. Dashed red lines show premia without adjustment, solid blue lines premia at equilibrium intensities

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 (22) 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 (22). 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 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. In particular, coal is a fuel with an emission factor that is high relative to its spot price. In other words, burning coal is a cheap way to produce emissions compared to burning natural gas, and in particular compared to burning oil. As a result, coal-using sectors exhibit a particularly large equilibrium adjustment in returns. To replicate a \$100 per ton 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. At the same time, oil-using sectors such as aviation and fishing in Figure 3 show minimal adjustment, and equilibrium premia resemble those without adjustment.

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

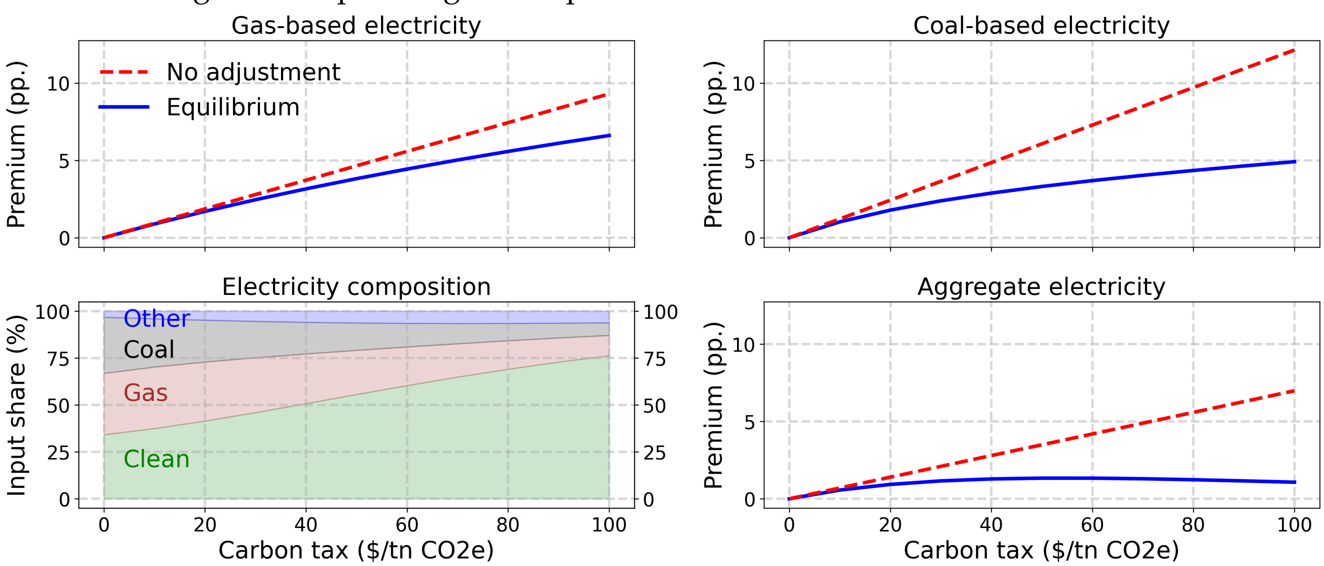

Note: Replicating premia and composition of production in electricity sector. Top line shows replicating premia in percent per year as a function of the carbon tax rate for gas- and coal-powered electricity sectors, bottom right panel for aggregate electricity sector. Dashed red lines show premia without 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.

The bottom left panel of Figure 4 summarizes substitution of the electricity sector across fuel types. As expected in models such as ours, a lot of substitution is possible by reducing dirty power generation and increasing instead clean power, in particular wind and solar. The bottom right panel shows replicating premia for the aggregates 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 carbon tax of about \$50 per ton. Intuitively, when the potential for substitution is high, endogenous cleaning of the electricity sector can generate lower equilibrium premia.

Table 3 shows how different margins of adjustment in the electricity sector contribute to the overall reduction in emissions, or the case of a \$100-per-ton 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 as well as 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 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 3: Sectoral Contributions to Emissions Changes

Electricity Type	Firm capital	Division Share	Intensity	Total	Emission share (baseline)
Coal	4	-64	-6	-66	76
Gas	8	-22	-2	-16	23
Other	0	1	0	1	1
Total	12	-85	-8	-82	100

Note: The table decomposes the -80% reduction in emissions within the electricity sector under the \$100 carbon tax replicating premia. Column **Firm Capital** describes the emission contribution of each variety of electricity if we only allow the firm to reduce the level of capital, but not reallocate resources across divisions or reduce the emission intensity. Column **Division Share** allows further adjustment by letting firms adjust the relative size of their divisions. Last, column **Intensity** shows the extra reduction of letting adjust the emission intensity.

5 Investor taste for clean assets

We now ask how a schedule of capital taxes can come about not because of regulation but because households have a taste for clean investment. We distinguish two types of taste. *Nonpeunciary cost* means that agents experience disutility when their portfolio contains equity in dirty companies. *Nonpecuniary benefit*, in contrast, means that agents experience pride when the portfolio is particularly clean. We further allow agents to focus their attention on particular sectors: for example, households may get utility from holding clean electricity. In this section, equity markets thus distinguish firms not only by intensity but also by sector, and a portfolio is a collection of nonnegative equity (or capital) holdings such that $k_{t,n}(\varepsilon)$ is equity in firms in sector n that produces with intensity ε.

Household sector: a large family of investors. Our main interest is in the effects of cross sectional heterogeneity in taste on emissions. We thus shut down any dynamics of the wealth distribution that could be implied by differences in portfolio performance. To this end, we use the common device of a large family that pools resources. We now describe the optimization problem of the family. It implies, in particular, asset pricing equations and portfolios chosen by family members with different tastes for clean investing. The supply side of the model remains the same as in the economy with capital taxes above, so the definition of equilibrium changes only minimally. For simplicity, in this section we assume that all capital goods have the same price as the final consumption good.

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 capital more when it is used to generate more emissions. Felicity for a green household who consumes c_t^g and holds a portfolio with $s_{t,n}^g(\varepsilon)$ invested in equity of firms in sector n that produce at emission intensity ε equity is

$$u\left(c_t^g - \gamma \sum_n \int (\varepsilon - \varepsilon_n^*) R_{t,n}(\varepsilon) s_{t,n}^g(\varepsilon) d\varepsilon\right), \tag{23}$$

where the scalar $\gamma > 0$ describes the strength of the green investment motive and ε_n^* regulates whether green taste for firms in sector n is positive or negative.

In particular, if $\varepsilon_n^* = 0$ for all n, then any emissions generate nonpecuniary costs, or disutility, for the green investors. We have set up preferences such that this cost is proportional to the payoff from a position, that is, return times investment, and not, say, simply the invested amount $s_{t,n}^g$. In other words, disutility arises if funds available for spending at date t are part of income from date t dirty production. We view this timing convention as plausible and convenient, but it is not essential for the results of this section. What matters is that there is some nonpecuniary effect of investment on utility. Our functional form further makes consumption and nonpecuniary benefits of green investing perfect substitutes, and thereby avoids wealth effects, zeroing in only on heterogeneity by taste.

Positive $\varepsilon_n^* > 0$ in some sector n implies positive nonpecuniary benefit from holding firms $\varepsilon < \varepsilon_n^*$ that are cleaner than a threshold. In principle, the threshold could be the same for all sectors: with equal ε_n^*s, green investors would perceive a benefit from any clean capital, say because they take pride in holding it. An economy-wide threshold then works intuitively like a subsidy on all capital. We do not consider this case in what follows. Instead, we focus on $\varepsilon_n^* > 0$ in only a few sectors n to capture investors' ambition to "clean up" dirty sectors, such as electricity in our quantitative example below. Household members with positive benefit may prefer a relatively clean (but still dirty) firm in their preferred sector n to a clean firm in another sector. We assume that there are always some sectors n such that $\varepsilon_n^* = 0$.

The remaining share of $\kappa^t = 1 - \kappa^g$ households are *traditional* households who obtain felicity from consumption $u\left(c_t^t\right)$. The family pools resources and chooses an aggregate amount s_t. However, individual members must invest shares of capital. Every household is responsible for investing the same per capita amount s_t. The equity constraints for green and traditional member households are

$$\kappa^i \sum_n \int s^i_{t,n}(\varepsilon)\, d\varepsilon = \kappa^i s_t \quad , i = g, t \tag{24}$$

Family members' nonpecuniary cost or benefit in (23) depends only on investments that they themselves are responsible for. We also impose short sale constraints $s^i_{t,n}(\varepsilon) \geq 0$ on all family members and types of capital.

The objective of the family is the equally weighted sum of individual members' utilities. The family must further respect the aggregate budget constraint at date t

$$\kappa^g c^g_t + \kappa^t c^t_t + s_t = \sum_{i=g,t} \kappa^i \sum_n \int R_{t,n}(\varepsilon)\, s^i_{t,n}(\varepsilon)\, d\varepsilon + w_t + T_t. \tag{25}$$

In other words, overall consumption of both types plus total savings must equal payoffs from all portfolios chosen by members in the previous period plus wages and transfers. The family maximizes utility by choosing capital and consumption subject to the sequence of members' equity constraints (24) and budget constraints (25).

Portfolio choice and savings. Let λ_t denote the multiplier on the date t budget constraint (25). Since the objective is equally weighted welfare and it is optimal to treat family members of the same type i symmetrically, λ_t also equals individual households' common marginal utility. Let μ^i_t denote the multiplier on type i's date t equity constraint in (24), for $i = g, t$. As is standard in models with short-sale constraints, the relevant asset-pricing equations are Euler inequalities. In particular, the first-order conditions for equity of a firm in sector n that operates at intensity ε imply

$$\mu^t_t \geq \beta \lambda_{t+1} R_{t+1,n}(\varepsilon) \qquad \mu^g_t \geq \beta \lambda_{t+1} R_{t+1,n}(\varepsilon)\left(1 + \gamma\left(\varepsilon^*_n - \varepsilon\right)\right). \tag{26}$$

The marginal cost of investing in any firm – the shadow value of a unit of equity μ^i_t – must be larger than the marginal benefit. The conditions hold with equality if family member i invests a positive amount in a firm in sector n that operates at intensity ε.

For traditional investors, the benefits from investment are only pecuniary and hence given by the return multiplied by the present value of resources next period $\beta\lambda_{t+1}$. For green investors, there is an additional nonpecuniary benefit or cost. 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 is purely myopic: individual green households need to be compensated for disutility only for one period. This approach allows us to focus on the effect of taste heterogeneity on the cross section of capital allocations.

The family's optimal savings and consumption decision depends on the average return on total household wealth R^w_{t+1} implied by individual households' decisions, that is, the value-weighted average returns on all firms. Let $\bar{\varepsilon}_{t,g}$ denote the average intensity when equity is weighted by green investors' portfolio weights and let $\bar{\varepsilon}^*_{t,g}$ denote the average positive taste, using sector weights in the equilibrium portfolio of green investors. Combining first-order conditions, we then have

$$\lambda_t = \beta\lambda_{t+1}\left(R^w_{t+1} + \gamma\kappa^g\left(\bar{\varepsilon}^*_{t,g} - \bar{\varepsilon}_{t,g}\right)\right). \tag{27}$$

The average effective return (including the average nonpecuniary component) must equal an aggregate MRS $\lambda_t/\beta\lambda_{t+1}$. In steady state, it must be equal to the discount rate β^{-1}. The average pecuniary return R^w_{t+1} is higher (lower) than the discount rate if green investors' nonpecuniary cost (benefit) dominates *for the average firm held by green investors*. Nonpecuniary costs require returns higher than the discount rate to compensate green investors so green taste works like an overall tax on capital. Green investors who obtain nonpecuniary benefits, in contrast, requires less compensation. In this case, taste even works like a capital subsidy when the firm is sufficiently clean.

Green representative agent. When $\kappa^g = 1$, green investors hold all equity so a green asset pricing equation—the second inequality in (26)—must hold with equality for all firms. As a result, all effective returns, including the nonpecuniary component, are equal to the clean rate

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

Returns are higher on dirtier firms. Using an approximation of net returns as in (9), the return is linear in the emission intensity. The role of ε^*_n is to shift down the entire return schedule, including the clean rate, of particular sectors, thus making investments in those sectors more attractive.

A green representative agent who incurs only nonpecuniary cost generates the same return schedule as a carbon tax. Indeed, with $\varepsilon^*_{t,n} = 0$, we obtain the same cross-sectional returns as with a carbon tax $\tau^c_t = \gamma$ or a capital tax with rate schedule $\tau_t(\varepsilon) = \gamma\varepsilon$. In general, the allocation need not be exactly the same as in the case of taxes, since holding dirty capital induces utility costs. However, the steady state will be identical: the steady state clean rate in both cases is equal to the discount rate β^{-1}. We also note that $\gamma\varepsilon$ is the amount of consumption per unit of capital employed at intensity ε that the agent would like to sacrifice to not hold dirty capital.

Heterogeneous agent portfolios. More generally, the investor optimality conditions (26) provide a simple characterization of equilibrium portfolios. First, all firms held by traditional investors must earn the same return, and that return must be larger than the return on any firm held exclusively by green investors. Indeed, the first equation in (26) says that the ratio $\mu^t_t/\beta\lambda_{t+1}$ serves as an upper bound on all returns. Since traditional households hold some capital, the upper bound must bind for all the firms they fund in equilibrium. Second, by a similar argument, all firms held by green investors must earn the same *effective* return. Since the nonpecuniary component of green investors' effective return is decreasing in emission intensity, pecuniary returns among firms held by green investors must then increase in emission intensity.

The two properties together imply that traditional investors only hold firms within a sector that are sufficiently dirty, whereas green investors only hold firms that are sufficiently clean.

The equilibrium return schedule in any given sector is piecewise linear, with slopes of segments smaller at higher emission intensities: the slope is γ until a threshold intensity $\bar{\varepsilon}_n$, say, and zero above the threshold. Traditional (green) investors only hold firms with intensities above (below) the threshold. Intuitively, green investors view dirtier firms as relatively less desirable and hold them only if they are compensated by higher returns. The schedule thus responds to their taste over the range of equity they own in equilibrium. Importantly, this effect is about the relative valuation on the part of green investors, and holds regardless of whether they incur nonpecuniary benefits or costs.

The location of the thresholds $\bar{\varepsilon}_n$ as well as what happens at the threshold itself depends on the relative wealth of green and traditional investors, captured in our setting by the parameter κ^g. Since green investors must hold all capital below the threshold, the threshold is increasing in κ^g. When there are only a few green investors, the segment of the schedule that is responsive to emission intensity is shorter, and a larger share of capital in the economy earns the same return. In particular, the dirtiest firms must earn the same return as most other firms. Technically, another difference to the representative agent case is that the demand for equity is no longer perfectly elastic at the returns provided by the schedule. Instead, the demand for clean or dirty equity now reflects the wealth of the natural clientele for that equity. In equilibrium, the demand for equity meets the supply of equity from the firm side, which is also responsive to the return schedule.

Nonpecuniary benefits versus costs. Whether green investors' preferences reflect nonpecuniary benefits or costs determines the intercept of the return schedule. Consider first the case of costs in all sectors, that is $\varepsilon_n^* = 0$ for all n. There is then an economy-wide clean rate $R_t(0)$ that is earned by the cleanest firms with zero emissions. All other firms earn higher returns. The typical return schedule thus looks like the schedules in the right panel of Figure 5. The schedules increase linearly up to the threshold and remain constant from then on. In particular, dirty firms with intensities above the threshold require higher returns than the clean rate even though they are owned by traditional investors who do not care about the environment. While traditional investors would be happy to hold those firms also at lower returns, high equilibrium returns on the dirtiest firms are necessary for markets to clear. Traditional investors must be enticed to hold the dirtiest firms, so they do not prefer to hold the slightly cleaner firms held by green investors that pay a premium over the clean rate.

Suppose instead that green investors obtain nonpecuniary benefits from investing in a single "preferred" sector n. The cleanest firms in that sector now earn a return that is not only below any other return in sector n, but also below a common return earned in all other sectors. In the preferred sector, we again have a threshold intensity beyond which the premium is constant. A new feature with sector-specific taste is that the threshold depends on the equilibrium size and intensity distribution of the preferred sector relative to the economy as a whole and not only on the wealth distribution. In particular, if there are very few green investors relative to the supply of clean capital in the preferred sector, there may be a jump in the emission intensity at the threshold. Green investors buy the cleanest firms, with intensities up to the threshold, and

Figure 5: Nonpecuniary cost

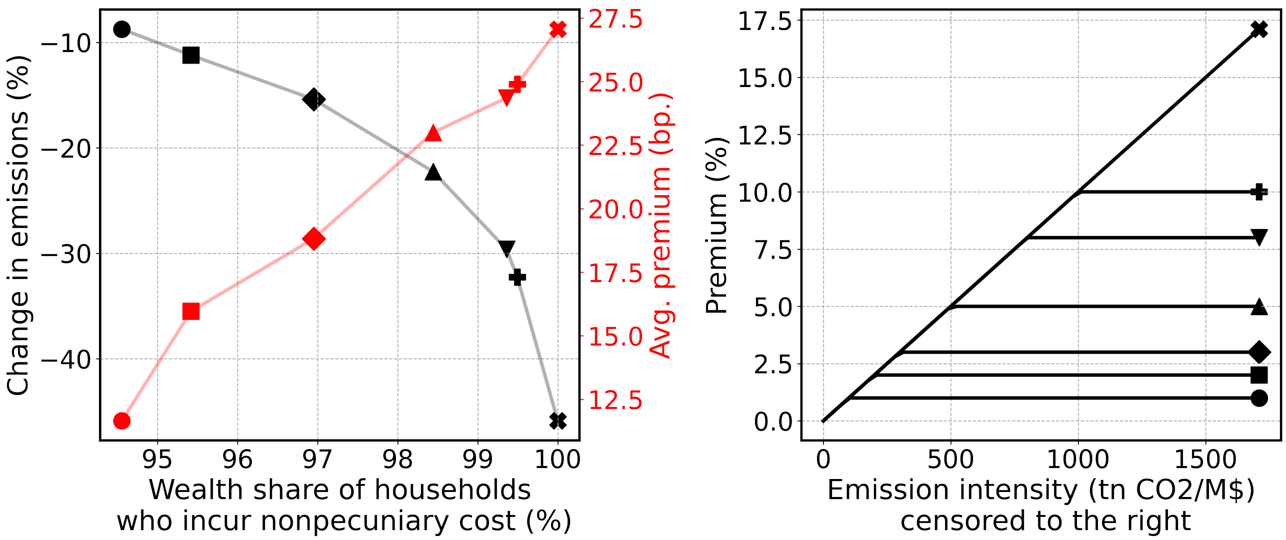

Note: Each dot in the figure represent an equilibrium counterfactual under different shares of green households κ^g, measured in the horizontal axis, holding constant the strength of green investment preference γ. For each equilibrium, we show the reduction of emissions with respect to no premium on the left axis and the average premium on the right axis.

obtain a nonpecuniary benefit which lowers the premium. A slightly more dirty firm beyond the threshold is held by traditional investors and hence pays a return that is discretely larger.

Our goal in this section is to clarify the potential of a small group of investors who obtain nonpecuniary benefits to reduce emissions, as opposed to providing an exhaustive characterization. We thus restrict attention to economies where the schedule of premia is continuous at the threshold. We display the set of schedules we consider for electricity in the right panel of Figure 6. Different schedules start at different intercepts, but all share an increasing piece with the same slope and flatten out once they reach zero. The varying intercept reflects the nonpecuniary benefit investors earn from holding the cleanest electricity firms. The key difference to the nonpecuniary cost case is that its intercept is below zero. Moreover, when green investors prefer a sector that is both relatively small and relatively dirty, we can have relatively dirty firms in that sector pay returns that respond to emission intensity.

Figure 6: Nonpecuniary benefit

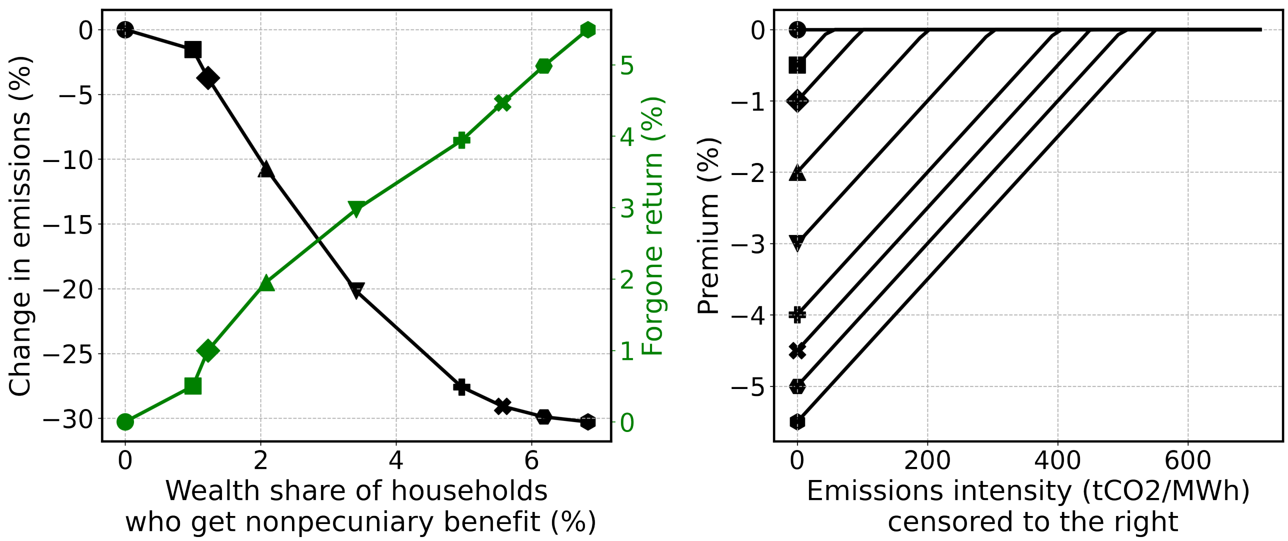

Note: Each dot in the figure represent an equilibrium counterfactual under different shares of green households κ^g, measured in the horizontal axis, holding constant the strength of green investment preference γ. For each equilibrium, we show the reduction of emissions with respect to no premium in the left axis and the average premium on the right axis.

Firm behavior. The firm problem in the economy with green tastes is the same as before. In particular, firms face return schedules that depend on the emission intensity they operate at and take those schedules into account when they choose their production plans and hence the markets ε in which to raise funds. 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 one-technology firms only: 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 the family of households as well as firms optimize and markets clear. The return schedule is determined jointly with the distribution of capital. From the household side, we have the exogenous distribution of equity by types, which says 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 as long as the Euler equations (26) hold. On the firm side, a return schedule gives rise to a collection of production plans and hence a distribution of equity by intensity. In equilibrium, the distributions and returns are consistent.

Quantifying the impact of taste heterogeneity. Our goal is to assess how small amounts of heterogeneity matter for the distribution of returns as well as for 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 carbon tax of \$100 per ton. We study two sets of counterfactuals that explore the impact of negative and positive taste, respectively. The first assumes that all intercepts are zero, or $\varepsilon_n^* = 0$. It thus explores robustness of a green representative agent who incurs nonpecuniary cost. 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 cost. Figure 5 summarizes how emission reduction depends on the share of green investor households. Along the horizontal axis, we measure the share κ^g, focusing on the range above 95%; that is, an overwhelming majority of households have a strong green investment motive. The far right point is $\kappa^g = 1$, a green representative agent who implements a carbon tax of \$100 per ton. The emission reduction, measured along the left vertical axis, is then the 45% reduction implied by a \$100 per ton carbon tax. We show a number of other economies for different values of κ^g, with emission reductions connected by the solid black line. The corresponding schedules of premia are marked in the right panel with the same symbol. The red symbols further show the average premium, in basis points, measured along the right vertical axis. The main result in the figure is that **emissions are extremely sensitive to a small drop in the share of green investors below 100%.** Even a one percent lower share lowers the emission reduction by 15pp or about one third. A five percent lower share means the emission reduction is 30pp or two thirds smaller.

The result is driven by the extreme skewness of the distributions of both emissions and emission intensities. A carbon tax creates large emission reductions because it creates incentives for very dirty firms to lower their intensities. Replicating return premia do the same thing. It is important, however that the returns increase steeply with intensity also in the right tail of the intensity distribution. 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. We thus lose the provision of incentives especially for the dirtiest firms. Since emissions are also skewed, there is a large cost in terms of total emissions.[7]

Nonpecuniary benefits from electricity. Figure 6 considers economies with a small share of green investors, again measured along the horizontal axis. Here the benchmark is the leftmost point with $\kappa^g = 0$, it corresponds to a representative agent in the absence of any carbon tax. In contrast to the previous exercise, green investors now obtain a nonpecuniary benefit from holding electricity firms. We thus need to choose two preference parameters, the slope γ and the intercept ε_n^*. We study a set of economies with continuous return schedules: we vary κ^g and ε^n together, so the wealth of green investors is just large enough to purchase all electricity capital below the threshold. This is how we construct the schedules in the right panel of Figure 6. The left panel shows as green symbols the foregone return per unit of wealth for the green investors, measured in percent along the right vertical axis.

The main takeaway from the economies with nonpecuniary benefit is that **a small wealth share of green investors is enough to generate large emission reductions if it focuses on a dirty sector with a lot of substitution potential**, here electricity. Indeed, with a two-percent wealth share, emissions are already more than 10% lower than at baseline. With 6% green investors, emissions fall by more than one-half of what a $100-per-ton carbon tax could achieve. The key effect is that electricity firms have strong incentives to expand clean production: the premium 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 cost case above. Nevertheless, losses are on the order of a few percentage points.

[7] Put differently, we can think of the return premium like a carbon tax rate times an intensity. In the flat part of the schedule, the intensity keeps increasing, so firms behave *as if* they face a declining carbon tax rate.

References

Aguiar, Angel, Maksym Chepeliev, Erwin Corong, and Dominique van der Mensbrugghe (2022) "The global trade analysis project (GTAP) data base: Version 11," *Journal of Global Economic Analysis*, 7 (2).

——— (2023) "The Global Trade Analysis Project (GTAP) Data Base: Version 11," *Journal of Global Economic Analysis*, 7 (2), 10.21642/JGEA.070201AF.

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.

Berk, Jonathan and Jules H van Binsbergen (2021) "The Impact of Impact Investing," *Available at SSRN 3909166*.

Bonnefon, Jean-François, Augustin Landier, Parinitha Sastry, and David Thesmar (2025) "The moral preferences of investors: Experimental evidence," *Journal of Financial Economics*, 163.

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, 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.

Gormsen, Niels Joachim, Kilian Huber, and Sangmin Oh (2024) "Climate Capitalists," *Available at SSRN 4366445*.

Goulder, Lawrence and Marc Hafstead (2017) "Confronting the Climate Challenge U.S. Policy Options," *Columbia University Press*.

Heeb, F., J. F. K¨olbel, 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.

Nordhaus, William (2013) "Integrated Economic and Climate Modeling," *Handbook of Computable General Equilibrium Modeling, Elsevier*, 1, 1069–1131.

Papoutsi, 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).

Pastor, Lubos, Robert F. Stambaugh, and Lucian A. Taylor (2024) "Sustainable investing," Working paper, Chicago Booth.

Pedersen, Lasse (2024) "Carbon Pricing versus Green Finance," 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.

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.

A Appendix

A.1 Input-Output structure

Here we detail the input-output structure of the quantitative model used in Section 4. Much of the production structure follows the multisector model developed by Yuan et al. (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 4 lists all sectors in our model.

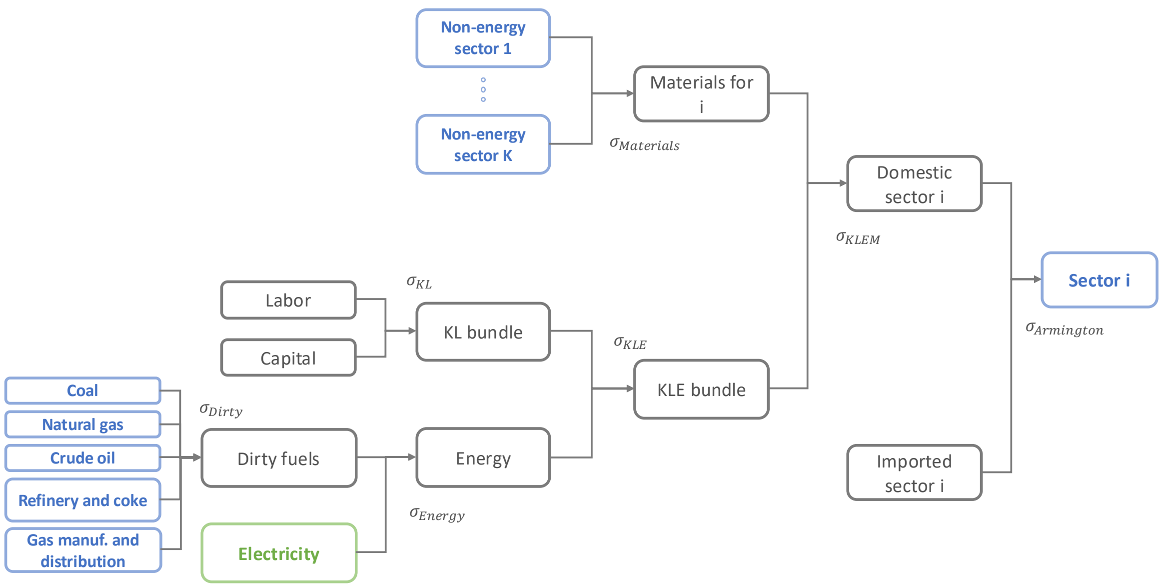

Figure 7: Input-output structure of a generic sector

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 ocrner 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 5 below.

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

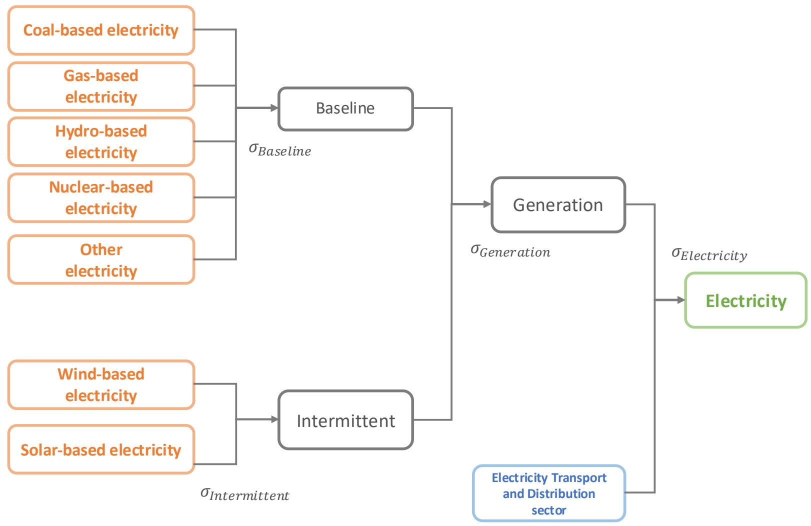

Figure 8: Input-output structure of electricity

of sector's goods. 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 other[8]. 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 A.4 below, we include an extra production layer in each electric variety to account for firm level 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.

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

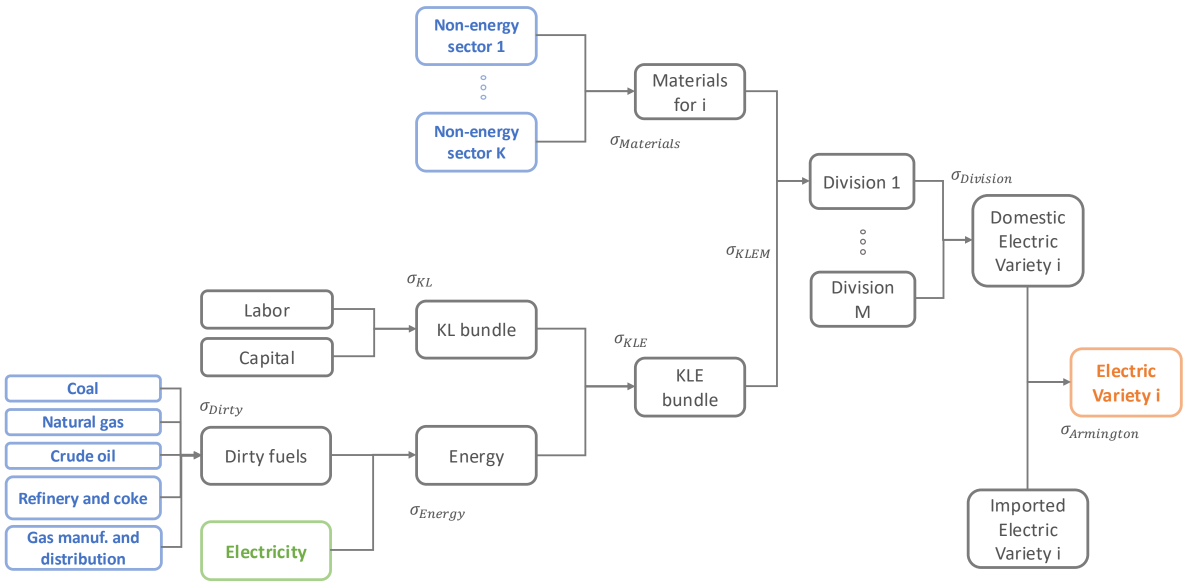

Figure 9: Input-output structure of a generic electric variety

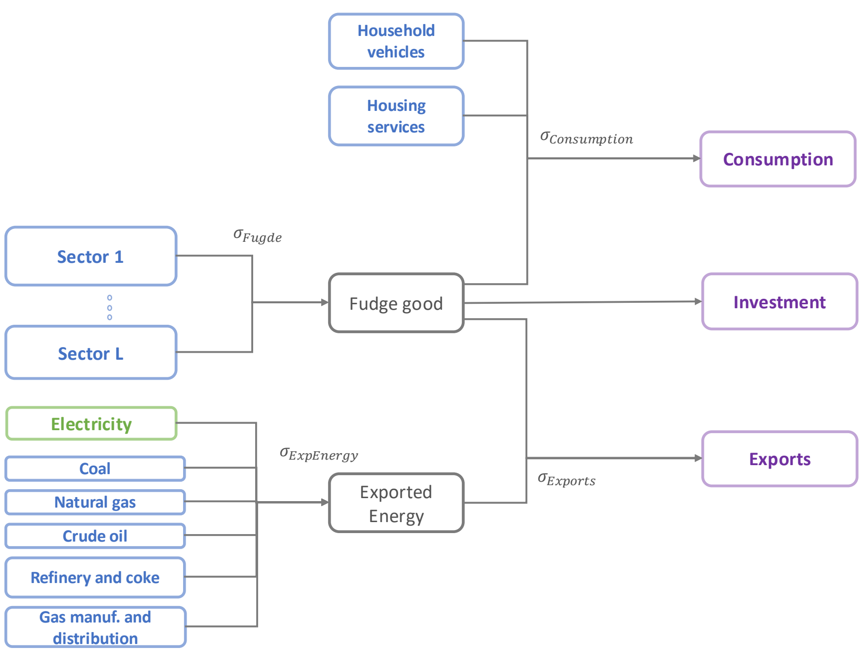

Figure 10: Input-output structure of final demand

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 4: Sectors and Aggregators

Sectors					Aggregators
Primary	**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 manufactur	Dairy products	Accommodation	Hydro-BL	
Oil seeds	Refinery	Processed rice	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	Dwellings		
		Other Metals			
		Metal products			
		Computer, electronic and optics			
		Electrical equipment			
		Other Machinery and equipment			
		Motor vehicles and parts for firms			
		Motor vehicles and parts for households			
		Other Transport equipment			
		Other Manufactures			

A.2 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.

A.3 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 et al., 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 are 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 CO_2.

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 CO_2 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 quantiative 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 quantiative 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.

A.4 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 5: Elasticities of substitution

Parameter	Comment	Value
Generic sector (Figure 7)		
$\sigma_{Armington}$	except energy-related goods (0.3)	2.5
σ_{KLEM}		0
σ_{KLE}		0.4
$\sigma_{Materials}$		0
σ_{KL}		1
σ_{Energy}		0.5
σ_{Dirty}		1
Electricity sector (Figure 8)		
$\sigma_{Baseline}$		10
$\sigma_{Intermittent}$		10
$\sigma_{Generation}$		10
$\sigma_{Electricity}$		10
$\sigma_{Divisions}$		10
Electric variety (Figure 9)		
$\sigma_{Armington}$		0.3
$\sigma_{Division}$		6
σ_{KLEM}		0
σ_{KLE}	exception NuclearBL and HydroBL (0)	0.4
$\sigma_{Materials}$		0
σ_{KL}	exception NuclearBL and HydroBL (0)	1
σ_{Energy}		0.5
σ_{Dirty}		1
Final demand (Figure 10)		
$\sigma_{Consumption}$		1
σ_{Fudge}		1
$\sigma_{ExpEnergy}$		0
$\sigma_{Exports}$		0