Mirrleesian Carbon Taxation*

Andrea Chiavari Alexandre Kohlhas

November 2025

Abstract

Alleviating the economic damages from climate change is, to first order, a problem of efficiently limiting firms' emissions. We analyze a neoclassical general-equilibrium model in which fossil-energy use generates climate damages. The model crucially incorporates substantial cross-firm heterogeneity in emission efficiency that we document using a novel firm-level dataset spanning 150 countries. Firms choose energy and other inputs and self-report emissions that are otherwise privately observed. Our central result is a simple formula for the marginal externality damage of emissions—the optimal carbon tax—that accounts for firms' incentives to distort reported emissions. The optimal tax varies markedly across firms as a function of marginal damages, output, and emission efficiencies, and exceeds common uniform-tax benchmarks, on average. Quantifying the model shows large welfare gains from internalizing reporting incentives: the constrained optimal tax recovers 3/4 of the potential welfare benefits; a uniform tax that does not internalize reporting incentives, by contrast, recovers almost none of them.

JEL codes: E13, E17, E6, H23, Q28, Q4, Q5.

Keywords: Climate change, optimal policy, Mirrleesian optimal taxes.

*First draft: November 2025. Chiavari: University of Oxford (email: andrea.chiavari@economics.ox.ac.uk). Kohlhas: University of Oxford (email: alexandre.kohlhas@economics.ox.ac.uk). We thank John Hassler, Marek Kapička, Rick van der Ploeg, Per Krusell, and seminar and conference participants at various institutions for their feedback. Andrea Chiavari gratefully acknowledges financial support from Christ Church Research Center, and Alexandre Kohlhas from Handelsbank Stiftelsen.

1 Introduction

A defining challenge of our time is that of climate change. Given the central role that firms' carbon emissions play in its development, it is important to a theory of carbon taxation that is effective. There is reason to believe that such a theory needs to be richer than the benchmark model of Piguvian taxation. Building on the work of Pigou (1912, 1920), the dominant model of carbon taxation over the past three decades has been that of a tax on firms' unpriced carbon emissions, to ensure that firms internalize the true social cost of their actions. This approach has, in turn, led to several real-world proposal (e.g., Nordhaus, 1977; Economists' Statement on Climate Change, 1997; Economists' Statement on Carbon Dividends, 2019). The practical implementation of carbon taxation, however, faces several challenges.

An important obstacle—long recognized by by policymakers (e.g., European Commission, 2023) and practitioners (e.g., ECA, 2015 alike—is that firms' emissions are primarily self-reported and based on firm-specific information. Combined, with the presence of a carbon tax, this creates an incentive for firms' to distort their carbon disclosures, hindering the effectiveness of carbon taxation.[1] In part as a consequence of these concerns and the high-profile court cases that have followed, trust in the effectiveness of carbon taxation is low among the general public (e.g., Dechezleprêtre *et al.*, 2025)—preventing its widespread adoption.

Our aim with this paper is to fill the gap between the macroeconomic theory that address firms' climate externality through carbon taxation and its practical implementation plagued by the presence of imperfect information about firms' emissions. To do so, we propose a modification of the standard Pigouvian approach that recognizes the need to incentivize firms' truthful disclosures about their carbon emissions. Within the context of a workhorse macroeconomic-climate framework, we show that a simple adjustment to a benchmark uniform tax accounts for firms' incentives to distort their carbon disclosures. This constrained optimal tax varies markedly across firms as a function of emission efficiencies, and exceeds common uniform-tax benchmarks, on average. Quantifying our tax formula shows substantial welfare gains from internalizing firmâs reporting incentives. In particular, we show that the constrained optimal tax recovers three-quarters of the potential welfare benefits from carbon taxation; a uniform tax that does not internalize reporting incentives, by contrast, recovers little.

To motivate the assumptions of our subsequent model, we perform an empirical analysis based on the ICE ESG Company Data, which contains reported CO_2e emissions for a comprehensive set of public and private firms from 150 countries over 15 years. We complement this data set with Compustat Fundamentals Annual to retrieve firm-level outcomes and we validate the ICE ESG Company Data using S&P Truecost—a more commonly used source of

[1] See discussion in Section 2.

reported CO_2e emissions that covers only publicly traded firms.

To justify our focus on curbing firm emissions as a climate policy lever, we begin by showing that firms account for the lion's share of global emissions—at least 50%—a point already stressed by . We then examine emission intensities (tCO_2e/\$m revenue) and confirm the findings in Shapiro and Walker (2018), Lyubich *et al.* (2018), and Capelle *et al.* (2023): emission intensities are widely dispersed, even after controlling for country-sector-time fixed effects, and this dispersion is large relative to other familiar sources of heterogeneity such as TFP. Finally, by comparing firms' reported emissions that are third-party verified with those that are not, we find that non-verified firms systematically report lower emissions even when compared to very similar verified firms. We interpret this as supporting practitioners' and policymakers' view that the vast heterogeneity in firms' emission intensities—which is private information—leaves ample scope for misreporting.

Motivated by the evidence above, we extend a standard dynamic integrated-assessment model in the spirit of Golosov *et al.* (2014) and Barrage and Nordhaus (2024) along two dimensions: we introduce firm-level heterogeneity in productivities and carbon intensities, and we treat carbon intensities, and thus emissions, as private information of the firms. In particular, the model features a representative household that solves a standard consumption-saving problem and supplies capital and labor to firms, from which it receives profits. Firms, which differ in productivity and carbon intensity, produce using capital, labor, and energy—with energy produced indirectly from output in a roundabout fashion—, and invest in abatement to curb emissions. While the baseline model allows for abatement but only a single energy type, as in Barrage and Nordhaus (2024), our extensions explore the implications of permitting substitution between green and brown energy, as in Golosov *et al.* (2014).

Emissions affect aggregate productivity through a damage function, which individual agents do not internalize, representing the environmental externality. In this respect, we follow Golosov *et al.* (2014) and approximate the feedback between the economy and the climate with a low-dimensional process in emissions only. In a series of extensions we consider a more complex model with the joint evolution of emissions and temperature, where temperature affects aggregate productivity as in Miftakhova *et al.* (2020)—an approach that has gained traction in the natural sciences (e.g., Allen *et al.*, 2009; Matthews *et al.*, 2009, 2012) and in macroeconomics (e.g., Dietz *et al.*, 2021; Fernández-Villaverde *et al.*, 2024).

Our first set of theoretical results show how our new elements challenge, motivated by the data, lead to conclusions that challenge the conventional wisdom. In particular, we show that the standard homogenues carbon tax implied by a planner that was to have full information introduces profit differential between equally productive firms with different carbon intensities, as this tax scheme leads to a larger bill the higher your carbon intensity and thus emissions.

In a world where these carbon intensities are private information of the firm this implies that the standard carbon tax is incentive incompatible, generating strong incentives for firms to misreport and pretend to have lower carbon intensities than they really have. We show that the cost of neglecting firms incentive to misreport can be arbitrary large, depending on the hteregeneity in carbon intensities across firms, and thus the cleanest type firms misreporting can try to impersonate.

Motivated by this failure of the standard approach, we propose a Mirrleesian approach to carbon taxation, in which the planner constrains its tax to deliver allocations that are ex post incentive-compatible. Our environment differs from the standard household taxation problem in Mirrlees (1976). Because emissions do not matter per se for firms, in any incentive-incompatible allocation each firm has the same dominant strategy: to misreport its type to impersonate those associated with the highest profit. Consequently, the planner—conditional on observable productivity—must equalize profits across firms with different carbon intensities, producing bunching along that dimension in equilibrium.

Since the incentive is unidirectional—high-carbon-intensity firms want to impersonate low-carbon-intensity ones—, the optimal carbon-tax schedule declines with reported carbon intensity to remove this incentive. Crucially, the scheme is incentive-compatible because, conditional on productivity, all firms pay the same total carbon tax as a share of revenue: low-intensity firms face low emissions and higher tax rates, while high-intensity firms face lower tax rates but higher emissions. We show that this heterogeneous tax does not require more information to be implemented than the standard homogeneous carbon tax. In fact, our incentive-compatible heterogeneous tax formula is the homogeneous carbon tax in Golosov *et al.* (2014) multiplied by a Mirrleesian adjustment—a function of firms' carbon intensities that induces truthful revelation under the scheme.

Lastly, we characterize the fundamental trade-off any planner faces when setting carbon taxes. On the one hand, planners want to curb emissions to raise aggregate productivity. On the other hand, efforts to reduce emissions require distorting input use in ways that generate input misallocation and lower aggregate productivity. We call this the efficiency-misallocation trade-off. Fundamentally, the optimal carbon tax is the one that optimally balances this trade-off while, in the Mirrleesian case also ensuring incentive compatibility.

Having established these theoretical results, we explore the model quantitatively to assess how far the Mirrleesian carbon tax goes in achieving first-best welfare gains and how costly it is to ignore incentive-compatibility considerations. We calibrate the novel heterogeneity we introduced using the data presented in the empirical analysis. For the first best we obtain a homogeneous carbon tax of \$300/tCO_2e, in line with the existing literature (e.g., Tol, 2023), with the total carbon bill ranging from almost 0% of revenues for low carbon intensity firms

to almost 7% for high ones. By contrast, our Mirrleesian carbon-tax schedule implies a tax of about \$1,500/tCO_2e for the median carbon-intensity firm, rising to roughly \$40,000/tCO_2e for firms two standard deviation below the median, and declining to about \$60/tCO_2e for firms two standard deviation above the median. Despite these rate differences, the total tax bill is the same across firms with similar productivity and amounts to less than 1% of revenues. Looking at the welfare implications, we find that ignoring incentive-compatibility considerations achieves less than 10% of the welfare gains attainable in the first best, whereas the Mirrleesian carbon tax captures nearly 80% of those gains.

Decomposing welfare improvements into misallocation, higher aggregate productivity from lower emissions, and changes in capital and energy use, we find that the effects of energy-use changes and the misallocation costs of carbon taxation are relatively small. This is because only energy use is distorted by taxes and energy comprises a small share of total output. Changes in capital make a positive contribution, but most welfare gains stem from higher aggregate productivity resulting from emissions reductions. Looking at the aggregate decline in emissions, we find that most of the reduction comes from abatement by high-emission firms. This decomposition also explains when the Mirrleesian planner performs well relative to the first best: performance is strong when the first-best covariance between energy use and carbon intensity is low—i.e., when productivity is an important determinant of energy choice. This result highlights the importance of jointly accounting for productivity and carbon-intensity heterogeneity in the analysis.

We conclude our quantitative analysis by showing that a range of robustness exercises and extensions common in the literature do not alter our conclusions. First, we examine transition dynamics and a more complex climate specification in which the climate is approximated by both emissions and temperature, following Miftakhova *et al.* (2020)—an approach increasingly used in the natural sciences (e.g., Allen *et al.*, 2009; Matthews *et al.*, 2009, 2012) and in macroeconomics (e.g., Dietz *et al.*, 2021; Fernández-Villaverde *et al.*, 2024). Second, we a sensitivity analysis of climate-parameter uncertainty, focusing in particular on the abatement-cost elasticity and the climate-damage elasticity. Third, we introduce renewable and fossil energy as imperfect substitutes, following Golosov *et al.* (2014). Fourth, we vary the precision of the planner's signal about firms' carbon intensities. Overall, our main message is robust across these permutations: ignoring incentive-compatibility considerations is consistently very costly, and a Mirrleesian carbon tax under realistic parameterizations goes a long way toward capturing the first-best gains.

Related Literature. In addition to the works cited above, this paper relates to three strands of research. We review these in order of proximity below.

First, our paper builds on the work that followed Pigou's seminal contributions to the *optimal taxation of externalities* (Pigou, 1912, 1920). The general theory behind Piguvian taxation is extended in Kang (2020) and Pai and Strack (2022). Important applications to climate change include Weitzman (1974), Nordhaus (1977, 1993), Acemoglu *et al.* (2012, 2016), Golosov *et al.* (2014), and Barrage and Nordhaus (2024), among others. We complement this work by showing that, when emission-intensity types are highly heterogeneous and partially private information to firms, the standard Pigouvian approach results in incentive-incompatible allocations with distorted disclosures. The optimal tax formula that we derive documents how to optimally adjust the standard Piguovian carbon tax for such concerns.

Second, our work rests on important insights from the literature on *optimal taxation under private information*. The literature on household income taxation in the presence of private information was pioneered by Mirrlees (1971) and Diamond (1998), among others, and extended to dynamic contexts in Golosov *et al.* (2006), Kocherlakota (2010), andStantcheva (2020). The literature on the optimal regulation of firms with private information is reviewed in Tirole (1988) and Laffont (1994), and found a recent dynamic application to firm R&D investments in Akcigit *et al.* (2022). Our contribution within this context is to show that integrating the insights from this literature can restore the effectiveness of the standard Pigouvian approach to carbon taxation: a simple adjustment—based on firms' reported carbon intensities, which under our scheme firms report truthfully—recovers most of the benefits of the textbook Pigouvian approach while ensuring compliance by firms.

Finally, our work is related to that which followed Allingham and Sandmo (1972) contribution on *tax evasion* and *misreporting incentives*. Recent applications to business taxation include Di Nola *et al.* (2021), Fernandez-Bastidas (2023), and Bhandari *et al.* (2024). We contribute to this literature by showing that its insights about the complex interaction between disclosure incentives and taxation extend to Pigouvian taxes when agent-types are partially private information. But, unlike this work, we do not address this challenge through improvements in auditing technology; instead, we propose a modification of the otherwise standard Pigouvian approach that recognize its impact on firms' disclosure incentives. In this sense, our approach presents a natural evolution of the Pigouvian approach to carbon taxation—one that internalizes firms' disclosure incentives.

2 Motivating Evidence

We present new evidence on the importance of firms' emissions in aggregate emissions, the heterogeneity in firms' emissions intensities, and their private-information nature. To start, we use micro data from *Intercontinental Exchange's (ICE's)* ESG company database. The

ICE-ESG database provides the most comprehensive dataset on private and public firms' emissions, spans the period 2009-2024 for 46,928 firms across 151 countries, and reports firm-year estimates of carbon emissions broken into three components: scope 1 (direct emissions from production), scope 2 (indirect emissions from energy use), and scope 3 (all other indirect emissions).[2] We focus on scope 1 emissions, as these are the model-consistent measure of in-house production emissions that firms can directly control and avoid double-counting when aggregating. We validate our results using publicly listed firms' emissions data from *S&P Trucost*. Appendix A.1 provides more information on sample construction and our definition of emissions.

We begin by illustrating the importance of firms' carbon emissions for overall emissions into the atmosphere. Figure 1 shows the share of overall emissions composed of carbon emissions from firms in the ICE-ESG sample. We plot these estimates separately for the US, Europe, China and the World and measure them in terms of tons of CO_2 equivalent emissions.[3] Data on country- and region-wide emissions are taken from the *Global Carbon Budget report* (2024 estimates). For the US, the country for which ICE has the best coverage, firms' direct emissions comprise more than 80% of overall carbon emissions into the atmosphere.[4] For Europe, China and the World, the ratio is somewhat smaller—between 40-60% of overall emissions—partly reflecting lower firm coverage by ICE outside the US. We conclude that firms' direct emissions comprise the lion's share of carbon emissions into the atmosphere.[5] Addressing climate change is, to a large extent, a problem of addressing firms' direct, production-based emissions.

Direct emissions are the by-product of firms' creation of goods and services for the economy. The size of overall carbon emissions attributable to firms, visible in Figure 1, thus raises the question of whether—and to what extent—firms differ in the efficiency of their carbon use. Figure 2 addresses this question by plotting histograms of *emission intensities* in the ICE-ESG and TruCost samples. We measure emission intensity at the firm level by tons of CO_2 equivalent emissions per million $ of revenue. Panel (a) shows substantial heterogeneity in emission intensities in the raw data.[6] The larger heterogeneity visible in the ICE-ESG sample reflects in part the presence of non-publicly-listed firms.

Clearly, part of this heterogeneity is due to differences in output and the production processes used in different sectors and countries at different point in time. To address this issue,

[2] When referring to *carbon emissions*, we, in all cases, mean *carbon equivalent emissions*. Figure A.2 further reports the average share of scope 1, 2 and 3 emissions for a firm in the ICE-ESG database.

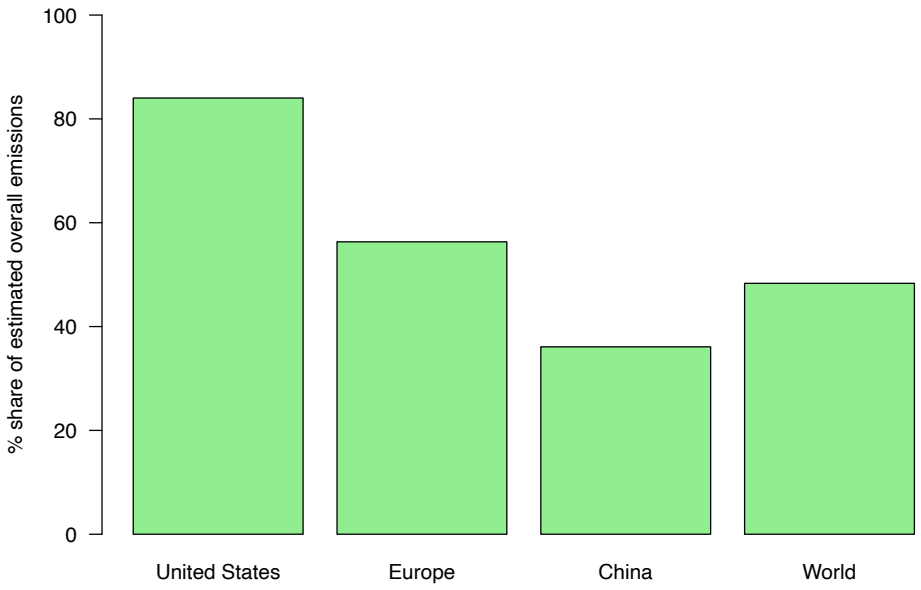
[3] For reference, total US, European, and Chinese firms' ICE emissions comprise around 60% of estimated World carbon emissions into the atmosphere (see also the Global Carbon Budget, 2024 report).

[4] The equivalent share for firms in the TruCost sample is around 25% (Figure A.1 in the Appendix), as this data set comprises a substantially smaller share of firms.

[5] This is in line with several technical reports on the topic (e.g., Report, 2017; IPCC, 2022).

[6] Shapiro and Walker (2018), Lyubich *et al.* (2018), and Capelle *et al.* (2023), among others, document similar heterogeneity using the EPA's *National Emissions Inventory* and *TruCost*, respectively.

Figure 1: Estimated Share of Overall Emissions



Note: Data from the ICE-ESG sample. The figure shows overall scope 1 firm emissions in 2024 (measured by tons of CO_2 equivalent emissions) divided by overall country/region-wide emissions. Estimates for country/region-wide emissions are taken from the *Global Carbon Budget* website (2024 estimates).

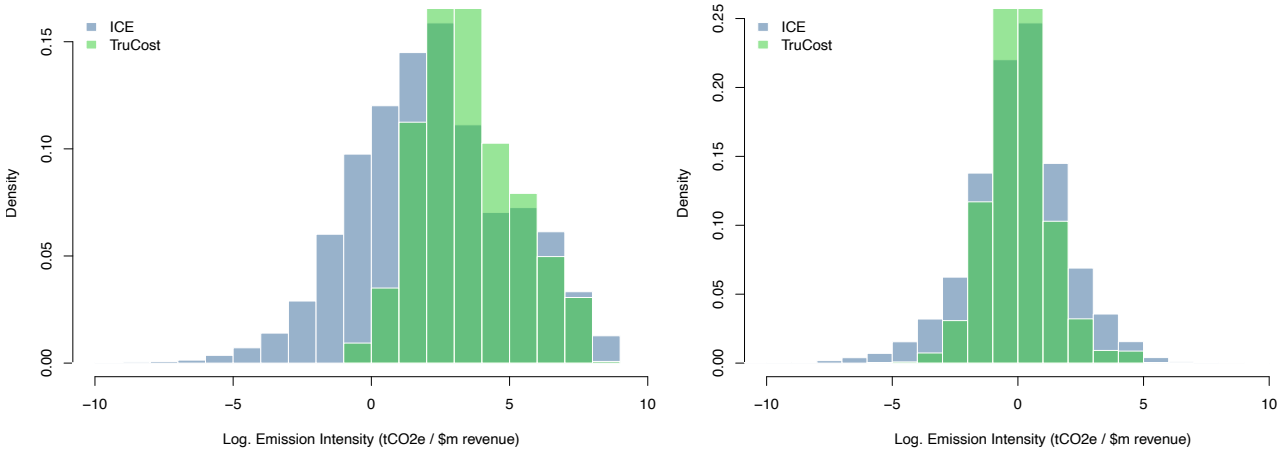
Panel (b) in Figure 2 residualizes emission intensities for sector, country, and time fixed effects. Although a substantial share of the variation can be attributed to these factors—the standard deviation of the distribution decreases by 30%—the bulk of the heterogeneity remains.[7] Most of the differences in the emission intensity by which firms produce their revenue are attributable to factors associated with a firm within a given sector, country, and time period. Figure 3 gauges the overall size of this heterogeneity by comparing it to the heterogeneity that exists in firm-level TFP in the US. All else equal, the figure shows a substantially wider distribution for emission intensities than TFP: the standard deviation of emission intensities is almost 3 times larger than that for TFP. Obviously, this comparison is not informative of the size of the aggregate consequences of this heterogeneity; but what it does highlight is that there are indeed large-scale differences in the emission efficiencies of firms when compared to other, standard measures of firm heterogeneity.

We show that large-scale heterogeneity in emission intensity is also present in sectors with comparatively homogeneous output, further alleviating concerns that output differences drive the observed heterogeneity.[8] We perform the same analysis for the US energy, airline and

[7] Figures A.4 and A.5 document heterogeneity in emission intensities across sectors and countries. Consistent with earlier work, the most carbon intensive sectors are energy production, agriculture, and resource extraction. Emerging economies are also, all else equal, more carbon intensive than developed economies.

[8] Petersen (2025) finds large heterogeneity in emission intensities within narrowly defined industries in the

Figure 2: Heterogeneity in Emission Intensities



Panel (a): Overall Sample

Panel (b): Residualized Sample

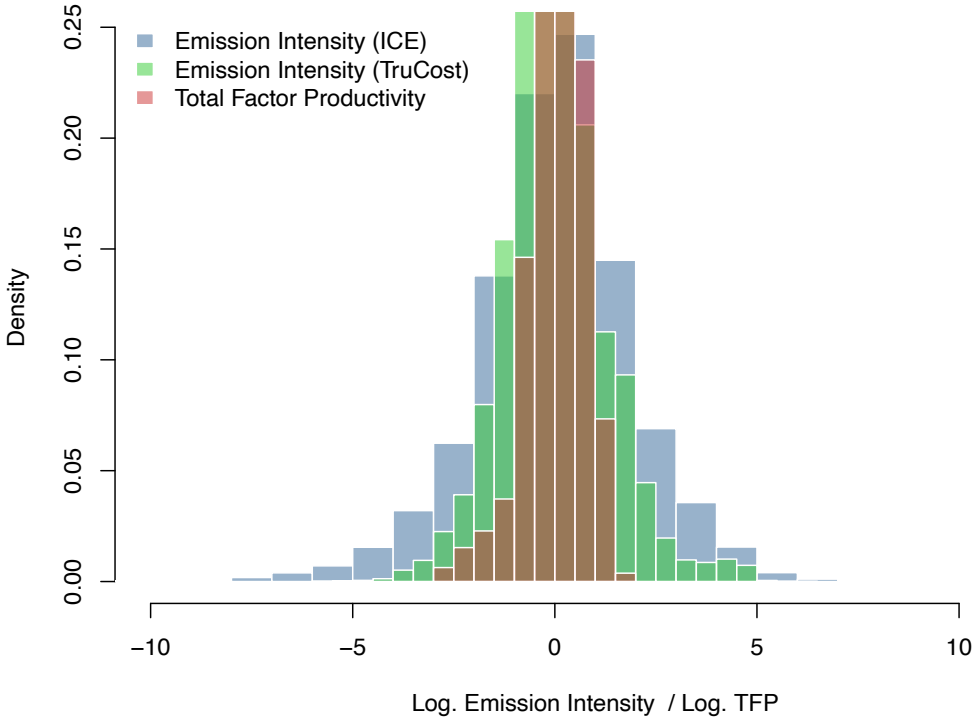
Note: Data from the ICE-ESG and TruCost samples. The figure shows the histogram of log. (scope 1) emission intensities. Emission intensities are measured by tons of CO_2 equivalent emissions divided by revenue (millions of $). Panel (a) shows the raw data, while Panel (b) residualizes the data wrt. sector, time and country fixed effects. We use ICE or NAICS-2 sectoral definitions, respectively. We consider only firms with high-quality reporting and winsorize emission intensities at the top and bottom 1% level (Appendix A.1).

steel industries in Figure A.3 in the Appendix. We measure output by megajoules of energy produced, thousand round-trip flights flown, and tons of steel produced. All histograms show substantial variation within sectors, as well as clear differences between them.

An important feature of the emissions data used in Figure 1-3, and those available from other sources, is that they are predominantly self-reported by firms (e.g., ICE, 2024; SEC, 2024). Consultancies and outside organizations can validate the processes by which estimates for scope 1 emissions have been attained; but the overall inputs into the process and the final outcome are, in all cases, up to a firm's discretion (World Resources Institute, 2004; IFRS Foundation, 2023, SEC, 2024). In this sense, a firm's overall emission is, ultimately, *private information* to the firm. As argued in the introduction, the private-information nature of the data used to estimate firms' emissions has raised concerns about misreporting among policymakers and practitioners alike (e.g., ECA, 2015; European Commission, 2023). An important corollary of this private-information setting is that, if firms' emissions are to be taxed to internalize climate externalities, addressing firms' incentives to skew their estimates (so-called "greenwashing") becomes potentially important. We restrict Figures 2 and 3 to observations that ICE classifies as "high quality"—that is, third-party-validated disclosures or records with sufficiently detailed data documentation. In Table A.2 in the Appendix, we

German manufacturing sector.

Figure 3: Emission Intensities and TFP



Note: Data from the ICE-ESG and Compustat samples. The figure shows the histogram of log. (scope 1) emission intensities and log. total factor productivity. The top and bottom 1 percent of (i) emission intensities (tons of CO2 equivalent emissions divided by revenue in millions of $); and (ii) total factor productivity estimates have been removed. TFP is estimated as in Ottonello and Winberry (2020). We consider only firms with high-quality reporting (Appendix A.1). NAICS-2 sectoral definitions are used in the estimation of TFP.

show that this restriction matters: firms without third-party validation tend to report lower emissions-intensities, all else equal, suggesting some scope for under-reporting.

Finally, a natural question that arises is whether differences in emission intensity are related to other measures of firm-level efficiency, such as TFP. Capelle *et al.* (2023) provide some evidence in this direction. Table A.3 in the Appendix shows that, all else equal, more productive firms are more carbon efficient; however, the overall contribution of this relationship to explaining emission intensities is small ($R^2 \approx 5\%$). TFP is a poor predictor of emission intensities. This will be important later for the quantification of our baseline framework.

In summary, the results in this section show that firms' carbon emissions often account for the lion's share of total carbon emissions into the atmosphere and that there is substantial heterogeneity in the emission intensity of firms, even within narrowly-defined sectors. Motivated by these findings and the self-reported nature of firms' emissions, in the next section, we extend a workhorse heterogenous-firm model to allow for heterogeneity in emission intensity and asymmetric information. We then proceed by quantifying the impact of heterogeneous

emission intensities for the optimal taxation of firms in general equilibrium.

3 Baseline Model Framework

We start by developing our baseline economy with two key features; that firms are heterogeneous in their emission intensity and that firms' emission intensities are firm-specific information, in part unknown to the social planner and other firms in the economy.

3.1 Environment

Preferences and Endowments: We consider an infinite-horizon economy with discrete time. The economy is populated by a representative competitive household with preferences over consumption streams, described by the utility function:

$$\mathcal{U} = \sum_{t=0}^{\infty} \beta^t \cdot u(C_t), \tag{1}$$

where $\beta \in (0,1)$ is the discount factor, $u(\cdot)$ is the standard concave period utility function, and C_t is consumption at time t. The household is endowed with N units of labor, which it supplies inelastically to the labor market every period, and an initial stock of capital, $K_0 > 0$. The capital stock accumulates according to the recurssion:

$$K_{t+1} = (1 - \delta_K) \cdot K_t + I_t, \tag{2}$$

where $\delta_K \in (0,1)$ is the capital depreciation rate and I_t is the household's investment.

Technology, Climate, and Markets: Output is produced by a final-goods sector comprised of a continuum of heterogeneous competitive firms, owned by households. Representative firm (i,j) is characterized by its *productivity*, a_i, and its *carbon intensity*, z_j. Pairs $\{a_i, z_j\} \in \mathcal{A} \times \mathcal{Z} \subseteq \mathbb{R}^2$ are jointly distributed with mass ω_{ij} so that $\sum_{ij} \omega_{ij} = 1$. Firm (i,j) produces quantity y_{ijt} of output at time t according to the production technology:

$$y_{ijt} = A\left(G_t\right) \cdot \exp^{a_i} \cdot F\left(n_{ijt}, k_{ijt}, e_{ijt}\right), \tag{3}$$

where n_{ijt} and k_{ijt} are the units of labor and capital employed by the firm, respectively, e_{ijt} is the firm's energy use, and $F(\cdot)$ is a standard neoclassical production function with decreasing returns-to-scale. Economy-wide productivity, $A(G_t)$, depends negatively on the stock of carbon in the atmosphere, G_t, so that $A_G < 0$ and increases in carbon damage the production-capabilities of the economy. Given an initial concentration of carbon $G_0 > 0$, the

stock of carbon in the atmosphere, in turn, evolves in accordance with:

$$G_{t+1} = (1 - \delta_G) \cdot G_t + \sum_{ij} \omega_{ij} \cdot g_{ijt}, \tag{4}$$

where g_{ijt} are firm (i, j)'s emissions and $\delta_G \in [0, 1)$ is the rate of depreciation of carbon in the atmosphere. Firm (i, j) releases emissions into the atmosphere due to its energy use:

$$g_{ijt} = \exp^{z_j - x_{ijt}} \cdot e_{ijt}, \tag{5}$$

where x_{ijt} is the firm's investment into *abatement* of its emissions at cost $Q(x_{ijt})$ with $Q' > 0$, $Q'' > 0$ and $Q(0) = 0$. Energy is available to firms in infinite supply at price P.[9] Firms rent capital and labor in competitive markets with factor prices given by R_t and W_t, respectively.

Information Structure: A central friction in our economy is that a firm's carbon-intensity type z_j is *private information* to the firm. When choosing its energy use, a firm of type j knows z_j, and thus the emissions g_{ijt} that arise as a result of its energy use; however, this information is not shared with other firms in the economy or the social planner. Clearly, this assumption presents a natural limit case of an information friction that exists between firms and the planner about firms' carbon emissions, and we later study other cases in which the planner instead obtains noisy signals of the carbon-intensity type of firms (Section 5.4).

Discussion of the Environment: The above environment includes three features critical to the carbon taxation of firms. First, our baseline model encodes a direct relationship between a firm's carbon emissions, g_{ijt}, and its energy use, e_{ijt}, which varies between firms, due to differences in their carbon-intensity type z_j (Equation 5). As we shall see below, variations in z_j translate in equilibrium one-for-one into variations in the emission efficiency measure—emissions per million $ of revenue—shown in Figure 2. Differences in carbon-intensity type can, in turn, arise due to differences in the type and blend of the energy used (e.g., different types of coal versus natural gas or renewables), differences in the filtration technology employed (e.g., the use of filters that limit the scope of nitrous oxide leakage), or even differences in the basic burner design used (e.g., burners that lead to soot formation rather than carbon emission), among others.[10] There are indeed many technological reasons for why a given

[9] This is equivalent to energy being produced by a competitive sector that produces energy from final goods at rate P—i.e., a so-called "round-about production function" for energy.

[10] For evidence that fuel type and blend affect emission intensities, see Jaramillo *et al.* (2007) and Brandt *et al.* (2014); that N_2O controls in nitric and adipic acid production yield large reductions, see Shimizu *et al.* (2000) and Pérez-Ramırez *et al.* (2003); that particulate filters reduce black carbon, see Zhang *et al.* (2023); and that burner design strongly influences soot formation and emissions, see Nemitallah *et al.* (2019) and Chu

quantity of energy use can lead to different carbon-equivalent emissions, and we above capture such differences with a firm's carbon-intensity type, z_j. Through this lens, a firm's investment into abatement x_{ijt} presents an upgrade of the technology it employs that makes its energy use less polluting. We extend our framework to explicitly allow for a choice between the use of *renweable energy*, which does not emit carbon, and *fossil-based energy* in Section 5.4. However, we here want to emphasize that—because our framework allows for differences in carbon intensities and abatement—our framework already largely accounts for such differences.

Second, notice that in our baseline model there are two sources of firm-level heterogeneity: there are firm-level differences in carbon intensity z_j and firm-level differences in productivity a_i. We show that accounting for both types of heterogeneity within one framework is crucial for the optimal carbon tax on firms. Indeed, the optimal allocation of energy use across firms depends centrally on the link between firms' productivity and their carbon intensity.

Finally, our baseline model views the climate as well-approximated by one variable: the current stock of carbon in the atmosphere, G_t. Following Golosov *et al.* (2014) and others, we argue that is a reasonable approximation given available medium-complexity climate models used in the natural sciences, which imply that the climate is well-described by current carbon concentration. Section 5.4 analyzes the case in which aggregate productivity, A_t, instead depends negatively on the global-mean temperature T_t, which in turn depends in a non-stationary manner on the stock of carbon in atmosphere, G_t, in addition to its own lagged value, T_{t-1}. This follows the approach taken in Miftakhova *et al.* (2020), among others, and creates a lagged effect of carbon emissions on productivity. This approach has recently found some popularity in the natural sciences (e.g., Allen *et al.*, 2009; Matthews *et al.*, 2009, 2012), as well as in macroeconomics (e.g., Dietz *et al.*, 2021; Fernández-Villaverde *et al.*, 2024). We stress that our main results below do not hinge on the speed over which economic damages from climate change unfold. Our assumption that it is only the production-side of the economy that is affected by emissions is made to keep our analysis as comparable as possible to that of Nordhaus (2007), Golosov *et al.* (2014), Barrage and Nordhaus (2024), and others.

4 Equilibrium and Planner Allocations

We next turn to the equilibrium and social planner allocations for our economy, the central differences between them being the extent to which firms' energy uses internalize the emission externality, $A(G_t)$. We start with the decentralized equilibrium.

et al. (2023).

4.1 Decentralized Equilibrium

Household Problem: The representative household chooses consumption C_t and capital K_{t+1} to maximize its utility in (1) subject to the per-period budget constraint:

$$C_t + K_{t+1} = W_t \cdot N + (R_t + 1 - \delta_K) \cdot K_t + \sum_{ij} \omega_{ij} \cdot \pi_{ijt}, \tag{6}$$

where π_{ijt} denotes the period-t profits of firm (i, j). The household takes the wage, W_t, and the rental rate of capital, R_t, as given when solving its problem. The solution yields:

$$u'(C_t) = \beta \cdot (R_{t+1} + 1 - \delta_K) \cdot u'(C_{t+1}). \tag{7}$$

Firm Problem: The problem of firm (i, j) is similarly standard, as firms in the baseline equilibrium do not internalize the damages their emissions have on production. Firm (i, j) chooses labor n_{ijt}, capital k_{ijt}, energy e_{ijt}, and abatement x_{ijt} to maximize its profits,

$$\pi_{ijt} = y_{ijt} - W_t \cdot n_{ijt} - R_t \cdot k_{ijt} - P \cdot e_{ijt} - Q(x_{ijt}), \tag{8}$$

subject to the production constraint in (3). The firm thus optimally equates the marginal revenue product for labor n_{ijt} and capital k_{ijt} with the real wage, W_t, and the rental rate of capital, R_t, respectively; the relevant conditions for energy, e_{ijt}, and abatement, x_{ijt}, are:[11]

$$\text{mrpe}_{ijt} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_e(n_{ijt}, k_{ijt}, e_{ijt}) = P, \qquad Q'(x_{ijt}) = 0. \tag{10}$$

Equation (10) presents a main departure point for our analysis: a profit-maximizing firm (i, j) optimally equates the marginal revenue product of energy, mrpe_{ijt}, to its price, P, and chooses zero abatement, $x_{iijt} = 0$. We are now ready to define the equilibrium of our economy:

Equilibrium Definition: An competitive equilibrium consists of household allocations $\{C_t, K_{t+1}\}_{t=0}^{\infty}$, firm allocations $\{\{n_{ijt}, k_{ijt}, e_{ijt}, x_{ijt}\}_{i,j}\}_{t=0}^{\infty}$, factor prices $\{W_t, R_t\}_{t=0}^{\infty}$, and firm-led emissions $\{\{g_{ijt}\}_{i.j}, G_t\}_{t=0}^{\infty}$ such that:

- Given factor prices, W_t and R_t, the household and firm allocations solve the household and firm problem, respectively (i.e, Equations 6-10 hold).

[11] Thus,

$$\text{mrpn}_{ijt} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_n(n_{ijt}, k_{ijt}, e_{ijt}) = W_t, \qquad \text{mrpk}_{ijt} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_k(n_{ijt}, k_{ijt}, e_{ijt}) = R_t. \tag{9}$$

- Given allocations, the labor, capital, and goods market clears; i.e., for all t: $\sum_{ij} \omega_{ij} \cdot n_{ijt} = N$, $\sum_{ij} \omega_{ij} \cdot k_{ijt} = K_t$, and:

$$C_t + K_{t+1} = (1 - \delta_K) \cdot K_t + \sum_{ij} \omega_{ij} \cdot y_{ijt} - \sum_{ij} \omega_{ij} \cdot (P \cdot e_{ijt} + Q(x_{ijt})) . \tag{11}$$

- Emissions evolve in accordance with Equations (4) and (5).

The equilibrium features a tight link between a firm's carbon-intensity type, z_j, and the emission-efficiency measure, $\log(g_{ijt}/y_{ijt})$, shown in Figure 2. Indeed, it follows directly from Equation (10) that in the case in which $F(\cdot)$ admits the Cobb-Douglas form: $\log(g_{ijt}/y_{ijt}) = \log(P/\upsilon) + z_j$.[12] Variations in z_j translate one-for-one into variations in emission intensity, and are the ultimate source behind this heterogeneity. This tight relationship between z_j and the emission-intensity measure used in the data extends more generally to any neoclassical production function. We summarize this result and our preceding discussion in Lemma 1:

Lemma 1. *In the decentralized equilibrium (de):*

i. Firm (i, j)'s energy use, e^{de}_{ijt}, and abatement, x^{de}_{ijt}, satisfy:

$$\text{mrpe}\left(n^{de}_{ijt}, k^{de}_{ijt}, e^{de}_{ijt}\right) = P, \qquad Q'\left(x^{de}_{ijt}\right) = 0. \tag{12}$$

ii. Variations in carbon intensity, z_j, translate into variations in emission efficiency:

$$\log(g^{de}_{ijt}/y^{de}_{ijt}) = \log(P/\upsilon) + z_j + o(1). \tag{13}$$

Finally, notice that a key feature of the decentralized equilibrium is that firms do not internalize the economy-wide damages from carbon emissions, $A(G_t)$, when making their energy and abatement choices in Equation (12). A firm directly equates its marginal benefit from energy use, the marginal revenue product of energy, with its marginal private cost, equal to the market price. This, in turn, follows from the profit function in (8) being independent of emissions and carbon-intensity type, $\partial \pi_{ijt}/\partial g_{ijt} = \partial \pi_{ijt}/\partial z_j = 0$. The decentralized equilibrium, as such, features a conflict between profit maximization and the overall productivity of the economy, leading, as we will see next, to a suboptimal allocation of all factors of production.

4.2 The Full-Information Planning Problem

We proceed with the planning problem in which the social planner has full-information about firms' carbon-intensity types, $\{z_j\}$, and thus emissions, $\{g_{ijt}\}$. Although our ultimate interest

[12] $\log\left(\frac{g_{ijt}}{y_{ijt}}\right) = \log\left(P \cdot \frac{\partial \log e_{ijt}}{\partial \log F_{ijt}}\right) + z_j = \log\left(\frac{P}{\nu}\right) + z_j$.

lies with the solution for the planning problem in which firms have private information about z_j, it will be useful to start with the full-information case, as the comparison between the two cases will allow us to more sharply characterize the costs of information frictions.

The full-information social planner solves the problem:

$$\max_{\{C_t, K_{t+1}, \{k_{ijt}, n_{ijt}, e_{ijt}, x_{ijt}\}_{i,j}\}_{t=0}^{\infty}} \mathcal{U} = \sum_{t=0}^{\infty} \beta^t \cdot u\left(C_t\right) \tag{14}$$

subject to Equations (3), (4), (5), $\sum_{ij} \omega_{ij} \cdot n_{ijt} = N$, $\sum_{ij} \omega_{ij} \cdot k_{ijt} = K_t$, Equation (11), as well as the usual non-negativity constraints. Let $\Lambda_{k,t}$ denote the constraint multiplier on the market-clearing condition for capital at time t. The first-order condition for K_{t+1} and C_t then once more yields the standard consumption Euler equation:

$$u'\left(C_t\right) = \beta \cdot \left(\Lambda_{k,t+1} + 1 - \delta_K\right) \cdot u'\left(C_{t+1}\right). \tag{15}$$

Let $\Lambda_{n,t}$ analogously denote the constraint multiplier on the market-clearing condition for labor at time t. The full-information social planner then further optimally equates the marginal revenue product for labor n_{ijt} and capital k_{ijt} to the multipliers $\Lambda_{n,t}$ and $\Lambda_{k,t}$, respectively, while the relevant conditions for energy, e_{ijt}, and abatement, x_{ijt}, are:[13]

$$\text{mrpe}_{ijt} = P + \Lambda_G \cdot \exp^{z_j - x_{ijt}}, \quad Q'\left(x_{ijt}\right) = \Lambda_G \cdot \exp^{z_j - x_{ijt}} \cdot e_{ijt}, \tag{17}$$

where $\Lambda_{G,t}$ denotes the constraint multiplier on the law-of-motion for emissions in (4):

$$\Lambda_{G,t} = -\frac{A'(G_t)}{A(G_t)} \cdot \mathcal{Y}_t + \beta \cdot \frac{u'(C_{t+1})}{u'(C_t)} \cdot (1 - \delta_G) \cdot \Lambda_{G,t+1}, \tag{18}$$

and $\mathcal{Y}_t \equiv \sum_{ij} \omega_{ij} \cdot y_{ij}$ is gross output in the economy.

Equations (15)-(18) characterize the full-information planner allocation and can be directly compared to Equations (7), (9)-(10), which describe the decentralized equilibrium. The central difference between the two sets of equations arises because $\Lambda_{G,t} > 0$ in Equation (17). This term, which captures the relevant notion of the *social cost of carbon* for our economy, equals from Equation (18) the present discounted value of marginal externality damages, $-A'(G_t)/A(G_t) \cdot \mathcal{Y}_t$—as in Golosov *et al.* (2014) and Barrage and Nordhaus (2024).

Definition 1. *Let* $\Lambda_{G,t} = \sum_{h=0}^{\infty} \beta^h \cdot \frac{u'(C_{t+h+1})}{u'(C_{t+h})} \cdot (1-\delta_G) \cdot \left(-\frac{A'(G_{t+h})}{A(G_{t+h})}\right) \cdot \mathcal{Y}_{t+h} > 0$ *be the* ***social cost of carbon****. Firm* (i,j)*'s* ***marginal social cost of carbon*** *is* $\lambda_{G,ijt} \equiv \Lambda_G \cdot \exp^{z_j - x_{ijt}} > 0$,

[13] Thus,

$$\text{mrpn}_{ijt} = \Lambda_n, \qquad \text{mrpk}_{ijt} = \Lambda_k. \tag{16}$$

while its ***marginal social return on abatement*** *is* $\Delta_{ijt} \equiv \lambda_{G,ijt} \cdot e_{ijt} > 0$.

This social cost of carbon is then multiplied in Equation (17) by $\exp^{z_j - x_{ijt}}$ to create firm (i,j)'s *marginal social cost of carbon*, which as a consequence depends on the firm's carbon-intensity type, z_j, and its abatement, x_{ijt}. The left-hand side of Equation (17) shows that the full-information planner optimally equates firm (i,j)'s marginal revenue product of energy, mrpe_{ijt}, to its overall societal cost, equal to the sum of the market price, $P > 0$, and its marginal social cost of carbon, $\lambda_{G,ijt} > 0$. The decentralized equilibrium, by contrast, equates the marginal revenue product of energy only to its price (Equation 12). The right-hand side of Equation (17) then further shows that the full-information planner optimally also equates the marginal cost of abatement, $Q'(x_{ijt})$ to its *marginal social return*, Δ_{ijt}, which scales with firm (i,j)'s energy use, e_{ijt}, in addition to its marginal social cost of carbon, $\lambda_{G,ijt}$.

Because of the internalization of the emission exeternality in Equation (17)—i.e., the positive social cost of carbon ($\Lambda_{G,t} > 0$)—the full-information planner allocation differs from that of the decentralized equilibrium. Indeed, it follows from Equations (15)-(18) that:

Lemma 2. *In the full-information planner (first-best, fb) allocation:*

i. Firm (i,j)*'s energy use,* e^{fb}_{ijt}*, and abatement,* x^{fb}_{ijt}*, satisfy:*

$$\text{mrpe}(n^{fb}_{ijt}, k^{fb}_{ijt}, e^{fb}_{ijt}) = P + \lambda_{G,ijt}, \qquad Q'\left(x^{fb}_{ijt}\right) = \Delta_{G,ijt}. \tag{19}$$

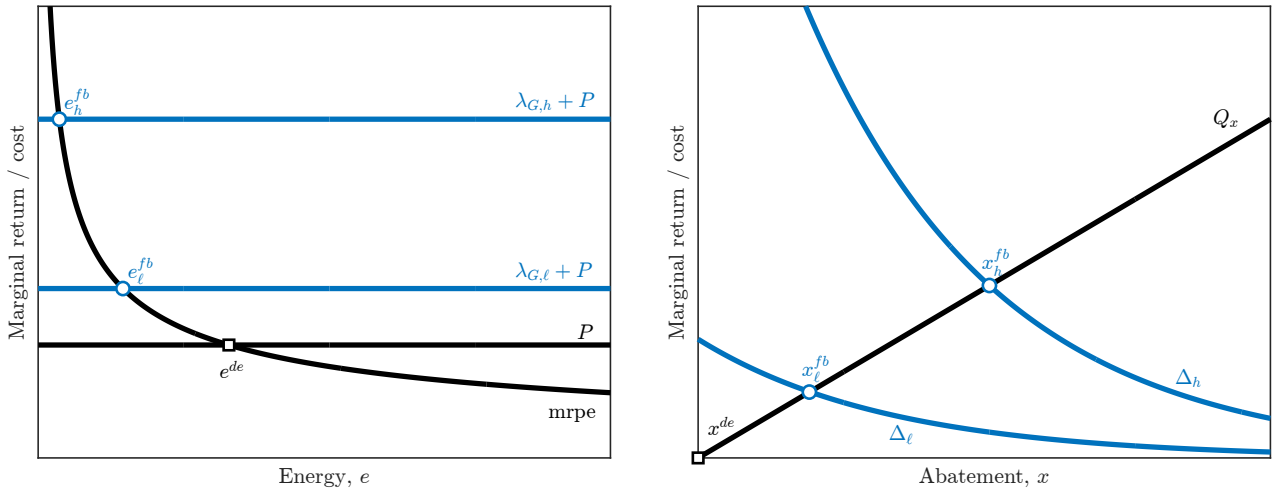
ii. The planner depresses energy and increases abatement heterogeneously across firms. In particular, for any i*, and any* j', j *with* $z_{j'} > z_j$*:* $e^{fb}_{ij't} < e^{fb}_{ijt} < e^{de}_{ijt}$ *and* $x^{fb}_{ij't} > x^{fb}_{ijt} > x^{de}_{ijt}$.

Since the marginal social cost of carbon is positive for all firms, $\lambda_{G,ijt} > 0$, the full-information planner depresses all firm's energy uses relative to what firms would choose to maximize their profits, $e^{fb}_{ijt} < e^{de}_{ijt}$ (Equation 12 and 19).[14] The social planner internalizes the emission externality and optimally drives down firms' energy use, and thus overall emissions to increase productivity in the economy. The social planner likewise also commands firms to, in part, abate their emissions, $x^{fb}_{ijt} > x^{de}_{ijt} = 0$. A larger carbon-intensity type, z_j, or a larger externality damage, $A'(G_t)/A(G_t) \cdot \mathcal{Y}_t$, all else equal, causes a firm to exhibit a larger marginal social cost of carbon, which in turn drives a larger wedge between the full-information allocation and that featured in the decentralized equilibrium. Figure 4 shows graphically how the allocation of energy and abatement are modified by the full-information planner.

Crucially, notice that the marginal social cost of carbon in Equation (19) is *type-specific*: each firm (i,j) features its *own* marginal social cost of carbon. Because the planner optimally

[14] This further follows from the assumption of a neoclassical production technology $F(\cdot)$.

Figure 4: The Full-Information Social Planner Allocation



Panel (a): Energy

Panel (b): Abatement

Note: The figure shows the determination of the equilibrium (in black) and full-information optimal (in blue) use of energy (Panel a) and abatement (Panel b). Equations (12) and (19) determine the optimal use of energy, e_{ij}, and abatement, x_{ij}, in the two cases. Panel (a) and (b) further considers the case in which a firms has a high (low) marginal social cost of capital, $\lambda_G^h > \lambda_G^l$, and hence a high (low) marginal benefit of abatement.

equates each firm's marginal revenue product of energy to the sum of the market price of energy and the firm-specific marginal social cost of carbon, it follows that the planner introduces *factor misallocation* across firms, to optimally account for the emission externality. This feature of the full-information planner allocation will be important for later.

Finally, because firms' marginal revenue products and the marginal cost of abatement are equated to type-specific objects, it follows that the planner also creates *profit differentials*, even among equally productive firms. That is, take two firms, one of type (i, j) and one of type (i, j'), where $j' \neq j$. Then, it follows from (15)-(19) that the implied profits under the planner allocation will generically differ between the two firms, $\pi_{ij}^{fb} \neq \pi_{ij'}^{fb}$, due to differences in their marginal social cost of carbon. In particular, the full-information planner allocates more energy to less carbon-intensive firms, pushing their allocation closer to that of the decentralized equilibrium and thus allowing them to attain higher profits. Indeed, it follows that:

Corollary 1 (**Incentive-incompatibility**). *In the full-information planner allocation the implied profits of firm (i, j) are larger than those of firm (i, j'), $\pi_{ijt}^{fb} > \pi_{ij't}^{fb}$, when $z_j < z_{j'}$.*

As we will see, it is straightforward to implement the full-information planner allocation with a linear carbon tax that is identical across firms. For now, however, notice that the existence of any profit differential among equally productive firms under such a tax translates into a direct incentive for high carbon-intensive firms—those with a large value of z_j—to misreport

their carbon intensities. Doing so would cause them to attain higher profits. Misreporting would, as such, be part of firms' profit-maximizing strategy and undo the benefits of taxation.

A Naive Planner: Lastly, suppose a *naive planner (nv)* was to act as the full-information planner presented above but without having information on firms' carbon intensities, $\{z_j\}$. In such an economy, all firms would find it optimal to report having the lowest carbon intensity, i.e., $\underline{z} \equiv \min_{z \in \mathcal{Z}} z_j = \arg\max \pi_{ijt}$. A firm would always have a *dominant reporting strategy* given by reporting to be of type $(i, \underline{j})$, where $\underline{j}$ is the index of $\underline{z}$. The first-order conditions for energy e_{ijt} and abatement x_{ijt} in such an economy would equal:

$$\text{mrpe}(n^{nv}_{ijt}, k^{nv}_{ijt}, e^{nv}_{ijt}) = P + \lambda_{G,i\underline{j}t}, \quad Q'(x^{nv}_{ijt}) = \Delta_{G,i\underline{j}t}, \tag{20}$$

where $\lambda_{G,i\underline{j}t} = \Lambda_G \cdot \exp^{\underline{z} - x_{ijt}}$ is the *misreported* marginal social cost of carbon and $\Delta_{G,i\underline{j}t} = \lambda_{G,i\underline{j}t} \cdot e_{ijt}$ is the *misreported* marginal return of abatement. Clearly, both conditions imply that all firms would use more energy and do less abatement than in the full-information case, $e^{nv}_{ijt} > e^{fb}_{ijt}$ and $x^{nv}_{ijt} < x^{fb}_{ijt}$, implying a smaller internalization of the emission externality. Depending on the dispersion of emission intensities, ignoring misreporting incentives could, as a result, substantially hinder the benefits of carbon taxation. By contrast, a planner who accounts for incentives cold restore the above Pigouvian logic and further welfare.

4.3 The Mirrleesian Planning Problem

We turn to our main problem of interest: the problem faced by a social planner that is uncertain about a firm's carbon-intensity type, z_j. This Mirrleesian planner faces the same problem as the full-information planner, but in addition faces two constraints: an incentive-compatibility constraint and an individual-rationality constraint. Since a firm of type (i, j) can pretend to be of type (i, j'), where $j \neq j'$, in the planner's allocation, each firm-type's profits must exceed those from a different emission-intensity type (*incentive compatibility*):[15]

$$\pi(n_{ijt}, k_{ijt}, e_{ijt}, x_{ijt} \mid a_i, z_j) \geq \pi(n_{ij't}, k_{ij't}, e_{ij't}, x_{ij't} \mid a_i, z_j) \quad \forall j' \neq j, \tag{21}$$

where the right-hand side of (21) is the profit of firm (i, j) when it receives the allocation intended for firm (i, j'). Moreover, to be willing to participate in the mechanism, a firm of

[15]Because we are assuming that productivity a_i is known, or alternatively that all inputs are known and thus productivity can be recovered from the production function accordingly, no firm of type (i, j) can pretend to be of type (i', j) and thus no incentive compatibility constraints are needed on this dimension.

type (i, j) must obtain non-negative profits from it (*individual rationality*):

$$\pi(n_{ijt}, k_{ijt}, e_{ijt}, x_{ijt} \mid a_i, z_j) \geq 0. \tag{22}$$

The problem faced by this planner resembles the classical problem faced by all Mirrleesian planners in the household taxation literature (e.g., Mirrlees, 1976). The central difference being that a firm's unobserved type is not directly payoff-relevant. Recall that, conditional on inputs and abatement, a firm's profits in (8) do not depend on the carbon-intensity type, $\partial \pi_{ijt}/\partial z_j = 0$.[16] This contrasts with the household case in which the household's unobserved type (i.e., the effort-type) matters for the household's payoff through its income.

A consequence of this feature—and the monotone ranking of firm profits in carbon-intensity types z_j in the full-information allocation (Corollary 1)—is that each firm has a dominant strategy: to misreport its carbon-intensity type to the planner and claim to be of type $\underline{z} \equiv \min_{z \in \mathcal{Z}} z_j$, as is the case under the naive planner. It follows that a firm's profits in the Mirrleesian case have to be independent of the carbon-intensity type, $\pi_{ijt} = \pi_{it}$ for all j. We show in Appendix B.1 that such profit equailization—the fact that the incentive-compatibility constraint binds for all j—in turn implies that the Mirrleesian planner *optimally bunches allocations* across carbon-intensity types within productivity type i (Lemma A.X). The planner chooses an allocation such that $\{k_{ijt}, n_{ijt}, e_{ijt}, x_{ijt}\} = \{k_{ij't}, n_{ij't}, e_{ij't}, x_{ij't}\}$ for all $j' \neq j$.[17]

A simple example can help with the intuition behind this result. Suppose energy is the only input into production and that there is no abatement. Suppose further that for a firm of type i, there are two energy allocations, $e_{1,it}$ and $e_{2,it}$, which satisfy the incentive-compatibility constraint, $\pi(e_{1,it}|a_i, z_j) = \pi(e_{2,it}|a_i, z_j)$, with $e_{1,it} < e_{2,it}$. With energy as the sole input, consumption equals value-added in the economy, which in turn equals total profits. Thus, any allocation that uses only $e_{1,it}$ and $e_{2,it}$ is incentive compatible and delivers the same welfare. Yet, because emissions enter through Equation (3), and thus create a negative externality, $A(G_t)$, the Mirrleesian planner prefers the lower energy level, $e_{1,it}$, for all type-i firms.

Using this symmetry of the allocation, it is straightforward to derive the solution to the Mirrleesian problem. To characterize its solution, it will, however, be useful to first define:

Definition 2. *Let* $\overline{\lambda}_{G,it} \equiv \mathbb{E}[\lambda_{G,ijt} \mid a_i] > 0$ *be type* i*'s* ***average social cost of carbon*** *and* $\overline{\Delta}_{G,it} \equiv \overline{\lambda}_{G,it} \cdot e_{it} > 0$ *be type* i*'s* ***average social return on abatement****.*

Let $\Lambda_{k,t}$ and $\Lambda_{n,t}$ once more denote the constraint multipliers on the market-clearing condition for capital and labor. The Mirrleesian planner, as the full-information planner, equates

[16] Thus, the *single-crossing condition*—often used as a sufficient condition for the existence of a truth-telling mechanism, because it allows one to focus on local (nearest) deviations—is not satisfied in our setting.

[17] The Inada condition on the production technology, $F(\cdot)$, ensures that the individual-rationality constraint in Equation (22) never binds.

the marginal revenue product for capital k_{ijt} and labor n_{ijt} to the multipliers $\Lambda_{k,t}$ and $\Lambda_{n,t}$, respectively.[18] The planner's conditions for K_{t+1} and C_t also again yield the standard Euler equation for consumption in (15), while the social cost of carbon Λ_G equals that in (18). The difference between the full-information and the Mirrleesian allocation only manifests itself in the firm-specific conditions for energy and abatement:

Lemma 3. *In the Mirrleesian planner (mr) allocation:*

i. Firm (i,j)'s energy use, e^{mr}_{ijt}, and abatement, x^{mr}_{ijt}, satisfy:

$$\text{mrpe}(n^{mr}_{it}, k^{mr}_{it}, e^{mr}_{it}) = P + \overline{\lambda}_{G,it}, \qquad Q'\left(x^{mr}_{it}\right) = \overline{\Delta}_{G,it}, \quad \forall j. \tag{24}$$

ii. The Mirrleesian planner depresses energy and increases abatement, while bunching firms across carbon intensities. In particular, for any i,j: $e^{mr}_{it} < e^{de}_{ijt}$, and $x^{mr}_{it} > x^{de}_{ijt}$.

As the full-information planner, the Mirrleesian planner depresses energy uses for all firms relative to the decentralized equilibrium and abates emissions. However, because the Mirrleesian planner does not know a firm's carbon-intensity type z_j, she cannot equate firm (i,j)'s marginal revenue product of energy to its marginal social cost of carbon, $\lambda_{G,ijt}$. Instead, the Mirrleesian planner optimally averages and equates a firm (i,j)'s marginal revenue product of energy to group-i's *average* marginal social cost of carbon. A similar averaging occurs for abatement. In this sense, the differences between the type-specific marginal costs and returns, $\lambda_{G,ijt}$ and $\Delta_{G,ijt}$, and their average values, $\overline{\lambda}_{G,it}$ and $\overline{\Delta}_{G,it}$, capture the bite that the information friction has on the economy. Figure 5 illustrates this point graphically

Finally, notice that because of the bunching of allocations within productivity types, the Mirrleesian allocation features, all else equal, less factor misallocation than the full-information case. This is although the Mirrleesian planner clearly cannot be as effective as the full-information planner in addressing the economic damages from climate change.

4.4 Implementation and Carbon Taxation

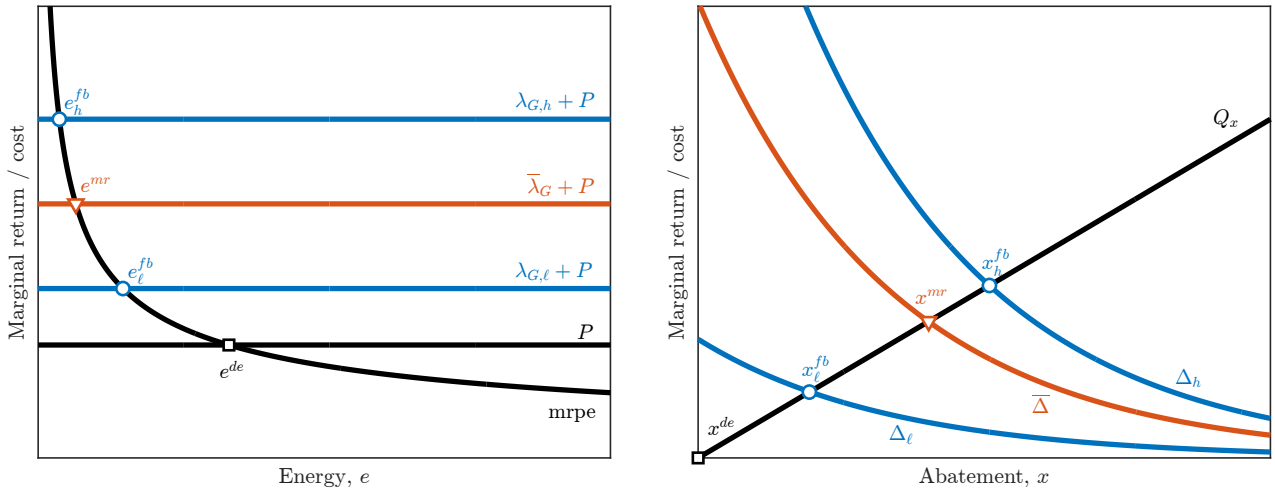
The Mirrleesian allocation, as well as its full-information counterpart, can straightforwardly be implemented by a simple, linear carbon tax in the decentralized equilibrium, combined with appropriate lump-sum rebates. To see this, let $\tau_{ijt}(g)$ denote the carbon-tax payments made by firm (i,j) when its emissions equal g. Firm (i,j)'s post-tax profits become:

$$\pi_{ijt} = y_{ijt} - W_t \cdot n_{ijt} - R_t \cdot k_{ijt} - P \cdot e_{ijt} - Q(x_{ijt}) - \tau_{ijt}(g_{ijt}), \tag{25}$$

[18] Thus,

$$\text{mrpn}_{ijt} = \Lambda_n, \qquad \text{mrpk}_{ijt} = \Lambda_k. \tag{23}$$

Figure 5: The Mirrleesian Social Planner Allocation



Panel (a): Energy

Panel (b): Abatement

Note: The figure shows the determination of the equilibrium (in black), full-information optimal (in blue), and Mirrleesian optimal (in orange) use of energy (Panel a) and abatement (Panel b). Equations (10), (17), and (24) determine the optimal use of energy, e_{ij}, and abatement, x_{ij}, in the three cases, respectively. Panel (a) and (b) further considers the case in which a firms has a high (low) marginal social cost of capital, $\lambda_G^h > \lambda_G^l$, and hence a high (low) marginal benefit of abatement. The average value is denoted by $\overline{\lambda}_G$.

The maximization of firm (i, j)'s profits subject to the production feasibility constraint (3) now results in the following *tax-adjusted* conditions for energy e_{ijt} and abatement x_{ijt}:

$$\text{mrpe}_{ijt} = P + \tau'_{ijt}\left(g_{ijt}\right) \cdot \exp^{z_j - x_{ijt}}, \quad Q'\left(x_{ijt}\right) = \tau'_{ijt}\left(g_{ijt}\right) \cdot \exp^{z_j - x_{ijt}} \cdot e_{ijt}; \tag{26}$$

while the conditions for labor n_{ijt} and capital k_{ijt}, and the Euler equation for consumption C_t, are identical to those in Section 4.1. Comparing Equation (26) with Equations (17) and (24), respectively, which form the basis for Lemma 2 and 3, immediately shows that:

Proposition 1. *Let the carbon tax schedule for firm (i, j), $\tau_{ijt}\left(g\right)$, equal:*

$$\tau_{ijt}(g) = \Lambda_{G,t} \cdot g \equiv \tau_t^{fb}(g) \quad \textit{or} \quad \tau_{ijt}(g) = \Lambda_{G,t} \cdot \frac{\mathbb{E}[\exp^{z_j} | a_i]}{\exp^{z_j}} \cdot g \equiv \tau_{ijt}^{mr}(g), \tag{27}$$

and let tax proceeds be rebated lump-sum to the representative consumer. Then, the decentralized equilibrium allocation coincides with: (i) the full-information social planner allocation when $\tau_{ijt} = \tau_{ijt}^{fb}$; and (ii) the Mirrleesian social planner allocation when $\tau_{ijt} = \tau_{ijt}^{mr}$.

Proposition 1 shows that a simple, linear carbon tax, equal to the social cost of carbon $\Lambda_{G,t}$, implements the full-information planner allocation. The proposition, as a result, recov-

ers the main insight of Golosov *et al.* (2014), but extends it to the case with heterogeneity in productivity and carbon intensity. The result arises because the marginal damage to the production-side of the economy from any additional unit of carbon is identical across firms: total factor productivity, $A(G_t)$, is a function of aggregate emissions G_t alone; firm-level emission are perfect substitutes. This, in turn, implies that the planner optimally sets a *homogeneous* tax across firms, equal to the social cost of carbon $\Lambda_{G,t}$, to adjust for the externality.

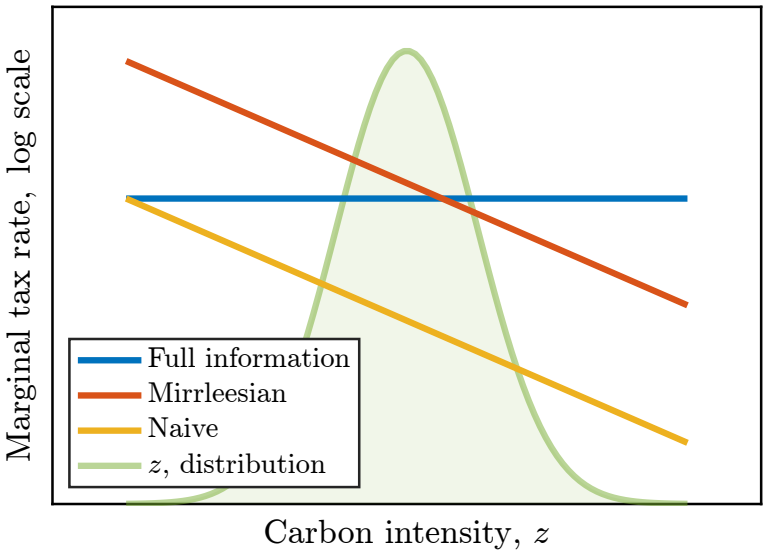
When heterogeneity is coupled with an information friction, however, the optimal carbon tax changes. Surprisingly, Proposition 1 shows that the Mirrleesian planner's allocation can still be implemented by a linear carbon tax; but crucially one that is *heterogeneous* across firms: the optimal marginal tax, $\Lambda_{G,t} \cdot \left(\frac{\mathbb{E}[\exp^{z_j}|a_i]}{\exp^{\bar{z}_j}}\right)$, is constant *within a firm* but depends on the carbon-intensity type of the firm, z_j, and, through the conditional expectation term, may also vary with the productivity type, a_i. This dependence arises because of the incentive-compatibility constraint in (21). In order to equalize profits within productivity groups—and thus prevent misreporting—the social planner must set a *higher marginal tax* for less carbon-intensive firms. This, at first glance, counterintuitive results arises because more carbon-intensive firms have a stronger incentive to misreport their carbon intensity. The misreporting incentive is one-sided: firms always prefer to report a lower z_j.

Figure 6 illustrates the two log-linear marginal tax rates as a function of carbon-intensity types, z_j. It follows directly from Proposition 1 that the Mirrleesian rate exceeds the full-information rate for firms whose carbon intensity is below their group average, and falls below it for firms whose carbon intensity is above the group average. Importantly, the two taxes are conceptually similar: the Mirrleesian schedule merely proportionally adjusts the full-information tax by $\frac{\mathbb{E}[\exp^{z_j}|a_i]}{\exp^{\bar{z}_j}}$ to ensure incentive compatibility. This adjustment determines the distinctive shape shown in Figure 6. In this sense, the Mirrleesian tax is not substantially harder to implement than its full-information counterpart. Given a firm's carbon-intensity type—which firms have no incentive to misreport under the tax scheme—the tax is still linear and a simple modification of the uniform Pigouvian tax. However, as we will see in the next section, this simple adjustment goes a long way to recover the potential benefits of carbon taxation in environments in which firms can misreport their emissions.

Finally, Figure 6 presents the evolution of the tax rate that would prevail with the naive planner, who acts as the full-information planner but without having information on carbon-intensity types, encountered earlier. It follows directly from comparing Equation (20) with Equation (26) that the tax that implements the naive allocation is also linear within firms:

$$\tau_{jt}^{nv} \equiv \Lambda_{G,t} \cdot \exp^{\bar{z}-z_j} \cdot g_t. \tag{28}$$

Figure 6: Optimal Marginal Tax Rates



Note: The figure shows the marginal tax rates that implement the full-information social planner allocation (in blue), the Mirrleesian planner allocation (in red), and the naive tax schedule (in orange). The first two are taken from Proposition 1, while the latter equals: $(\tau_t')^{\text{naive}} \equiv \Lambda_{G,t} \cdot \exp^{z-z_j} \neq (\tau_t')^{\text{fb}}$.

The naive tax rate *declines* in firms' carbon-intensity types, as the benefits of misreporting, which are reflected in the effective tax, are increasing in z_j, and is always lower than the marginal rates charged by the full-information and the Mirrleesian planner.

Before proceeding with the potential benefits of a Mirrleesian approach to taxation and the quantification of our framework, we should mention that other, more elaborate policies exist that implement the planner's allocation. A tax on energy combined with a subsidy on abatement can, for example, also attain this aim. Appendix D studies this case in detail. [19]

4.5 A Welfare Decomposition

The above implementation of the planner allocations is informative about the extent to which firms' misreporting incentives alter the optimal carbon tax on firms. Building on our analysis of the decentralized equilibrium and social planner allocations, in this subsection, we turn to studying the trade-off inherent to the planner allocations, resulting in the above tax schedules. To do so, we will make use of the following decomposition of consumption, which helps showcase the bite firms' misreporting incentives have on the economy:

[19] In particular, we focus on carbon taxes for two reasons: (i) it is a single instrument, whereas the alternative requires coordinating two or more instruments; and (ii) a carbon tax is robust to the exact functional form of the emissions function, whereas the energy tax-abatement subsidy combination works only under the (common but restrictive) assumption that emissions depend solely on energy, rather than on capital or other inputs.

Proposition 2. *Let $F(n,k,e) = n^{\alpha}\cdot k^{\gamma}\cdot e^{\upsilon}$ with $\alpha,\gamma,\upsilon > 0$ and $\zeta \equiv 1-\alpha-\gamma-\nu \in (0,1)$, and let $Q(x) = (1+\theta)^{-1}\cdot x^{1+\theta}$ with $\theta > 0$. Then, in both the decentralized equilibrium (de), the full-information (mp), the Mirrleesian planner (mp), and the naive planner (nv) allocations:*

$$C_t = \Omega_t \cdot \Theta_t \cdot \left(A(G_t)\cdot \exp^{a}\cdot K_t^{\alpha}\cdot N^{\gamma}\cdot E_t^{\upsilon}\right), \tag{29}$$

where:

- *a is the firm-level productivity average, $a \equiv \log\left(\sum_{ij}\omega_{ij}\cdot\exp^{\frac{a_i}{\zeta}}\right)^{\zeta}$;*
- *$\Omega_t \in (0,1]$ is a measure of factor misallocation:*

$$\Omega_t \equiv \left(\sum_{ij}\varpi_{ij}^{\Omega}\cdot\left(1+\frac{\tau'_{ijt}\cdot\exp^{z_j-x_{ijt}}}{P}\right)^{-\frac{\nu}{\zeta}}\right)^{1-\alpha-\gamma}\left(\sum_{ij}\varpi_{ij}^{\Omega}\cdot\left(1+\frac{\tau'_{ijt}\cdot\exp^{z_j-x_{ijt}}}{P}\right)^{-\frac{1-\alpha-\gamma}{\zeta}}\right)^{-\nu},$$

  with $\varpi_{ij}^{\Omega} \equiv \omega_{ij}\cdot\exp^{\frac{a_i}{\zeta}}\cdot\left(\sum_{ij}\omega_{ij}\exp^{\frac{a_i}{\zeta}}\right)^{-1}$ and $\tau'_{ijt} \in \left\{0,\ \frac{d\tau_t^{fb}}{dg},\ \frac{d\tau_{jt}^{nv}}{dg},\ \frac{d\tau_{ijt}^{mr}}{dg}\right\}$;
- *$\Theta_t \equiv \left(P\cdot E_t + \sum_{ij}\omega_{ij}\cdot Q(x_{ijt})\right)\cdot\mathcal{Y}_t^{-1} > 0$ is the intermediate input share.*

Proposition 2 uses the results from Proposition 1 to characterize consumption, the determinant of household welfare, as the outcomes of the decentralized economy with carbon taxes. The original decentralized equilibrium without carbon taxation corresponds to the case in which $\tau'_{ijt} = 0$, while the allocations of the full-information planner, the Mirrleesian planner, and the naive planner correspond to $\tau'_{ijt} = d\tau_t^{\ell}/dg,\ \ell \in \{fb, nv, mr\}$, respectively.

The proposition shows that since carbon taxes are zero in the original decentralized equilibrium, $\Omega_t^{de} = 1$: there is no factor misallocation in the original decentralized equilibrium; however, because emissions are universally higher in this case, $G_t^{de} > G_t^{nv} > G_t^{mr} > G_t^{fb}$, aggregate productivity is also depressed, $A_t^{de} < A_t^{nv} < A_t^{mr} < A_t^{fb}$. The naive planner delivers an outcome close to that of the decentralized equilibrium. By imposing only minimal distortions, it avoids further misallocation, but at the cost of only modest emissions reductions and thus limited productivity gains. In contrast, the full-information planner optimally balances the *efficiency-misallocation trade-off*. By setting $\tau'_{ijt} = d\tau_t^{fb}/dg$, the planner optimally internalizes the emission externality, and thus attains the highest level of productivity and consumption; yet the planner does so at the cost of creating the most misallocation, $\Omega_t^{fb} < \Omega_t^{nv} < 1$.

Compared to the full-information planner, the Mirrleesian planner, in addition, faces the incentive-compatibility constraint in (21). This constraint, in turn, worsens the possible trade-off between efficiency and misallocation. Consequently, the social planner optimally lessens the misallocation that occurs relative to that under full-information, but at the cost of larger

emissions and thus lower productivity. In effect, the four allocations considered—the decentralized equilibrium, the full-information planner, the naive planner, and the Mirrleesian planner—can be seen as tracing out different points on an *efficiency-misallocation curve*, one that naturally adjusts for the intermediate and accumulated input use. We will later return to the decomposition in Proposition 2 when evaluating the benefits from Mirrleesian taxation.

5 Quantification and Extensions

We have demonstrated that firms' reporting incentives change the optimal carbon-tax schedule on firms. Our results show that accounting for firms' misreporting incentives, all else equal, leads to higher marginal taxes on lower carbon intensity firms. In this section, we use our analysis to shed light on the quantitative implications of our results for the optimal carbon taxes on US firms. Does accounting for firms' reporting incentives quantitatively alter the environmental and welfare benefits of carbon taxation? And if so, to what extent and through which channels? To answer these questions, we proceed to parametrize and quantify our baseline economy using our data and empirical findings from Section 2.

5.1 Functional Forms and Parametrization

The aim of our functional-form assumptions and parameterizations is to ensure that our model quantitatively replicates key features of the US macroeconomy and firm-level outcomes, and captures the rich heterogeneity in emission efficiencies documented in the data.

Functional Forms: To this end, we retain the specifications of the production technology and the abatement cost curve used in Section 4.5: we let $F(n,k,e) = n^{\alpha} \cdot k^{\gamma} \cdot e^{\upsilon}$, where $\alpha, \gamma, \upsilon > 0$ and $\alpha + \gamma + \nu < 1$, and set $Q(x) = (1+\theta)^{-1} \cdot x^{1+\theta}$ with $\theta > 0$. The per-period utility function is of the logarithmic form, $u(C_t) = \log(C_t)$. Similar to Golosov *et al.* (2014), we assume the aggregate damage function for emissions is $A(G) = \exp^{-\phi \cdot G}$, where ϕ governs the semi-elasticity of the damage function, central to our analysis,and thus the implied aggregate productivity losses from carbon emissions. Finally, we assume that firm productivities and carbon intensities are drawn from normal distributions, $a_i \sim \mathcal{N}(\mu_a, \sigma_a^2)$ and $z_j \sim \mathcal{N}(\mu_z, \sigma_z^2)$, and allow for a nonzero correlation ρ_{az} between the two types. We will later truncate these distributions, as well as study changes in the variances of the distributions (σ_a^2 and σ_z^2), to analyze how changes in the extent of firm-level heterogeneity affect our main results.

Parametrization: We assume that a model period corresponds to one year. We choose standard parameters for the discount factor ($\beta = 0.96$) and the energy share ($\upsilon = 0.08$). The

latter being, for example, consistent with the data reported by the U.S. Energy Information Administration (2024). Assuming decreasing returns to scale of $\alpha + \gamma + \nu = 0.8$—well within the range use in the literature e.g., Basu and Kimball (1997)—and a standard split of GDP between capital (0.4) and labor (0.6), we obtain a production elasticity for capital of $\alpha = 0.29$ and for labor of $\gamma = 0.43$.[20] We assume a 10% depreciation rate for capital ($\delta_K = 0.10$) and normalize average productivity e^{μ_a} to 1 (i.e., $\mu_a = 0$). We set the standard deviation of productivity, σ_a to 0.6, consistent with Compustat and Bloom *et al.* (2018), among others.

The remaining parameters describe the climate portion of the economy. The semi-elasticity of the damage function, ϕ, and the mean carbon efficiency, μ_z, both affect the overall cost of climate change, as they govern how one unit of energy translates into productivity losses. Consequently, we normalize the average carbon intensity e^{μ_z} to 1 (i.e., set $\mu_z = 0$) and calibrate $\phi = 3.5 \cdot 10^{-4}$ to yield a 2% long-run aggregate productivity loss; the latter being consistent with the standard IAM practice (e.g., Nordhaus and Moffat, 2017) and substantially more conservative than recent estimates (Bilal and Känzig, 2024; Caggese *et al.*, 2025). Section 5.4 studies how our results depend on this calibration choice. We set the standard deviation of carbon efficiencies, σ_z, to 1.6, consistent with the ICE estimates in Panel (b) in Figure 2,[21] and set the carbon depreciation rate, δ_G, to 0.02; the latter being close to common IPCC (2021) estimates. Section 5.4 shows that using more conservative values does not affect our results. Finally, we set the inverse elasticity of abatement to carbon prices, θ, to 2, in line with the estimates in Känzig (2023) on the effects of carbon tax shocks on green technologies. We also assume that firm-level productivity and carbon intensity are uncorrelated ($\rho_{az} = 0$), based on the evidence discussed in Section 2 showing that productivity has little predictive power for emission intensity. Table C.1 in Appendix C summarizes the parameters.

5.2 Quantification and Decomposition

We quantify the economy-wide consequences of firms' reporting incentive. To do so, we compare results based on the Mirrleesian optimal tax schedule with those from the naive tax schedule, which does not internalize firms' reporting incentives (Proposition 1 and Figure 7). We also here compare estimates to those from the full-information case, in which the planner has full-information about firms' carbon intensities. The latter, namely, provides a quantification of the first-best benefits from carbon taxation. To start with, we focus on

[20] These numbers convert GDP shares to gross-output shares and account for decreasing returns to scale: $\alpha = 0.4 \cdot 0.8 \cdot (1 - \nu)$ and $\gamma = 0.6 \cdot 0.8 \cdot (1 - \nu)$, with $\nu = 0.08$.

[21] In practice, we truncate the normal distributions of carbon intensities (and productivities) at five standard deviations. This balances the desire to approximate the normal distribution well with the need to ensure that $\min z_j$ is well defined and finite, and lies well within the emission intensities reported by the firms as shown in Figure 2.

Table I: Quantitative Effects of Carbon Taxation

	$\mathcal{Y}$	C	G	E	X	K
Naive planner	+0.09	+0.09	−2.96	+0.09	+0.05	+0.09
Mirrleesian planner	+1.47	+1.78	−59.74	−2.01	+4.63	+1.47
Full-informed planner	+2.00	+2.20	−72.22	−0.41	+3.22	+2.00

Note. All numbers are reported in %-deviation from the decentralized equilibrium and refer to steady-state comparisons. The decentralized levels are: $\mathcal{Y} = 4.00$, $C = 2.87$, $G = 57.65$, $E = 0.32$, $X = 0$, and $K = 8.15$.

steady-state comparisons. Section 5.4 shows that our results carry over when accounting for the full transition dynamics of the model. In order to compute estimates under the naive tax schedule, we need to take a stance on the maximum dispersion in carbon intensities, as this controls the scope of misreporting. We truncate the distribution of emission intensities at the 5 standard-deviation level. This balances the desire to approximate the distribution well with the need to ensure that $\min z_j$ is well defined. Section 5.4 shows that our results extend to other assumptions about the extent of misreporting. Table I presents the main results.

Several points are worth highlighting. First, consistent with our theoretical results, the full-information planner delivers the highest welfare benefit, measured by consumption increases, and the largest emission reduction (+2.20% and -72.22%, respectively). This is followed by the Mirrleesian planner (+1.78% and -59.74%), with the naive planner substantially further behind (+0.09% and -2.96%). Overall, the potential benefits from either Mirrleesian or full-information taxation are large—close to 2% of steady-state consumption and around a 2/3 decline in emissions. The consumption benefits are close to size of the overall decline in productivity due to climate change. Second, while the Mirrleesian planner captures the lion's share of the potential, first-best welfare and environmental gains (around 80% of them), the naive planner captures virtually none of them (around 3-4%). Combined, our quantitative results thus suggest that Mirrleesian taxation can close-to replicate the benefits of first-best carbon taxation. Accounting for firms' reporting incentives does not substantially constrain the potential welfare and environmental benefits of carbon taxation. By contrast, however, ignoring firms' reporting incentives does render the standard Pigouvian approach largely ineffective at reducing emissions, with substantial consumption losses as a result.

In the next two subsections, we unpack these findings: we first analyze the sources of welfare benefits from carbon taxation in Table I, and then examine the drivers of differences in emission reductions. Our aim is to both discern what causes the above benefits of carbon taxation—abatement or the reallocation of energy—as well as to analyze the circumstances that drive Mirrleesian taxation to substantially outperform the naive case.

Table II: Welfare Decomposition of Consumption Changes

	Ω	$\mathcal{O}$	$\mathcal{A}$	$\mathcal{K}^{\alpha}$	$\mathcal{E}^{\nu}$	C
Naive planner	+0.00	+0.00	+0.06	+0.03	+0.00	+0.09
Mirrleesian planner	−0.00	+0.30	+1.21	+0.42	−0.16	+1.78
Full-informed planner	−0.01	+0.20	+1.47	+0.57	−0.03	+2.20

Note. All numbers are reported in percent deviation from the decentralized equilibrium steady state. The decentralized equilibrium levels are: $\Omega = 1$, $\Theta = 0.71$, $K = 8.15$, $E = 0.32$, $C = 2.87$.

Welfare Decomposition: We start by breaking down the causes of the large differences in welfare, measured by consumption increases, in Table I. To do so, we exploit the following decomposition of consumption, building on our result in Proposition 2:

$$\frac{C_t}{C_t^{de}} = \Omega_t \cdot \underbrace{\left(\frac{\Theta_t}{1-\nu-\frac{\alpha \cdot I_t}{R_t \cdot K_t}}\right)}_{\equiv \mathcal{O}_t} \cdot \underbrace{\left(\frac{A(G_t)}{A(G_t^{de})}\right)}_{\equiv \mathcal{A}_t} \cdot \underbrace{\left(\frac{K_t}{K_t^{de}}\right)^{\alpha}}_{\equiv \mathcal{K}_t^{\alpha}} \cdot \underbrace{\left(\frac{E_t}{E_t^{de}}\right)^{\nu}}_{\equiv \mathcal{E}_t^{\nu}}. \tag{30}$$

Equation (30) shows that consumption differences between any of the planner's allocation and the decentralized equilibrium can be decomposed into five components: (i) factor misallocation, Ω; (ii) differences in the intermediate input share, $\mathcal{O}_t$, (iii) differences in productivity damages, $\mathcal{A}_t$; (iv) differences in capital use, $\mathcal{K}_t$; and lastly (v) differences in aggregate energy use, $\mathcal{E}_t$. Table II presents the decomposition, focusing on the steady state of the economy.

First, notice that the bulk of the consumption increase achievable by the Mirrleesian and full-information planner are driven by productivity increases, which in turn are caused by declines in aggregate emissions. The use of additional capital, as it becomes more productive, contributes some to the consumption increase, yet much less than increase productivity. Indeed, productivity increases drive close to 2/3-3/4 of the overall consumption response. Second, despite the increase in productivity and the accompanying decline in emissions, energy use declines by less than 1/4 of a percent. The Mirrleesian and full-information allocations combine large declines in emissions with small declines in energy use, which is why these allocations can create large consumption increases. The next subsection investigates the sources behind the large reduction in emissions. Finally, recall that the Mirrleesian and full-information planners introduce misallocation into the economy. The marginal revenue products of energy is not equated to its private cost. Table II, however, shows that the quantitative costs of this misallocation are small—less than 1bp. The benefits of productivity increases are large and the costs of misallocation small. Table II further shows that the planner allocations also tend

to make the economy more intermediate-input intensive: the small suppression of energy use is counteracted by the higher expenditure on abatement.

In sum, the results in Table II demonstrate that the main quantitative driver behind the documented consumption increases are the large declines in emissions that we saw in Table I. These, in turn, drive sizable increases in productivity and, to some extent, are amplified by the increased use of capital. The misallocation introduced by the planners' interventions is, by contrast, quantitatively small, and so too is the decline in overall energy use. That said, Table II also raises the question of how the full-information and Mirrleesian planners combine a *large decrease* in emissions with a *small decline* in energy. The combination is here the catalyst behind the benefits from climate taxation. It is to this question that we turn to next.

Emission Decomposition: Conceptually, emissions can decline for three reasons: because of declines in overall energy use, because of increased overall abatement, or because of a better allocation of energy and abatement to high emission firms. Using the law-of-motion for emissions in (4), the following equation breaks-down changes in aggregate emissions relative to the decentralized equilibrium into these three components (Appendix D):

$$\frac{G_t}{G_t^{de}} = \underbrace{\underbrace{\left(\frac{\sum_{ij}\varpi_{ijt}\cdot\exp^{z_j-x_{ijt}+X_t}}{\sum_{ij}\varpi_{ijt}\cdot\exp^{z_j}}\right)}_{\equiv\mathcal{Z}_x}\cdot\underbrace{\left(\frac{\sum_{ij}\varpi_{ijt}\cdot\exp^{z_j}}{\sum_{ij}\varpi^{de}_{ijt}\cdot\exp^{z_j}}\right)}_{\equiv\mathcal{Z}_e}}_{\equiv\mathcal{Z}}\cdot\underbrace{\left(\exp^{-X_t}\right)}_{\equiv\mathcal{X}}\cdot\underbrace{\left(\frac{E_t}{E_t^{de}}\right)}_{\equiv\mathcal{E}}, \tag{31}$$

where $\varpi_{ijt} \equiv \omega_{ij}\cdot\frac{e_{ijt}}{E_t}$. Equation (31) shows that emission differences between any of the planner's allocations and the decentralized equilibrium can be decomposed into: (i) changes in an allocative term, $\mathcal{Z}$; (ii) changes in an aggregate abatement term, $\mathcal{X}$; and (iii) changes in aggregate energy use, $\mathcal{E}$. We have already seen that the third term is not the reason behind decline in emissions. The allocative term $\mathcal{Z}$ can then be further decomposed into two different components: $\mathcal{Z}_x$ and $\mathcal{Z}_e$, where $\mathcal{Z}_e$ captures the role of *energy reallocation* to low-z firms, while $\mathcal{Z}_x$ captures the effect of allocating a higher level of abatement—relative to the aggregate level—to high-energy and/or high-z firms. Table III presents the results.

The table shows that the *abatement component*, which captures the allocation of abatement across firms, is by far the dominant force behind the estimated emissions reductions. In the full-information and Mirrleesian cases, it contributes 82-90% of the overall emission decrease. The *energy-reallocation component*, by contrast, is 0 both in the Mirrleesian and naive cases, as these allocations are bunched across carbon-intensity types, like in the decentralized equilibrium. As a result, there is no between-firm energy reallocation effect. Conversely, in the full-information case, the social planner does reallocate energy toward low-z firms. Table

Table III: Decomposition of Changes in Aggregate Emissions

	$\mathcal{Z}_x$	$\mathcal{Z}_e$	$\mathcal{X}$	$\mathcal{E}$	G
Naive planner	-3.00	$+0.00$	-0.05	$+0.09$	-2.96
Mirrleesian planner	-53.59	$+0.00$	-4.25	-1.89	-59.74
Full-informed planner	-58.66	-11.98	-3.16	-0.42	-72.22

Note. All numbers are reported in percent deviation form the decentralized equilibrium steady state. The decentralized equilibrium levels are: $\sum_{ij} \varpi_{ijt}^{de} \cdot \exp^{z_j} = 3.60$, $X = 0$, and $E = 0.32$.

III shows that this effect is quantitatively meaningful, although not nearly the size of the abatement component, and accounts for most of the difference between the Mirrleesian and the full-information cases. Finally, we note that the rise in aggregate abatement is meaningful, albeit not a driving quantitative force in the decline in emissions.

We conclude that, in both the Mirrleesian and the full-information (first-best) case, most of the welfare and environmental benefits from carbon taxation derive from $\mathcal{Z}_x$— that is, from the planner incentivizing abatement by high-energy and/or high-carbon-intensity firms. By contrast, most of the differences in welfare and environmental outcomes between the Mirrleesian and the full-information cases stem from $\mathcal{Z}_e$ –that is, from the additional ability of the full-information taxes to reallocate energy to low-carbon-intensity firms.

To understand what drives the overall size of the *abatement component*, $\mathcal{Z}_x$, we exploit the following characterization (Appendix D):

$$\log \mathcal{Z}_x \approx -\mathbb{C}ov\left(\frac{g_{ijt}}{G_t}, \frac{\exp^{x_{ijt}}}{\exp^{X_t}}\right) < 0. \tag{32}$$

This result shows that the abatement component is large when relative emissions, g_{ijt}/G_t, co-move strongly with relative abatement efforts across firms, $\exp^{x_{ijt}} / \exp^{X_t}$, which is indeed the case under both the full-information and the Mirrleesian allocation. The first-order conditions for abatement (Equations 19 and 24) directly show that:

$$\left(x_{ijt}^{fb}\right)^{\theta} = \Lambda_{G,t} \cdot g_{ijt}^{fb} \quad \text{and} \quad \left(x_{ijt}^{mr}\right)^{\theta} = \Lambda_{G,t} \cdot g_{ijt}^{mr} \cdot \frac{\mathbb{E}[\exp^{z_j}]}{\exp^{z_j}}. \tag{33}$$

Thus, abatement is a direct function of emissions in both planner allocations. Combined with the fact that $\mathbb{E}[\exp^{z_j}]/\exp^{z_j}$ is, on average, small for our calibrated distribution of z_j—the implied Jensen inequality term is small—this explains the high correlation between relative emissions and relative abatement in the two cases, and hence the large contribution of the abatement component to the overall reduction in emissions that we find. Both the full-

information planner and its Mirrleesian counterpart enforce a strong correlation at the firm level between emissions and abatement. AppendixD documents the allocation of emissions and abatement efforts across productivities and carbon intensity-types.

Although our findings clarify the forces behind the reduction in aggregate emissions, it does not explain why the Mirrleesian planner—despite the incentive-compatibility constraint—performs nearly as well as the full-information planner. As explained above, the difference between her and the full-information planner is entirely accounted for by the *energy-reallocation component*, $\mathcal{Z}_e$. To address this issue, we examine what determines the strength of the energy-reallocation effect. A few simple derivations show that:

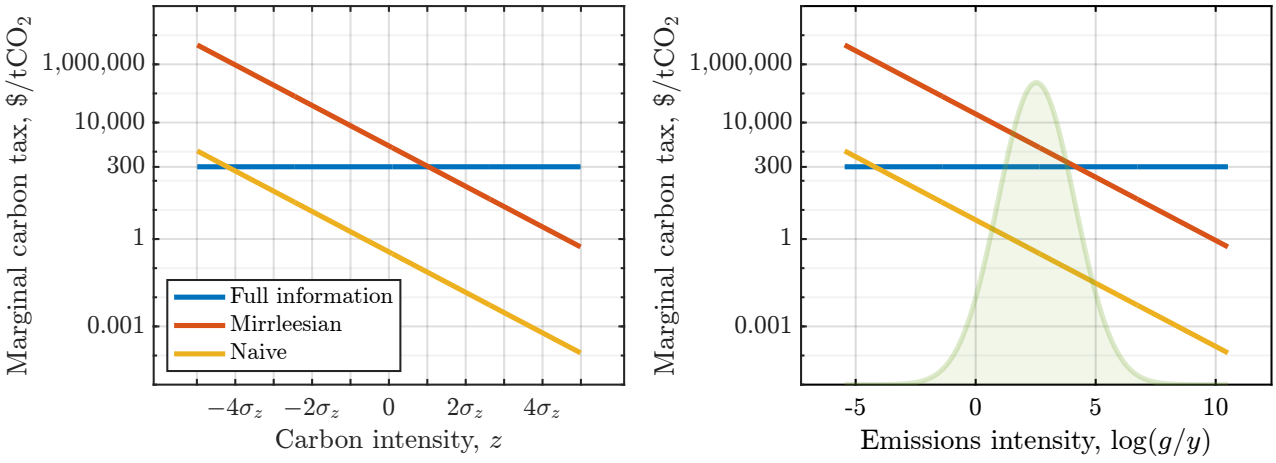
$$\mathcal{Z}_e^{fb} - \mathcal{Z}_e^{mr} = \mathbb{C}ov\left(\frac{e_{ijt}^{fb}}{E_t^{fb}},\ \frac{\exp^{z_j}}{\mathbb{E}\left[\exp^{z_j}\right]}\right) < 0. \tag{34}$$

This equation decomposes differences in $\mathcal{Z}_e$ between the two planner allocations into the covariance between firm-level energy shares, e_{ijt}^{fb}/E_t^{fb}, in the full-information allocation and the relative carbon-efficiency type of the firm, $\exp^{z_j}/\mathbb{E}_\omega\left[\exp^{z_j}\right]$. Notice that this covariance is always negative, as firm-level energy shares comove negatively with relative carbon intensities (Lemma 2). As a result, the gap between $\mathcal{Z}_e^{fb}$ and $\mathcal{Z}_e^{mr}$ is, all else equal, larger when heterogeneity in energy use under the full-information allocation is mainly driven by carbon intensity differences rather than productivity differences. By contrast, when variation in energy shares is mainly due to productivity differences, the covariance in Equation (34) is close to zero. The Mirrleesian planner, in this case, performs nearly as well as the full-information planner.

Based on this decomposition, we conclude that the reason the Mirrleesian planner is able to reap most of the benefits from carbon taxation is that—in the first-best allocation—the productivity level of a firm is *quantitatively* an important driver of its energy use. Although carbon intensities clearly matter, under our baseline calibration, productivities are also a quantitatively important determinant of firms' optimal energy use, highlighting the importance of considering both sources of heterogeneity in the analysis. Appendix D shows graphically the similarities in energy use in the Mirrleesian and full-information cases.

In summary, in this section, we have demonstrated that disregarding firms' reporting incentives renders the standard Pigouvian approach to carbon taxation largely ineffective in reducing emissions—with substantial consumption losses as a result. By contrast, accounting for firms' reporting incentives does not substantially constrain the potential welfare and environmental benefits of carbon taxation. The Mirrleesian optimal tax attains approximately 80% of the potential benefits. The root cause behind the comparative performance of the Mirrleesian approach is that the first-best carbon tax attains its welfare improvement largely by reallocating abatement efforts—a task the Mirrleesian approach does almost equally well

Figure 7: Marginal Carbon Tax Rate



Note: This figure shows the marginal tax rates as a function of carbon (left-hand side) and emission intensities (right-hand side). The right-hand side panel further depicts the distribution of firm emission intensities in the decentralized equilibrium. We use 2024 US GDP numbers to convert model estimates into USD ($) values.

when productivity optimally determines a large fraction of the variation in energy uses.

5.3 Carbon Taxes in Practice

We next turn to the implied carbon taxes in the three cases—the first-best, full-information case, the Mirrleesian case, and the naive case—that underlie these results. Common estimates of the Piguvian tax on carbon—equivalent to that which implements the first-best, full-information allocation—is around $280 per ton of CO_2 equivalent emissions (tCO_2e).[22] Figure 7 shows that our baseline calibration produces average estimates that are close to this value. The left-hand side panel depicts the marginal tax rate by carbon-intensity type implied by Proposition 1, while the right-hand side panel shows the marginal tax rate over observable emission intensities (Section 2). The marginal tax that implements the full-information allocation is, in both cases, around $300 per tCO_2e and constant across the intensity measures, corroborating the soundness of our baseline calibration strategy.

In line with Proposition 1, the optimal Mirrleesian tax is log-linear in carbon-intensity type, z, and monotonically declining. The left-hand side panel shows that for a firm that is 2 standard deviations *below* the mean carbon-intensity type, its value is around $40,000 per tCO_2e; for a firm that is 2 standard deviations *above*, it is only $65 per tCO_2e. At the ± 1 standard deviation level, the range is between $7,000-300 per tCO_2e. There is substantial heterogeneity in the Mirrleesian optimal marginal tax. This heterogeneity carries over to the case in which we consider the Mirrleesian tax as a function of emission intensities. The

[22] See, for example, Ricke *et al.* (2018), Pindyck (2019), and Tol (2023), among others.

full-information marginal tax is, by contrast, constant across the two panels.

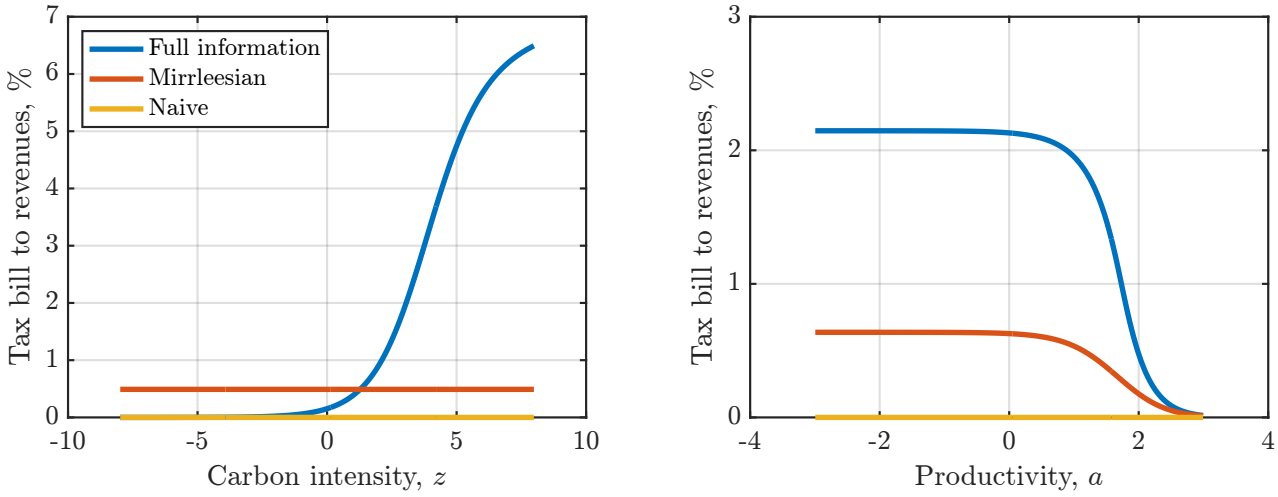
A fear one might have is that by enacting high marginal tax rates on carbon-efficient firms—such as those illustrated in Figure 7—one may heavily penalize the profits of carbon-efficient firms, potentially to the detriment of their future innovations. Figure 8 reports the overall tax bill spent on carbon taxes as a share of a firm's revenue across firms with different carbon intensities (left-hand panel) and productivities (right-hand panel). Although the marginal carbon tax can be high for an individual firm, the overall tax bill in the Mirrleesian case stays well below 1% of revenue. The marginal taxes can be high, but the average value paid is low. Clearly, changes to the incentive structure of firms are the driver of our main results.

The left-hand side panel of Figure 7 further shows a fundamental difference between the full-information planner and the Mirrleesian planner. The full information planner, unconstrained by incentive-compatibility considerations, sets carbon taxes so that the tax bill as a share of revenue rises steeply with a firm's carbon intensity. This highlights two features: (i) the planner penalizes heavily more carbon-intensive firms, as this is the most cost-effective way to reduce emissions; and (ii) this, in turn, creates a strong incentive for carbon-intensive firms to misreport their carbon-intensity type. By contrast, the Mirrleesian planner internalizes this incentive and, to eliminate it, charges all firms the same tax-bill share of revenue.

Finally, the right-hand side panel of Figure 7 depicts the difference between the full-information tax bill and the Mirrleesian tax bill along the productivity dimension. Although the marginal carbon tax is independent of productivity, the tax bill's share of revenue declines with productivity. In the presence of carbon taxes, more productive firms emit less per unit of revenue, as they have the strongest incentives to abate. As commented on earlier, the marginal benefit from abatement scales with firm size (i.e., productivity). Consequently, both planners find it most cost-effective to shift the burden from high- to low-productivity firms. This is, in part, related to our discussion in the previous section on the reason behind the comparative performance of the Mirrleesian planner.

We conclude that the Mirrleesian tax schedule amends the full-information tax schedule by making the marginal carbon tax decline strongly in a firm's carbon-intensity type. The marginal tax for most firms declines from around \$40,000 to \$ 65 per tCO_2e, and is on average above standard estimates of the optimal carbon tax (approx. $\$1,500$ per tCO_2e). The average tax paid by firms is, however, small, less than 0.75 % of revenue, on average. The Mirrleesian, thus, combines strong incentives with a low average cost to firms.

Figure 8: Tax Bill Share of Revenues



Note: The figure shows the overall carbon tax bill paid in the different allocations compared to firm revenue.

5.4 Quantitative Refinements

Transition Dynamics: A potential concern with the above findings is that they present the truly *long-run benefits* of climate taxation. We have thus far detailed results by comparing the steady states of the different planner allocations with the steady state of the decentralized equilibrium. To explore the benefits of carbon taxation along the *transition*—where by transition we mean over the period from the carbon tax implementation until the economy reaches its new steady state—we extend our baseline economy to allow for a richer description of the dynamic evolution of the climate and consider the full transition path of the economy.

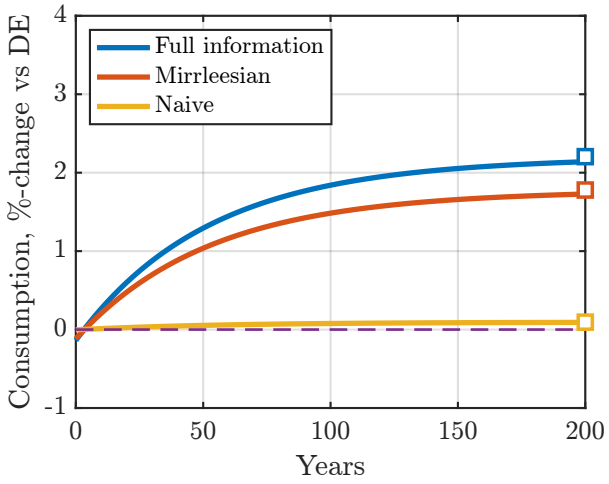
We follow Miftakhova *et al.* (2020) and consider an alternative climate block that has become increasingly popular in both the natural sciences (e.g., Allen *et al.*, 2009; Matthews *et al.*, 2009, 2012) and in macroeconomics (e.g., Dietz *et al.*, 2021; Fernández-Villaverde *et al.*, 2024). We assume that economy-wide productivity depends on temperature, $A = \exp^{-\phi \cdot T}$, rather than directly on emissions, and that temperature, T, evolves dynamically with the stock of accumulated emissions in the atmosphere in accordance with:

$$T_{t+1} = (1 - \xi_T) \cdot T_t + \xi_T \cdot (\mathcal{T} + \Xi_G \cdot G_t), \tag{35}$$

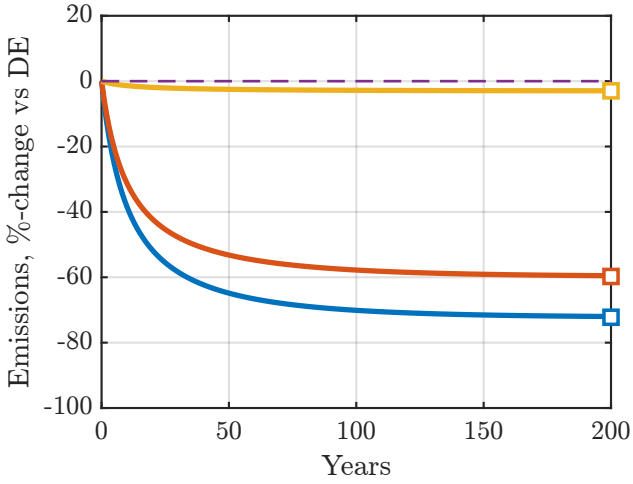
where $1 - \xi_T$ is the persistence in global temperatures, $\mathcal{T}$ is its historical mean, and Ξ_G is the conversion factor mapping emissions into temperature. Consistent with the historical global average temperature, we set $\mathcal{T} = 14°\text{C}$, and set $\Xi_G = 7.34 \cdot 10^{-3}$ to generate a 2°C increase in temperature in the decentralized equilibrium within the first 100 years. The latter is in line with standard "business-as-usual" scenarios. We then recalibrate $\phi = 1.3 \cdot 10^{-3}$ to match a 2%

Figure 9: Transition Dynamics of Consumption and Emissions

Panel (a): Baseline Model

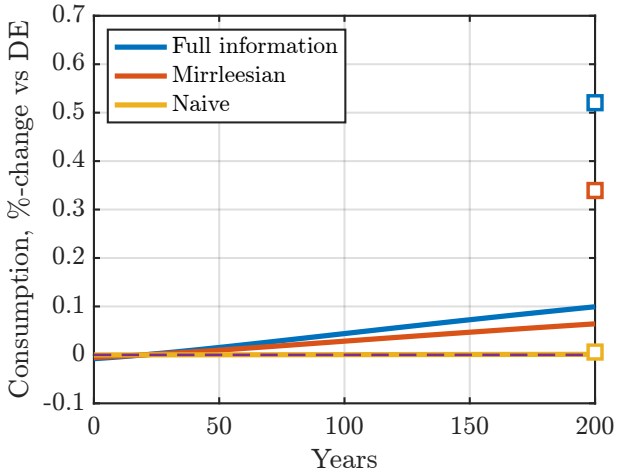


Panel (a.1): Consumption

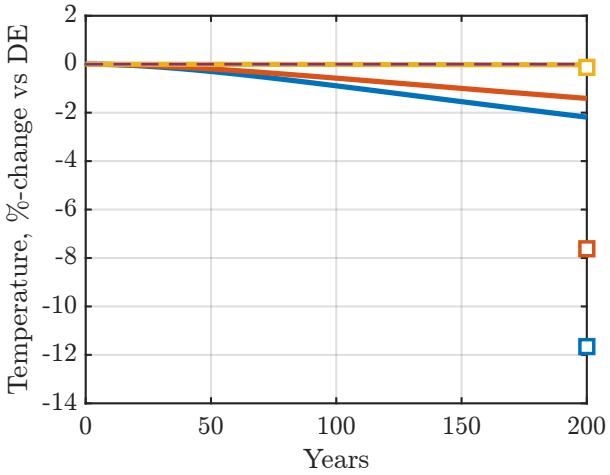


Panel (a.2): Emissions

Panel (b): Alternative Climate Model



Panel (b.1): Consumption



Panel (b.2): Temperature

*Note.*The figure shows the transition dynamics for our baseline model (Panel a) and the extended climate model (Panel b). Appendix D describes the solution for the transition path of the extended model.

aggregate productivity loss from a 2°C increase in temperature relative to the historical mean, consistent with our baseline calibration. We calibrate temperature persistence $1 - \xi = 0.98$, in line with the estimates in Bruns *et al.* (2020). Finally, we set $\delta_G = 1 \cdot 10^{-3}$ to approximate a (near) unit-root emissions process, as is often assumed in the natural sciences.

Figure 9 shows the dynamics of consumption and climate along the transition. We detail these for both our baseline economy (Panel a) and for the economy with the extended climate

block (Panel b). In our baseline economy, as in the steady-state analysis, the Mirrleesian planner allocations follow closely the full-information outcomes.[23] Relative to the decentralized equilibrium, the full-information planner achieves a net-present value consumption-equivalent welfare gain of 0.59% and a 51.18% reduction in the net-present value of emissions.[24] The Mirrleesian planner, by contrast, delivers a 0.47% net-present value welfare gain and a 41.88% reduction in the net-present value of emissions. Thus, the Mirrleesian planner once more captures the lion's share of the full-information gains (around 80%). In contrast, the naive planner yields only a small welfare gain and a 1.93% reduction in emissions. Overall, our results along the transition—although with net-present value level effects for each planner somewhat muted by the cumulative nature of climate change—are consistent with the main insights from our steady-state analysis of the baseline model.

Panel (b) shows that these insights also carry over to the version with the extended climate block.[25] The Mirrleesian planner once more captures the bulk of the possible full-information benefits from climate taxation (around 65% in net-present value terms), while naive taxation results in negligible welfare and environmental benefits. The striking difference between the two climate models, however, manifests in the horizon over which the effects take place. In the baseline model, most of the steady-state effects materialize within 100 years. By contrast, the effects are both smaller (even in the long-run) and take substantially longer to materialize in the extended model. As such, in the extended model the net-present value benefits of climate taxes are much smaller. Relative to the decentralized equilibrium, the full-information planner achieves a net-present value welfare gain of $3 \cdot 10^{-3}\%$ and a -0.13% reduction in net-present value of temperature; in the long run, it achieves a 0.52% welfare increase and a -11.66% temperature reduction. The Mirrleesian planner instead achieves a net-present value welfare gain of $2 \cdot 10^{-3}\%$ and a -0.08% reduction in temperature. In the long run it achieves a 0.34% consumption increase and a -7.62% temperature reduction.

On balance, our results thus confirm the *comparative* insights from the baseline model: the Mirrleesian planner recovers most of the benefits from climate taxation. The *absolute* magnitudes of those benefits yet rest crucially on the exact characterization of the climate process used—of which there is substantial debate in the natural sciences—and the horizon

[23] For this experiment we need to set discipline on two additional objects necessary to perform this exercise: the starting stock of capital and emissions $\{K_0, G_0\}$. To this end, we calibrate K_0 as representing the steady state level of capital for an economy with a stock of emission G_0, and we set G_0 such that most of the effects of climate change unfold within the next 100 years, as in most common scenarios. However, we find that the initial point does not fundamentally drive our results.

[24] We here use the household discount factor, β, to discount consumption and emission flows.

[25] As we show in Appendix **??** in this model we need to set only two initial conditions, i.e., $\{K_0, T_0\}$, as the temperature emission system cha be reduced to a single AR(2) temperature process. We set K_0 as representing the steady state level of capital for an economy with temperature T_0 and set $T_0 = 15°\text{C}$, which approximates presents days where temperature has already increased by approximately 1°C above historical mean.

that one whishes to study: dynamics and the considerations of short- versus long-run matters. Our conclusions are, in this sense, in line with the findings of much of the literature that studies the macroeconomic consequences of climate change (e.g., Nordhaus, 2019).

Climate Parameter Uncertainty: Our baseline model assesses the consequences of Mirrleesian climate taxation using estimates of the environmental damages to production, ϕ, and the inverse elasticity of abatement, θ. As with the processes linking temperature and emissions, both parameters are, however, subject to considerable parameter uncertainty. Figure 10 explores the robustness of our results to variations in these parameters.

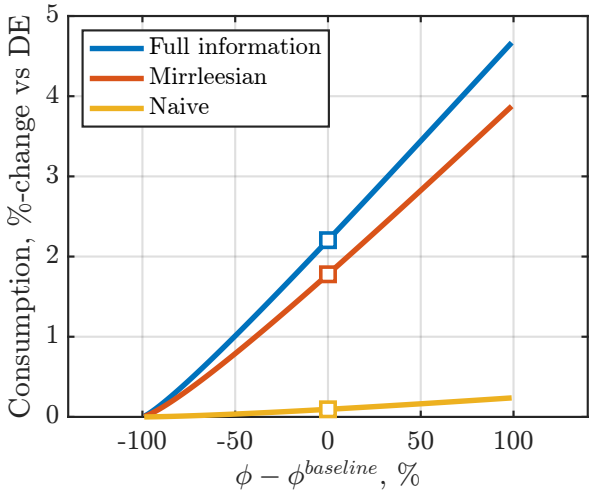
Our baseline calibration conservatively sets ϕ to deliver a 2% long-run productivity loss in the decentralized equilibrium. Although this value is consistent with traditional estimates of the cost of climate change (e.g., Nordhaus, 2007), some modern estimates, however, project substantially larger losses (e.g, Bilal and Känzig, 2024, Caggese *et al.*, 2025). Panel (a) in in Figure 10 shows that—although the benefits of both full-information and Mirrleesian taxation increase substantially as we push ϕ closer to the higher losses estimated in recent approaches—the overall take-away remains: Mirrleesian taxation recovers a large share of the potential welfare benefits, while naive taxation does not. Crucially, the comparative performance of Mirrleesian taxation does not rest on the assumed cost of climate change.

Our conclusions likewise do not hinge on the assumed value for θ. Our baseline calibration of the inverse tax elasticity of abatement is informed by Känzig (2023) estimate of the response of green technologies to carbon tax shocks. While not an ideal measure of abatement behavior, it provides a useful benchmark given the scarcity of precise estimates. Panel (b) of Figure 10 shows that our results are relatively insensitive to the choice of θ. Two offsetting forces drive this result. On one hand, a larger θ reduces the abatement response to carbon taxes; but, on the other hand, it also lowers the overall cost of abatement, $Q(x) = \frac{1}{1+\theta} \cdot x^{1+\theta}$. These effects largely cancel each other in equilibrium by moving in opposite directions, so that the ultimate benefits of climate taxation do not depend much on the exact value of θ. We view this fact as a strength of our framework, especially given the limited evidence on θ in the literature.

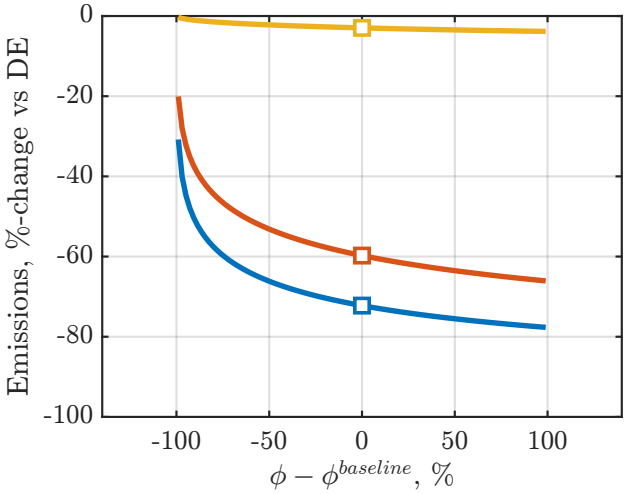
Renewable versus Fossil Energy: One reason for why firms, even within narrowly-defined sectors, may exhibit differences in emission intensities is the availability of renewable versus fossil energy. Other technological differences—such as differences in filters used and burner designs, mentioned in Section 3—lead to differences in emission intensities, in addition to which arises from differences in the use of different fossil fuels. But, clearly, one reason behind the differences in emission intensities documented in Section 2 is the firm-level availability of renewable versus fossil energy. To explore the implications of the introduction of reneweable

Figure 10: Robustness to Environmental Damages and Abatement

Panel (a): Environmental Damages ϕ

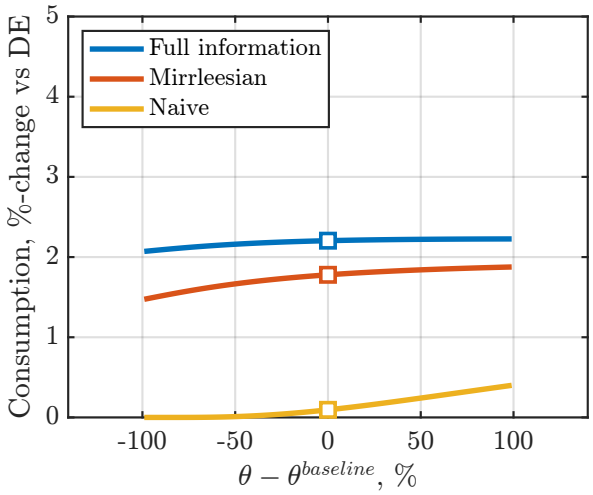


Panel (a.1): Consumption

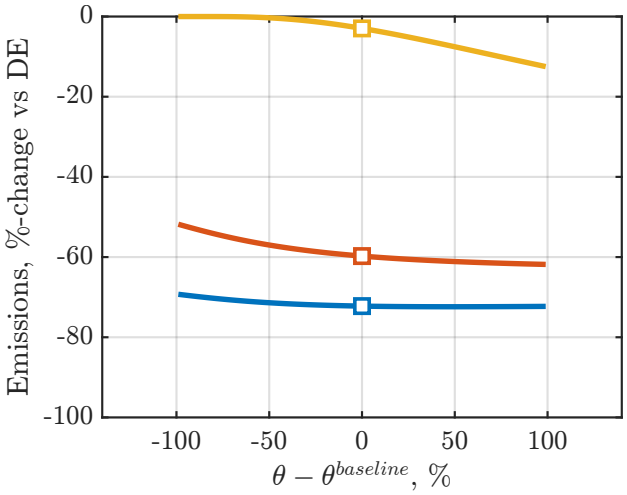


Panel (a.2): Emissions

Panel (b): Abatement Elasticity θ



Panel (b.1): Consumption



Panel (b.2): Emissions

Note: This figure shows the steady-state consequences of climate change for consumption and emissions for different parameter configurations for θ and ϕ. We consider $+/-$ 100% variations in either parameter.

energy sector into our model, we extend our baseline framework by assuming that energy, e_{ijt}, is comprised of renewable, $e_{\text{green},ijt}$, and fossil energy, $e_{\text{brown},ijt}$ in accordance with:

$$e_{ijt} = \left(e_{\text{green},ijt}^{\frac{\varepsilon-1}{\varepsilon}} + e_{\text{brown},ijt}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{36}$$

Table IV: Extension with Renweables and Fossil Fuels

	$\mathcal{Y}$	C	G	A	K	X
Naive planner	+0.00	+0.01	−0.04	+0.02	+0.02	+0.01
Mirrleesian planner	+1.44	+1.39	-47.17	+1.00	+1.45	+1.38
Full-information planner	+2.35	+2.40	-77.08	+1.64	+2.35	+1.98

Note: All numbers are reported in %-deviation from the decentralized equilibrium and refer to steady-state comparisons. The decentralized levels are: $\mathcal{Y} = 4.06$, $C = 2.91$, $G = 57.65$, $E = 0.32$, $X = 0$, and $K = 8.15$.

where both fossil and reneweable energy are produced as in the baseline model, with P being the price of fossil energy and $(1 + \tau_{\text{green}}) \cdot P$ the price of renewable energy with $\tau_{\text{green}} > 0$. We assume that only the use of fossil energy creates emissions in Equation (5). Although conceptually potentially problematic, the CES structure in (36) has proven to be a popular and tractable means to model firms' choice of renewable and fossil energy (e.g, Hassler *et al.*, 2016). We calibrate the elasticity of substitution, ϵ, to 4, within but close to the top of common empirical estimates (e.g., Papageorgiou *et al.*, 2017; Jo, 2025; Petersen, 2025), to account for the steady-state nature of our analysis. We also set the rewnweable price wedge, τ_{green}, such that 1-in-4 energy $ is spent in the decentralized steady state on renewable energy, consistent with the EIA's 2024 estimate for the US.[26] Table IV shows the steady-state results.

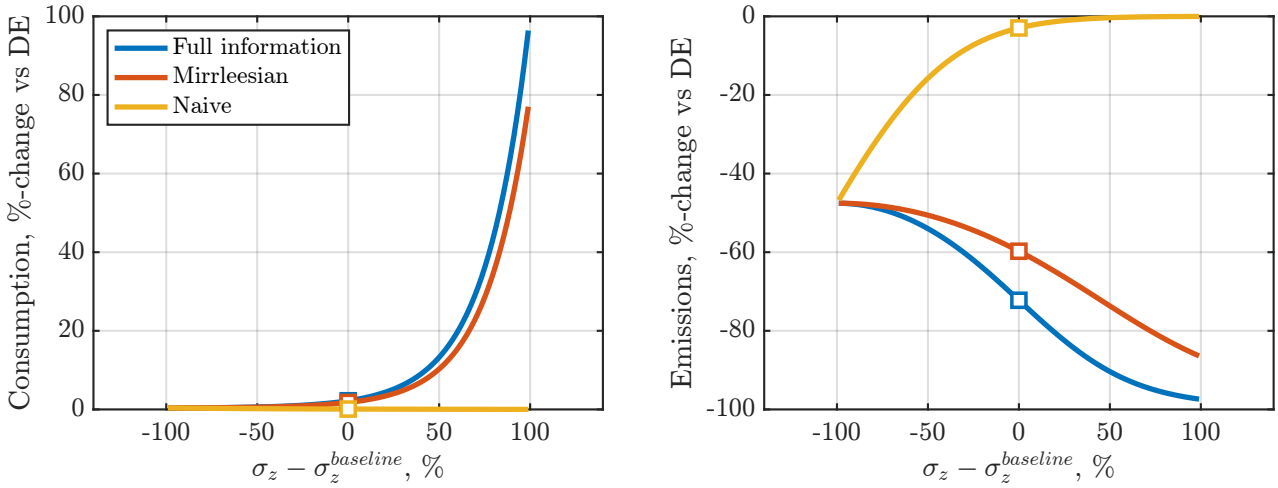
The results in Table IV, on balance, show that the insights from the baseline model extend to the case with a renewable sector. The Mirrleesian tax once more recovers around 2/3 of the possible decline in emissions—and although it captures somewhat less of the potential rise in consumption and welfare than before (slightly more than 1/2), the Mirrlesian tax still drastically outperforms the naive benchmark. Compared to the full-information, first-best allocation, the Mirrlesian tax does not create as forceful as switch from fossil energy to renewable energy, as should be expected given its limitations. But, because the reallocation of abatement efforts is still potent force under the Mirrleesian tax, the Mirrleesian planner still manages to recover most of the environmental and welfare benefits from carbon taxation. We conclude that our results do not hinge on the absence of renewable sector.

6 Final Remarks

Addressing the economic damages from climate change is, to a first-order approximation, a problem of curtailing firms' direct carbon emissions (Section 2). In response to this challenge, considerable efforts have recently gone into the design and implementation of (Pigouvian)

[26] See: www.reuters.com/business/energy/us-power-use-reach-record-highs-2025-2026-eia-says-2025-05-06/.

Figure 11: Tax Bill Share of Revenues by Carbon Intensity and Productivity



carbon taxes that correct firms' climate externality. As acknowledged by practitioners and policymakers alike, a central challenge with carbon taxation is, however, the largely self-reported nature of emissions disclosures. Carbon disclosures are in most cases based on firm-level information that is difficult or costly to verify for policymakers. Combined with a sizable carbon tax, this creates a strong incentive for firms to misreport their carbon disclosures—consistent with recent legal evidence.

In this paper, we formulate a neoclassical general-equilibrium model in which firms' fossil-energy use generates climate damages to production. The model crucially incorporates substantial cross-firm heterogeneity in carbon efficiency that we documented using a novel firm-level dataset spanning 150 countries. Firms choose energy and other inputs and self-report emissions that are otherwise privately observed. Our main result is a simple formula for the optimal carbon tax that accounts for firms' incentives to distort their carbon disclosures. The optimal Mirrleesian tax varies markedly across firms as a function of a firm's carbon intensity (based on its carbon disclosure), marginal climate damages, and output.

We evaluate this formula quantitatively, using the US distribution of carbon intensities, and find that optimal carbon taxes vary substantially across firms. The information friction related to firms' emissions makes firm-level taxes optimally *heterogeneous across firms* but *constant within a firm*. We find that the optimal carbon tax for a firm that is 1 standard deviation below the mean carbon-intensity level is around \$5,000-7,000 per tCO_2e; a firm that is 1 standard deviation above the mean carbon-intensity level, by contrast, faces an optimal tax slightly less than \$300 per tCO_2e. The median carbon tax paid is \$1,500 per tCO_2e. Despite this heterogeneity, the overall Mirrleesian tax bill faced by firms in equilibrium never exceeds 1% of firm revenue. Instead, we show that it is changes in firm-level disclosure incentives that

drive the bulk of the benefits of carbon taxation that we document.

Our tax contrasts with other carbon-tax estimates that do not internalize firms' reporting incentives. Our model—presuming full-information about firms' carbon emissions—provides an optimal carbon tax, equal to the social cost of carbon, that is homogenous across firms and close to \$300 per tCO_2e. Although there are substantial uncertainties about the true social cost of carbon, this estimate is close to the median estimate used in the literature. Our results, thus, do not rest on the presumption of a large social cost of carbon; accounting for firms' reporting incentives increases most carbon taxes paid by firms and by itself leads to substantial heterogeneity in optimal carbon tax rates.[27]

Crucially, we document that the optimal Mirrleesian tax recovers most of the potential benefits from full-information carbon taxation. In our baseline calibration, accounting for firms' reporting incentives allows us to recover over 80% of the possible welfare and environmental benefits from carbon taxation. A naive approach—which implements the full-information tax without accounting for firms' reporting incentives—in contrast recovers less than 10%. We show that this comparative performance of the optimal Mirrleesian tax rests on the ability of a well-constructed tax to incentive abatement efforts by carbon-intensive firms. As a result, we document that the comparative performance of the optimal Mirrleesian tax extends to several variations of our model framework. In all cases, the optimal Mirrleesian tax recovers more than 2/3 of the potential welfare and environmental benefits from carbon taxation; the naive approach by contrast recovers substantially less—and some times close to zero.

An important feature of our proposed tax formula is its simplicity. The optimal carbon tax paid by a firm corresponds to a simple, proportional adjustment of the standard social cost of carbon estimate, based on a firm's location in the carbon-intensity distribution. Given our formula, a firm has no incentive to misreport its location in the distribution. As such, our proposed carbon tax retains many of the advantages in terms of simplicity and implementation of standard carbon-tax proposals. Although it should be clear from our discussions throughout the text that several extensions of the present setting are desirable, one advantage of the tractability of our framework is that such extensions should come at a low cost.

Finally, beyond the analysis in this paper, our approach suggests that accounting for firms' reporting incentives can meaningfully alter the optimal Pigouvian taxation of firms. We see important scope for extending these ideas to other circumstances in which firms' actions have externalities on the broader economy and their efforts are hard to monitor—such as is the case with, for example, industrial policy. Overall, we view the research in this paper as a useful step towards a unified framework for the Mirrleesian taxation of firms, and the current

[27] We may also relate our estimates to the price of emission rights in the European Union Emission Trading System, covering large CO_2 emitters in the EU. The price is currently around \$100 per tC02 and equal across all of firms, and thus well-below our estimate for the optimal carbon tax rates paid by most firms.

model framework as a useful basis for extensions that incorporate dynamic investments in abatement, cross-jurisdictional leakage, and other reporting and verification environments.

A Motivating Evidence

A.1 Data Construction

ICE-ESG Database: Our main analysis uses data from the ICE's ESG company database. The total data set covers the period 2009-2024 for 402,956 firm-years spanning 46,928 firms across 151 countries. To construct the sample that we use in Figures 2 and 3 and Appendix A, we discard ICE-model-generated emission estimates (derived from company data) and focus only on firms' self-reported emissions. We also discard observations that are not marked by ICE as having "high-quality reporting data" behind them (i.e, observations for which the ICE's disclosure-quality category is above 3). Figure 1 uses the full sample of observations. Direct emission intensities are measured by reported scope 1 emissions (tons of CO_2 equivalent emissions) per million $ of revenue (`code: Reported Emissions Intensity Scope 1 (tCO2e/$m Revenue)`). Firms' direct emissions are measured by the numerator in this expression (`code: Reported Emissions Scope 1 (tCO2e)`). We remove all observations that are related to the top and bottom 1% of the emission intensity distribution.

S&P TruCost Database: We supplement the data from ICE with S&P's TruCost Database on publicly-listed firms emissions. The overall sample covers the period 2009-2019 for 50,013 firm-years spanning 9,227 firms across 39 different countries. Emission intensities are again measured by self-reported scope 1 emissions (tons of CO_2 equivalent emissions) per million $ of revenue (`code: di319407`). Overall emissions are measured by the numerator of this expression (`code: di319413`). We once more remove all observations that are related to the top and bottom 1% of the emission intensity distribution.

Compustat Database: Finally, we use the following variables from Compustat Fundamentals Annual: revenue (`code: sale`), capital (`code: ppent, ppegt`), investment (`code: capx`), employment (`code: emp`), and industry classification (`code: naics and sic`). We deflate nominal variables where appropriate with US CPI (`code: CPIAUCLS` from FRED) and compute (revenue-based) total factor productivity as in Ottonello and Winberry (2020).

Descriptive Statistics: Table A.1 reports descriptive statistics for the various data sets.

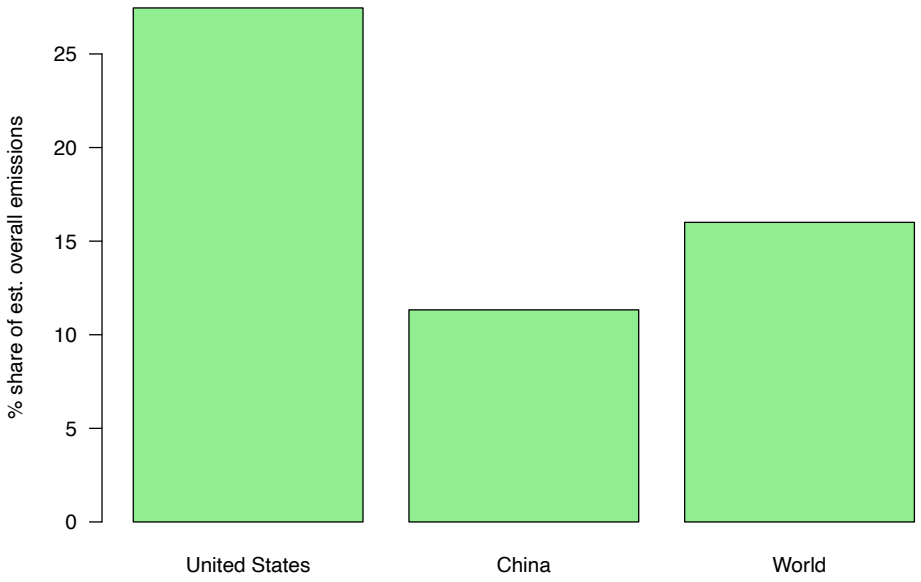
Table A.1: Descriptive statistics: ICE-TruCost-Compustat

Variable Name	Obs.	Mean	Std.	Median
Scope 1 emission (ICE)	32,599	1,461,794	5,139,210	29,200
Scope 1 emission (TruCost)	49,011	541,617	2,042,981	26,061
Scope 1 emission intensity (ICE)	32,599	181.31	568.25	8.13
Scope 1 emission intensity (TruCost)	49,012	131.05	357.26	21.64
Revenue (Compustat)	xx,xxx	xx,xxx	xx,xxx	xx,xxx
Capital (Compustat)	xx,xxx	xx,xxx	xx,xxx	xx,xxx
Investment (Compustat)	xx,xxx	xx,xxx	xx,xxx	xx,xxx
Assets (Compustat)	xx,xxx	xx,xxx	xx,xxx	xx,xxx
Employment (Compustat)	xx,xxx	xx,xxx	xx,xxx	xx,xxx
Tfp (Compustat)	xx,xxx	xx,xxx	xx,xxx	xx,xxx

Notes: The table reports descriptive statistics for the sample of firms in the ICE, TruCost, or Compustat database. The units of the first two rows are tons of CO_2 equivalent emissions (tCO2e). The two subsequent rows are tCO2e per million $ in revenue. For all rows, observations have been removed that are in top and bottom 1 percent of the variable's distribution.

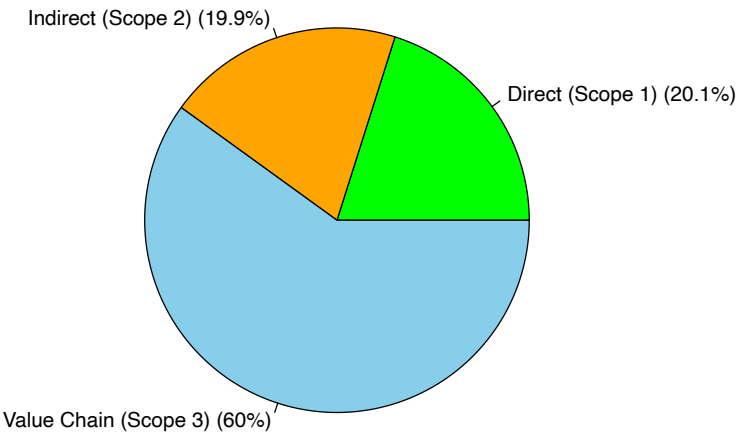
A.2 Additional Evidence

Figure A.1: Estimated Shares of Overall Emissions—TruCost



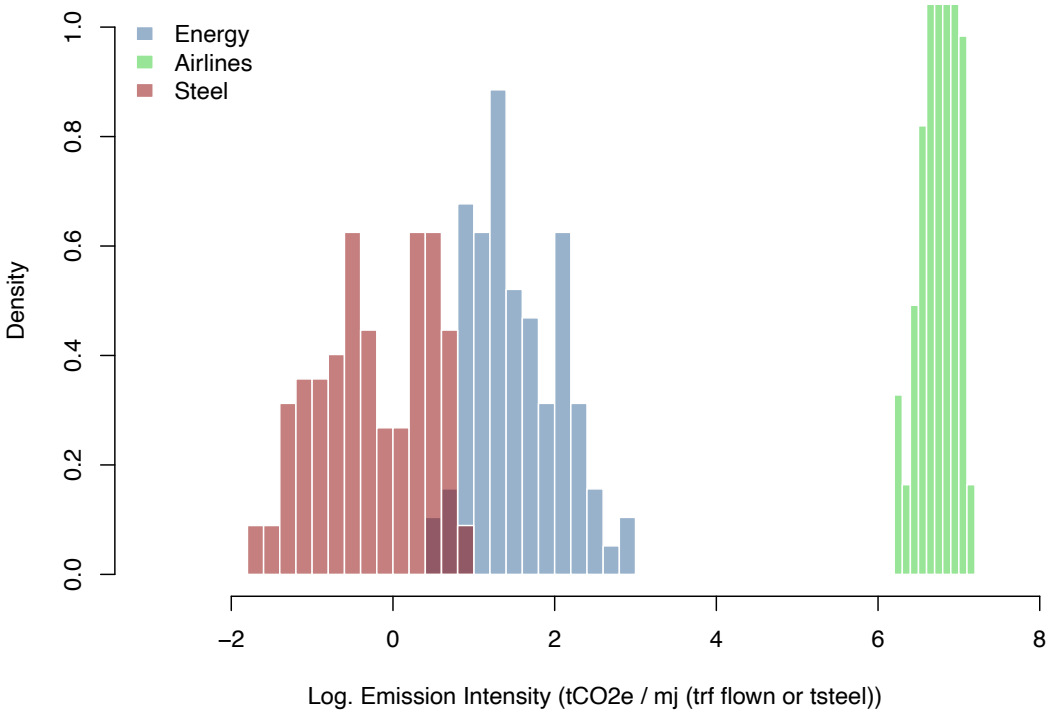
Note: Data from the TruCost sample. The figure shows overall scope 1 emissions in 2019 (measured by tons of CO_2 equivalent emissions) divided by overall country/region-wide emissions. Estimates for country/region-wide emissions are taken from the Global Carbon Budget website (20219 estimates).

Figure A.2: Breakdown of ICE Emissions



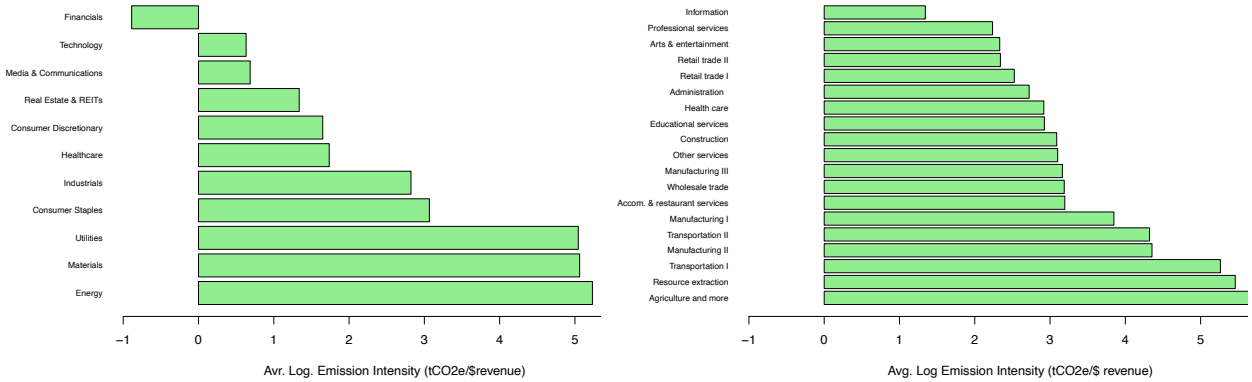
Note: Data from the ICE-ESG sample. The figure shows the breakdown of the average firm's emissions into scope 1, 2 and 3 emissions. We consider only firms with validated or high-quality reporting (Appendix A.1).

Figure A.3: Heterogeneity in Emission Intensities—Output Measures



Note: Data from the ICE-ESG sample. The figure shows the distribution of *log. emission intensities* within the US energy, airlines and steel sectors. Emission intensities are measured by tons of CO_2 equivalent emissions (tCO2e) divided by output. We measure output by: mega jules of energy produced (mj), thousands of return flights flown (trf) and tons of steel produced (tsteel), respectively. We consider only firms with validated or high-quality reporting (Appendix A.1). We winsorize emission intensities at the top/bottom 1% level.

Figure A.4: Sectoral Heterogeneity in Emission Intensities

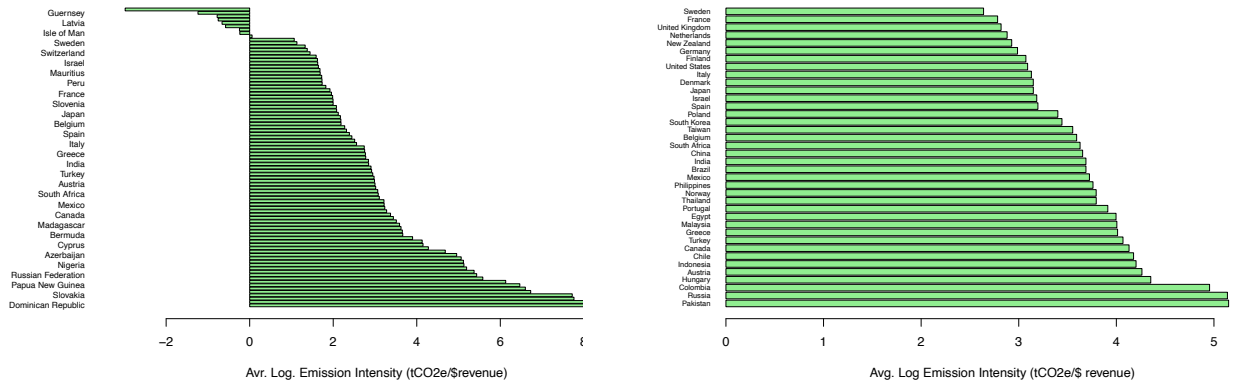


Panel (a): ICE-ESG Sample

Panel (b): TruCost Sample

Note: Data from the ICE-ESG and TruCost samples. The figure shows the avr. log (scope 1) emission intensity within ICE and TruCost sectors (NAICS level 2). Emission intensities are measured by tons of CO_2 equivalent emissions (tCO2e) divided by revenue (millions of $). We consider only firms with validated or high-quality reporting (Appendix A.1) and winsorize all emission intensities at the top/bottom 1% level.

Figure A.5: Country Heterogeneity in Emission Intensities



Panel (a): ICE-ESG Sample

Panel (b): TruCost Sample

Note: Data from the ICE-ESG and TruCost samples. The figure shows the avr. log (scope 1) emission intensity within countries. Emission intensities are measured by tons of CO_2 equivalent emissions (tCO2e) divided by revenue (millions of $). We consider only firms with validated or high-quality reporting (Appendix A.1) and winsorize all emission intensities at the top/bottom 1% level.

Table A.2: Emission Disclosures and Emission Intensity

	Log. emission intensity	
	(1)	(2)
No third-party validation	-0.193^{***}	-0.214^{***}
	(0.031)	(0.023)
Time fixed effects	×	✓
Sector fixed effects	×	✓
Country fixed effects	×	✓
Observations	32,268	32,268

Notes: Panel least-squares estimates from the ICE-ESG sample. The top and bottom 1 percent of emission intensities (tons of scope 1 CO_2 equivalent emissions, tCO2e, divided by revenue in millions of \$) are removed. ICE sectoral definitions and verification status are used. We consider only firms with high-quality reporting (Appendix A.1). Robust clustered (sectors) standard errors in parentheses. Sample: 2010–2024.

Table A.3: Carbon Intensity and Productivity

	Log. emission intensity	
	(1)	(2)
Log. tfp	-0.556***	-0.567***
	(0.150)	(0.153)
Time fixed effects	×	✓
Sector fixed effects	×	✓
Country fixed effects	×	✓
Observations	12,508	12,508
R-squared	0.047	0.058

Notes: Panel least-squares estimates from the merged ICE-ESG-Compustat sample. The top and bottom 1 percent of emission intensities (tons of scope 1 CO_2 equivalent emissions, tCO2e, divided by revenue in millions of \$) and total factor productivity are removed. Tfp is estimated as in Ottonello and Winberry (2020). We consider only firms with high-quality reporting (Appendix A.1). NAICS-4 sectoral definitions are used throughout. Robust clustered (sectors) standard errors in parentheses. Sample: 2010–2024.

Table C.1: Parameters and Targets

Description	Parameter	Value	Target
Economic block			
Discount rate	β	0.96	
Capital elasticity	α	0.29	
Labor elasticity	γ	0.43	Standard
Energy elasticity	ν	0.08	
Capital depreciation rate	δ_K	0.10	
Mean productivity	μ_a	0	Normalization
Std productivity	σ_a	0.6	Compustat data estimates
Climate block			
Carbon depreciation rate	δ_G	0.02	Golosov *et al.* (2014)
Abatement cost inverse elasticity	θ	2	Känzig (2023)
Damage semi-elasticity	ϕ	$3.5{\cdot}10^{-4}$	Social cost of carbon estimates
Corr between a and z distributions	ρ_{az}	0	ICE ESG data estimates
Mean carbon intensity	μ_z	0	Normalization
Std carbon intensity	σ_z	1.6	ICE ESG data estimates

B Equilibrium and Planner Allocations

B.1 The Mirrleesian Case

`tbd...`

C Quantitative Results

`tbd...`

References

Acemoglu, D., Aghion, P., Bursztyn, L. and Hemous, D. (2012). The environment and directed technical change. *American economic review*, **102** (1), 131–166.

—, Akcigit, U., Hanley, D. and Kerr, W. (2016). Transition to clean technology. *Journal of political economy*, **124** (1), 52–104.

Akcigit, U., Hanley, D. and Stantcheva, S. (2022). Optimal taxation and r&d policies. *Econometrica*, **90** (2), 645–684.

Allen, M. R., Frame, D. J., Huntingford, C., Jones, C. D., Lowe, J. A., Meinshausen, M. and Meinshausen, N. (2009). Warming caused by cumulative carbon emissions towards the trillionth tonne. *Nature*, **458** (7242), 1163–1166.

Allingham, M. G. and Sandmo, A. (1972). Income tax evasion: A theoretical analysis. *Journal of public economics*, **1** (3-4), 323–338.

Barrage, L. and Nordhaus, W. (2024). Policies, projections, and the social cost of carbon: Results from the dice-2023 model. *Proceedings of the National Academy of Sciences*, **121** (13), e2312030121.

Basu, S. and Kimball, M. S. (1997). Cyclical productivity with unobserved input variation.

Bhandari, A., Evans, D., McGrattan, E. R. and Yao, Y. (2024). *Business Income Underreporting and Public Finance*. Tech. rep., working paper, University of Minnesota.

Bilal, A. and Känzig, D. R. (2024). *The macroeconomic impact of climate change: Global vs. local temperature*. Tech. rep., National Bureau of Economic Research.

Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I. and Terry, S. J. (2018). Really uncertain business cycles. *Econometrica*, **86** (3), 1031–1065.

Brandt, A. R., Heath, G., Kort, E., O'Sullivan, F., Pétron, G., Jordaan, S. M., Tans, P., Wilcox, J., Gopstein, A., Arent, D. *et al.* (2014). Methane leaks from north american natural gas systems. *Science*, **343** (6172), 733–735.

Bruns, S. B., Csereklyei, Z. and Stern, D. I. (2020). A multicointegration model of global climate change. *Journal of Econometrics*, **214** (1), 175–197.

Caggese, A., Chiavari, A., Goraya, S. and Villegas-Sanchez, C. (2025). *Climate change, firms, and aggregate productivity*. Centre for Economic Policy Research.

Capelle, M. D., Kirti, M. D., Pierri, M. N. and Bauer, M. G. V. (2023). Mitigating climate change at the firm level: Mind the laggards. *International Monetary Fund WP*.

Chu, H., Qi, J., Feng, S., Dong, W., Hong, R., Qiu, B. and Han, W. (2023). Soot formation in high-pressure combustion: Status and challenges. *Fuel*, **345**, 128236.

Dechezleprêtre, A., Fabre, A., Kruse, T., Planterose, B., Sanchez Chico, A. and Stantcheva, S. (2025). Fighting climate change: International attitudes toward climate policies. *American Economic Review*, **115** (4), 1258–1300.

Di Nola, A., Kocharkov, G., Scholl, A. and Tkhir, A.-M. (2021). The aggregate consequences of tax evasion. *Review of Economic Dynamics*, **40**, 198–227.

Diamond, P. A. (1998). Optimal income taxation: an example with a u-shaped pattern of optimal marginal tax rates. *American Economic Review*, pp. 83–95.

Dietz, S., Van Der Ploeg, F., Rezai, A. and Venmans, F. (2021). Are economists getting climate dynamics right and does it matter? *Journal of the Association of Environmental and Resource Economists*, **8** (5), 895–921.

ECA (2015). *The integrity and implementation of the EU ETS*. Special Report 06/2015, European Court of Auditors, Luxembourg, eN 2015 No 06.

Economists' Statement on Carbon Dividends (2019). Economists' statement on carbon dividends. *The Wall Street Journal*, full-page statement organized by the Climate Leadership Council; 3,500+ signatories.

Economists' Statement on Climate Change (1997). Economists' statement on climate change. Public statement coordinated by Redefining Progress, signed by 2,600+ economists, including 19 Nobel laureates.

European Commission (2023). Esrs e1 climate change. *European Commission Working paper*.

Fernandez-Bastidas, R. (2023). Entrepreneurship and tax evasion. *Economic Modelling*, **128**, 106488.

Fernández-Villaverde, J., Gillingham, K. T. and Scheidegger, S. (2024). Climate change through the lens of macroeconomic modeling. *Annual Review of Economics*, **17**.

Golosov, M., Hassler, J., Krusell, P. and Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, **82** (1), 41–88.

—, Tsyvinski, A., Werning, I., Diamond, P. and Judd, K. L. (2006). New dynamic public finance: A user's guide [with comments and discussion]. *NBER macroeconomics annual*, **21**, 317–387.

Hassler, J., Krusell, P. and Smith Jr, A. A. (2016). Environmental macroeconomics. In *Handbook of macroeconomics*, vol. 2, Elsevier, pp. 1893–2008.

ICE (2024). Sustainability report and documentation. *Intercontinental Exchange WP*.

IFRS Foundation (2023). Ifrs s2: Climate-related disclosures. *IFRS Working Paper*.

IPCC (2021). *Climate Change 2021: The Physical Science Basis*. Cambridge, UK and New York, NY, USA: Cambridge University Press.

IPCC (2022). *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, UK and New York, NY, USA: Cambridge University Press.

Jaramillo, P., Griffin, W. M. and Matthews, H. S. (2007). Comparative life-cycle air emissions of coal, domestic natural gas, lng, and sng for electricity generation. *Environmental science & technology*, **41** (17), 6290–6296.

Jo, A. (2025). Substitution between clean and dirty energy with biased technical change. *International Economic Review*, **66** (2), 883–902.

Kang, Z. Y. (2020). Optimal indirect regulation of externalities. *Available at SSRN 3586050*.

Känzig, D. R. (2023). *The unequal economic consequences of carbon pricing*. Tech. rep., National Bureau of Economic Research.

Kocherlakota, N. R. (2010). The new dynamic public finance. In *The New Dynamic Public Finance*, Princeton University Press.

Laffont, J.-J. (1994). The new economics of regulation ten years after. *Econometrica: journal of the Econometric Society*, pp. 507–537.

Lyubich, E., Shapiro, J. S. and Walker, R. (2018). Regulating mismeasured pollution: Implications of firm heterogeneity for environmental policy. In *AEA papers and proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 108, pp. 136–142.

Matthews, H. D., Gillett, N. P., Stott, P. A. and Zickfeld, K. (2009). The proportionality of global warming to cumulative carbon emissions. *Nature*, **459** (7248), 829–832.

—, Solomon, S. and Pierrehumbert, R. (2012). Cumulative carbon as a policy framework for achieving climate stabilization. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, **370** (1974), 4365–4379.

Miftakhova, A., Judd, K. L., Lontzek, T. S. and Schmedders, K. (2020). Statistical approximation of high-dimensional climate models. *Journal of Econometrics*, **214** (1), 67–80.

Mirrlees, J. A. (1971). An exploration in the theory of optimum income taxation. *The review of economic studies*, **38** (2), 175–208.

— (1976). Optimal tax theory: A synthesis. *Journal of public Economics*, **6** (4), 327–358.

Nemitallah, M. A., Abdelhafez, A. A., Ali, A., Mansir, I. and Habib, M. A. (2019). Frontiers in combustion techniques and burner designs for emissions control and co2 capture: A review. *International Journal of Energy Research*, **43** (14), 7790–7822.

Nordhaus, W. (2019). Climate change: The ultimate challenge for economics. *American Economic Review*, **109** (6), 1991–2014.

Nordhaus, W. D. (1977). Economic growth and climate: the carbon dioxide problem. *The American Economic Review*, **67** (1), 341–346.

— (1993). Optimal greenhouse-gas reductions and tax policy in the" dice" model. *The American Economic Review*, **83** (2), 313–317.

— (2007). A review of the stern review on the economics of climate change. *Journal of economic literature*, **45** (3), 686–702.

— and Moffat, A. (2017). A survey of global impacts of climate change: replication, survey methods, and a statistical analysis.

Ottonello, P. and Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, **88** (6), 2473–2502.

Pai, M. and Strack, P. (2022). Taxing externalities without hurting the poor.

Papageorgiou, C., Saam, M. and Schulte, P. (2017). Substitution between clean and dirty energy inputs: A macroeconomic perspective. *Review of Economics and Statistics*, **99** (2), 281–290.

Pérez-Ramirez, J., Kapteijn, F., Schöffel, K. and Moulijn, J. (2003). Formation and control of n2o in nitric acid production: where do we stand today? *Applied Catalysis B: Environmental*, **44** (2), 117–151.

Petersen, T. (2025). Can the plants turn green? *Mimeo*.

Pigou, A. (1920). *The economics of welfare*. Routledge.

Pigou, A. C. (1912). *Wealth and welfare*. Macmillan and Company, limited.

Pindyck, R. S. (2019). The social cost of carbon revisited. *Journal of Environmental Economics and Management*, **94**, 140–160.

Report, C. C. M. (2017). *The Carbon Majors Database: CDP Carbon Majors Report 2017*. Tech. rep., CDP, London, cDP report.

Ricke, K., Drouet, L., Caldeira, K. and Tavoni, M. (2018). Country-level social cost of carbon. *Nature Climate Change*, **8** (10), 895–900.

SEC (2024). The enhancement and standardization of climate-related disclosures: Final rules. *Security and Exchange Commission: Final Rule*.

Shapiro, J. S. and Walker, R. (2018). Why is pollution from us manufacturing declining? the roles of environmental regulation, productivity, and trade. *American economic review*, **108** (12), 3814–3854.

Shimizu, A., Tanaka, K. and Fujimori, M. (2000). Abatement technologies for n2o emissions in the adipic acid industry. *Chemosphere-global change science*, **2** (3-4), 425–434.

Stantcheva, S. (2020). Dynamic taxation. *Annual Review of Economics*, **12** (1), 801–831.

Tirole, J. (1988). *The theory of industrial organization*. MIT press.

Tol, R. S. (2023). Social cost of carbon estimates have increased over time. *Nature climate change*, **13** (6), 532–536.

U.S. Energy Information Administration (2024). U.s. energy spending increased by more than 20% in 2022. Today in Energy, principal contributors: Melissa Alejandro and Brett Marohl.

Weitzman, M. L. (1974). Prices vs. quantities. *The Review of Economic Studies*, **41** (4), 477–491.

World Resources Institute (2004). *The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard (Revised Edition)*.

Zhang, Z., Dong, R., Lan, G., Yuan, T. and Tan, D. (2023). Diesel particulate filter regeneration mechanism of modern automobile engines and methods of reducing pm emissions: A review. *Environmental Science and Pollution Research*, **30** (14), 39338–39376.