Mirrleesian Carbon Taxation*

Andrea Chiavari Alexandre Kohlhas

July 2026

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 intensity 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 optimal carbon tax that modifies marginal externality damages to account for firms' incentives to distort reported emissions. Unlike standard carbon-tax proposals, the optimal tax varies markedly across firms as a function of emission intensities 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.

Keywords: Climate Change, Optimal Policy, Mirrleesian Optimal Taxes.

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

*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 Adrien Bilal, John Hassler, Marek Kapička, Per Krusell, Andrea Lanteri, Rick van der Ploeg, Hannah Roemer, Ludvig Sinander, Sasha Teytelbaum, and seminar and conference participants at Bologna, Bonn, CERGE-EI, Cyprus, Duke, EIEF, Expectations and Macro (Seoul), Mannheim, Minneapolis Fed, Oxford, SED 2026, Toronto Climate Change, T2M, and Yale. We gratefully acknowledge support from Christ Church Research Center and the Hedelius Foundation.

1 Introduction

A defining challenge of our time is climate change. Given the central role that firms' carbon emissions play in its development, it is important to have 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 Pigouvian 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 (e.g., Nordhaus, 1977, 2007). This approach has, in turn, led to several real-world carbon tax proposals (e.g., Statement on Carbon Dividends, 2019).[1] Nevertheless, the practical implementation of carbon taxation faces several challenges.

An important obstacle—long recognized by policymakers (e.g., European Commission, 2023a) and practitioners (e.g., ECA, 2015) alike—is that firms' carbon disclosures are based on *firm-specific information* and primarily *self-reported*.[2] 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.[3] In part as a consequence of these concerns and several high-profile court cases that have substantiated their relevance, trust in the effectiveness of carbon taxation is low among the general public (e.g., Dechezleprêtre et al., 2025).[4] Our aim with this paper is to fill the gap between the macroeconomic theory that addresses firms' climate externality through carbon taxation and its practical implementation, plagued by imperfect knowledge about firms' emissions.

To attain it, we propose a Mirrleesian modification of the standard Pigouvian approach that recognizes the need to incentivize firms' truthful disclosures. Within the context of a workhorse dynamic integrated-assessment model, we show that a simple, firm-specific adjustment to a uniform carbon tax optimally internalizes firms' disclosure incentives. Our constrained optimal tax varies markedly across firms as a function of a firm's emission intensity, marginal damages, and output, and exceeds common uniform-tax benchmarks, on average. Quantifying our tax formula, we find substantial welfare

[1] See, for example, Stiglitz et al. (2017), Parry et al. (2021), and Shultz and Baker (2017).

[2] See also GHG Protocols (2004), International Sustainability Standards Board (2023), ECA (2015), and SEC Carbon Disclosures (2024), among others.

[3] We discuss firms' carbon disclosures in the EU, US, and the UK in Section 2 and Appendix A.3.

[4] In Europe, recent enforcement actions have highlighted failures in emissions reporting and verification, including the Brikel/Bobov Dol EU ETS investigation (2021), the Scottish Environment Protection Agency v. Diageo plc (2022), and the Scottish Environment Protection Agency v. ExxonMobil Chemical Ltd (Mossmorran) (2010). We discuss these three representative legal cases in detail in Section 2, with additional institutional background and documentation provided in Appendix A. Even observed non-compliance often does not result in sizable fines (Calel et al., 2025). As acknowledged by policymakers, due to difficulties verifying carbon disclosures, these presumably represent the tip of the iceberg.

and environmental gains from internalizing reporting incentives. The constrained optimal tax recovers around 3/4 of the potential social benefits; a uniform tax that does not internalize reporting incentives, by contrast, recovers less than 1/10.

To motivate our analysis, we provide new evidence on firms' reported emissions. Using a comprehensive dataset on public and private firms' carbon-equivalent emissions (CO_2e) from 150 countries over 15 years—*the ICE ESG Company Database*—we show that firms' *direct, production-based* (scope 1) emissions account for over half of overall emissions into the atmosphere—a lower bound, given that our estimate covers only ICE-ESG firms. Addressing climate change is, to a large extent, a problem of curtailing firms' emissions.[5]

Crucially, we document substantial *heterogeneity in firms' emission intensity*, defined as emissions per unit of output or revenue. Consistent with Lyubich et al. (2018), emission intensities are widely dispersed, even after controlling for country-sector-time fixed effects, and the spread is large relative to other familiar sources of firm heterogeneity such as TFP (Section 2 and Appendix A.3 discuss in detail the technological reasons behind this heterogeneity).[6] Finally, we exploit a novel feature of the ICE ESG database: we compare reported emissions for firms whose disclosures are *third-party verified* (that is, for which the internal emissions calculations are checked by external consultants) with those that are not. We find that non-verified firms report, on average, 20% lower emissions than otherwise similar verified firms, with the gap being larger in sectors with greater heterogeneity in emissions intensity. Our findings are therefore consistent with policymaker concerns about systematic optimistic reporting of carbon-equivalent emissions.

Our core contribution is to develop a theory of carbon taxation consistent with firms' incentives. We consider a workhorse dynamic integrated-assessment model in the spirit of Golosov et al. (2014) and Barrage and Nordhaus (2024). The distinguishing feature of our model is that firms are *heterogeneous*, both in their *productivity* and their *emission intensity*. The latter measures the extent to which a firm's energy use converts into carbon-equivalent emissions. Importantly, a firm's emission intensity, and hence its emissions, are partially *private information* to the firm, unknown to the social planner or other firms in the economy. This information friction creates a conflict between the social planner's desired disclosures for a firm and the firm's own optimal disclosure. The rest of the model is standard. In the model, firms produce output using capital, labor, and energy—with energy produced indirectly from output in a roundabout fashion—and invest in abatement to curb emissions. Emissions reduce economy-wide productivity through a

[5]The ICE database covers substantially more firms than the more commonly used S&P Trucost dataset, which is why our estimates are larger than those reported in e.g., Bolton and Kacperczyk (2023).

[6]Other papers documenting pervasive heterogeneity in emissions intensity include Capelle et al. (2023), De Lyon and Dechezleprêtre (2025), and Fontaine et al. (2026).

standard damage function (i.e., an *environmental externality*) that firms do not internalize. Within this environment, we derive two main results.

Our first main result shows that the combination of firm-specific information and heterogeneity in emission intensity challenges conventional carbon tax proposals. We show that a standard, homogeneous carbon tax—as implied in our economy by the Pigouvian optimal tax when the planner has full information about firms' emission intensities—introduces profit differentials between equally productive firms with different emission intensities. The tax scheme leads to larger tax bills for firms with higher emission intensity. In an economy in which emission intensities are partially private information, homogeneous carbon taxes are therefore *incentive-incompatible*; they generate strong incentives for firms to misreport their emissions. We show that the social cost of neglecting firms' incentives can be arbitrarily large, depending on the extent of heterogeneity.

The second main result builds on this challenge. We derive the optimal carbon tax when the social planner is constrained to choose allocations that are ex-post incentive-compatible. We refer to our constrained optimal tax as the *Mirrleesian carbon tax*. On the one hand, the Mirrleesian planner desires to curb emissions, to raise economy-wide productivity; but, on the other hand, efforts to reduce emissions also distort input use, creating misallocation that lowers productivity. The optimal full-information tax balances this *efficiency–misallocation trade*-off with a constant, homogeneous tax on emissions equal to the social cost of carbon (e.g., Golosov et al., 2014). By contrast, in the Mirrleesian case, the optimal constrained tax instead proportionally adjusts the full-information tax as a function of a firm's location in the carbon-intensity distribution, to ensure incentive compatibility. The Mirrleesian tax is, thus, *linear within firms but heterogeneous across firms*.

Indeed, a central feature of the Mirrleesian carbon tax is that it *declines* with a firm's emission intensity. A firm's incentive to misreport is unidirectional—high-emission-intensity firms wish to pretend to be less carbon intensive. Combined with the Mirrleesian tax schedule resulting in allocations—conditional on productivity—in which all firms pay the same total carbon tax as a share of revenue, this removes any incentive to misreport. We show that our constrained optimal tax varies markedly across firms as a function of their emission intensity and exceeds common uniform benchmarks, on average.

To explore the quantitative benefits of Mirrleesian carbon taxation, we calibrate the model to the within-sector, within-time heterogeneity in emission intensities for the verified U.S. ICE ESG sample. Absent information frictions, we estimate the optimal carbon tax to be \$164 per ton of CO_2 equivalent emissions (tCO_2e), in line with recent estimates of the social cost of carbon (e.g., Tol, 2023). The average associated tax bill ranges from close to 0% of revenues for low emission-intensity firms to almost 2% for high emission-

intensity firms. By contrast, the Mirrleesian carbon tax schedule implies substantially higher marginal rates. The optimal tax is around \$700 per tCO_2e for the median emission-intensity firm, rising to roughly \$3,500 per tCO_2e for firms 1 standard deviation below the median, and declining to about \$150 per tCO_2e for firms 1 standard deviation above. Notwithstanding this variation in marginal rates, total tax payments are *identical across firms with the same productivity* and on average amount to less than 0.5% of revenue. The schedule thus pairs heterogeneous marginal rates with equalized, modest total payments.

Crucially, the Mirrleesian adjustment extracts *most* of the potential benefits from carbon taxation. Compared to the (first-best) full-information tax, the Mirrleesian tax recovers 3/4 of the potential welfare gains and around 80% of the potential decline in emissions. A tax equal to its optimal value under full information—to which firms optimally respond by misreporting—by contrast, recovers less than 10% of the social and environmental benefits attainable in the first best. We conclude that the Mirrleesian tax schedule extracts most of the potential benefits from carbon taxation.

We show that the reason behind this comparative performance is that most of the benefits from carbon taxation arise due to higher aggregate productivity driven by reduced emissions. Changes in capital and energy use, as well as the misallocation costs of carbon taxation, are in contrast small. We document through an exact decomposition of firms' emissions that the Mirrleesian tax attains a similar decline in overall emissions to that in the first-best, full-information case whenever the first-best correlation between firms' energy use and carbon intensity is close to zero—i.e., whenever productivity (as in our model) is an important determinant of first-best energy choices. Our quantitative results thus underscore the importance of jointly accounting for productivity dynamics and carbon-intensity heterogeneity within one unified framework.

We conclude our analysis by showing that the relative performance of Mirrleesian taxation extends across a range of robustness exercises and extensions. First, we examine transition dynamics and allow for a more complex climate specification in which the climate system is characterized 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 undertake a sensitivity analysis of the parameters that govern the climate block of the economy, focusing in particular on the abatement-cost elasticity and the climate-damage elasticity, for both of which there is substantial parameter uncertainty. Third, we introduce renewable and fossil energy as imperfect substitutes, to account for one important driver of the heterogeneity in emission intensity. This follows the approach in Hassler et al. (2016), among others. And, lastly, we vary the accuracy

of the planner's information about firms' emission intensities. Overall, our main result remains the same: under realistic parameterizations, ignoring incentive-compatibility considerations is consistently socially and environmentally costly; by contrast, a Mirrleesian approach that instead internalizes such considerations in all cases extract at least 50-60% of the potential benefits of carbon taxation.

Finally, a premise of our analysis is that policymakers often have only crude estimates of firms' emission intensity. Greenstone et al. (2025) document a multi-decade attempt to gain more accurate information about firms' emissions in India and the costs associated with it.[7] Recognizing this informational gap and its impact on the efficacy of carbon taxation, policymakers across developed and developing economies have called for a combination of (i) *tougher compliance enforcement* (EC; ECA; WB) and (ii) *better disclosure standards plus independent assurance* (IFRS; IAASB; UNFCCC)—in effect attempting to decrease the information gap between policymakers and firms.[8] However, as also recognized by policymakers, the cost of doing so is often substantial—both in terms of the added regulatory burden placed on firms, resources spent on measurement devices, and the time spent constructing carbon estimates. In this paper, we instead propose a different approach, one that accepts an informational gap between policymakers and firms, and instead adjusts the carbon tax schedule to account for any incentive concerns. Our analysis suggests that information frictions should be treated as a first-order design concern for carbon taxation, much like uncertainty about social damages or discounting is currently done.

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 Pigouvian 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), Van den Bremer and Van der Ploeg (2021), 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 allo-

[7] A system referred to as continuous emissions measurement systems (CEMS) can in principle be fitted on every firm burner. But, as documented by European Commission (2023b) almost all installations instead use a calculation-based methodology. Indeed, only 145 installations (1.7%) in 22 countries reported using continuous emissions measurement systems (CEMS) in 2025, nine installations fewer than in 2021.

[8] These are: the European Commission (EC), the European Court of Auditors (ECA), the World Bank (WB), the International Financial Reporting Standards (IFRS), the International Auditing and Assurance Standards Board (IAASB), and the United Nations Framework Convention on Climate Change (UNFCCC).

cations with distorted disclosures. The optimal tax formula that we derive documents how to optimally adjust the standard carbon tax for such concerns. In complementary and closely-related work, Cicala et al. (2022) study the design of allocation mechanisms when firms hold private information about their emissions. While Cicala et al. (2022) elicit firms' private information through an allocation mechanism relying on costly third-party emissions verification, we show that when such verification is prohibitively costly or not credible, a simple modification of the standard Pigouvian tax can recover much of the efficiency gains from climate taxation.

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), and Stantcheva (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 utilizing 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 tax scheme firms report truthfully—recovers most of the benefits of the textbook Pigouvian approach while ensuring compliance.

Lastly, 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 recognizes 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' carbon disclosure incentives.

2 Motivating evidence

In this section, we provide evidence on the *scope* and *heterogeneity* in firms' carbon-equivalent emissions. We then examine the drivers of this heterogeneity and how disclosure practices can translate into systematic misreporting at the firm level.

2.1 Firm emissions and heterogeneity

We present new evidence on firms' carbon-equivalent emissions. To start, we use micro data from *Intercontinental Exchange's (ICE's)* ESG company database. The ICE-ESG database provides one of the most comprehensive dataset on private and public firms' emissions in existence and spans the period 2009-2024 for 46,928 firms across 151 countries. The database reports, for an individual firm-year, estimates of CO_2-equivalent (CO_2e) direct emissions from production (i.e., *scope 1* emissions). We further validate our results using publicly-listed firms' CO_2e emissions data from *S&P Trucost*. Appendix A.1 provides more information on sample construction and the ICE-ESG emission measures.[9]

We begin by illustrating the importance of firms' carbon-equivalent emissions for overall emissions. Figure 1 shows the share of global emissions accounted for by firms' emissions in the ICE-ESG sample. We plot estimates separately for the top 20 emitting countries. Data on world-wide emissions are taken from the *Global Carbon Budget report* (2024 estimates). Overall, firms' emissions account for around half of overall emissions into the atmosphere—with U.S. and Chinese firms contributing by far most to the total. Importantly, notice that these estimates, although substantial, present only a lower bound on the importance of firms' emissions, as they only cover firms covered in the ICE-ESG sample. We conclude that firms' direct emissions comprise the lion's share of carbon-equivalent emissions into the atmosphere.[10] Addressing climate change is, to a large extent, a problem of addressing firms' direct, production-based emissions.

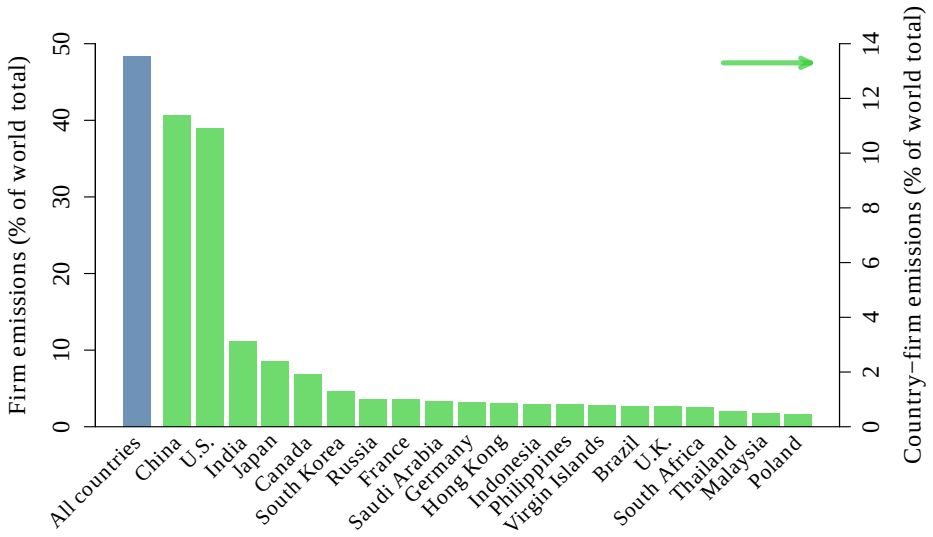
Such emissions are the by-product of firms' supply of goods and services to the economy. The size of overall emissions attributable to firms, visible in Figure 1, thus raises the question of whether—and to what extent—firms differ in the efficiency of their carbon use. Figure 2 addresses this question by plotting the distribution of *emission intensities* in the ICE-ESG and Trucost samples. We measure emission intensity at the firm level by tons of CO_2e (tCO_2e) emissions per million \$ of revenue created. Panel (a) shows substantial heterogeneity in emission intensities in the raw data—similar to that documented in Lyubich et al. (2018), among others.[11] The larger heterogeneity visible in the ICE-ESG sample when compared to Trucost reflects in part the presence of non-publicly-listed firms.

[9] When referring to *carbon emissions*, we in all cases mean *carbon-equivalent emissions*. The ICE-ESG database provides emission estimates broken down into three components: scope 1 (direct emissions from production), scope 2 (indirect emissions from energy use), and scope 3 (all other indirect emissions). We focus on scope 1 emissions throughout our analysis to avoid double-counting emission flows. Figure A.1 in the Appendix reports the average share of scope 1, 2, and 3 emissions for a firm in the ICE-ESG database.

[10] This is in line with several technical reports on the topic (e.g., Report, 2017; IPCC, 2022). Figure A.2 in the Appendix shows the equivalent figure to that in Figure 1 for the S&P Trucost sample.

[11] Capelle et al. (2023), De Lyon and Dechezleprêtre (2025), and Fontaine et al. (2026) make a related point.

Figure 1: Estimated share of overall emissions



Note: Data from the ICE-ESG sample. The figure shows overall firm emissions in 2024 (measured by tons of CO_2e emissions) divided by total world emissions. Estimates for world emissions are taken from the *Global Carbon Budget* website (2024 estimates). Top-20 country-contributions are shown on the right-side axis.

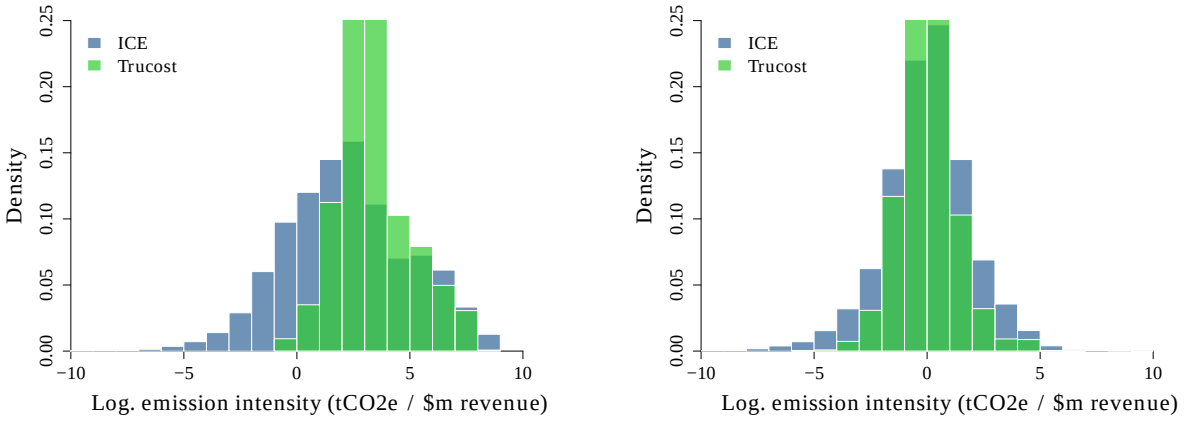
Clearly, part of this heterogeneity is due to differences in output and the production processes used in different sectors and countries at different points in time. To address this issue, Panel (b) in Figure 2 residualizes emission intensities for sector, country, and time fixed effects. Although a 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.[12] Most of the differences in emission intensity are attributable to factors associated with a firm *within* a given sector, country, and time period. Crucially, we show that large-scale heterogeneity is also present in narrowly-defined sectors with relatively homogeneous output, further alleviating concerns that output differences drive the observed heterogeneity.[13] We perform the same analysis for the U.S. energy, airline, and steel industries in Figure A.3 in the Appendix. We measure output by megajoules produced, thousand round-trip flights flown, and tons of steel created. All distributions show substantial variation within sectors, as well as clear differences between them.

A natural question that arises is whether the documented heterogeneity is large compared to other sources of firm heterogeneity. Figure 3 gauges the overall size of this heterogeneity for the case of the U.S. by comparing it to the within-sector heterogeneity that

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

[13] Petersen (2025) finds large heterogeneity in emission intensities within narrowly defined industries in the German manufacturing sector.

Figure 2: Heterogeneity in emission intensities



(a) Overall variation **(b)** Residualized variation

Note: Data from the ICE-ESG and TruCost samples. The figure shows the histogram of log. emission intensities. Emission intensities are measured by tons of CO_2e 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 trim emission intensities at the top and bottom 1% level (Appendix A.1).

exists in firm-level TFP. All else equal, the figure shows a substantially wider distribution for emission intensities than TFP: the level and the dispersion of the unit-free measure—the log within-industry 90/10 difference (Lyubich et al., 2018)—is almost 2-3 times larger than that for TFP.[14] Obviously, this comparison is not informative of the size of the aggregate consequences of this heterogeneity—our model will instead be informative of this—but what it does highlight is that there are indeed large-scale differences in the emission efficiencies when compared to other, standard measures of firm heterogeneity.

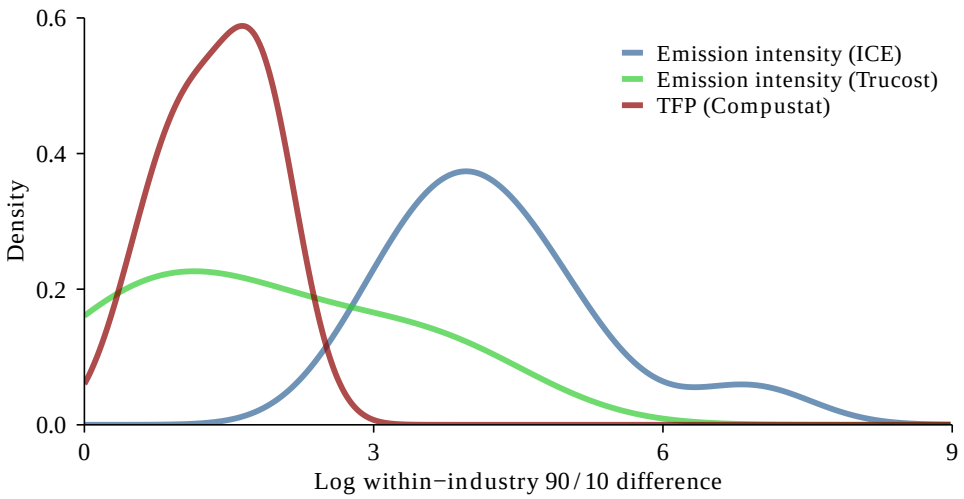
Finally, a related question is whether differences in emission intensity are associated with other measures of firm-level efficiency, such as TFP. Capelle et al. (2023) provide some evidence in this direction. Table A.2 in Appendix shows that, all else equal, more productive firms are also more carbon efficient; however, the overall contribution of this relationship to 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.

2.2 Sources of heterogeneity and disclosures

The above evidence documents substantial heterogeneity in emission intensities. As argued in the engineering and climate literatures, and embedded in the disclosure stan-

[14]Notice that Compustat and Trucost rely on the same sample of firms: U.S. publicly-traded firms.

Figure 3: Emission intensities and TFP



Note: Data from the ICE-ESG, Trucost, and Compustat samples for the U.S.. The figure shows the estimated kernel densities of the log within-industry 90/10 difference of emission intensities and total factor productivity. We use NAICS 2 industry codes for the Trucost and Compustat samples and ICE industry codes for the ICE-ESG sample. The top and bottom 1 percent of emission intensities and total factor productivity estimates have been removed. Total factor productivity is estimated as in Ottonello and Winberry (2020). We consider only firms with high-quality reporting in the ICE sample (Appendix A.1).

dards of major ETS markets, the ultimate source of this heterogeneity can be broken down into 3 components: (i) differences in the *amount of fuel* used in production (e.g., tons of coal or m^3 of natural gas); (ii) differences in the *fuel mix* used (e.g., different types of coal or natural gas); and (iii) differences in the *technological adjustment* (e.g., differences in the oxidation rates of burners, filtration technology employed, or carbon-equivalent capture).[15] Combined, these factors can cause two firms, who otherwise produce a similar type and quantity of output, to feature remarkably dissimilar emission intensities. We discuss how each of these factors lead to heterogeneity in emission intensities in Appendix A.3.

An important feature of the emission data used in Figures 1-3 and those available from other sources is that they are *self-reported* (e.g., ICE, 2024; SEC Carbon Disclosures, 2024). Because the factors highlighted above turn on the specific blend of fuels used, the condition of individual burners, and the filtration technology installed at each installation, among others, a firm's true emission intensity depends on a vector of operational details that outside parties cannot easily assess. Consultancies and outside organizations are sometimes hired to validate the computations by which emission estimates have been at-

[15]See, for example, European Commission Directorate-General for Climate Action (2025) and UK Department for Energy Security and Net Zero (2025).

tained, but the overall inputs into the process, and thus the final outcome, are in all cases up to a firm's discretion (GHG Protocols, 2004; International Sustainability Standards Board, 2023; SEC Carbon Disclosures, 2024). This private-information nature of the data used to compute firms' emissions has, as mentioned above, raised concerns about systematic misreporting among policymakers and practitioners alike (e.g., ECA, 2015; European Commission, 2023a). Such concerns have been compounded by recent enforcement cases that have exposed two distinct modes of disclosure failure—duplicitous disclosures and non-disclosure. Three cases have been particularly influential and are all in the context of the ETSs that operate in Europe:

- *The case of Brikel (EU ETS, 2021-2023):* Two coal-fired power plants in Bulgaria, linked by a common owner, allegedly reported coal emission factors of between 76.2 and 90.5 over 2018–2020, used to compute overall emissions, although they were allegedly burning coal with European Environment Agency reference value of 104.2, in addition to undisclosed industrial waste with a higher emission content. The Organized Crime and Corruption Reporting Project (2021) estimates that the under-declaration saved the two plants at least €26.6 million in EU ETS allowance costs over the period. Both plants' disclosures had been certified as accurate by a third-party verifier. As of early 2026, the investigation is ongoing; Brikel was briefly closed in 2022 but continued to operate during proceedings, incurring administrative fines of only a few thousand euros, a small fraction of the estimated ETS savings.

- *The case of Diageo plc (UK ETS, 2022):* Between 2013 and 2019, three Diageo whiskey-production sites in Scotland carried out regulated combustion activity without the necessary permits, with the consequence that emissions from these sites went unreported for over 6 years. Closely-situated Diageo sites held such permits. Under the U.K. ETS only emission sources over a certain size require permits, and thus payments for carbon allowances. The breach was identified through an externally-triggered audit in 2019, and in March 2022 the Scottish Environment Protection Agency imposed a fine of £1.2 million—against group profits of £3.7 billion in 2021.[16]

- *The case of ExxonMobil (EU ETS, 2010):* In its 2008 disclosure under the EU ETS, ExxonMobil omitted three emission sources at its Mossmorran ethylene plant in Fife, altogether, understating reportable CO2e by approximately 33,000 tonnes per year. The Scottish Environment Protection Agency, warned of the under-reporting

[16] Furthermore, notice that, according to Calel et al. (2025), the observed overall non-compliance in the EU ETS should have resulted in €13 billion in total fines, but only €2.1 billion appear to have been collected.

Table 1: Emission disclosures and emission intensity

	Log. emission intensity		
	(1)	(2)	(3)
No third-party validation	-0.193^{***}	-0.214^{***}	-0.205^{***}
	(0.031)	(0.023)	(0.023)
No third-party validation $\times$ log(90/10)	.	.	-0.051^{***}
			(0.017)
Fixed effects			
Time		✓	✓
Sector		✓	✓
Country		✓	✓
Observations	32,268	32,268	30,298

Notes: Least-squares estimates from the ICE-ESG sample. The top and bottom 1 percent of emission intensities (tons of CO_2e emissions divided by revenue in millions of \$) are removed. ICE sectoral definitions and verification status are used. Robust clustered (sectors) standard errors in parentheses. Sample: 2010–2024.

by a whistleblower, imposed a fine of £2.77 million in September 2010.[17]

Appendix A.3 discusses each case in detail and explains how such failures could arise within existing disclosure practices and regulations governing ETS protocols.

Although the overall scope of misreporting is obviously unknown, a clear advantage of the ICE-ESG database is that it contains information about which firms have their disclosures validated (i.e., have their emission calculations checked). Indeed, we restricted Figures 1 and 3 to observations that ICE classifies as "higher quality"—that is, have their disclosures third-party validated. Table 1 shows that this restriction matters: firms without third-party validation tend to report systematically lower and more optimistic emission intensities, all else equal, suggesting scope for systematic under-reporting.[18] Moreover, firms that are in sectors with more dispersed emission intensities appear to underreport more, consistent with these sectors having larger "scope for misreporting."

In summary, in this section, we have documented that firms' emissions account for

[17] Notice that the UK implementation of the EU ETS already embeds a per-tonne under-reporting penalty well above the allowance market price. The mandatory civil penalty for undeclared emissions is €100 per tCO2e, irrespective of intent—several multiples of the prevailing allowance price over most of the sample period, and the rate at which SEPA's 2010 £2.77m fine against ExxonMobil Chemical was mechanically computed. That the Brikel and Diageo cases nevertheless document multi-year disclosure failure indicates that the effective deterrent depends jointly on the penalty rate and the probability of detection.

[18] Notice that self-selection in validation seems less of a concern: if clean firms verify, verified firms should report lower intensities, not higher, as we find in Table 1.

the lion's share of overall 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 macro-climate model to allow for heterogeneity in emission intensity and asymmetric information. We then proceed by characterizing and quantifying the impact of the combination of heterogeneous emission intensities and information frictions for the optimal taxation of firms in general equilibrium.

3 Baseline model framework

We begin 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, 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 recursion:

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

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, who live for L periods, owned by households. Each firm is endowed at entry with a fixed *productivity* a_i and *carbon intensity* z_j, where pairs $\{a_i, z_j\} \in \mathcal{A} \times \mathcal{Z} \subseteq \mathbb{R}^2$ are jointly distributed with mass ω_{ij}, satisfying $\sum_{ij} \omega_{ij} = 1$. We assume $\mathcal{A} \times \mathcal{Z}$ is compact and index a firm of type (i,j) with age $\ell \leq L$ by

(i, j, ℓ). Upon exit at age L, firms are replaced by an entering cohort of age $\ell = 1$ of equal size, implying a stationary age distribution with mass $1/L$ at each ℓ. We let $w_{ij\ell} \equiv \omega_{ij}/L$ denote the joint p.d.f. over types and ages, so that $\sum_{ij\ell} w_{ij\ell} = 1$. Firm (i, j, ℓ) produces quantity $y_{ij\ell,t}$ of output at time t according to the standard production technology:

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

where $n_{ij\ell,t}$ and $k_{ij\ell,t}$ are the units of labor and capital employed by the firm, respectively, $e_{ij\ell,t}$ 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 ($A_G < 0$), so that increases in carbon concentration damage the productive-capabilities of the economy. Given an initial concentration $G_0 > 0$, the stock of carbon in the atmosphere evolves in accordance with:

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

where $g_{ij\ell,t}$ are firm (i, j, ℓ)'s emissions and $\delta_G \in [0, 1)$ is the rate of depreciation of carbon in the atmosphere. Firm (i, j, ℓ) produces emissions due to its energy use:

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

where $\exp^{z_j - x_{ij\ell,t}}$ is the *effective carbon intensity* and $x_{ij\ell,t}$ is the firm's expenditure on the *abatement* of its emissions at cost $Q\left(x_{ij\ell,t}\right)$ with $Q' > 0$, $Q'' > 0$ and $Q(0) = 0$. Energy is available to firms in infinite supply at price $P > 0$.[19] Finally, firms rent capital and labor in competitive markets with factor prices equal to R_t and W_t, respectively.

Information structure. A central friction in our economy is that a firm's carbon intensity z_j is *private information* to the firm. When choosing its energy use, a firm of type (i, j, ℓ) knows z_j and thus the emissions $g_{ij\ell,t}$ that arise as a result of its energy use and abatement. However, this information is not shared with other firms in the economy or the social planner.[20] Clearly, this assumption presents a natural limit case of an information friction that exists between firms and the social planner about firms' carbon emissions, and we later study other cases in which the planner instead has partial (but incomplete)

[19] 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.

[20] Consistent with the literature on information frictions and asymmetric information, we nevertheless assume that the distribution of z_j is known to all in the economy.

information about firms' carbon-intensity type (Section 5.4).

Discussion of the environment. The above environment includes four features central to the carbon taxation of firms. First, in the above environment, firms differ along two dimensions, in addition to their age: *carbon intensity* z_j and *productivity* a_i. Below, we show that jointly accounting for both sources of heterogeneity is crucial for the effective design of corrective policy. In particular, the effectiveness of corrective policy interventions will depend centrally on the relative importance of productivity and carbon intensity in shaping the first-best allocation of energy and abatement across firms.

Second, our baseline framework links firms' emissions to their energy use through an effective carbon intensity $\exp^{z_j - x_{ij\ell,t}}$ that varies across firms (Equation 5). As we shall see below, this heterogeneity maps one-for-one into differences in the emission intensity measure—emissions per unit of revenue—shown in Figure 2. Cross-firm variation in emission, in turn, reflects both differences in the amount of energy use and in the composition of fuels and technologies used (Section 2.2 and Appendix A.3), captured here by the term $\exp^{z_j - x_{ij\ell,t}}$. While abatement is the primary margin through which firms reduce their emissions intensity in our baseline model, we extend the framework in Section 5.4 to also allow firms to choose between *renewable energy*, which is carbon-free, and *fossil-based energy*. This follows the approach in Golosov et al. (2014).

Third, our environment allows firms to live for multiple periods, enabling us to study the static and dynamic incentives faced by a social planner in designing optimal corrective policy. Specifically, it allows us to analyze how current and future disclosures interact, and potentially modify the design of optimal corrective policy.

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 this is a reasonable approximation given available medium-complexity climate models used in the natural sciences. Section 5.4 analyzes the case in which economy-wide productivity instead depends negatively on the global-mean temperature T_t, which in turn depends in a near 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, which 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 nevertheless stress that our main results do not hinge on the speed over which economic damages from climate change unfold or whether they arise from carbon concentration or are modulated through temperature. Moreover, our assumption that it is only the production-side of the econ-

omy that is affected by emissions is made to keep our analysis comparable to those in 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 difference between them being the extent to which firms' energy use internalize the emission externality, $A(G_t)$. We start with the decentralized equilibrium.

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:

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

where $\pi_{ij\ell,t}$ denotes the period-t profits of firm (i, j, ℓ). 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, ℓ) is similarly standard, as firms in the decentralized equilibrium do not internalize the damages their emissions have on the productive capabilities of the economy. Instead, firm (i, j, ℓ) chooses labor $n_{ij\ell,t}$, capital $k_{ij\ell,t}$, energy $e_{ij\ell,t}$, and abatement $x_{ij\ell,t}$ to maximize its present discounted value of profits:

$$v_{ij\ell,t} = \pi_{ij\ell,t} + \mathbb{1}_{\{\ell<L\}} \cdot \beta \cdot \frac{u'(C_{t+1})}{u'(C_t)} \cdot v_{ij,\ell+1,t+1}, \tag{8}$$

where per-period profits are given by:

$$\pi_{ij\ell,t} = y_{ij\ell,t} - W_t \cdot n_{ij\ell,t} - R_t \cdot k_{ij\ell,t} - P \cdot e_{ij\ell,t} - Q(x_{ij\ell,t}), \tag{9}$$

subject to (3). The firm thus optimally equates the *marginal revenue product* of labor and capital with the real wage W_t and the rental rate of capital R_t, respectively:

$$mrpn_{ij\ell,t} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_n(\cdot) = W_t \quad \text{and} \quad mrpk_{ij\ell,t} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_k(\cdot) = R_t, \tag{10}$$

while the relevant conditions for energy $e_{ij\ell t}$ and abatement $x_{ij\ell t}$ are:

$$mrpe_{ij\ell,t} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_e\left(n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}\right) = P \quad \text{and} \quad Q'\left(x_{ij\ell,t}\right) = 0. \tag{11}$$

Equation (11) presents a main departure point for our analysis: a profit-maximizing firm (i, j, ℓ) optimally equates the marginal revenue product of energy $mrpe_{ij\ell,t}$ to its price P and chooses zero abatement, $x_{ij\ell t} = 0$. Importantly, notice that Equations (10) and (11) show that a firm's choice of inputs are also independent of its carbon intensity z_j and age ℓ. The reason is that neither carbon intensity nor age alter the firm's productive capabilities in Equation (3).

We are now ready to define the equilibrium of our economy:

Equilibrium definition and characterization. A competitive equilibrium consists of household allocations $\{C_t, K_{t+1}\}_{t=0}^{\infty}$, firm allocations $\left\{\{n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t}\}_{i,j,\ell}\right\}_{t=0}^{\infty}$, factor prices $\{W_t, R_t\}_{t=0}^{\infty}$, and firm-driven emissions $\{\{g_{ij\ell,t}\}_{i,j,\ell}, G_t\}_{t=0}^{\infty}$ 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)-(11) hold).
- Given allocations, the labor, capital, and goods market clears; i.e., for all t:

$$\sum_{ij\ell} w_{ij\ell} \cdot n_{ij\ell,t} = N \quad \text{and} \quad \sum_{ij\ell} w_{ij\ell} \cdot k_{ij\ell,t} = K_t, \tag{12}$$

  and:

$$C_t + K_{t+1} = (1 - \delta_K) \cdot K_t + \sum_{ij\ell} w_{ij\ell} \cdot \left(y_{ij\ell,t} - P \cdot e_{ij\ell,t} - Q(x_{ij\ell,t})\right). \tag{13}$$

- Emissions $\{g_{ij\ell,t}\}$ and G_t evolve in accordance with Equations (4) and (5).

The decentralized equilibrium features a tight link between a firm's carbon-intensity type z_j and the emission-efficiency measure $\log(g_{ij\ell,t}/y_{ij\ell,t})$ shown in Figure 2. Indeed, it follows directly from Equation (11) that, in the case in which $F(\cdot)$ is Cobb-Douglas, we have that: $\log(g_{ij\ell,t}/y_{ij\ell,t}) = \log(\nu/P) + z_j$, where $\nu \equiv \partial \log y_{ij\ell,t}/\partial \log e_{ij\ell,t}$.[21] Variations in z_j translate one-for-one into variations in emission intensity, and are the ultimate source behind this heterogeneity in the model. 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 discussions above in Lemma 1:

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

[21] In particular, we have that: $\log\left(\frac{g_{ij\ell,t}}{y_{ij\ell,t}}\right) = \log\left(P \cdot \frac{\partial \log e_{ij\ell,t}}{\partial \log y_{ij\ell,t}}\right) + z_j = \log\left(\frac{\nu}{P}\right) + z_j$.

i. Firm (i, j, ℓ)'s energy use $e_{i,t}^{DE}$ and abatement $x_{i,t}^{DE}$ satisfy:

$$mrpe_{i,t}^{DE} = P \quad \textit{and} \quad Q'\left(x_{i,t}^{DE}\right) = 0, \quad \forall j, \ell, \tag{14}$$

and are independent of the firm's carbon-intensity type j and age ℓ.

ii. Variations in carbon intensity z_j translate one-for-one into variations in emission efficiency:

$$\log\left(g_{ij,t}^{DE}/y_{i,t}^{DE}\right) = \log(\nu/P) + z_j + \mathcal{O}(1), \tag{15}$$

where $\partial \log y_{i,t}/\partial \log e_{i,t} = \nu + \mathcal{O}(1)$.

Finally, notice that a key feature of the decentralized equilibrium is that firms do not internalize the damages from carbon emissions, $A(G_t)$, when making their energy and abatement choices. 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 (9) being independent of emissions and carbon intensity, $\partial \pi_{i,t}/\partial g_{ij,t} = \partial \pi_{i,t}/\partial z_j = 0$, leading a firm's choice of all inputs to be independent of its carbon-intensity type. 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 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 $\{z_j\}$, and thus their emissions $\{g_{ij\ell,t}\}$. Although our ultimate interest lies with the solution to the planning problem in which firms have private information, it will be useful to start with the full-information case, as the comparison between the two cases will allow us to sharply characterize the costs of information frictions. The full-information social planner solves the problem:

$$\max_{\{C_t, K_{t+1}, \{k_{ij\ell,t}, n_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t}\}_{i,j,\ell}\}_{t=0}^{\infty}} \mathcal{U} = \sum_{t=0}^{\infty} \beta^t \cdot u\left(C_t\right) \tag{16}$$

subject to Equations (3), (4), (5), (12), and (13), as well as the usual non-negativity constraints (Appendix B.1). 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:

$$u'\left(C_t\right) = \beta \cdot u'\left(C_{t+1}\right) \cdot \left(\Lambda_{K,t+1} + 1 - \delta_K\right). \tag{17}$$

Let $\Lambda_{N,t}$ analogously denote the constraint multiplier on the market-clearing condition for labor at time t. The full-information planner then further optimally equates the marginal revenue product for labor and capital to the multipliers $\Lambda_{N,t}$ and $\Lambda_{K,t}$, respectively, while the relevant conditions for energy $e_{ij\ell,t}$ and abatement $x_{ij\ell,t}$ are:[22]

$$mrpe_{ij\ell,t} = P + \Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}} \quad \text{and} \quad Q'\left(x_{ij\ell,t}\right) = \Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t}, \tag{19}$$

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

$$\Lambda_{G,t} = \beta \cdot \frac{u'(C_{t+1})}{u'(C_t)} \cdot \left(-\frac{A'(G_{t+1})}{A(G_{t+1})} \cdot \mathcal{Y}_{t+1} + \Lambda_{G,t+1} \cdot (1 - \delta_G)\right), \tag{20}$$

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

Equations (17)-(20) characterize the full-information planner allocation and can be directly compared to Equations (7), (10)-(11), which describe the decentralized equilibrium. Notice that these equations once more provide solutions $\{n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t}\}$ that are independent of firm age ℓ, as age does not alter the productive capabilities of firms.[23] The central difference between the two sets of equations arises because $\Lambda_{G,t} > 0$ in Equation (19). This term, which captures the relevant notion of the *social cost of carbon* for our economy, equals from Equation (20) the present discounted value of marginal externality damages, $-\frac{A'(G_{t+1})}{A(G_{t+1})} \cdot \mathcal{Y}_{t+1}$, similar to findings of Golosov et al. (2014).

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

The social cost of carbon is then multiplied in Equation (19) by $\exp^{z_j - x_{ij,t}}$ to create firm (i,j,ℓ)'s *marginal social cost of energy*, which consequently depends on the firm's carbon intensity z_j and its abatement $x_{ij,t}$. The left-hand side of Equation (19) shows that the full-information planner optimally equates firm (i,j,ℓ)'s marginal revenue product of energy,

[22] The relevant conditions for labor $n_{ij\ell,t}$ and capital $k_{ij\ell,t}$ mirror those from the decentalized equilibrium:

$$mrpn_{ij\ell,t} = \Lambda_{N,t} \quad \text{and} \quad mrpk_{ij\ell,t} = \Lambda_{K,t}. \tag{18}$$

[23] We thus suppress the use of the subscript ℓ for the rest of this subsection when it does not cause confusion.

$mrpe_{ij\ell,t}$, to its overall societal cost, equal to the sum of the market price $P > 0$ and its marginal social cost of energy $\lambda_{G,ij,t} > 0$. The decentralized equilibrium, by contrast, only equates it to the market price (Equation 14). The right-hand side of Equation (19) then further shows that the full-information planner also optimally equates the marginal cost of abatement $Q'\left(x_{ij,t}\right)$ to its *marginal social return* $\Delta_{ij,t}$, which scales with firm (i,j,ℓ)'s energy use $e_{ij,t}$ in addition to its marginal social cost of energy $\lambda_{G,ij,t}$.

Thus, because of the internalization of the emission externality—driven by 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 (17)-(20) that:

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

i. Firm (i,j,ℓ)'s energy use $e^{FB}_{ij,t}$ and abatement $x^{FB}_{ij,t}$ satisfy:

$$mrpe^{FB}_{ij,t} = P + \lambda_{G,ij,t} \quad \textit{and} \quad Q'\left(x^{FB}_{ij,t}\right) = \Delta_{G,ij,t}, \quad \forall \ell, \tag{21}$$

and are independent of the firm's age ℓ.

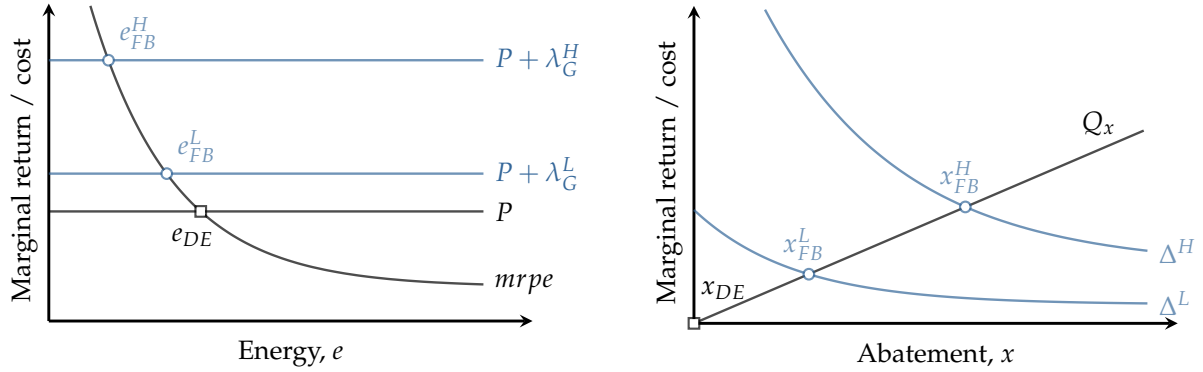
ii. The planner depresses energy use and increases abatement efforts heterogeneously across firms. In particular, for any i and any (j',j) with $z_{j'} > z_j$:

$$e^{FB}_{ij',t} < e^{FB}_{ij,t} < e^{DE}_{i,t} \quad \textit{and} \quad x^{FB}_{ij',t} > x^{FB}_{ij,t} > x^{DE}_{i,t} = 0. \tag{22}$$

Since the marginal social cost of energy is positive, the full-information planner depresses all firms' energy uses relative to what firms would choose to maximize their value, $e^{FB}_{ij,t} < e^{DE}_{ij,t}$ (Equation 14 and 21). 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, $x^{FB}_{ij,t} > x^{DE}_{i,t} = 0$. A larger carbon intensity z_j, or a larger externality damage $-A'(G_{t+1})/A(G_{t+1}) \cdot \mathcal{Y}_{t+1}$, all else equal, causes a firm to exhibit a larger marginal social cost of energy, 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 energy and abatement allocations are modified by the full-information planner relative to what a firm would choose in equilibrium.

Crucially, notice that the marginal social cost of energy, $\lambda_{G,ij,t} > 0$, is *type-specific*: each type (i,j) firm features its *own* marginal social cost of energy. Because the planner optimally equates each firm's marginal revenue product of energy to the sum of its market

Figure 4: The full-information social planner allocation



(a) Energy **(b)** Abatement

Note: The figure illustrates the determination of energy use (Panel a) and abatement (Panel b) under the decentralized equilibrium (black) and the full-information optimum (blue). The figure considers the case in which firms have either a high or low marginal social cost of energy, $\lambda^H_G > \lambda^L_G$, implying respectively a high or low marginal return to abatement.

price and the firm-specific marginal social cost of energy, it follows that the planner introduces *factor misallocation* across firms, to optimally account for the emission externality. This feature of the full-information 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 and value 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 Equations (17)-(21) that the implied profits under the planner allocation will generically differ between the two firms, $\pi^{FB}_{ij\ell,t} \neq \pi^{FB}_{ij'\ell,t}$ for all ℓ, triggering, all else equal, different firm values, $v^{FB}_{ij\ell,t} \neq v^{FB}_{ij'\ell,t}$. 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 such firms to attain higher profits and values. Indeed, it follows that:

Corollary 1 (Incentive-incompatibility). *In the full-information allocation, the implied value of firm (i, j, ℓ) is larger than those of firm (i, j', ℓ), $v^{FB}_{ij\ell,t} > v^{FB}_{ij'\ell,t}$ for all (i, ℓ) when $z_j < z_{j'}$.*

As we will show, the full-information allocation can be implemented with a uniform linear carbon tax. In the absence of perfect information, however, any value differential across otherwise similar firms generates incentives for high carbon-intensive firms—those with larger values of z_j—to misreport their emissions intensity in order to increase firm value. Such behavior in fact becomes part of firms' optimal disclosure strategy and undermines the effectiveness of the tax, as we demonstrate below.

A naïve planner. Suppose a *naïve planner (NV)* was to act as the full-information planner presented above but without having direct 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., $\underline{z} \equiv \min_{z \in \mathcal{Z}} z_j = \arg\max v_{ij\ell,t}$.[24] A firm would always have a *dominant reporting strategy* given by type $(i, \underline{j}, \ell)$, where $\underline{j}$ is the index of $\underline{z}$. The first-order conditions for energy and abatement would then become:

$$mrpe^{NV}_{ij,t} = P + \lambda_{G,i\underline{j},t}, \quad Q'(x^{NV}_{ij,t}) = \Delta_{G,i\underline{j},t}, \tag{23}$$

where $\lambda_{G,i\underline{j},t} = \Lambda_{G,t} \cdot \exp^{\underline{z} - x_{ij,t}}$ is the *misreported marginal social cost of energy* and $\Delta_{G,i\underline{j},t} = \lambda_{G,i\underline{j},t} \cdot e_{ij,t}$ is the *misreported marginal return of abatement*. Clearly, both conditions imply that firms would use more energy and do less abatement than in the full-information case (Equation 21), implying a smaller internalization of the emission externality. Depending on the dispersion of emission intensities, ignoring misreporting incentives could, as a result, substantially reduce the above-outlined full-information benefits.

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 firms' carbon intensity $\{z_j\}$. This Mirrleesian planner faces the same problem as the full-information planner but, in addition, faces two constraints: (i) an *incentive-compatibility constraint* and (ii) an *individual-rationality constraint*.

Since a firm of type (i, j, ℓ) can always pretend to be of type (i, j', ℓ), where $j \neq j'$, in the planner's allocation, each firm-type's value must exceed those from another carbon-intensity type (*incentive compatibility*):[25]

$$v(\iota_{ij\ell,t}, v_{ij,\ell+1,t+1} \mid a_i, z_j, \ell, W_t, R_t) \geq v(\iota_{ij'\ell,t}, v_{ij',\ell+1,t+1} \mid a_i, z_j, \ell, W_t, R_t) \quad \forall j' \neq j, \tag{24}$$

where $\iota_{ij\ell,t} \equiv \left[n_{ij\ell,t}\ k_{ij\ell,t}\ e_{ij\ell,t}\ x_{ij\ell,t}\right]$ denotes the vector with the firm's input choices, and the right-hand side of (24) is the value of firm (i, j, ℓ) when it receives the allocation intended for firm (i, j', ℓ). Moreover, to be willing to participate in the mechanism, a firm of

[24] Recall that Section 2 showed that firms who do not have their emission calculations checked report systematically lower emissions, and that firms who operate in sectors with more dispersed emission intensities "mistate" more.

[25] 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, ℓ) can pretend to be of type (i', j, ℓ) and thus no incentive compatibility constraints are needed on this dimension.

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

$$v(\iota_{ij\ell,t}, v_{ij,\ell+1,t+1} \mid a_i, z_j, \ell, W_t, R_t) \geq 0. \tag{25}$$

The problem faced by the planner subject to these additional constraints 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 (9) do not depend on the firm's carbon-intensity type, $\partial \pi_{ij\ell,t} / \partial z_j = 0$.[26] This contrasts with the household case in which the household's unobserved type (i.e., the effort-type) matters *directly* for the household's payoff through its earned income.

We add two further conditions to the problem. An important feature of our framework is that, although the use of energy has an environmental externality, the decentralized use of labor and capital is, all else equal, efficient. The full-information planner equates the marginal revenue product of these inputs to their respective prices. This, in turn, implies that the Mirrleesian planner has an incentive to distort labor and capital markets to relax the above incentive and rationality constraints. A "Theory of the Second Best" logic applies to our framework (Lipsey and Lancaster, 1956)—a logic that may be challenging to implement in practice. Rather than allowing for such allocations, we focus in the following on the more constrained (and arguably more realistic) allocations in which:

$$W_t = \Lambda_{N,t} \quad \text{and} \quad R_t = \Lambda_{K,t}, \tag{26}$$

where $\Lambda_{N,t}$ and $\Lambda_{K,t}$ once more denote the multipliers on the market-clearing conditions for labor and capital. That is, we assume that the (implied) wages and interest rates used by firms in Equations (24) and (25) to compute their values equal the associated multipliers, thus ensuring that the planner leaves labor and capital markets undistorted. As we will show in Section 5.2, imposing these constraints does, in either case, not matter much quantitatively for the relative benefits of the Mirrleesian approach.

The Mirrleesian planner thus solves the problem:

$$\max_{\{C_t, K_{t+1}, W_t, R_t, \{k_{ij\ell,t}, n_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t}\}_{i,j,\ell}\}_{t=0}^{\infty}} \mathcal{U} = \sum_{t=0}^{\infty} \beta^t \cdot u\left(C_t\right) \tag{27}$$

subject to the aforementioned Equations (3), (4), (5), (12), and (13), in addition to Equations (24), (25), and (26) above, as well as the usual non-negativity constraints.

[26] 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.

At first pass, the Mirrleesian planner's problem may seem daunting. The non-convexity of the constraint set, caused by the incentive constraints, precludes the use of standard methods. The direct pay-off irrelevance of a firm's carbon-intensity type, $\partial \pi_{ij\ell,t}/\partial z_j = 0$, combined with the monotone ranking of firm profits in z_j in the full-information allocation (Corollary 1), however, provides a way forward. A consequence of the two is that each firm has a dominant strategy: to (mis)report its carbon-intensity type and claim to be of type $\underline{z} \equiv \min_{z \in \mathcal{Z}} z_j$, as is the case under the naïve planner. It follows that a firm's profits in the Mirrleesian case must be independent of its carbon-intensity type, $\pi_{ij\ell,t} = \pi_{i\ell,t}$ for all j, and so too its value, $v_{ij\ell,t} = v_{i\ell,t}$ for all j. We show in Appendix B.1 that such value equalization—the fact that the incentive-compatibility constraint binds for all j—in turn implies that the planner chooses a symmetric allocation such that $\{k_{ij\ell,t}, n_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t}\} = \{k_{i,t}, n_{i,t}, e_{i,t}, x_{i,t}\}$ for all (j, ℓ).[27]

Crucially, notice that this allocation is also independent of firm age ℓ. We show in Appendix B.1 that the latter follows from the technologies available to a firm being independent of its age, and the fact that we consider a planning problem under *full commitment*. The latter ensures that the Mirrleesian planner does not choose allocations to first incentivize disclosures that it will later use to modify firms' input allocations. Although we view the case *without commitment* to be interesting, we view the case *with commitment* to ultimately be the case of greater applied interest.

A simple example can help with the intuition behind the symmetry of the optimal allocation. Suppose energy is the only input into production, that firms share the same productivity, live for only one period, and that there is no abatement. Suppose further that for a firm of type j, there are two energy allocations, $e_{1,jt}$ and $e_{2,jt}$, which satisfy incentive-compatibility, $\pi(e_{1,jt}|z_j) = \pi(e_{2,jt}|z_j)$, with $e_{1,jt} < e_{2,jt}$. This is generically the case. 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 $e_{1,jt}$ and $e_{2,jt}$ is incentive compatible and delivers the same welfare absent the externality. Yet, because emissions enter through Equation (3), and thus create a negative externality, $\partial A(G_t)/\partial G_t < 0$, the Mirrleesian planner prefers the lower energy level $e_{1,jt}$ for all type-j firms.

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

Definition 2. *Let* $\overline{\lambda}_{G,i,t} \equiv \mathbb{E}[\lambda_{G,ij,t} \mid a_i] > 0$ *be type* i*'s* ***average social cost of carbon*** *and* $\overline{\Delta}_{G,i,t} \equiv \overline{\lambda}_{G,i,t} \cdot e_{i,t} > 0$ *be type* i*'s* ***average social return on abatement****.*

[27] The Inada condition on the production technology, $F(\cdot)$, ensures that the individual-rationality constraint in Equation (25) never binds.

Recall that $\Lambda_{K,t}$ and $\Lambda_{N,t}$ denote the constraint multipliers on the market-clearing condition for capital and labor. The Mirrleesian planner, as the full-information planner, equates the marginal revenue product for capital $k_{i,t}$ and labor $n_{i,t}$ to the multipliers $\Lambda_{K,t}$ and $\Lambda_{N,t}$, respectively.[28] The planner's conditions for K_{t+1} and C_t also again yield the standard Euler equation for consumption in (17), while the social cost of carbon $\Lambda_{G,t}$ equals that in (20). 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, ℓ)'s energy use $e^{MR}_{i,t}$ and abatement $x^{MR}_{i,t}$ satisfy:

$$mrpe^{MR}_{i,t} = P + \overline{\lambda}_{G,i,t}, \qquad Q'\left(x^{MR}_{i,t}\right) = \overline{\Delta}_{G,i,t}, \quad \forall j, \ell, \tag{29}$$

and are independent of the firm's carbon-intensity type j and age ℓ.

ii. The Mirrleesian planner depresses energy and increases abatement, while bunching firms across carbon intensities and age. In particular, for any i, j, ℓ:

$$e^{MR}_{i,t} < e^{DE}_{ij\ell,t} \quad \textit{and} \quad x^{MR}_{i,t} > x^{DE}_{ij\ell,t}.$$

Figure 5 illustrates how the Mirrleesian planner allocations differ from those considered above. 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 z_j, she cannot equate a firm's marginal revenue product of energy to its marginal social cost of energy. Instead, the Mirrleesian planner optimally averages and equates a firm of type (i, j)'s marginal revenue product of energy to group-i's *average* marginal social cost of energy. 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.[29]

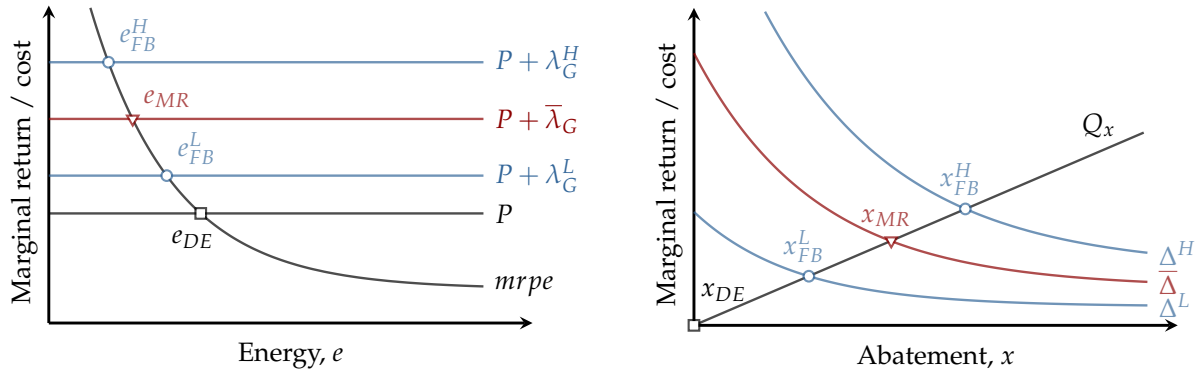
Finally, notice two features of the Mirrleesian allocation. First, because of the bunching of allocations within productivity types, the Mirrleesian allocation features, all else equal, less *factor misallocation* than the full-information case. This is although the Mirrleesian

[28] Thus,

$$mrpn_{i,t} = \Lambda_{N,t}, \quad \text{and} \quad mrpk_{i,t} = \Lambda_{K,t}. \tag{28}$$

[29] Notice that the average marginal social cost of energy $\bar{\lambda}_{G,i,t}$ is, in effect, the social planner's "best guess" (Bayesian optimal) for the individual marginal costs $\lambda_{G,ij,t}$.

Figure 5: The Mirrleesian social planner allocation



(a) Energy **(b)** Abatement

Note: The figure shows the determination of energy (Panel a) and abatement (Panel b) in the decentralized equilibrium (black), the full-information optimum (blue), and the Mirrleesian optimum (orange). The figure considers the case in which firms have either a high or low marginal social cost of capital, $\lambda^H_G > \lambda^L_G$, implying respectively a high or low marginal benefit of abatement. The average value is denoted by $\overline{\lambda}_G$.

planner clearly cannot be as effective as the full-information planner in addressing the economic damages from climate change. Second, despite the fact that a firm's carbon-intensity type is fixed across time, the planner never *experiments* with allocations across firm age, to learn a firm's carbon-intensity type and then change the allocation later on. Instead, as we show through a backwards induction argument in Appendix B.1, the full-commitment planner always sets allocations independent of firm age ℓ.

4.4 Implementation and carbon taxation

Both the Mirrleesian allocation and its full-information counterpart can straightforwardly be implemented by a simple, linear carbon tax in the decentralized equilibrium. To see this, let $\tau_{ij\ell,t}(g)$ denote the carbon-tax payments made by firm (i, j, ℓ) when its emissions equal g. Firm (i, j, ℓ)'s post-tax profits then become:

$$\pi_{ij\ell,t} = y_{ij\ell,t} - W_t \cdot n_{ij\ell,t} - R_t \cdot k_{ij\ell,t} - P \cdot e_{ij\ell,t} - Q(x_{ij\ell,t}) - \tau_{ij\ell,t}\left(g_{ij\ell,t}\right). \tag{30}$$

The maximization of firm (i, j, ℓ)'s value subject to the profit expression in (30) now results in the following *tax-adjusted* conditions for energy $e_{ij\ell,t}$ and abatement $x_{ij\ell,t}$:

$$mrpe_{ij\ell,t} = P + \tau'_{ij\ell,t}\left(g_{ij\ell,t}\right) \cdot \exp^{z_j - x_{ij\ell,t}}, \quad Q'\left(x_{ij\ell,t}\right) = \tau'_{ij\ell,t}\left(g_{ij\ell,t}\right) \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t}; \tag{31}$$

while the conditions for labor $n_{ij\ell,t}$ and capital $k_{ij\ell,t}$, and the Euler equation for consumption C_t, are identical to those in Section 4.1. Comparing Equation (31) with the corre-

sponding equations in Lemma 2 and 3 then immediately shows that:

Proposition 1. *Let the carbon tax schedule for firm* (i, j, ℓ)*,* $\tau_{ij\ell,t}(g)$*, equal:*

$$\tau_{ij\ell,t}(g) = \Lambda_{G,t} \cdot g \equiv \tau_t^{FB}(g) \quad or \quad \tau_{ij\ell,t}(g) = \Lambda_{G,t} \cdot \frac{\mathbb{E}[\exp^{z_j} | a_i]}{\exp^{z_j}} \cdot g \equiv \tau_{ij,t}^{MR}(g), \tag{32}$$

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 $\tau_{ij\ell,t}(g) = \tau_t^{FB}(g)$ *and firms disclose truthfully without explicit justification:; and (ii) the Mirrleesian social planner allocation when* $\tau_{ij\ell,t}(g) = \tau_{ij,t}^{MR}(g)$*.*

Proposition 1 shows that a simple, linear carbon tax, equal to the social cost of carbon $\Lambda_{G,t}$ implements the full-information allocation when firms truthfully disclose. The proposition, as a result, recovers 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 production from any additional unit of carbon is identical across firms: $A(G_t)$ 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 information frictions, however, the optimal carbon tax becomes *type-specific*.[30] Indeed, 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, $\Lambda_{G,t} \cdot (\mathbb{E}[\exp^{z_j} | a_i] / \exp^{z_j})$, is still constant but depends on the carbon-intensity type of the firm, and through the conditional expectation term may also vary with the productivity type. This dependence arises because of the incentive-compatibility constraint in (24). To equalize profits within productivity groups—and thus prevent misreporting—the planner must set a *higher marginal tax* for less carbon-intensive firms. This, at first glance, counterintuitive result arises because more carbon-intensive firms have a stronger incentive to misreport emissions.

Figure 6 illustrates the two log-linear marginal tax rates in Proposition 1 as a function of the carbon-intensity type. It follows from Proposition 1 that the Mirrleesian rate exceeds the full-information rate for firms whose carbon intensity is below their group average, while it falls below it for firms whose carbon intensity is above the group average. Importantly, the two taxes are conceptually similar: the Mirrleesian schedule merely pro-

[30] This is a direct implication—and, arguably, an advantage—of constraining the planner not to distort labor and capital markets through the additional constraint in Equation (26). Absent this restriction, the planner would generally choose a (unrealistic) corrective policy that could not be implemented through a single carbon tax alone.

portionally adjusts the full-information tax by $\frac{\mathbb{E}[\exp^{z_j}|a_i]}{\exp^{z_j}}$ to ensure incentive compatibility. This ratio-based (or "predictive-relative-to-actual") adjustment determines the distinctive shape in Figure 6. In this sense, the Mirrleesian tax is no 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 in recovering the potential benefits of carbon taxation in environments where firms can misreport emissions.

As we will argue, the reason for this is in part that—although Mirrleesian marginal tax rates $\tau_{ij,t}^{\text{MR}}$ vary systematically across the carbon-intensity distribution—the overall tax share of revenue $\tau_{ij,t}^{MR}\left(g_{ij,t}^{MR}\right)/y_{i,t}^{MR}$ is constant across z_j in the Mirrleesian case.[31] This follows from the profit and value equalization implied by incentive compatibility. Indeed, we will show that the overall tax bill will mainly remain "small".

Finally, Figure 6 also shows the evolution of the effective tax rate that would prevail with the naïve planner encountered earlier, who acts as the full-information planner but without having accurate information on carbon-intensity types. It follows from comparing Equation (23) with Equation (31) that the effective tax implied by the naïve planner allocation is also linear within firms and equal to:

$$\tau_{j,t}^{NV}(g) \equiv \Lambda_{G,t} \cdot \exp^{\bar{z}-z_j} \cdot g_t. \tag{33}$$

Notice that the naïve tax rate *declines* in z_j, 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 to decompose and quantify the benefits of a Mirrleesian approach to taxation, we should mention that other, more elaborate policies exist that also implement the planners' allocations. A tax on energy combined with a subsidy on abatement can, for example, also attain this aim in the Mirrleesian case (Appendix B.2).[32]

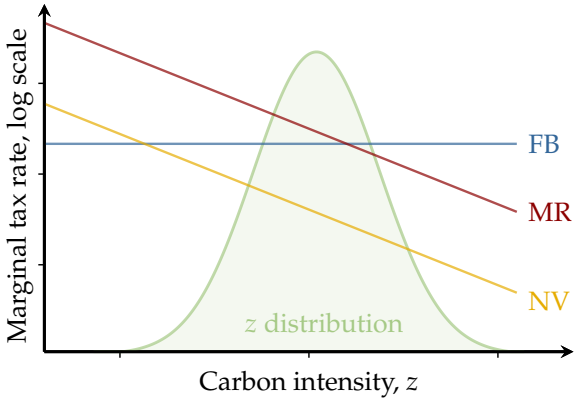
[31] A few simple derivations show that the tax share equals:

$$\frac{\tau_{ij,t}^{MR}\left(g_{ij,t}^{MR}\right)}{y_{i,t}^{MR}} = \nu \frac{\Lambda_{G,t}\mathbb{E}\left[\exp(z_j) \mid a_i\right]\exp\left(-x_{i,t}^{MR}\right)}{P + \Lambda_{G,t}\mathbb{E}\left[\exp(z_j) \mid a_i\right]\exp\left(-x_{i,t}^{MR}\right)},$$

which is independent of z_j.

[32] 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.

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 naïve tax schedule (in orange).

4.5 A welfare decomposition

The above implementation is informative about the extent to which firms' misreporting incentives change the optimal carbon tax on firms. In this subsection, building on our earlier analysis, we turn to studying the trade-offs inherent to the planner allocations. 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:

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 each equilibrium* $\mathcal{P} \in \{DE, FB, MP, NV\}$*:*

$$C_t^{\mathcal{P}} = \Omega_t^{\mathcal{P}} \cdot \Theta_t^{\mathcal{P}} \cdot A(G_t^{\mathcal{P}}) \cdot \exp^a \cdot (K_t^{\mathcal{P}})^{\alpha} \cdot (N^{\mathcal{P}})^{\gamma} \cdot (E_t^{\mathcal{P}})^{\upsilon}, \tag{34}$$

where $\Omega_t^{\mathcal{P}} \in (0,1]$ *measures factor misallocation (Equation A50),* Θ_t *is the intermediate input share (Equation A52), and* $\exp^a$ *denotes the relevant firm-level productivity average (Equation A50). All objects in Equation (34) are a function of the distribution of taxes* $\{\tau_{ij\ell,t}^{\mathcal{P}}(g_{ij\ell,t})\}$

Proposition 2 uses Proposition 1 to characterize consumption, the determinant of household welfare, as the outcome of the decentralized economy with carbon taxes. The decentralized equilibrium without carbon taxation corresponds to the case where $\tau_{ij\ell,t}(g_{ij\ell,t}) = 0$; the full-information planner, the Mirrleesian planner, and the naïve planner allocations correspond to $\tau_{ij\ell,t}(g_{ij\ell,t}) \in \{\tau_t^{FB}(g_{ij\ell,t}), \tau_{j,t}^{MR}(g_{ij\ell,t}), \tau_{ij,t}^{NV}(g_{ij\ell,t})\}$, respectively.

Because carbon taxes are zero in the original decentralized equilibrium, $\Omega_t^{DE} = 1$ (Equation A50): 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 naïve planner delivers an outcome close to that of the decentralized equilibrium. By imposing only minimal distortions, it avoids the bulk of 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 $\tau_{ij\ell,t}(g_{ij\ell,t}) = \tau_t^{FB}(g_{ij\ell,t})$, the planner optimally internalizes the emission externality, and thus attains the highest level of productivity and consumption possible; yet the planner does so at the cost of creating the most misallocation, $\Omega_t^{FB} < \Omega_t^{NV} < 1$.

Finally, compared to the full-information planner, the Mirrleesian planner, in addition, faces the key incentive-compatibility constraint in (24). This constraint, in turn, worsens the possible trade-off between efficiency and misallocation. Consequently, the Mirrleesian 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 allocation, the naïve planner allocation, and the Mirrleesian planner allocation—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' reporting incentives, all else equal, leads to higher marginal taxes for the median carbon-intensity firm. In this section, we use our analysis to shed light on the quantitative implications of our results for the optimal carbon taxes on 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.

5.1 Functional forms and parametrization

The aim of our functional-form assumptions and parameterization is to ensure that our model quantitatively replicates key features of firm-level outcomes and the macroeconomy, as well as captures the rich heterogeneity in emissions documented earlier.

Functional forms. To this end, we retain the specification 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 and thus the implied aggregate productivity loss 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 non-zero correlation ρ_{az} between the two types. We later truncate both distributions to ensure that their supports are compact, consistent with our above assumptions.

Parametrization. We assume that one model period corresponds to a decade. Following Golosov et al. (2014), we set the discount factor β to 0.985^{10}, a value that lies between standard macroeconomic calibrations and the lower discount rates advocated by Stern (2008). We set the energy share of gross output to 4%. Assuming decreasing returns to scale of $\alpha + \gamma + \upsilon = 0.90$—well within the range estimated in the literature (e.g., Basu and Fernald, 1997)—and a standard division of income between capital (40%) and labor (60%), we obtain output elasticities of $\alpha = 0.346$ for capital, $\gamma = 0.518$ for labor, and $\upsilon = 0.036$ for energy.[33] We assume an annual depreciation rate of 7%, implying a model-period depreciation rate of $\delta_K \approx 0.52$. Finally, we normalize average productivity to one, $e^{\mu_a} = 1$ (equivalently, $\mu_a = 0$), and set the standard deviation of firm-level productivity σ_a to 0.6, consistent with the dispersion observed in our Compustat. To ensure the compactness of the set, we truncate the distribution at 3 standard deviations.

The remaining parameters govern the climate side of the economy. The semi-elasticity of the damage function ϕ and mean carbon intensity μ_z jointly determine the aggregate cost of climate change by governing how energy use translates into productivity losses. We therefore normalize mean carbon intensity to one, $e^{\mu_z} = 1$ (equivalently, $\mu_z = 0$), and calibrate $\phi = 4.9 \times 10^{-4}$ to generate a 2% long-run loss in aggregate productivity, consistent with standard integrated assessment models (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 the sensitivity of our results to this choice. We set the standard deviation of carbon intensity σ_z to 1.6, consistent with the higher-quality ICE sample for the U.S.,[34] and truncate the distribution at 3 standard deviations, as we do for

[33] These values convert income shares into gross-output elasticities under decreasing returns to scale: $\alpha = 0.40 \times 0.90 \times (1 - 0.04)$, $\gamma = 0.60 \times 0.90 \times (1 - 0.04)$, and $\upsilon = 0.04 \times 0.90$.

[34] Notice that, for the U.S., the original decentralized equilibrium is "incentive compatible"—firms have no direct mis-reporting incentives—as the U.S. does not operate a carbon tax.

Table 2: Decomposition of consumption changes

	Ω	$\mathcal{O}$	$\mathcal{A}$	$\mathcal{K}^{\alpha}$	$\mathcal{E}^{\nu}$	$\mathcal{C}$
Naïve planner	$+0.00$	$+0.00$	$+0.09$	$+0.05$	$+0.00$	$+0.15$
Mirrleesian planner	-0.00	$+0.08$	$+0.95$	$+0.49$	-0.02	$+1.51$
Full-informed planner	-0.00	$+0.06$	$+1.22$	$+0.65$	$+0.01$	$+1.95$

Note: All numbers are reported in %-deviation from the decentralized equilibrium. The decentralized levels are: $\Omega^{DE} = 1$, $\Theta^{DE} = 0.70$, $A^{DE} = 0.98$, $K^{DE} = 1.76$, $E^{DE} = 0.12$, and $C^{DE} = 2.42$.

productivity.[35] We set the carbon depreciation rate δ_G to 0.001, in line with the estimates in IPCC (2021) that imply a highly persistent climate process. We assume that firm-level productivity and carbon intensity are uncorrelated ($\rho_{az} = 0$), consistent with the evidence discussed in Section 2. Finally, we set the inverse elasticity of abatement with respect to carbon prices θ to 2, in line with the estimates of Känzig (2023) on the effects of carbon tax shocks on green technologies. Table C.1 in Appendix C summarizes the parameters.

5.2 Quantification and decomposition

Welfare 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 first-best case, where the planner has full-information about firms' carbon intensities. We also compare our results to those from the naïve planner, which does not internalize firms' reporting incentives. To start with, we focus on steady-state comparisons. Section 5.4 shows that our results carry over to when we account for the full transition dynamics of the model. Table 2 presents the main results, where to facilitate comparison we decompose consumption differences relative to those in the decentralized equilibrium into five factors, using the breakdown in Proposition 2:

$$\underbrace{\frac{C^{\mathcal{P}}}{C^{DE}}}_{\equiv \mathcal{C}} \equiv \Omega^{\mathcal{P}} \cdot \underbrace{\left(\frac{\Theta^{\mathcal{P}}}{1-\nu-\alpha\delta/(r+\delta)}\right)}_{\equiv \mathcal{O}} \cdot \underbrace{\left(\frac{A^{\mathcal{P}}}{A^{DE}}\right)}_{\equiv \mathcal{A}} \cdot \underbrace{\left(\frac{K^{\mathcal{P}}}{K^{DE}}\right)^{\alpha}}_{\equiv \mathcal{K}^{\alpha}} \cdot \underbrace{\left(\frac{E^{\mathcal{P}}}{E^{DE}}\right)^{\nu}}_{\equiv \mathcal{E}^{\nu}}, \tag{35}$$

where $\mathcal{P} = \{FB, MP, NV\}$. The five factors are: differences in (i) factor misallocation Ω (equal to 1 in the decentralized equilibrium), (ii) the intermediate input share $\mathcal{O}$, (iii) aggregate productivity $\mathcal{A}$, (iv) capital use $\mathcal{K}$, and lastly (v) energy use $\mathcal{E}$.

Several points are worth highlighting. First, consistent with our theoretical results,

[35] This implies a minimum carbon intensity of -4.79, which lies well within the support of the demeaned distribution of energy intensities shown in Panel (b) of Figure 2.

the full-information planner delivers the largest consumption increase and the largest increase in productivity (+1.95% and +1.22%, respectively). This is followed by the Mirrleesian planner (+1.51% and +0.95%) with the naïve planner substantially further behind (+0.15% and +0.09%). Overall, the potential benefits from Mirrleesian or full-information carbon taxation are thus large—close to 2% of steady-state consumption and close to the overall size of the productivity loss due to climate change itself.

Second, notice that the Mirrleesian planner captures the lion's share of the potential, first-best welfare benefits (around 78% of them), while the naïve planner captures substantially less of them (around 8%). Combined, our quantitative results thus suggest that Mirrleesian taxation can recover much of the theoretical benefits of first-best carbon taxation, even with the additional constraints in Equation (26). Accounting for firms' reporting incentives does not substantially constrain the potential welfare benefits. However, ignoring firms' reporting incentives altogether does render the standard Pigouvian approach largely ineffective at addressing climate change.

Third, Table 2 shows that the bulk of the increase in consumption achievable by the Mirrleesian and full-information planner is driven by productivity (c. 67%), which in turn is caused by declines in aggregate emissions. The use of additional capital, as it becomes more productive, contributes some to the overall consumption increase, yet much less than the increase productivity. Furthermore, despite the increase in productivity and the accompanying decline in emissions, overall energy use hardly moves. The Mirrleesian and full-information allocations combine large declines in emissions with small declines in energy, which is why these allocations can create large consumption increases in the first place. Appendix C.2 studies the firm-level drivers behind these economy-wide shifts. Finally, recall that the Mirrleesian and full-information planners introduce misallocation into the economy and alter an economy's intermediate input share. Table 2, however, shows that these forces are quantitatively small—in each case less than 1bp. This is because energy accounts for a small share of output.

In sum, the results in Table 2 demonstrate that most of the potential benefits of carbon taxation can be recovered in the Mirrleesian case, and that most of the benefits from carbon taxation arise due to emission reductions increasing productivity. In the next subsection, we unpack these findings: we first analyze the sources of the emission reductions, and then examine why the Mirrleesian planner is able to almost attain most of the first-best benefits. Our aim is to both discern what causes the above benefits of carbon taxation—abatement or the reallocation of energy use–and to analyze the circumstances that drive Mirrleesian taxation to substantially outperform the naïve case.

Table 3: Decomposition of changes in aggregate emissions

	$\mathcal{Z}_x$	$\mathcal{Z}_e$	$\mathcal{X}$	$\mathcal{E}$	$\mathcal{G}$
Naïve planner	-4.61	$+0.00$	-0.05	$+0.14$	-4.52
Mirrleesian planner	-45.47	$+0.00$	-0.87	-0.59	-46.93
Full-informed planner	-54.12	-5.92	-0.61	$+0.35$	-60.29

Note: All numbers are reported in percent and are based on steady-state values.

Emission decomposition. Conceptually, emissions can decline for three reasons: because of declines in overall energy use, because of increased abatement, or because of a better allocation of energy and abatement among 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 3 components (Appendix B.3):

$$\underbrace{\frac{G^{\mathcal{P}}}{G^{DE}}}_{\equiv \mathcal{G}} = \underbrace{\overbrace{\left(\frac{\sum_{ij\ell} \omega^{\mathcal{P}}_{ij\ell} \cdot \exp^{z_j - x^{\mathcal{P}}_{ij} + X^{\mathcal{P}}}}{\sum_{ij} \omega^{\mathcal{P}}_{ij\ell} \cdot \exp^{z_j}}\right)}^{\equiv \mathcal{Z}_x} \cdot \overbrace{\left(\frac{\sum_{ij\ell} \omega^{\mathcal{P}}_{ij\ell} \cdot \exp^{z_j}}{\sum_{ij\ell} \omega^{DE}_{ijt} \cdot \exp^{z_j}}\right)}^{\equiv \mathcal{Z}_e}}_{\equiv \mathcal{Z}} \cdot \underbrace{\left(\exp^{-X^{\mathcal{P}}}\right)}_{\equiv \mathcal{X}} \cdot \underbrace{\left(\frac{E^{\mathcal{P}}}{E^{DE}}\right)}_{\equiv \mathcal{E}}, \quad (36)$$

where $\omega^{\mathcal{P}}_{ij\ell} \equiv \omega_{ij\ell} \cdot \frac{e^{\mathcal{P}}_{ij\ell,t}}{E^{\mathcal{P}}}$ and $\mathcal{P} = \{FB, MP, NP\}$. Changes in aggregate abatement and energy use are denoted by $\mathcal{X}$ and $\mathcal{E}$, respectively. We have already seen that the latter term is not the reason behind the decline in emissions. The allocative term $\mathcal{Z} \equiv \mathcal{Z}_x \cdot \mathcal{Z}_e$ is the product of two different components: $\mathcal{Z}_x$ and $\mathcal{Z}_e$, where $\mathcal{Z}_e$ captures the role of *energy allocation 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 3 presents the results of the decomposition in Equation (36).

The table shows that the *abatement component*, which captures the allocation of abatement across firms, is by far the dominant force behind the estimated emission reduction. In the full-information and Mirrleesian cases, it contributes around 90% of the overall emission decrease. The *energy allocation component* is instead 0 both in the Mirrleesian and naïve cases, as these allocations are bunched across carbon-intensity types, like in the decentralized equilibrium. As a result, there is no between-firm energy allocation effect. Conversely, in the full-information case, the social planner does allocate energy towards low-z firms. Table 3 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 there is no rise in aggregate abatement; only in the extent of its reallocation.

We conclude that, in both the Mirrleesian and the full-information 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 planner to reallocate energy to low-carbon-intensity firms.

Allocation of abatement. To understand what drives the overall size of the *abatement component*, $\mathcal{Z}_x$, we exploit the following characterization (Appendix B.3):

$$\log \mathcal{Z}_x \approx -\mathbb{C}ov\left(\frac{g_{ij\ell,t}}{G_t}, \frac{\exp^{x_{ij\ell,t}}}{\exp^{X_t}}\right) < 0. \tag{37}$$

This result shows that the abatement component is large when relative emissions, $g_{ij\ell,t}/G_t$, comove strongly with relative abatement efforts across firms, $\exp^{x_{ij\ell,t}} / \exp^{X_t}$, which is indeed the case under both the full-information and the Mirrleesian allocation. The first-order conditions for abatement (Equations 21 and 29) directly show that:

$$\left(x^{FB}_{ij\ell,t}\right)^{\theta} = \Lambda_{G,t} \cdot g^{FB}_{ij\ell,t} \quad \text{and} \quad \left(x^{MP}_{ij\ell,t}\right)^{\theta} = \Lambda_{G,t} \cdot g^{MP}_{ij\ell,t} \cdot \frac{\mathbb{E}[\exp^{z_j} \mid a_i]}{\exp^{z_j}}. \tag{38}$$

Thus, abatement is a direct function of emissions in both planner allocations. Combined with the fact that $\mathbb{E}[\exp^{z_j} | a_i] / \exp^{z_j}$ is, on average, close to 1 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 emission reduction. Both the full-information planner and its Mirrleesian counterpart create a strong correlation at the firm level between emissions and abatement. Appendix C.2 shows allocations of emission and abatement efforts across productivities and carbon intensity-types.

Allocation of energy. Although our findings show the forces behind the emission reduction, they do not explain why the Mirrleesian planner performs nearly as well as the full-information planner. As explained above, the difference between her and the full-information planner is primarily driven by the *energy-reallocation component*, $\mathcal{Z}_e$. To address why this component is small, we next examine what determines the strength of the

energy-reallocation effect. A few simple derivations show that (Appendix B.3):

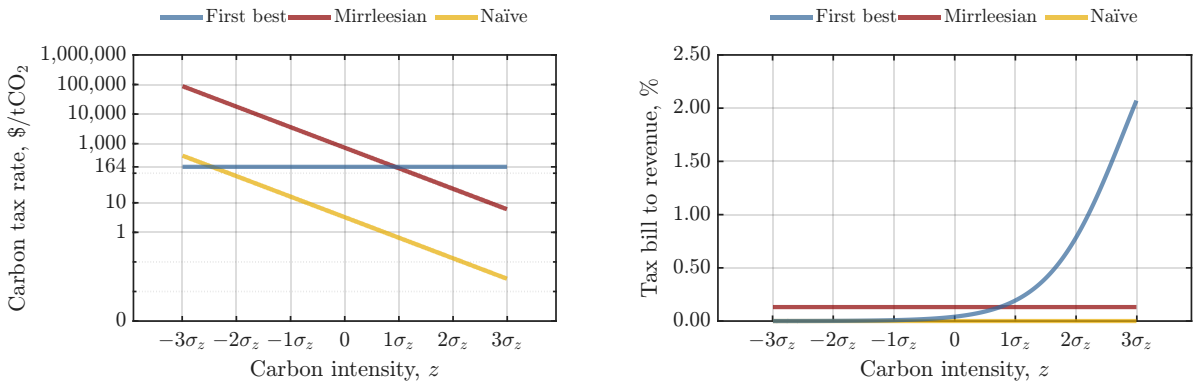
$$\mathcal{Z}_e^{FB} - \mathcal{Z}_e^{MP} = \mathbb{C}ov\left(\frac{e_{ij\ell,t}^{FB}}{E_t^{FB}},\ \frac{\exp^{z_j}}{\mathbb{E}\left[\exp^{z_j}\right]}\right) < 0. \tag{39}$$

This equation decomposes differences in $\mathcal{Z}_e$ between the two planner allocations into the covariance between firm-level energy shares in the full-information allocation, $e_{ij\ell,t}^{FB}/E_t^{FB}$, and the relative carbon-efficiency type of the firm, $\exp^{z_j}/\mathbb{E}\left[\exp^{z_j}\right]$. 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 firms' energy use in the full-information allocation is mainly driven by carbon intensity differences rather than productivity differences. By contrast, when the variation in energy shares is mainly due to productivity differences, the covariance in Equation (39) is close to zero, and the Mirrleesian planner in this case performs nearly as well as the full-information planner.

Thus, the reason the Mirrleesian planner can reap most of the first-best benefits from carbon taxation is that—in the first-best allocation—firm productivity is *quantitatively* the dominant driver of firms' optimal energy use. Although carbon intensities clearly matter, under our baseline calibration, firm-level productivity is a quantitatively important determinant of firms' optimal energy use, highlighting the importance of considering both sources of heterogeneity in the analysis. Appendix C.2 shows graphically the similarities in energy use in the Mirrleesian and full-information cases.

In summary, in this section, we have shown that abstracting from firms' reporting incentives can render the standard (naïve) Pigouvian approach to carbon taxation largely ineffective in reducing emissions—with substantial consumption losses as a result. Clearly, the extent of firms' misreporting is, ultimately, close to unknowable; yet one advantage of our quantitative-theoretical framework is that we can bound the potential losses from misreporting. Conversely, accounting for firms' reporting incentives does not substantially constrain the potential welfare and environmental benefits of carbon taxation. The Mirrleesian optimal tax attains close to 3/4 of the potential, first-best benefits. The root cause of this comparative performance of the Mirrleesian approach is that the first-best carbon tax attains its welfare improvement largely by allocating abatement efforts across firms—a task the Mirrleesian approach does almost equally well when firm productivity optimally determines a large share of the variation in firms' energy use.

Figure 7: Calibrated carbon tax rates



(a) Marginal carbon tax rate

(b) Overall carbon tax bill

Note: This figure shows the marginal tax rate, Panel (a), and the overall tax bill, Panel (b), as a function of firms' carbon intensities. We use 2024 US GDP numbers to convert model estimates into USD values.

5.3 Calibrated carbon taxes

Before closing with several extensions that show how our above results carry over to more elaborate frameworks, we briefly turn to the implied carbon taxes that implement the different planner allocations. Common estimates of the social cost of carbon—equivalent to the tax that implements the first-best, full-information allocation—is around \$100-250 per tCO_2e.[36] Panel (a) in Figure 7 shows that our baseline calibration produces estimates that are close to this value. The marginal tax that implements the full-information allocation is around \$164 per tCO_2e and constant across the carbon-intensity distribution, corroborating the soundness of our baseline calibration.

Furthermore, notice that, in line with Proposition 1, the optimal Mirrleesian tax is log-linear in z and monotonically declining. Indeed, Panel (a) shows that for a firm that is 2 standard deviations *below* the mean carbon-intensity type, its value is around \$15,000 per tCO_2e; for a firm that is 2 standard deviations *above*, it is around \$30 per tCO_2e. At the ± 1 standard deviation level, the range is between \$3,500-150 per tCO_2e, with a median value of around \$800 per tCO_2e. Unlike the first-best full-information tax, there is substantial heterogeneity in the Mirrleesian optimal carbon tax, and its median value is larger.

A fear one might have is that by enacting such high marginal tax rates on select carbon-efficient firms one may heavily penalize their profits, potentially to the detriment of future innovation. Panel (b) in Figure 7 reports the overall tax bill spent on carbon taxes as a share of a firm's revenue. Although the marginal carbon tax can be high for an individual

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

firm, the overall tax bill in the Mirrleesian case stays well below 0.5% of revenue. Clearly, changes to the incentive structure faced by firms is the main driver of our results rather than the size of the climate tax bills themselves.

Panel (b) further highlights a fundamental difference between the full-information planner and the Mirrleesian planner allocations. 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 illustrates 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 emissions. By contrast, the Mirrleesian planner internalizes this incentive and, to eliminate it, charges all firms the same tax-bill share of revenue (conditional on their productivity type; Section 4.4).

Finally, Appendix C.3 shows the differences between the first-best 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 noted earlier, the marginal benefit from abatement scales with firm size (i.e., productivity). Consequently, both the first-best and the Mirrleesian planner 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 in a firm's carbon-intensity type. The marginal tax for most firms declines from around \$15,000 to \$30 per tCO_2e, and is on average equal to approximately \$800 per tCO_2e, above standard estimates of the social cost of carbon. The average tax paid by firms is, however, small, less than 0.5% of revenue, on average. The Mirrleesian, thus, combines strong incentives with a low average costs.

5.4 Quantitative refinements

Before concluding, we briefly turn to several alternative versions of our benchmark framework. The overall aim of these is to demonstrate that our main result—the relative performance of Mirrleesian carbon taxation—is robust to alternative assumptions about the climate block of the economy, the energy market, and the information structure. We recalibrate (where appropriate) all alternative models to target the same moments as our benchmark model. Table 4 and Appendix C summarize the results.

Table 4: Overview of quantitative refinements

	Consumption				*Emissions*			
	FB	MR	NV	MR/FB	FB	MR	NV	MR/FB
Baseline								
Steady state	+1.95	+1.51	+0.15	77%	−60.3	−46.9	−4.5	78%
Extensions/robustness								
Transition, baseline climate	+0.08	+0.06	+0.01	72%	−5.09	−3.94	−0.36	77%
Transition, temperature block	+0.01	+0.01	+0.00	67%	−3.27	−2.20	−0.15	67%
Renewable and fossil energy	+2.34	+1.36	+0.15	58%	−66.8	−52.7	−5.4	79%
Higher (50%) planner information, η	+1.95	+1.64	+0.71	84%	−60.29	−50.35	−21.82	83%
Higher (50%) damage semi-elasticity, ϕ	+3.09	+2.44	+0.26	79%	−64.30	−51.30	−5.32	80%
Lower (50%) abatement elasticity, θ	+1.85	+1.25	+0.01	67%	−57.46	−39.44	−0.20	69%
Lower (50%) carbon intensity dispersion, σ_z	+0.51	+0.47	+0.19	91%	−39.95	−36.68	−14.96	92%

Notes: FB, MR, and NV denote the full-information, Mirrleesian, and naïve planner allocations and report %-deviations of consumption and emissions from the decentralized equilibrium. MR/FB reports the share of the first-best gains in consumption or emission reductions recovered by the Mirrleesian allocation.

Transition dynamics and alternative climate model. A potential concern with the above findings is that they characterize the *long-run benefits* of climate taxation by comparing steady-state allocations. To assess the benefits of carbon taxation along the *transition*, we study below the full transition dynamics of both the baseline model and an extension that incorporates a richer description of climate dynamics.

In particular, 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 emissions in the atmosphere in accordance with:

$$T_{t+1} = (1 - \xi_T) \cdot T_t + \xi_T \cdot (\mathcal{T} + \Xi_G \cdot G_t), \tag{40}$$

where $1 - \xi_T$ is the persistence in global temperature, $\mathcal{T}$ is its historical mean, and Ξ_G is the conversion factor mapping emissions into temperature. The calibration of the initial conditions for the baseline model, combined with the calibration of the alternative climate model, is described in Appendix D.1. Table 4 summarizes our findings while Figures D.1 and D.2 show the dynamics of consumption and the climate along the transition.

In our baseline economy along the transition, as in its steady state, the Mirrleesian planner allocations follow closely the full-information outcomes. Relative to the decentralized equilibrium, the full-information (Mirrleesian) planner, for example, achieves a

reduction in the net-present value of emissions of -5.09% (-3.94%).[37] The Mirrleesian planner once more captures the bulk of the full-information welfare and environmental benefits (around 70%). In contrast, the naïve planner yields only a minor benefit.

These insights carry over to the version with the extended climate block. The Mirrleesian planner once more captures the majority of the potential full-information benefits from climate taxation (around 67% in net-present value terms), while naïve taxation results in negligible welfare and environmental gains. In net-present value terms, the first-best (Mirrleesian) planner attains an emission reduction of -3.27% (-2.20%) in the extended model. In both cases, the bulk of the benefits arise after around 100 years.

On balance, our results thus reinforce the insights from the baseline model: the Mirrleesian planner consistently exhausts most of the potential gains from climate taxation. While the absolute magnitude of these gains depends on the specification of the climate block—as emphasized throughout the macro-climate literature (e.g., Nordhaus, 2019)—the relative performance of the Mirrleesian planner is largely unaffected. The reason is that the climate block affects a planner's problems primarily through the social cost of carbon, proportionally scaling both the first-best and the Mirrleesian optimal tax schedules. Our main conclusions are therefore robust to alternative climate dynamics.

Renewable and fossil energy. One reason why firms, even within narrowly-defined sectors, exhibit differences in emission intensities is the availability of renewable versus fossil energy in their area of operation. Other technological differences—such as differences in filters used and burner designs, mentioned in Section 2—lead to differences in emission intensities, in addition to those which arise from differences in the use of different fuels. To explore the implications of the introduction of a reneweable energy sector into our model, we enrich our baseline framework by assuming that energy $e_{ij\ell,t}$ is comprised of renewable $e_{\text{green},ij\ell,t}$ and fossil energy $e_{\text{brown},ij\ell,t}$ in accordance with:

$$e_{ij\ell,t} = \left(e_{\text{green},ij\ell,t}^{\frac{\varepsilon-1}{\varepsilon}} + e_{\text{brown},ij\ell,t}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}, \tag{41}$$

where both fossil and renewable energy are produced as in the baseline model. We assume that only the use of fossil energy creates emissions. Although conceptual issues with the CES structure are known, it has proven to be a both popular and tractable means to model firms' choices of renewable versus fossil energy (e.g, Hassler et al., 2016). Appendix D.2 provides details of the calibration.

Table 4 and Table D.1 show that the main insights from the baseline model extend

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

to the economy with a renewable energy sector. The Mirrleesian tax once again recovers around 2/3 of the potential reduction in emissions and approximately 60 percent of the first-best welfare gains, substantially outperforming the naïve benchmark. Relative to the full-information allocation, the Mirrleesian planner induces a smaller reallocation from fossil to renewable energy, reflecting the incentive-compatibility constraints it faces. Nevertheless, because it continues to allocate abatement efforts efficiently across firms, it captures most of the environmental and welfare gains from carbon taxation.

Partial information. In our baseline analysis, the information friction between the planner and firms takes a particularly simple form: each firm knows its carbon intensity z_j, while the planner knows only the distribution from which z_j is drawn. To account for the possibility of additional planner information, we enrich the baseline model with an additional planner signal. In particular, the signal is such that social planner observes the share $\eta \in [0, 1]$ of a firm's carbon intensity z_j without error; the planner is only uncertain about the remaining $1 - \eta$ share. Varying η from 0 to 1 thus traces out the entire range of possible information frictions. Appendix D.3 derives the implied wedges in the first-order conditions. Table 4 and Figure D.3 report the corresponding quantitative results.

Notice that, for all values of η, the Mirrleesian planner attains most of the potential welfare and environmental gains. Indeed, decreasing the information friction causes the Mirrleesian planner to attain a *larger* share of the possible welfare and environmental benefits. Clearly, increasing η also improves the share of the potential benefits that the naïve planner can attain. But even assuming a 50% reduction in the planner's uncertainty about a firm's carbon intensity, the naïve planner still attains substantially less than half of the welfare and environmental benefits of the Mirrleesian planner (Table 4). Accounting for firms' disclosure incentives remains important, even when policymakers possess substantial information about firms' carbon intensities.

Climate parameter uncertainty. Finally, our baseline model relies on estimates of the environmental damage parameter ϕ, the inverse elasticity of abatement θ, and the standard deviation of the carbon-intensity distribution σ_z. As with the processes governing emissions, these parameters are subject to considerable uncertainty. Table 4 and Figures D.4–D.6 examine the robustness of our results to alternative parameter values.

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). Table 4 and Table D.4 show 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 the same: Mirrleesian taxation recovers a large share of the potential, first-best benefits.

Likewise, our conclusions do not hinge on the precise value of θ used. 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. Table 4 and Figures D.5 show that our results are relatively insensitive to the choice of θ. Two offsetting forces drive this finding. 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, so that the overall benefits of climate taxation do not depend much on the value of θ. We view this fact as a strength of our framework, given the limited evidence on θ.

Lastly, in our baseline calibration, we discipline the standard deviation of the carbon-intensity distribution σ_z using data from an economy without a carbon tax—the U.S.—where firms have no incentive to misreport emissions. Table 4 and Figures D.6 examine the robustness of our results to alternative values of σ_z. On balance, the results show that our conclusions remain robust to plausible departures from the baseline calibration. Although a lower dispersion in carbon intensities reduces firms' incentives to misreport, even when σ_z is reduced by 50 percent the Mirrleesian planner continues to recover a large share of the first-best gains, whereas the naïve planner captures only a small fraction.

To summarize, calibrating our model framework to microdata, we estimate that the Mirrleesian approach attains at least 60% of the potential welfare and environmental benefits from climate taxation. This stands in stark contrast to a naïve approach, which does not internalize firms' reporting incentives. We estimate that the latter often attains only around 10% of the potential benefits. Crucially, we have documented that this relative performance of the Mirrleesian approach holds across a range of model extensions that vary the climate block of the economy and considers alternative parameter values, among others. We view this robustness as an important feature of our framework given the uncertainty that exists about the best climate specification to use in macro-climate contexts.

6 Final remarks

Addressing the economic damages from climate change is to first-order a problem of curtailing firms' direct carbon emissions. In response to this challenge, considerable efforts

have recently gone into the design and implementation of (Pigouvian) 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 emission disclosures. Carbon disclosures are in most cases based on firm-level information that is either difficult or costly to verify for policymakers (International Sustainability Standards Board, 2023; SEC Carbon Disclosures, 2024). Combined with a sizable carbon tax, this creates a strong incentive for firms to misreport their carbon disclosures—consistent with recent legal evidence and court cases.

In this paper, we have formulated 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 has been 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, marginal climate damages, and output.

We evaluated this formula quantitatively, using the distribution of carbon intensities, and found that optimal carbon taxes vary substantially across firms. The information friction makes firm-level taxes optimally *heterogeneous across firms* but *constant within a firm*. We found that the optimal carbon tax for a firm that is 1 standard deviation below the mean carbon-intensity level is around \$35,000 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 \$150 per tCO_2e. The median carbon tax paid is \$800 per tCO_2e. Despite this heterogeneity, the overall Mirrleesian tax bill faced by firms in equilibrium never exceeds 0.5% of firm revenue. Instead, we have shown 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 equal to \$164 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.[38]

[38] 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 \$90 per tC02 and equal

Crucially, we have documented 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 around 80% of the possible welfare and environmental benefits from carbon taxation. A naïve approach—which implements the full-information tax without accounting for firms' reporting incentives—in contrast recovers less than 10%. We have showed that this comparative performance rests on the ability of a well-constructed tax to incentivize abatement efforts by carbon-intensive firms. As a result, we have shown 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 naïve approach by contrast recovers substantially less—and sometimes close to zero.

An important feature of our proposed tax is its simplicity. The optimal carbon tax 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 that several extensions of the present setting are desirable, one advantage of our framework is that such extensions should come at a low cost.

Finally, beyond the analysis in this paper, our approach suggests that 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 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 model framework as a useful basis for extensions that incorporate dynamic investments in abatement, jurisdictional leakages, and other reporting and verification environments.

across all of firms, and thus well-below our estimate for the optimal carbon tax rates paid by most firms.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous, "The environment and directed technical change," *American economic review*, 2012, *102* (1), 131–166.

_ , Ufuk Akcigit, Douglas Hanley, and William Kerr, "Transition to clean technology," *Journal of political economy*, 2016, *124* (1), 52–104.

Akcigit, Ufuk, Douglas Hanley, and Stefanie Stantcheva, "Optimal taxation and R&D policies," *Econometrica*, 2022, *90* (2), 645–684.

Allen, Myles R, David J Frame, Chris Huntingford, Chris D Jones, Jason A Lowe, Malte Meinshausen, and Nicolai Meinshausen, "Warming caused by cumulative carbon emissions towards the trillionth tonne," *Nature*, 2009, *458* (7242), 1163–1166.

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

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

Basu, Susanto and John G Fernald, "Returns to scale in US production: Estimates and implications," *Journal of political economy*, 1997, *105* (2), 249–283.

Bhandari, Anmol, David Evans, Ellen R McGrattan, and Yuki Yao, "Business Income Underreporting and Public Finance," Technical Report, working paper, University of Minnesota 2024.

Bilal, Adrien and Diego R Känzig, "The macroeconomic impact of climate change: Global vs. local temperature," Technical Report, National Bureau of Economic Research 2024.

Bivol, "Waste Year," 2020.

Bolton, P and M Kacperczyk, "Firm commitments (NBER Working Paper No. w31244)," 2023.

Bruns, Stephan B, Zsuzsanna Csereklyei, and David I Stern, "A multicointegration model of global climate change," *Journal of Econometrics*, 2020, *214* (1), 175–197.

Caggese, Andrea, Andrea Chiavari, Sampreet Goraya, and Carolina Villegas-Sanchez, *Climate change, firms, and aggregate productivity*, Centre for Economic Policy Research, 2025.

Calel, Raphael, Antoine Dechezlepretre, and Frank Venmans, "Policing carbon markets," *Climate Policy*, 2025, pp. 1–19.

Capelle, Mr Damien, Mr Divya Kirti, Mr Nicola Pierri, and Mr German Villegas Bauer, "Mitigating climate change at the firm level: Mind the laggards," *International Monetary Fund WP*, 2023.

Cicala, Steve, David Hémous, and Morten G Olsen, "Adverse selection as a policy instrument: unraveling climate change," Technical Report, National Bureau of Economic Research 2022.

ClientEarth, "Lawyers Call for EU Crack Down on Illegal Waste Burning in Coal Plants," 2020.

Dechezleprêtre, Antoine, Adrien Fabre, Tobias Kruse, Bluebery Planterose, Ana Sanchez Chico, and Stefanie Stantcheva, "Fighting climate change: International attitudes toward climate policies," *American Economic Review*, 2025, *115* (4), 1258–1300.

den Bremer, Ton S Van and Frederick Van der Ploeg, "The risk-adjusted carbon price," *American Economic Review*, 2021, *111* (9), 2782–2810.

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

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

ECA, "The integrity and implementation of the EU ETS," Special Report 06/2015, European Court of Auditors, Luxembourg 2015. EN 2015 No 06.

European Commission, "Commission Implementing Regulation (EU) 2018/2066 of 19 December 2018 on the Monitoring and Reporting of Greenhouse Gas Emissions Pursuant to Directive 2003/87/EC," 2018. Official Journal of the European Union, L 334, 31 December 2018.

__ , "Commission Implementing Regulation (EU) 2018/2067 of 19 December 2018 on the Verification of Data and on the Accreditation of Verifiers pursuant to Directive 2003/87/EC of the European Parliament and of the Council," 2018. Article 32: Conditions for not carrying out site visits.

__ , "ESRS E1 Climate change," *European Commission Working paper*, 2023.

__ , "Report from the Commission to the European Parliament and the Council on the Functioning of the European Carbon Market in 2022 pursuant to Articles 10(5) and 21(2) of Directive 2003/87/EC (COM(2023) 654 final)," *European Commission Report*, 2023.

European Commission Directorate-General for Climate Action, "The Monitoring and Reporting Regulation – General Guidance for Installations," Technical Report MRR Guidance document No. 1, European Commission January 2025. Updated version, final draft, 14 January 2025.

Fernandez-Bastidas, Rocio, "Entrepreneurship and tax evasion," *Economic Modelling*, 2023, *128*, 106488.

Fernández-Villaverde, Jesús, Kenneth T Gillingham, and Simon Scheidegger, "Climate change through the lens of macroeconomic modeling," *Annual Review of Economics*, 2024, *17.*

Fontaine, François, Carole Marullaz, and Katheline Schubert, "Carbon Tax and Labor Reallocation: The Role of Firm Heterogeneity and Energy Efficiency," *Mimeo*, 2026.

GHG Protocols, *The Greenhouse Gas Protocol: A Corporate Accounting and Reporting Standard (Revised Edition)* 2004.

Golosov, Mikhail, Aleh Tsyvinski, Ivan Werning, Peter Diamond, and Kenneth L Judd, "New dynamic public finance: A user's guide [with comments and discussion]," *NBER macroeconomics annual*, 2006, *21*, 317–387.

__ **, John Hassler, Per Krusell, and Aleh Tsyvinski**, "Optimal taxes on fossil fuel in general equilibrium," *Econometrica*, 2014, *82* (1), 41–88.

Greenstone, Michael, Rohini Pande, Nicholas Ryan, and Anant Sudarshan, "Can pollution markets work in developing countries? Experimental evidence from India," *The Quarterly Journal of Economics*, 2025, *140* (2), 1003–1060.

Hassler, John, Per Krusell, and Anthony A Smith Jr, "Environmental macroeconomics," in "Handbook of macroeconomics," Vol. 2, Elsevier, 2016, pp. 1893–2008.

ICE, "Sustainability Report and Documentation," *Intercontinental Exchange WP*, 2024.

International Energy Agency, "Carbon Capture, Utilisation and Storage," 2024.

International Sustainability Standards Board, "IFRS S2 Climate-related Disclosures," IFRS Sustainability Disclosure Standard, IFRS Foundation June 2023.

IPCC, "2006 IPCC Guidelines for National Greenhouse Gas Inventories," Technical Report, Intergovernmental Panel on Climate Change 2006.

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

IPCC, *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, 2022.

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

Kang, Zi Yang, "Optimal indirect regulation of externalities," *Available at SSRN 3586050*, 2020.

Känzig, Diego R, "The unequal economic consequences of carbon pricing," Technical Report, National Bureau of Economic Research 2023.

Kocherlakota, Narayana R, "The new dynamic public finance," in "The New Dynamic Public Finance," Princeton University Press, 2010.

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

Lipsey, Richard G and Kelvin Lancaster, "The general theory of second best," *The review of economic studies*, 1956, *24* (1), 11–32.

Lyon, Josh De and Antoine Dechezleprêtre, "The great dispersion in energy productivity between firms," *OECD Science, Technology and Industry Working Papers*, 2025, *2025* (24).

Lyubich, Eva, Joseph S Shapiro, and Reed Walker, "Regulating mismeasured pollution: Implications of firm heterogeneity for environmental policy," in "AEA papers and proceedings," Vol. 108 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2018, pp. 136–142.

Matthews, H Damon, Nathan P Gillett, Peter A Stott, and Kirsten Zickfeld, "The proportionality of global warming to cumulative carbon emissions," *Nature*, 2009, *459* (7248), 829–832.

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

Miftakhova, Alena, Kenneth L Judd, Thomas S Lontzek, and Karl Schmedders, "Statistical approximation of high-dimensional climate models," *Journal of Econometrics*, 2020, *214* (1), 67–80.

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

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

Nola, Alessandro Di, Georgi Kocharkov, Almuth Scholl, and Anna-Mariia Tkhir, "The aggregate consequences of tax evasion," *Review of Economic Dynamics*, 2021, *40*, 198–227.

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

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

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

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

_ **and Andrew Moffat**, "A survey of global impacts of climate change: replication, survey methods, and a statistical analysis," 2017.

Organized Crime and Corruption Reporting Project, "Bulgarian Coal Magnate's Plants May Have Saved Around 30M Euros by Under-Declaring Emissions," 2021.

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

Pai, Mallesh and Philipp Strack, "Taxing externalities without hurting the poor," 2022.

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

Parry, Ian, Simon Black, and James Roaf, "Proposal for an international carbon price floor among large emitters," 2021.

Petersen, Thore, "Can the plants turn green?," *Mimeo*, 2025.

Pigou, Arthur, *The economics of welfare*, Routledge, 1920.

Pigou, Arthur Cecil, *Wealth and welfare*, Macmillan and Company, limited, 1912.

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

Quick, Jeffrey C., "Carbon Dioxide Emission Factors for U.S. Coal by Origin and Destination," *Environmental Science & Technology*, 2010, *44* (7).

Report, CDP Carbon Majors, "The Carbon Majors Database: CDP Carbon Majors Report 2017," Technical Report, CDP, London 2017. CDP report.

Ricke, Katharine, Laurent Drouet, Ken Caldeira, and Massimo Tavoni, "Country-level social cost of carbon," *Nature Climate Change*, 2018, *8* (10), 895–900.

SEC Carbon Disclosures, "The Enhancement and Standardization of Climate-Related Disclosures: Final Rules," *Security and Exchange Commission: Final Rule*, 2024.

Shultz, George P and James A Baker, "A conservative answer to climate change," *Wall Street Journal*, 2017, *8.*

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

Statement on Carbon Dividends, "Economists' Statement on Carbon Dividends," *The Wall Street Journal*, January 2019. Full-page statement organized by the Climate Leadership Council; 3,500+ signatories.

Stern, Nicholas, "The economics of climate change," *American economic review*, 2008, *98* (2), 1–37.

Stiglitz, Joseph E, Nicholas Stern, Maosheng Duan, Ottmar Edenhofer, Gaël Giraud, Geoffrey M Heal, Emilio Lèbre La Rovere, Adele Morris, Elisabeth Moyer, Mari Pangestu et al., "Report of the high-level commission on carbon prices," 2017.

Tirole, Jean, *The theory of industrial organization*, MIT press, 1988.

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

UK Department for Energy Security and Net Zero, "UK ETS Monitoring and Reporting Regulation (MRR): General Guidance for Installations," Guidance Document, UK Government July 2025. UKETS01.

U.S. Environmental Protection Agency, "Cost of Selective Catalytic Reduction (SCR) Application for NOx Control on Coal-Fired Boilers," Technical Report, U.S. Environmental Protection Agency 2002.

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

Online Appendix for "Mirrleesian Carbon Taxation"

Andrea Chiavari Alexandre Kohlhas

July 2026

Contents

A Motivating evidence **1**
A.1 Data construction . . . 1
A.2 Additional evidence . . . 3
A.3 Sources of heterogeneity, disclosure standards, and misreporting . . . 6

B Model framework and results **10**
B.1 Proofs and derivations . . . 10
B.2 Alternative implementations . . . 18
B.3 Aggregate emission decomposition . . . 19

C Quantitative results **21**
C.1 Parametrization . . . 21
C.2 Firm-level outcomes . . . 21
C.3 Carbon tax . . . 23

D Quantitative refinements **25**
D.1 Transitional dynamics and alternative climate model . . . 25
D.2 Green and brown energy inputs . . . 27
D.3 Alternative degrees of imperfect information . . . 27
D.4 Climate parameters uncertainty . . . 28

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

	Obs.	Mean	Std.	Median
ICE				
Scope 1 emission (All)	402,523	409,770	5,801,209	2,307
Scope 1 emission (High-quality)	32,599	1,461,794	5,139.210	29,200
Scope 1 emission intensity (High-quality)	32,599	181	568	8
Trucost				
Scope 1 emission	49,011	541,617	2,042,981	26,061
Scope 1 emission intensity	49,012	131	357	22

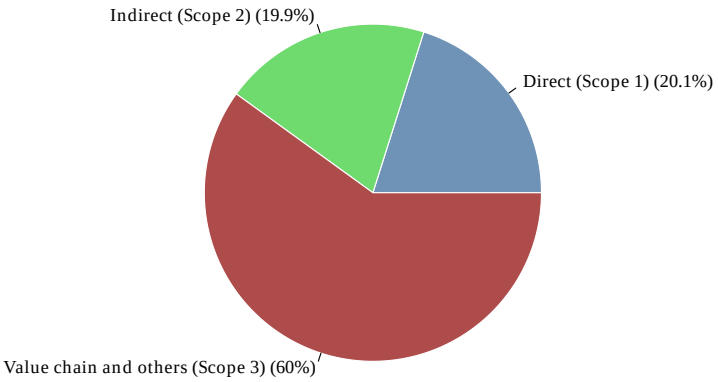
Notes: The table reports descriptive statistics for the sample of firms in the ICE and Trucost database. The ICE high-quality sample is restricted to observations with third-party-validated disclosures, identified by ICE's high-quality reporting flag; the Trucost extract used here does not contain an analogous observation-level validation indicator. The units of the first two rows are tons of CO_2e emissions (tCO2e). The two subsequent rows are tCO2e per million $ in revenue. For ICE, observations in the top and bottom 1 percent of the full-sample distribution are removed before selecting high-quality observations. For Trucost, observations in the top and bottom 1 percent of the Trucost distribution are removed.

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. The sample includes model-estimated emissions, disclosed emissions, and, among the disclosed emissions, those that have been validated by an independent third party. To construct the sample that we use in Figures 2, 3, A.3, A.4, and A.5 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. In our ICE data, this restriction identifies firm-year observations with third-party-validated emissions disclosures. Table 1 uses all disclosed emissions, but not model-estimated ones, as its purpose is to compare reporting differences between firms whose emissions receive third-party validation and those whose emissions do not. Figures 1 and A.1 use the full sample of observations, as they aim to capture representative aggregates. Direct emission intensities are measured by reported scope 1 emissions (tCO_2e) 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`

Figure A.1: 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 ICE observations with third-party-validated disclosures (Appendix A.1). We limit the sample to firms from Europe or North America, where scope 2 and 3 emissions are standardized.

`Scope 1 (tCO2e)`). We remove observations in 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 (tCO_2e) per million \$ of revenue (`code: di319407`). Overall, emissions are measured by the numerator of this expression (`code: di319413`). Unlike ICE, the Trucost extract used here does not contain an observation-level indicator for third-party validation, so the validation-based sample restriction is applied only to ICE. We again remove observations in 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) TFP as in Ottonello and Winberry (2020).

Descriptive statistics. Table A.1 reports descriptive statistics for the various data sets. The table highlights two features of the data that are used throughout the empirical sec-

tion. First, the ICE-ESG database contains substantially broader firm coverage than Trucost, which is why the aggregate emission shares in Figure 1 are larger than their Trucost counterpart in Figure A.2. Second, emission intensities are highly dispersed even after trimming the tails of the distribution and, in ICE, restricting to third-party-validated disclosures. This motivates the focus in the main text on heterogeneity in firms' carbon use.[39]

Figure A.1 describes the scope composition of reported emissions in the ICE-ESG sample. For the average firm, direct production emissions are only one component of total emissions reported for ESG purposes, with value-chain and other emissions accounting for the largest share. This decomposition is useful for interpreting the scope of the paper. The model and empirical analysis focus on Scope 1 emissions, as these are generated within the firm and are the natural object of interest for a production-based carbon tax. Scope 2 and Scope 3 emissions are also important, but using them to aggregate firm-level emissions would mechanically double-count emissions produced elsewhere in the economy and would shift the focus toward emissions over which firms have limited control.

A.2 Additional evidence

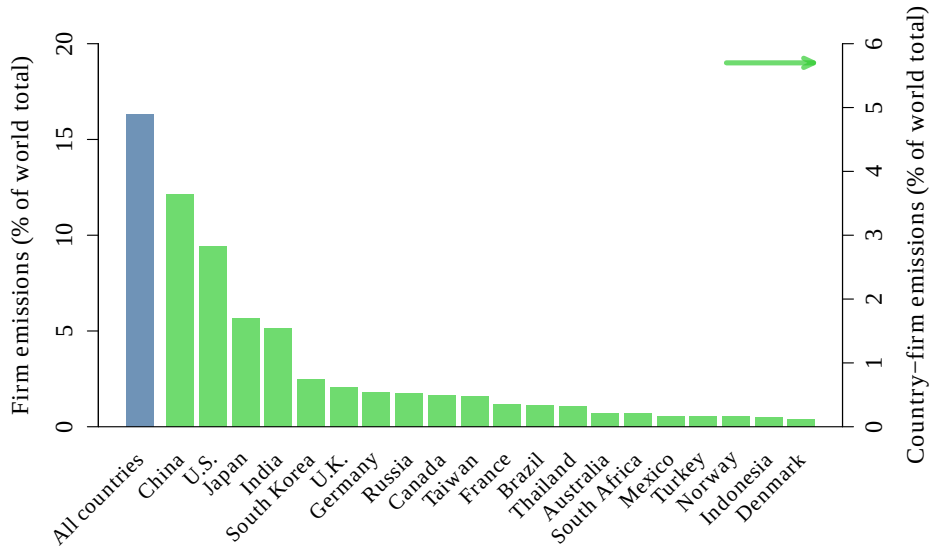
This subsection provides additional descriptive evidence supporting the two empirical claims in Section 2: firms account for a large share of aggregate emissions, and emission intensities are highly heterogeneous across and within economically meaningful groups.

Figure A.2 shows that publicly listed firms in Trucost account for a smaller share of world emissions than firms in ICE-ESG. This difference is expected, since Trucost covers a narrower set of firms; however, the magnitude remains substantial, and the ranking of large-emitting countries is similar. The comparison supports the interpretation in the main text (Figure 1): firms remain a substantial driver of global emissions despite the smaller sample, and, given the representativeness of the data, these measures should be viewed as conservative lower bounds.

Figure A.3 shows that large differences in emission intensity are present even when output is measured in physical units within narrowly defined sectors. Energy producers, airlines, and steel firms differ sharply from each other, but, more importantly for our analysis, each narrow sector also displays substantial within-sector dispersion. This evidence reinforces the interpretation of Figure 2: the dispersion in carbon efficiency is not driven by broadly defined sectors, nor by comparisons across firms producing different goods

[39] The large standard deviations for total scope 1 emissions reflect firm-size heterogeneity among large emitters. They are in tCO_2e and therefore are not directly comparable to the normalized emission-intensity rows. For ICE, trimming thresholds are computed in the full sample and then applied before selecting the high-quality subsample.

Figure A.2: Estimated share of overall emissions—Trucost



Note: Data from the Trucost sample. The figure shows overall firm emissions in 2024 (measured by tons of CO_2e emissions) divided by total world emissions. Estimates for world emissions are taken from the *Global Carbon Budget* website (2024 estimates). Top-20 country-contributions are shown on the right-side axis.

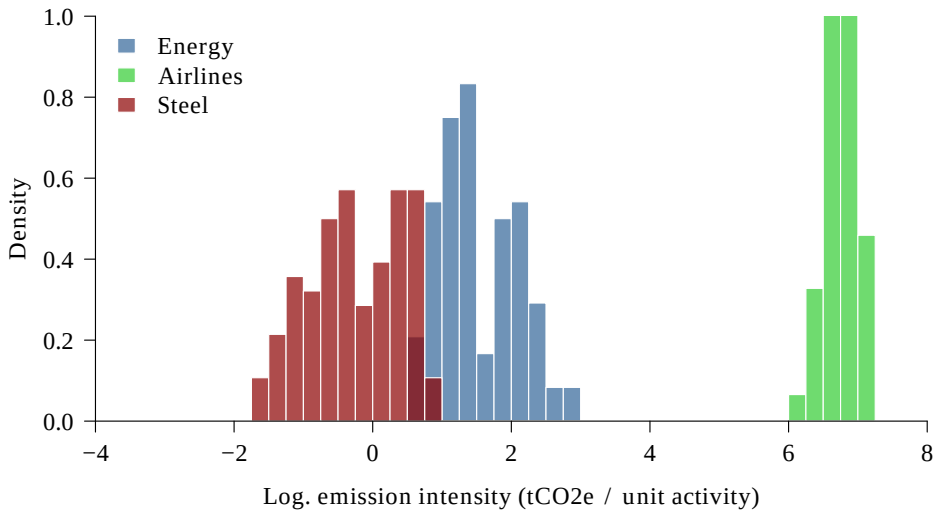
or charging different prices.

Figure A.4 documents the sectoral component of this dispersion. Emission intensities are highest in sectors such as energy, materials, utilities, agriculture, resource extraction, transportation, and manufacturing. These patterns are consistent with the conventional wisdom on sectoral emissions and also show why sector fixed effects are important in Figure 2. The fact that substantial dispersion remains after absorbing these sectoral differences motivates the model's firm-level heterogeneity in carbon intensity.

Figure A.5 shows that country-level differences are also sizable. The figure reports the most and least carbon-intensive countries in each dataset, omitting the middle of the distribution for readability. The cross-country ranking presents some differences across samples, reflecting differences in firm coverage, but both panels show that country composition is an important source of raw variation. This motivates the country fixed effects used in the main empirical analysis.

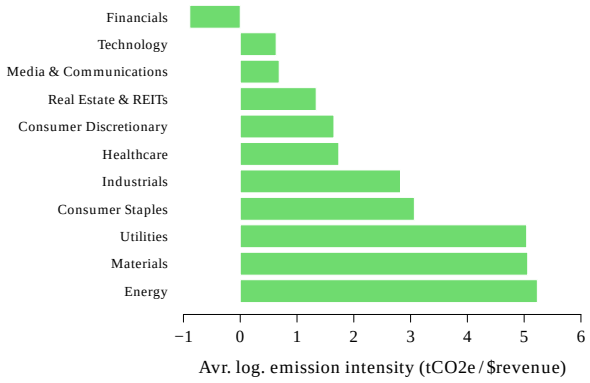
Finally, Table A.2 relates emission intensity to firm productivity in the merged ICE-ESG and Compustat sample. More productive firms are significantly less emission intensive, both unconditionally and after absorbing year, sector, and country fixed effects, although the standard errors are large. Quantitatively, however, productivity accounts for little of the overall dispersion: the R^2 remains close to 5 percent. Carbon intensity therefore captures a distinct dimension of firm heterogeneity, rather than simply reflecting standard productivity differences.

Figure A.3: Heterogeneity in emission intensities—output measures

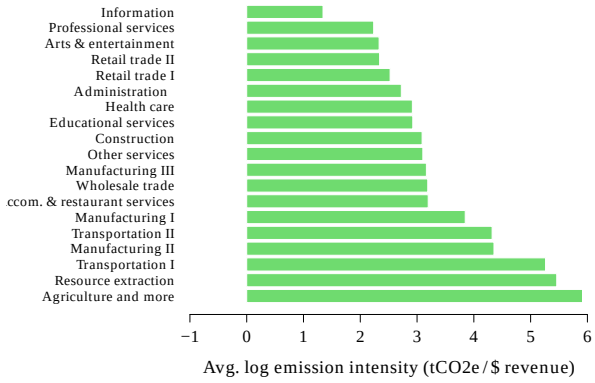


Note: Data from the ICE-ESG sample. The figure shows the distribution of log. emission intensities within the US energy, airlines and steel sectors. Emission intensities are measured by tons of CO_2e emissions divided by output. We measure output by: mega jules of energy produced, thousands of return flights flown and tons of steel produced, respectively. We consider only ICE observations with third-party-validated disclosures (Appendix A.1). We winsorize emission intensities at the top/bottom 1% level.

Figure A.4: Sectoral heterogeneity in emission intensities



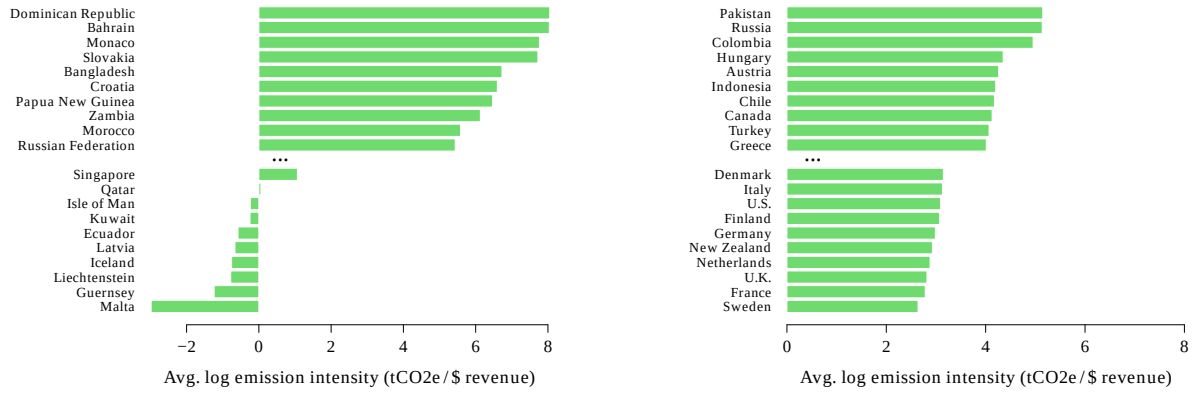
(a) ICE-ESG sample



(b) Trucost sample

Note: Data from the ICE-ESG and TruCost samples. The figure shows the avr. log emission intensity within ICE and Trucost sectors (NAICS level 2). Emission intensities are measured by tons of CO_2e emissions divided by revenue (millions of $). For ICE, we consider only observations with third-party-validated disclosures (Appendix A.1); Trucost does not provide an analogous validation flag in our extract. We winsorize all emission intensities at the top/bottom 1% level.

Figure A.5: Country heterogeneity in emission intensities



(a) ICE-ESG sample

(b) Trucost sample

Note: Data from the ICE-ESG and Trucost samples. The figure shows the average log. emission intensity for different countries for the top and bottom 10 countries by log. emission intensity. Labels for a subset of countries are given. Emission intensities are measured by tons of CO_2e emissions divided by revenue (millions of $). For ICE, we consider only observations with third-party-validated disclosures (Appendix A.1); Trucost does not provide an analogous validation flag in our extract. We winsorize emission intensities at the top/bottom 1% level.

A.3 Sources of heterogeneity, disclosure standards, and misreporting

Sources of heterogeneity and disclosure standards. Section 2 documents that emissions intensities are highly heterogeneous across firms even within narrowly defined sectors—a finding first highlighted by Lyubich et al. (2018). This is consistent with a large body of evidence from the engineering and sustainability literature and is reflected in the emissions calculation standards adopted by all major ETS markets, as well as by the carbon tax framework introduced in Chile. We now explain, within the European ETS frameworks, constituting some of the largest carbon price mechanisms, which practical elements—fuel composition and technology choices—drive the wide heterogeneity in emissions outcomes, and how they are disclosed under existing regulatory frameworks.

Emissions disclosure in major ETS markets must define an installation boundary, identify emission sources and source streams, specify the monitoring methodology, and document procedures for collecting, verifying, and retaining data (see European Commission Directorate-General for Climate Action, 2025; UK Department for Energy Security and Net Zero, 2025). Emissions are rarely measured automatically at the stack: in the EU ETS, only 145 installations, or 1.7 percent of installations across 22 countries, reported using continuous emissions monitoring systems (CEMS) in 2022 (European Commission,

Table A.2: Carbon intensity and productivity

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

Notes: Least-squares estimates from the merged Trucost-Compustat sample. The top and bottom 1 percent of emission intensities (tons of CO_2e emissions divided by revenue in millions of $) and total factor productivity are removed. TFP is estimated as in Ottonello and Winberry (2020). NAICS-4 sectoral definitions are used throughout. Robust clustered (sectors) standard errors in parentheses. Sample: 2010–2024.

2023b).[40] Instead, reported emissions are calculated from firm-specific inputs—fuel quantities, emission factors, process parameters, and abatement records—and occasionally validated by a third party.

A useful way to summarize this accounting framework is through the equation:

$$g = \underbrace{\left(\sum_i \omega_i^e \cdot (1 - x_i)\right)}_{\text{Technological factors}} \cdot \underbrace{\left(\sum_i \omega_i^e \cdot z_i\right)}_{\text{Average emission factor}} \cdot \underbrace{\left(\sum_i e_i\right)}_{\text{Energy amount}}, \tag{A1}$$

where g denotes reported emissions in CO_2e units, e_i is energy use from source stream i, $\omega_i^e \equiv e_i / \sum_j e_j$ is the corresponding energy share, z_i is the emissions factor of that source stream, and x_i summarizes source-stream-specific technological and process adjustments that reduce reportable emissions. The equation is a parsimonious version of the calculation-based reporting logic: reported emissions equal total energy use, multiplied by the average emissions intensity of the energy bundle, and adjusted for technologies or processes that reduce the share of potential emissions ultimately reported.

[40] CEMS require continuous measurement of gas concentrations and exhaust flows, missing-data substitution protocols, certified equipment, along with regular quality assurance and surveillance tests by accredited laboratories. They are infeasible where emissions are diffuse or lack a well-defined release point. As a result, they are costly, technically demanding, and require a well-defined emission point, making them typically applicable only to a few large plants with concentrated emission streams.

• *Energy amount.* The last term, $\sum_i e_i$, in Equation (A1) captures reported energy use. The installation must identify all relevant source streams (and report the activity data required to reconstruct it): physical quantities of fuels and materials, stock changes, and the net calorific values used to convert physical quantities into energy units. The index i runs over a fine partition of inputs—distinguishing coal ranks and batches, natural gas, refinery gas, fuel oil, coke, waste gas, biomass, mixed fossil-biomass fuels, and other process energy inputs. An installation's reported emissions therefore depend on whether the operator has disclosed the quantity and composition of every fuel used in production, before any emissions factor or technological adjustment is applied.

• *Average emission factor.* The middle term, $\sum_i \omega_i^e \cdot z_i$, in Equation (A1) is the average warming content of the energy bundle, where z_i maps one unit of energy from source stream i into CO_2e. Depending on the tier and source stream, z_i may be drawn from legislated default factors, supplier-guaranteed values, or installation-commissioned laboratory analysis. The margin is quantitatively meaningful even for relatively homogeneous fuels: IPCC default factors for stationary combustion range from 94.6 kg CO_2/GJ for bituminous coal to 101.0 for lignite (IPCC, 2006), and more disaggregated evidence shows that factors vary further by origin, destination, and sample composition (Quick, 2010). Heterogeneous inputs exhibit substantially wider dispersion. Biomass and waste-derived fuels can span the full range from zero to fossil-like emission factors, because ETS accounting assigns zero to sustainable biomass while non-biogenic fractions—such as plastics in refuse-derived fuels—carry emission factors comparable to coal (European Commission, 2018a).

• *Technological factors.* The first term, $\sum_i \omega_i^e \cdot (1 - x_i)$, in Equation (A1) is the energy-weighted share of potential emissions remaining after source-stream-specific technological and process adjustments. The correction x_i is not a single technology but a composite of abatement choices that can differ substantially across installations using identical fuel mixes. Carbon capture at the stack targets CO_2 directly, with commercial post-combustion systems designed around capture rates near 90 percent (International Energy Agency, 2024). Incomplete combustion and oxidation corrections reduce reportable emissions when carbon exits as ash, slag, or unburnt residue rather than CO_2. Process modifications—such as chemical additives, altered conversion rates, or waste-gas recovery—reduce emissions by changing the stoichiometry of combustion itself. For non-CO_2 greenhouse gases, selective catalytic reduction and similar post-combustion controls can reduce NO_x emissions by 80–90 percent (U.S. Environmental Protection Agency, 2002), while methane oxidation and flaring govern the CO_2e contribution of fugitive gas streams. The implication is that two installations operating with identical fuel quantities

and composition can report substantially different emissions depending on which subset of these technologies is installed—a feature which our model framework in Section 3 directly captures.

Reporting margins. The same identity from Equation (A1) clarifies the main documented margins of misreporting and how it can occur.

First, operators can omit energy sources entirely. In the Brikel case, discussed in the main text, the installation was documented burning waste streams—plastics and petroleum-derived industrial waste—alongside coal without the corresponding permit updates or public disclosure (ClientEarth, 2020; Bivol, 2020). Because such materials carry fossil-carbon content comparable to or exceeding that of coal, omitting them from reported source streams directly suppressed emissions. Second, firms can manipulate the emission factor. Brikel, allegedly burning lignite with reference factors between 96.1 and 104.2 kg CO_2/GJ, allegedly reported factors of 90.5, 79.5, and 76.2 in 2018, 2019, and 2020, respectively (Organized Crime and Corruption Reporting Project, 2021). The plant supported these figures with laboratory analyses that are difficult to challenge without independent physical access to the fuel sample.

Third, firms can evade the reporting boundary altogether. Since only installations above a threshold are often required to participate in the disclosures, misreporting activity data or fragmenting operations across sites can place an installation below the regulatory perimeter. Diageo operated three Scottish whisky-production sites carrying out regulated combustion activity without the required permits for multiple years. By never registering these sites, their emissions fell entirely outside the reporting population until the breach was identified. ExxonMobil's Mossmorran case shows the same logic applied within a registered installation: specific emission sources inside the plant were omitted from the 2008 EU ETS return, understating emissions by approximately 33,000t.

Finally, notice that the Brikel disclosures were certified by a third-party verifier, and the European Public Prosecutor's Office later opened a fraud investigation into a verification company alleged to have submitted false data and documentation. None of the other cases were uncovered through routine regulatory oversight: investigative journalists, NGOs, whistle-blowers, and later prosecutors played the central role in each. This pattern is consistent with the disclosure framework above. Regulators have access only to reported data—monitoring plans, audited records, and verified calculations—and do not conduct systematic on-site inspections of fuel composition or emission sources (European Commission, 2018b). Without independent physical access to the installation, detecting manipulation requires either an anomaly large enough to trigger scrutiny or an external

informant. The documented cases are therefore likely the tip of the iceberg.

B Model framework and results

B.1 Proofs and derivations

Proof of Lemma 1. We start with firm (i, j, ℓ)'s problem:

$$\max_{\{n_{ij\ell,\tau}, k_{ij\ell,\tau}, e_{ij\ell,\tau}, x_{ij\ell,\tau}\}_{\tau=t}^{\infty}} v_{ij\ell,t} = \pi_{ij\ell,t} + \mathbb{1}_{\{\ell<L\}} \cdot \beta \cdot \frac{u'(C_{t+1})}{u'(C_t)} \cdot v_{ij,\ell+1,t+1}, \tag{A2}$$

where:

$$\pi_{ij\ell,t} = A\left(G_t\right) \cdot \exp^{a_i} \cdot F\left(n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}\right) - W_t \cdot n_{ij\ell,t} - R_t \cdot k_{ij\ell,t} - P \cdot e_{ij\ell,t} - Q(x_{ij\ell,t}). \tag{A3}$$

Since the firm can costlessly reoptimize input choices at any point in time, it follows that the optimal choices maximize per-period profits and hence satisfy the expressions:

$$mrpn_{ij\ell,t} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_n\left(\cdot\right) = W_t \quad \text{and} \quad mrpk_{ij\ell,t} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_k\left(\cdot\right) = R_t, \tag{A4}$$

$$mrpe_{ij\ell,t} \equiv A(G_t) \cdot \exp^{a_i} \cdot F_e\left(\cdot\right) = P_t \quad \text{and} \quad Q'(x_{ij\ell,\tau}) = 0. \tag{A5}$$

Furthermore, notice that from the condition for energy:

$$\frac{\partial y_{ij\ell,t}}{\partial e_{ij\ell,t}} \cdot \frac{e_{ij\ell,t}}{y_{ij\ell,t}} = \frac{\partial \log y_{ij\ell,t}}{\partial \log e_{ij\ell,t}} = P \cdot \frac{e_{ij\ell,t}}{y_{ij\ell,t}}. \tag{A6}$$

Thus,

$$\log\left(\frac{g_{ij\ell,t}}{y_{ij\ell,t}}\right) = \log\left(\frac{\exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t}}{P \cdot \frac{\partial \log e_{ij\ell,t}}{\partial \log y_{ij\ell,t}} \cdot e_{ij\ell,t}}\right) = \log(z_j) + \log\left(\frac{\eta}{P}\right) + \mathcal{O}(1), \tag{A7}$$

since $x_{ij\ell,t} = 0$ from Equation (A5) and $\frac{\partial \log y_{ij\ell,t}}{\partial \log e_{ij\ell,t}} = \eta + \mathcal{O}(1)$. ■

Proof of Lemma 2. The full-information planner maximizes (16) subject to subject to Equations (3), (4), (5), (12), and (13). This yields the following Lagrangian:

$$\begin{aligned}\mathcal{L} = \sum_t \beta^t \cdot \Bigg[& u(C_t) + \Lambda_{C,t} \cdot \Bigg(- C_t - K_{t+1} + (1-\delta_K) \cdot K_t \\ & + \sum_{ij\ell} w_{ij\ell} \cdot \left(A(G_t) \cdot \exp^{a_i} \cdot F\left(n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}\right) - P \cdot e_{ij\ell,t} - Q(x_{ij\ell,t})\right) \Bigg) \\ & - \Lambda_{C,t} \cdot \Lambda_{G,t} \cdot \left((1-\delta_G) \cdot G_t + \sum_{ij\ell} w_{ij\ell} \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t} - G_{t+1} \right) \\ & + \Lambda_{C,t} \cdot \Lambda_{N,t} \cdot \left(1 - \sum_{ij\ell} w_{ij\ell} \cdot n_{ij\ell,t} \right) + \Lambda_{C,t} \cdot \Lambda_{K,t} \cdot \left(K_t - \sum_{ij\ell} w_{ij\ell} \cdot k_{ij\ell,t} \right) \Bigg]. \end{aligned} \tag{A8}$$

• The first-order conditions for C_t, K_{t+1}, and G_{t+1} are:

$$\frac{\partial \mathcal{L}}{\partial C_t} = u'(C_t) - \Lambda_{c,t} = 0 \tag{A9}$$

$$\frac{\partial \mathcal{L}}{\partial K_{t+1}} = -\beta^t \cdot \Lambda_{C,t} + \beta^{t+1} \cdot \Lambda_{C,t+1} \cdot ((1-\delta_K) + \Lambda_{K,t+1}) = 0 \tag{A10}$$

$$\frac{\partial \mathcal{L}}{\partial G_{t+1}} = \beta^t \cdot \Lambda_{C,t} \cdot \Lambda_{G,t} + \beta^{t+1} \cdot \Lambda_{C,t+1} \cdot \frac{A'(G_{t+1})}{A(G_{t+1})} \cdot \mathcal{Y}_{t+1} - \beta^{t+1} \cdot \Lambda_{C,t+1} \cdot \Lambda_{G,t+1} \cdot (1-\delta_G) = 0, \tag{A11}$$

which combined give the Euler equations for consumption and the social cost of carbon:

$$u'(C_t) = \beta \cdot u'(C_{t+1}) \cdot ((1-\delta_K) + \Lambda_{K,t+1}) \tag{A12}$$

$$\Lambda_{G,t} = \beta \cdot \frac{u'(C_{t+1})}{u'(C_t)} \cdot \left(-\frac{A'(G_{t+1})}{A(G_{t+1})} \cdot \mathcal{Y}_{t+1} + \Lambda_{G,t+1} \cdot (1-\delta_G) \right). \tag{A13}$$

• The first-order conditions for a firm's input $\left(n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t},\right)$ are:

$$\frac{\partial \mathcal{L}}{\partial n_{ij\ell,t}} = \beta^t \cdot \Lambda_{C,t} \cdot w_{ij\ell} \cdot mrpn_{ij\ell,t} - \beta^t \cdot \Lambda_{C,t} \cdot \Lambda_{N,t} \cdot w_{ij\ell} = 0 \tag{A14}$$

$$\frac{\partial \mathcal{L}}{\partial k_{ij\ell,t}} = \beta^t \cdot \Lambda_{C,t} \cdot w_{ij\ell} \cdot mrpk_{ij\ell,t} - \beta^t \cdot \Lambda_{C,t} \cdot \Lambda_{K,t} \cdot w_{ij\ell} = 0 \tag{A15}$$

$$\frac{\partial \mathcal{L}}{\partial e_{ij\ell,t}} = \beta^t \cdot \Lambda_{C,t} \cdot (w_{ij\ell} \cdot mrpe_{ij\ell,t} - w_{ij\ell} \cdot P) + \beta^t \cdot \Lambda_{C,t} \cdot \Lambda_{G,t} \cdot w_{ij\ell} \cdot \exp^{z_j - x_{ij\ell,t}} = 0 \tag{A16}$$

$$\frac{\partial \mathcal{L}}{\partial x_{ij\ell,t}} = -\beta^t \cdot \Lambda_{C,t} \cdot w_{ij\ell} \cdot Q'(x_{ij\ell,t}) - \beta^t \cdot \Lambda_{C,t} \cdot \Lambda_{G,t} \cdot w_{ij\ell} \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t} = 0. \tag{A17}$$

Thus,

$$mrpn_{ij\ell,t} = \Lambda_{N,t} \quad \text{and} \quad mrpk_{ij\ell,t} = \Lambda_{K,t}, \tag{A18}$$

and:

$$mrpe_{ij\ell,t} = P + \Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}} \quad \text{and} \quad Q'\left(x_{ij\ell,t}\right) = \Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t}. \tag{A19}$$

• The market-clearing conditions for labor and capital, the low-of-motion for emissions, and the aggregate resource constraint, in turn, follow from the relevant first-order conditions for $\Lambda_{C,t}, \Lambda_{G,t}, \Lambda_{N,t}$, and $\Lambda_{K,t}$. This completes the proof. ■

Proof of Corollary 1. The result follows directly from the second part of Lemma 1. Notice that, for any (i, ℓ, t), and any j', j with $z_{j'} > z_j$:

$$e^{FB}_{ij'\ell,t} < e^{FB}_{ij\ell,t} < e^{DE}_{ij\ell,t} \quad \text{and} \quad x^{FB}_{ij'\ell,t} > x^{FB}_{ij\ell,t} > x^{DE}_{ij\ell,t} = 0. \tag{A20}$$

As a result, for any t, $\pi^{FB}_{ij\ell,t} > \pi^{FB}_{ij'\ell,t}$ by revealed preference, and hence $v^{FB}_{ij\ell,t} > v^{FB}_{ij'\ell,t}$. ■

Proof of Lemma 3. The Mirrleesian planner solves the problem:

$$\max_{\{C_t, K_{t+1}, G_{t+1}, R_t, W_t, \{k_{ij\ell,t}, n_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t}, v_{ij,\ell+1,t+1}\}_{i,j,\ell}\}_{t=0}^{\infty}} \mathcal{U} = \sum_{t=0}^{\infty} \beta^t \cdot u\left(C_t\right) \tag{A21}$$

subject to Equations (3), (4), (5), (12), and (13). In addition, the Mirrleesian planner faces the incentive compatibility constraint, also detailed in the main text:

$$v_{ij\ell,t} = \sum_{\tau=0}^{L-\ell} m_{t,t+\tau}\pi(n_{ij\ell+\tau,t+\tau}, k_{ij\ell+\tau,t+\tau}, e_{ij\ell+\tau,t+\tau}, x_{ij\ell+\tau,t+\tau}) \geq \sum_{\tau=0}^{L-\ell} m_{t,t+\tau}\pi(n_{ij'\ell+\tau,t+\tau}, k_{ij'\ell+\tau,t+\tau}, e_{ij'\ell,t+\tau}, x_{ij'\ell+\tau,t+\tau}) = v_{ij'\ell,t}, \quad \forall j' \neq j. \tag{A22}$$

We abstract below from the individual rationality constraint, as this constraint is automatically slack due to the assumption of decreasing returns to scale. Finally, the planner also faces the following implementability constraint:

$$\Lambda_{N,t} = W_t \quad \text{and} \quad \Lambda_{K,t} = R_t, \tag{A23}$$

where $\Lambda_{N,t}$ and $\Lambda_{K,t}$ are the shadow costs of labor and capital.

Since profits do not depend directly on z_j, firm can freely (and credibly) report any carbon intensity $z_j \in \mathcal{Z}$ that gives them the highest profit conditional on their productivity type, a_i, and their age, ℓ. Thus, in equilibrium, incentive compatibility requires Equation (A22) to hold with equality implying that:

$$\sum_{\tau=0}^{L-\ell} m_{t,t+\tau}\pi(n_{ij\ell+\tau,t+\tau}, k_{ij\ell+\tau,t+\tau}, e_{ij\ell+\tau,t+\tau}, x_{ij\ell+\tau,t+\tau}) = v_{i\ell,t}, \tag{A24}$$

where $v_{i\ell,t} \in \mathbb{R}_+$ is the common value for firms of type $\{i, \ell\}$ at time t.

Now, starting from the last year of a firm's life, $\ell = L$ it follows that:

$$\pi(n_{ijL,t}, k_{ijL,t}, e_{ijL,t}, x_{ijL,t}) = v_{i\ell,t} \equiv \tilde{\pi}_{iL,t}. \tag{A25}$$

Thus, at $\ell = L - 1$, we have:

$$\begin{aligned} \pi(n_{ijL-1,t}, k_{ijL-1,t}, e_{ijL-1,t}, x_{ijL-1,t}) + m_{t,t+1}\pi(n_{ijL,t}, k_{ijL,t}, e_{ijL,t}, x_{ijL,t}) &= v_{iL-1,t} \\ \pi(n_{ijL-1,t}, k_{ijL-1,t}, e_{ijL-1,t}, x_{ijL-1,t}) + m_{t,t+1}v_{iL,t} &= v_{iL-1,t} \\ \pi(n_{ijL-1,t}, k_{ijL-1,t}, e_{ijL-1,t}, x_{ijL-1,t}) = v_{iL-1,t} - m_{t,t+1}v_{iL,t} &\equiv \tilde{\pi}_{iL-1,t}, \end{aligned} \tag{A26}$$

and so on until $\ell = 1$. This implies that an incentive compatible allocation satisfying

(A24) also satisfies equivalently the following set of per-period constraints:

$$\pi_{ij\ell,t} = \tilde{\pi}_{i\ell,t}, \tag{A27}$$

which we use below.

The Lagrangian solved by the Mirrleesian planner can therefore be stated as:

$$\begin{aligned}
\mathcal{L} = \sum_t \beta^t \cdot \Bigg[& u(C_t) + \Lambda_{C,t} \cdot \Bigg(- C_t - K_{t+1} + (1-\delta_K) \cdot K_t \\
& + \sum_{ij\ell} w_{ij\ell} \cdot \Big(A(G_t) \cdot e^{a_i} \cdot F(n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}) - P \cdot e_{ij\ell,t} - Q(x_{ij\ell,t}) \Big) \Bigg) \\
& - \Lambda_{C,t} \cdot \Lambda_{G,t} \cdot \left((1-\delta_G) \cdot G_t + \sum_{ij\ell} w_{ij\ell} \cdot e^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t} - G_{t+1} \right) \\
& + \Lambda_{C,t} \cdot \Lambda_{N,t} \cdot \left(1 - \sum_{ij\ell} w_{ij\ell} \cdot n_{ij\ell,t} \right) \\
& + \Lambda_{C,t} \cdot \Lambda_{K,t} \cdot \left(K_t - \sum_{ij\ell} w_{ij\ell} \cdot k_{ij\ell,t} \right) \\
& + \Lambda_{C,t} \cdot \sum_{ij\ell} w_{ij\ell} \Lambda_{\pi,ij\ell,t} \cdot \left(\pi_{ij\ell,t} - \tilde{\pi}_{i\ell,t} \right) \Bigg].
\end{aligned} \tag{A28}$$

The first-order conditions for a firm's inputs $(n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t})$ is thus:

$$mrpn_{ij\ell,t} = \Lambda_{N,t} \tag{A29}$$

$$mrpk_{ij\ell,t} = \Lambda_{K,t} \tag{A30}$$

$$mrpe_{ij\ell,t} = P + \frac{\Lambda_{G,t}}{1 + \Lambda_{\pi,ij\ell,t}} \cdot \exp^{z_j - x_{ij\ell,t}} \tag{A31}$$

$$Q'(x_{ij\ell,t}) = \frac{\Lambda_{G,t}}{1 + \Lambda_{\pi,ij\ell,t}} \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t}, \tag{A32}$$

where (A31) and (A32) can be combined as follows:

$$mrpe_{ij\ell,t} = \frac{Q'(x_{ij\ell,t})}{e_{ij\ell,t}}. \tag{A33}$$

Now, notice that Equations (A27), (A29), (A30), and (A33) jointly characterize the firm's choice of $(n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t})$ independently of z_j—i.e., $(n_{ij\ell,t}, k_{ij\ell,t}, e_{ij\ell,t}, x_{ij\ell,t}) =$

$(n_{i\ell,t}, k_{i\ell,t}, e_{i\ell,t}, x_{i\ell,t})$. There is bunching across types j.

Finally, the first-order condition with respect to $\tilde{\pi}_{i\ell,t}$ takes the form:

$$\sum_j w_{j|i\ell} \cdot \Lambda_{\pi,ij\ell,t} = 0, \tag{A34}$$

where $w_{j|i\ell} \equiv \frac{w_{ij\ell}}{\sum_j w_{ij\ell}}$. Combining (A31) and (A32) with (A34), and exploiting bunching across types j, yields:

$$mrpe_{i\ell,t} = P + \sum_j w_{j|i\ell} \cdot \Lambda_{G,t} \cdot \exp^{z_j - x_{i\ell,t}} \tag{A35}$$

$$Q'(x_{i\ell,t}) = \sum_j w_{j|i\ell} \cdot \Lambda_{G,t} \cdot \exp^{z_j - x_{i\ell,t}} \cdot e_{i\ell,t}. \tag{A36}$$

Since

$$w_{j|i\ell} \equiv \frac{w_{ij\ell}}{\sum_j w_{ij\ell}} = \frac{\omega_{ij} \cdot L^{-1}}{\sum_j \omega_{ij} \cdot L^{-1}} = \frac{\omega_{ij}}{\sum_j \omega_{ij}}, \tag{A37}$$

is independent of ℓ, it follows that the entire system of Equations (A29), (A30), (A35), and (A36), characterizing $(n_{i\ell,t}, k_{i\ell,t}, e_{i\ell,t}, x_{i\ell,t})$, does not directly depend on ℓ either—i.e., $(n_{i\ell,t}, k_{i\ell,t}, e_{i\ell,t}, x_{i\ell,t}) = (n_{i,t}, k_{i,t}, e_{i,t}, x_{i,t})$. There is bunching across types ℓ in addition to types j. This concludes the proof. ■

Proof of Proposition 1: Follows directly from Equation (31) and Lemma 2 and 3. ■

Proof of Proposition 2: We express all the first-order conditions across the different equilibria $\mathcal{P}$ as follows:

$$mrpk_{ij\ell,t} = R_t, \tag{A38}$$

$$mrpn_{ij\ell,t} = W_t, \tag{A39}$$

$$mrpe_{ij\ell,t} = (1 + \Delta^{\mathcal{P}}_{ij\ell,t}) \cdot P; \tag{A40}$$

where $\mathcal{P} \in \{DE, FB, MR, NV\}$ with

$$\Delta_{ij\ell,t}^{DE} = 0, \tag{A41}$$

$$\Delta_{ij\ell,t}^{FB} = \frac{\Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}}}{P}, \tag{A42}$$

$$\Delta_{ij\ell,t}^{MR} = \frac{\Lambda_{G,t} \cdot \mathbb{E}[\exp^{z_j - x_{ij\ell,t}} | a_i]}{P}, \tag{A43}$$

$$\Delta_{ij\ell,t}^{NV}, = \frac{\Lambda_{G,t} \cdot \exp^{\underline{z} - x_{ij\ell,t}}}{P}. \tag{A44}$$

Then, we can express each input choice as follows:

$$k_{ij\ell,t} = (A \cdot \exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{\gamma}{R_t}\right)^{\frac{1-\alpha-\nu}{\zeta}} \cdot \left(\frac{\alpha}{W_t}\right)^{\frac{\alpha}{\zeta}} \cdot \left(\frac{\nu}{(1+\Delta_{ij\ell,t}^{\mathcal{P}}) \cdot P}\right)^{\frac{\nu}{\zeta}}, \tag{A45}$$

$$n_{ij\ell,t} = (A \cdot \exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{\gamma}{R_t}\right)^{\frac{\gamma}{\zeta}} \cdot \left(\frac{\alpha}{W_t}\right)^{\frac{1-\gamma-\nu}{\zeta}} \cdot \left(\frac{\nu}{(1+\Delta_{ij\ell,t}^{\mathcal{P}}) \cdot P}\right)^{\frac{\nu}{\zeta}}, \tag{A46}$$

$$e_{ij\ell,t} = (A \cdot \exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{\gamma}{R_t}\right)^{\frac{\gamma}{\zeta}} \cdot \left(\frac{\alpha}{W_t}\right)^{\frac{\alpha}{\zeta}} \cdot \left(\frac{\nu}{(1+\Delta_{ij\ell,t}^{\mathcal{P}}) \cdot P}\right)^{\frac{1-\alpha-\gamma}{\zeta}}; \tag{A47}$$

implying the following expression for output:

$$y_{ij\ell,t} = (A \cdot \exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{\gamma}{R_t}\right)^{\frac{\gamma}{\zeta}} \cdot \left(\frac{\alpha}{W_t}\right)^{\frac{\alpha}{\zeta}} \cdot \left(\frac{\nu}{(1+\Delta_{ij\ell,t}^{\mathcal{P}}) \cdot P}\right)^{\frac{\nu}{\zeta}}. \tag{A48}$$

We can then express aggregate gross output as follows:

$$\mathcal{Y}_t \equiv \sum_{ij\ell} w_{ij\ell} y_{ij\ell,t} \tag{A49}$$

$$= A_t^{\frac{1}{\zeta}} \cdot \left(\frac{\gamma}{R_t}\right)^{\frac{\gamma}{\zeta}} \cdot \left(\frac{\alpha}{W_t}\right)^{\frac{\alpha}{\zeta}} \cdot \left(\frac{\nu}{P}\right)^{\frac{\nu}{\zeta}} \cdot \sum_{ij\ell} w_{ij\ell} \cdot (\exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{1}{1+\Delta^{\mathcal{P}}_{ij\ell,t}}\right)^{\frac{\nu}{\zeta}}$$

$$= \frac{A_t^{\frac{1}{\zeta}} \cdot \left(\frac{\gamma}{R_t}\right)^{\frac{\gamma}{\zeta}} \cdot \left(\frac{\alpha}{W_t}\right)^{\frac{\alpha}{\zeta}} \cdot \left(\frac{\nu}{P}\right)^{\frac{\nu}{\zeta}} \cdot \sum_{ij\ell} w_{ij\ell} \cdot (\exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{1}{1+\Delta^{\mathcal{P}}_{ij\ell,t}}\right)^{\frac{\nu}{\zeta}}}{\left(\sum_{ij\ell} n_{ij\ell,t}\right)^{\alpha} \cdot \left(\sum_{ij\ell} k_{ij\ell,t}\right)^{\gamma} \cdot \left(\sum_{ij\ell} e_{ij\ell,t}\right)^{\nu}} \cdot N_t^{\alpha} \cdot K_t^{\gamma} \cdot E_t^{\nu}$$

$$= A_t \cdot \frac{\sum_{ij\ell} w_{ij\ell} \cdot (\exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{1}{1+\Delta^{\mathcal{P}}_{ij\ell,t}}\right)^{\frac{\nu}{\zeta}}}{\left(\sum_{ij\ell} w_{ij\ell} \cdot (\exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{1}{1+\Delta^{\mathcal{P}}_{ij\ell,t}}\right)^{\frac{\nu}{\zeta}}\right)^{\alpha+\gamma} \cdot \left(\sum_{ij\ell} w_{ij\ell} \cdot (\exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{1}{1+\Delta^{\mathcal{P}}_{ij\ell,t}}\right)^{\frac{1-\alpha-\gamma}{\zeta}}\right)^{\nu}} \cdot N_t^{\alpha} \cdot K_t^{\gamma} \cdot E_t^{\nu}$$

$$= A_t \cdot \exp^{a} \cdot \underbrace{\frac{\sum_{ij\ell} w_{ij\ell} \cdot (\exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{1}{1+\Delta^{\mathcal{P}}_{ij\ell,t}}\right)^{\frac{\nu}{\zeta}} \cdot \exp^{-a}}{\left(\sum_{ij\ell} w_{ij\ell} \cdot (\exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{1}{1+\Delta^{\mathcal{P}}_{ij\ell,t}}\right)^{\frac{\nu}{\zeta}}\right)^{\alpha+\gamma} \cdot \left(\sum_{ij\ell} w_{ij\ell} \cdot (\exp^{a_i})^{\frac{1}{\zeta}} \cdot \left(\frac{1}{1+\Delta^{\mathcal{P}}_{ij\ell,t}}\right)^{\frac{1-\alpha-\gamma}{\zeta}}\right)^{\nu}}}_{\equiv \Omega_t^{\mathcal{P}}} \tag{A50}$$

$$\cdot N_t^{\alpha} \cdot K_t^{\gamma} \cdot E_t^{\nu},$$

where $\exp^{a} \equiv (\sum_{ij\ell} w_{ij\ell} \cdot \exp^{\frac{a_i}{\zeta}})^{\zeta}$.

This can be expressed in terms of consumption as follows:

$$C_t = \mathcal{Y}_t - P \cdot E_t - \sum_{ij\ell} w_{ij\ell,t} \cdot Q(x_{ij\ell,t}), \tag{A51}$$

$$= A_t \underbrace{\left(1 - \frac{P \cdot E_t + \sum_{ij\ell} w_{ij\ell,t} \cdot Q(x_{ij\ell,t})}{\mathcal{Y}_t}\right)}_{\equiv \Theta_t} \cdot \Omega_t^{\mathcal{P}} \cdot N_t^{\alpha} \cdot K_t^{\gamma} \cdot E_t^{\nu}. \tag{A52}$$

This concludes the proof. ■

B.2 Alternative implementations

Assume firms are charged a carbon tax $\tau_{ij\ell,t}(g_{ij\ell,t})$, a energy tax $\tau_{ij\ell,t}(e_{ij\ell,t})$, and receive an abatement subsidy $\tau_{ij\ell,t}(x_{ij\ell,t})$, then profits read as

$$\pi_{ij\ell,t} = y_{ijt,t} - W_t \cdot n_{ij\ell,t} - R_t \cdot k_{ij\ell,t} - P \cdot e_{ij\ell,t} - Q(x_{ij\ell,t}) - \tau_{ij\ell,t}(g_{ij\ell,t}) - \tau_{ij\ell,t}(e_{ij\ell,t}) + \tau_{ij\ell,t}(x_{ij\ell,t}). \tag{A53}$$

The first-order conditions of the firm are given by

$$mrpk_{ij\ell,t} = R_t, \tag{A54}$$

$$mrpn_{ij\ell,t} = W_t, \tag{A55}$$

$$mrpe_{ij\ell,t} = P + \tau'_{ij\ell,t}(g_{ij\ell,t}) \cdot \exp^{z_j - x_{ij\ell,t}} + \tau'_{ij\ell,t}(e_{ij\ell,t}), \tag{A56}$$

$$\tau'_{ij\ell,t}(g_{ij\ell,t}) \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t} + \tau'_{ij\ell,t}(x_{ij\ell,t}) = Q'(x_{ij\ell,t}); \tag{A57}$$

First best planner. Here we look at the implementation of the first best planner, i.e., the planner with full information on firms carbon types z_j, whose distorted first order conditions are given by the following set of equations:

$$\Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}} = \tau'_{ij\ell,t}(g_{ij\ell,t}) \cdot \exp^{z_j - x_{ij\ell,t}} + \tau'_{ij\ell,t}(e_{ij\ell,t}), \tag{A58}$$

$$\tau'_{ij\ell,t}(g_{ij\ell,t}) \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t} + \tau'_{ij\ell,t}(x_{ij\ell,t}) = \Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t}. \tag{A59}$$

Below we show that these first order condition can either be implemented by a carbon tax, or by a combination of an energy tax and abatement subsidy.

- *Carbon tax implementation:* To look at the implementation via a carbon tax only, we shut down the energy tax and the abatement subsidy, i.e., $\tau_{ij\ell,t}(e_{ij\ell,t}) = 0$, $\tau_{ij\ell,t}(x_{ij\ell,t}) = 0$. Then, the carbon tax implementing the first best is given by

$$\Lambda_{G,t} = \tau'_{ij\ell,t}(g_{ij\ell,t}). \tag{A60}$$

- *Energy tax and abatement subsidy implementation:* To look at the implementation via energy tax and abatement subsidy, we shut down the carbon tax, i.e., $\tau_{ij\ell,t}(g_{ij\ell,t}) = 0$. Then, the energy tax and abatement subsidy implementing the first best equilibrium

is given by

$$\Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}} = \tau'_{ij\ell,t}(e_{ij\ell,t}), \tag{A61}$$

$$\tau'_{ij\ell,t}(x_{ij\ell,t}) = \Lambda_{G,t} \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t}. \tag{A62}$$

Mirrleesian planner. Here we look at the implementation of the Mirrleesian planner whose distorted equilibrium conditions are:

$$\Lambda_{G,t} \cdot \mathbb{E}[\exp^{z_j} | a_i] \cdot \exp^{-x_{ij\ell,t}} = \tau'_{ij\ell,t}(g_{ij\ell,t}) \cdot \exp^{z_j - x_{ij\ell,t}} + \tau'_{ij\ell,t}(e_{ij\ell,t}), \tag{A63}$$

$$\tau'_{ij\ell,t}(g_{ij\ell,t}) \cdot \exp^{z_j - x_{ij\ell,t}} \cdot e_{ij\ell,t} + \tau'_{ij\ell,t}(x_{ij\ell,t}) = \Lambda_{G,t} \cdot \mathbb{E}[\exp^{z_j} | a_i] \cdot \exp^{-x_{ij\ell,t}} \cdot e_{ij\ell,t}. \tag{A64}$$

- *Carbon tax implementation:*

$$\Lambda_{G,t} \cdot \frac{\mathbb{E}[\exp^{z_j} | a_i]}{\exp^{z_j}} = \tau'_{ij\ell,t}(g_{ij\ell,t}). \tag{A65}$$

- *Energy tax and abatement subsidy implementation:*

$$\Lambda_{G,t} \cdot \mathbb{E}[\exp^{z_j} | a_i] \cdot \exp^{-x_{ij\ell,t}} = \tau'_{ij\ell,t}(e_{ij\ell,t}), \tag{A66}$$

$$\tau'_{ij\ell,t}(x_{ij\ell,t}) = \Lambda_{G,t} \cdot \mathbb{E}[\exp^{z_j} | a_i] \cdot \exp^{-x_{ij\ell,t}} \cdot e_{ij\ell,t}. \tag{A67}$$

B.3 Aggregate emission decomposition

Here, we present the derivation of the aggregate emissions decomposition as well as the derivations of each of its firm-level components in terms of covariances.

Aggregate decomposition. Start from the aggregate emission definition and substitute the definition of individual level emissions to get

$$G_t = \sum_{ij\ell} w_{ij\ell} g_{ij\ell,t}, \tag{A68}$$

$$= \left(\frac{\sum_{ij\ell} w_{ij\ell} \cdot \frac{e_{ij\ell,t}}{E_t} \cdot \exp^{z_j - x_{ij\ell,t} + X_t}}{\sum_{ij\ell} w_{ij\ell} \cdot \frac{e_{ij\ell,t}}{E_t} \cdot \exp^{z_j}} \right) \cdot \left(\sum_{ij\ell} w_{ij\ell} \cdot \frac{e_{ij\ell,t}}{E_t} \cdot \exp^{z_j} \right) \cdot \exp^{-X_t} \cdot E_t; \tag{A69}$$

similarly for the decentralized equilibrium case

$$G_t^{DE} = \left(\sum_{ij\ell} w_{ij\ell} \cdot \frac{e_{ij\ell,t}^{DE}}{E_t^{DE}} \cdot \exp^{z_j} \right) \cdot E_t^{DE}. \tag{A70}$$

$$\tag{A71}$$

Finally, taking the ratios of G_t and G_t^{DE} yileds the decomposition in Equation (36).

Abatement allocation component. Starting from the definition of the abatement component and rearranging yileds

$$\mathcal{Z}_x \equiv \frac{\sum_{ij\ell} w_{ij\ell} \cdot e_{ij\ell,t} \cdot \exp^{z_j - x_{ij\ell,t} + X_t}}{\sum_{ij\ell} w_{ij\ell} \cdot e_{ij\ell,t} \cdot \exp^{z_j}}, \tag{A72}$$

$$= \frac{1}{\mathbb{E}\left[\frac{g_{ij\ell,t}}{G_t} \cdot \exp^{x_{ij\ell,t} - X_t} \right]}; \tag{A73}$$

Now, the denominator can be approximated as follows

$$\mathbb{E}\left[\frac{g_{ij\ell,t}}{G_t} \cdot \exp^{x_{ij\ell,t} - X_t} \right] = \mathbb{E}\left[\frac{g_{ij\ell,t}}{G_t} \right] \mathbb{E}\left[\exp^{x_{ij\ell,t} - X_t} \right] + \mathbb{COV}\left(\frac{g_{ij\ell,t}}{G_t}, \exp^{x_{ij\ell,t} - X_t} \right), \tag{A74}$$

$$= \mathbb{E}\left[\exp^{x_{ij\ell,t} - \mathbb{E}[x_{ij\ell,t}]} \right] + \mathbb{COV}\left(\frac{g_{ij\ell,t}}{G_t}, \exp^{x_{ij\ell,t} - X_t} \right), \tag{A75}$$

$$\approx 1 + \mathbb{COV}\left(\frac{g_{ij\ell,t}}{G_t}, \exp^{x_{ij\ell,t} - X_t} \right); \tag{A76}$$

The last approximate equality follows from the fact that the first term arises from Jensen's inequality and is equal to one up to a first-order approximation. Therefore, it follows:

$$\log \mathcal{Z}_x = -\log\left(1 + \mathbb{COV}\left(\frac{g_{ij\ell,t}}{G_t}, \frac{\exp^{x_{ij\ell,t}}}{\exp^{X_t}} \right) \right), \tag{A77}$$

$$\approx -\mathbb{COV}\left(\frac{g_{ij\ell,t}}{G_t}, \frac{\exp^{x_{ij\ell,t}}}{\exp^{X_t}} \right). \tag{A78}$$

Energy reallocation component. Using the definition of the energy reallocation component, and some rearranging yields:

$$\mathcal{Z}_e^{FB} - \mathcal{Z}_e^{MR} = \frac{\sum_{ij\ell} w_{ij\ell} \cdot \frac{e^{FB}_{ij\ell,t}}{E^{FB}_t} \cdot \exp^{z_j}}{\sum_{ij\ell} w_{ij\ell} \cdot \frac{e^{DE}_{ij\ell,t}}{E^{DE}_t} \cdot \exp^{z_j}} - \frac{\sum_{ij\ell} w_{ij\ell} \cdot \frac{e^{MR}_{ij\ell,t}}{E^{MR}_t} \cdot \exp^{z_j}}{\sum_{ij\ell} w_{ij\ell} \cdot \frac{e^{DE}_{ij\ell,t}}{E^{FB}_t} \cdot \exp^{z_j}}, \tag{A79}$$

$$= \frac{\mathbb{E}\left[\frac{e^{FB}_{ij\ell,t}}{E^{FB}_t}\right]\mathbb{E}\left[\exp^{z_j}\right] + \mathbb{C}\left(\frac{e^{FB}_{ij\ell,t}}{E^{FB}_t}, \exp^{z_j}\right) - \mathbb{E}\left[\frac{e^{MR}_{ij\ell,t}}{E^{MR}_t}\right]\mathbb{E}\left[\exp^{z_j}\right] - \mathbb{C}\left(\frac{e^{MR}_{ij\ell,t}}{E^{MR}_t}, \exp^{z_j}\right)}{\mathbb{E}\left[\frac{e^{DE}_{ij\ell,t}}{E^{DE}_t}\right]\mathbb{E}\left[\exp^{z_j}\right] + \mathbb{C}\left(\frac{e^{DE}_{ij\ell,t}}{E^{DE}_t}, \exp^{z_j}\right)}, \tag{A80}$$

$$= \frac{\mathbb{E}\left[\exp^{z_j}\right] + \mathbb{C}\left(\frac{e^{FB}_{ij\ell,t}}{E^{FB}_t}, \exp^{z_j}\right) - \mathbb{E}\left[\exp^{z_j}\right] - \mathbb{C}\left(\frac{e^{MR}_{ij\ell,t}}{E^{MR}_t}, \exp^{z_j}\right)}{\mathbb{E}\left[\exp^{z_j}\right] + \mathbb{C}\left(\frac{e^{DE}_{ij\ell,t}}{E^{DE}_t}, \exp^{z_j}\right)}, \tag{A81}$$

$$= \frac{\mathbb{C}\left(\frac{e^{FB}_{ij\ell,t}}{E^{FB}_t}, \frac{\exp^{z_j}}{\mathbb{E}\left[\exp^{z_j}\right]}\right) - \mathbb{C}\left(\frac{e^{MR}_{ij\ell,t}}{E^{MR}_t}, \frac{\exp^{z_j}}{\mathbb{E}\left[\exp^{z_j}\right]}\right)}{1 + \mathbb{C}\left(\frac{e^{DE}_{ij\ell,t}}{E^{DE}_t}, \frac{\exp^{z_j}}{\mathbb{E}\left[\exp^{z_j}\right]}\right)}, \tag{A82}$$

$$= \mathbb{C}\left(\frac{e^{FB}_{ij\ell,t}}{E^{FB}_t}, \frac{\exp^{z_j}}{\mathbb{E}\left[\exp^{z_j}\right]}\right); \tag{A83}$$

since $\mathbb{C}\left(\frac{e^{DE}_{ij\ell,t}}{E^{DE}_t}, \frac{\exp^{z_j}}{\mathbb{E}\left[\exp^{z_j}\right]}\right) = \mathbb{C}\left(\frac{e^{MR}_{ij\ell,t}}{E^{DE}_t}, \frac{\exp^{z_j}}{\mathbb{E}\left[\exp^{z_j}\right]}\right) = 0$, as energy in these two equilibria is independent of carbon intensity.

C Quantitative results

C.1 Parametrization

Table C.1 reports the model parameters, their calibrated values, and the corresponding calibration targets discussed in Section 5.1.

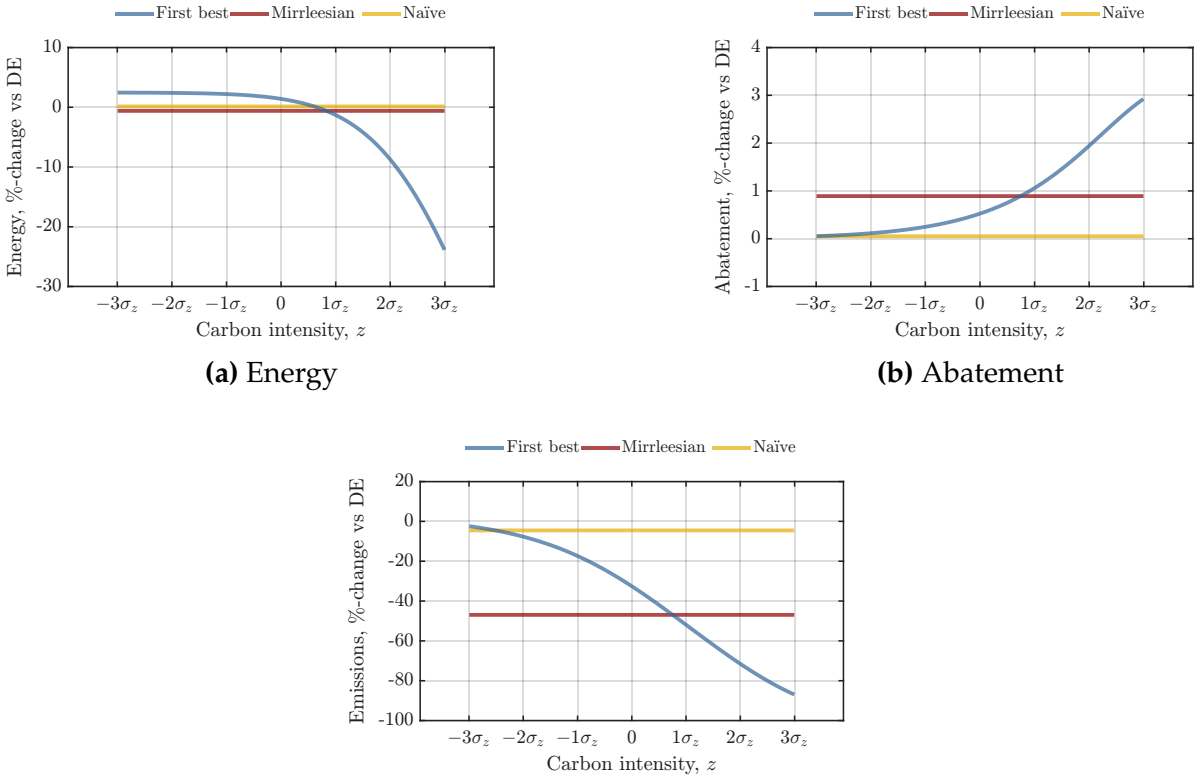
C.2 Firm-level outcomes

Figures C.1 and C.2 plot firm-level energy use, abatement, and emissions under the different planner allocations as percentage deviations from the decentralized equilibrium along the carbon intensity and productivity dimensions.

Table C.1: Parameters and targets

Description	Parameter	Value	Target
Economic block			
Discount rate	β	$0.985^{1}0$	
Capital elasticity	α	0.346	
Labor elasticity	γ	0.518	Standard
Energy elasticity	ν	0.036	
Capital depreciation rate	δ_K	$1-(1-0.07^{10})$	
Mean productivity	μ_a	0	Normalization
Std productivity	σ_a	0.6	Compustat data estimates
Climate block			
Carbon depreciation rate	δ_G	0.001	Golosov et al. (2014)
Abatement cost inverse elasticity	θ	2.0	Känzig (2023)
Damage semi-elasticity	ϕ	$4.9{\cdot}10^{-4}$	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

Figure C.1: Firm-level outcomes by carbon intensity



(a) Energy

(b) Abatement

(c) Emissions

Note: The figure reports firm-level energy use, abatement, and emissions under the full-information, Mirrleesian, and naïve planner allocations as percentage deviations from the decentralized equilibrium, by carbon-intensity type. Higher values of z correspond to more carbon-intensive firms.

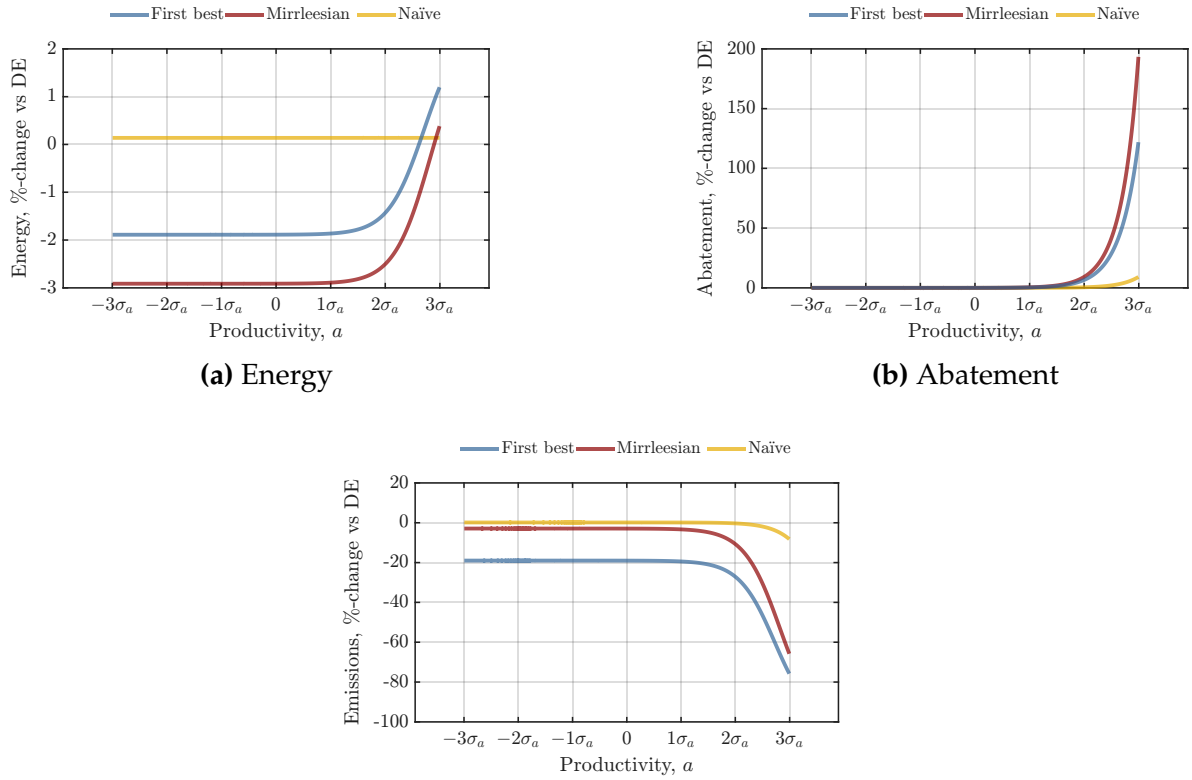
Figure C.1 illustrates how the planner allocations vary with firms' carbon intensity. Under full information, the planner separates firms by carbon intensity, imposing progressively higher carbon taxes on more carbon-intensive firms. As a result, these firms reduce energy use more, undertake greater abatement, and experience larger declines in emissions. By contrast, the Mirrleesian planner bunches firms along the carbon-intensity dimension to eliminate incentives to misreport. The naïve planner also induces bunching, as all firms optimally misreport the lowest carbon-intensity type. Because the resulting allocation fails to internalize the true variation in carbon intensities, it delivers only negligible reductions in emissions.

Figure C.2 illustrates how the planner allocations vary with firm productivity. Under full information, more productive firms use more energy because of their higher marginal revenue product of energy. Consequently, they also undertake more abatement and achieve larger reductions in emissions. The Mirrleesian allocation closely tracks these patterns across the productivity distribution, although it delivers somewhat smaller emission reductions on average because incentive compatibility prevents the planner from fully conditioning allocations on firms' carbon intensities. By contrast, the naïve planner fails to generate meaningful emission reductions because all firms optimally misreport the lowest carbon-intensity type, weakening the effective carbon price faced by high-emission firms.

C.3 Carbon tax

Figure C.3 reports the overall carbon tax bill as a share of revenues across the productivity distribution. Under both the full-information and Mirrleesian allocations, the tax burden declines with productivity, reflecting that more productive firms generate proportionally more revenue relative to their emissions because they undertake greater abatement. The Mirrleesian allocation closely tracks the first-best, although tax bills are slightly higher because the planner cannot fully condition allocations on carbon intensity and therefore achieves somewhat less abatement. By contrast, the naïve planner yields a substantially smaller tax burden because firms optimally misreport their carbon intensity, reducing the effective taxation of emissions.

Figure C.2: Firm-level outcomes by productivity



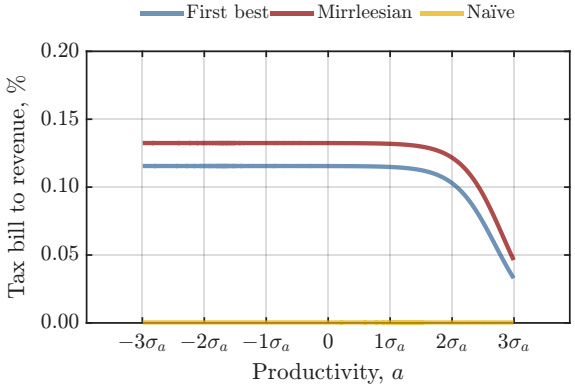
(a) Energy

(b) Abatement

(c) Emissions

Note: The figure reports firm-level energy use, abatement, and emissions under the full-information, Mirrleesian, and naïve planner allocations as percentage deviations from the decentralized equilibrium, by productivity type. Higher values of z correspond to more productive firms.

Figure C.3: Carbon tax bill as a share of revenues by productivity



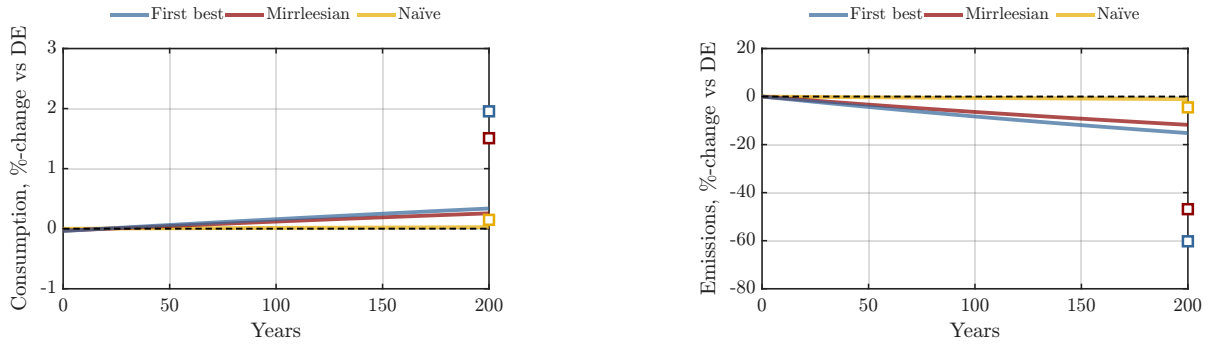
Note: The figure reports the overall carbon tax bill as a share of firm revenues under the full-information, Mirrleesian, and naïve planner allocations, by productivity type.

D Quantitative refinements

D.1 Transitional dynamics and alternative climate model

Baseline. To compute the transitional dynamics of the baseline model, we need to discipline two additional objects: the initial capital stock and the initial emissions stock, $\{K_0, G_0\}$. We therefore calibrate K_0 to the steady-state level of capital in an economy with emissions stock G_0, and set G_0 equal to 70 percent of the total emissions stock in the decentralized equilibrium, which we interpret as reflecting that 70 percent of the carbon budget has already been exhausted, in line with IPCC (2022). However, the precise initial condition does not materially affect our results.

Figure D.1: Transitional dynamics baseline model



(a) Consumption **(b)** Emissions

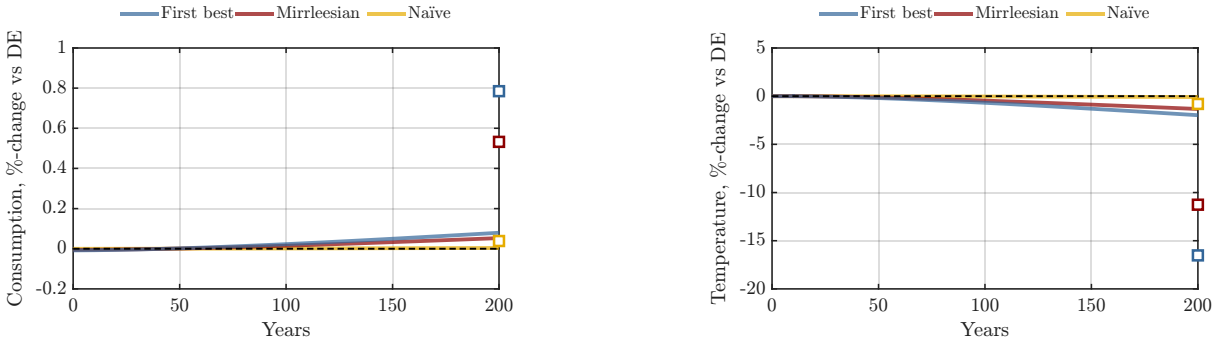
Note: The figure plots the first 200 years of the transition dynamics in the baseline model for consumption and emissions under the full-information, Mirrleesian, and naïve planner allocations, expressed as percentage deviations from the decentralized equilibrium. Square symbols denote the corresponding long-run steady states.

Figure D.1 reports the first 200 years of the transition dynamics under each planner, expressed as percentage deviations from the decentralized equilibrium. Two observations stand out. First, the transition is extremely slow because of the high persistence of atmospheric carbon, reflected in the low depreciation rate of emissions, consistent with the evidence. Second, consistent with the steady-state results in Table 4, the Mirrleesian planner closely tracks the first-best allocation throughout the transition, whereas the naïve planner captures only a small fraction of the potential gains from carbon taxation.

Alternative climate model. The additional parameters implied by the extended climate model are calibrated as follows. Consistent with the historical global average temperature, we set $\mathcal{T} = 14^{\circ}\text{C}$ and the temperature persistence parameter to $1 - \xi = 0.99^{10}$, in line

with the annual estimates of Bruns et al. (2020). We further choose $\Xi_G = 0.23$ to generate a 2°C increase in temperature over the first 100 years in the decentralized equilibrium, consistent with standard business-as-usual scenarios. We then recalibrate $\phi = 1.3 \cdot 10^{-3}$ to match a 2 percent decline in aggregate productivity following a 2°C increase in temperature relative to the historical mean, thereby preserving the baseline calibration. We set $\delta_G = 1 - (1 - 0.001)^{10}$, implying an almost unit-root process for emissions, as is common in the natural sciences. Finally, this specification requires three initial conditions, $\{K_0, G_0, T_0\}$. We set K_0 equal to the steady-state capital stock associated with $T_0 = 14$°C, which approximates the historical average temperature, and set G_0 equal to 70 percent of the decentralized-equilibrium emissions stock, as in the baseline model.

Figure D.2: Transitional dynamics alternative climate model



(a) Consumption **(b)** Temperature

Note: The figure plots the first 200 years of the transition dynamics in the augmented climate model for consumption and emissions under the full-information, Mirrleesian, and naïve planner allocations, expressed as percentage deviations from the decentralized equilibrium. Square symbols denote the corresponding long-run steady states.

Figure D.2 reports the first 200 years of the transition dynamics under each planner, expressed as percentage deviations from the decentralized equilibrium. As in the baseline model, two observations stand out. First, the transition is extremely slow because of the high persistence of atmospheric carbon and the inertia of global temperature, reflected in the low depreciation rate of emissions and the near-unit-root dynamics of temperature, consistent with the evidence. Second, consistent with the steady-state results in Table 4, the Mirrleesian planner closely tracks the first-best allocation throughout the transition, whereas the naïve planner captures only a small fraction of the potential gains from carbon taxation. Overall, the choice of climate model has little effect on the relative performance of the Mirrleesian planner. The climate block affects the allocation through the social cost of carbon, implying a proportional scaling of the first-best and Mirrleesian tax schedules and, therefore, little change in their relative performance.

D.2 Green and brown energy inputs

We calibrate the elasticity of substitution, ϵ, to 5, consistent with the recent estimates of Petersen (2025) and near the upper end of the empirical range reported by Papageorgiou et al. (2017) and Jo (2025), to reflect the long-run 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.

Table D.1: Green versus brown energy

	$\mathcal{C}$	$\mathcal{G}$	$\mathcal{E}_{\text{brown}}$	$\mathcal{E}_{\text{green}}$	$\mathcal{K}$	$\mathcal{X}$
Naive planner	$+0.15$	-5.36	$+0.15$	$+0.21$	$+0.17$	$+0.06$
Mirrleesian planner	$+1.36$	-52.68	-7.35	$+15.86$	$+1.63$	$+1.05$
Full-informed planner	$+2.34$	-66.78	-76.65	$+197.06$	$+2.12$	$+0.76$

Note: The table reports percentage deviations from the decentralized-equilibrium steady state in consumption, the emissions stock, brown energy, green energy, capital, and abatement under the full-information, Mirrleesian, and naïve planner allocations.

Table D.1 reports percentage deviations from the decentralized-equilibrium steady state in consumption, the emissions stock, brown energy, green energy, capital, and abatement under the full-information, Mirrleesian, and naïve planner allocations. Overall, allowing for substitution between brown and green energy has only a modest effect on our quantitative results. As in the baseline model, the Mirrleesian planner recovers a substantial share of the first-best gains, whereas the naïve planner captures almost none of them.

D.3 Alternative degrees of imperfect information

In this case, the wedge in the Mirrleesian planner's first-order conditions can be expressed as

$$\lambda^{MR}_{G,ij,t} = \Lambda_{G,t} \cdot \exp^{\eta \cdot z_j} \cdot \mathbb{E}[\exp^{(1-\eta)\cdot z_j} | a_i] \cdot \exp^{x_{ij,t}} \tag{A84}$$

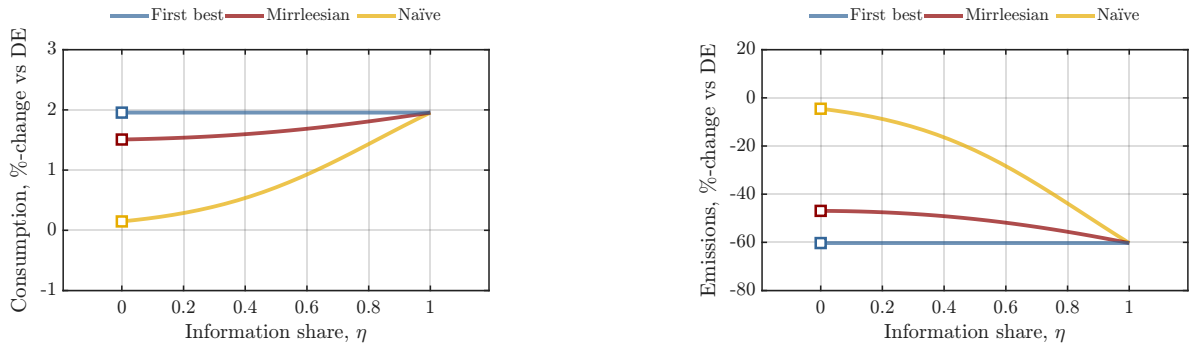
while the corresponding wedge for the naïve planner is

$$\lambda^{NV}_{G,ij,t} = \Lambda_{G,t} \cdot \exp^{\eta \cdot z_j} \cdot \exp^{(1-\eta)\cdot \underline{z}} \cdot \exp^{x_{ij,t}} . \tag{A85}$$

Clearly, as $\eta \to 1$, both expressions converge to the full-information case. The baseline specification corresponds to $\eta = 0$ under which the planner knows only the distribu-

tion of carbon intensities but has no firm-specific information about individual carbon-intensity types.

Figure D.3: Alternative degrees of imperfect information



(a) Consumption **(b)** Emissions

Note: The figure plots consumption and emissions under the full-information, Mirrleesian, and naïve planner allocations, expressed as percentage deviations from the decentralized equilibrium, for different values of the information parameter η. Square symbols denote the baseline steady states.

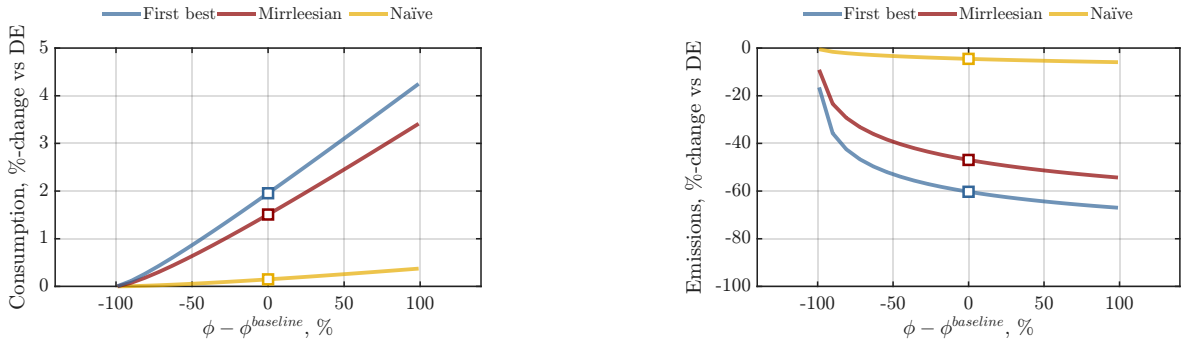
Figure D.3 plots consumption and emissions under the full-information, Mirrleesian, and naïve planner allocations, expressed as percentage deviations from the decentralized equilibrium, for different values of the information parameter η. The main message is that accounting for firms' disclosure incentives remains important regardless of the planner's information set. Across all values of η, the Mirrleesian planner recovers most of the first-best welfare gains, whereas the naïve planner performs substantially worse. For example, even when the planner observes 60 percent of firms' carbon-intensity information, ignoring misreporting incentives recovers only about half of the first-best welfare gains, while the Mirrleesian planner continues to recover most of them.

D.4 Climate parameters uncertainty

In this section, we perform comparative statics with respect to three climate-related parameters: the environmental damage parameter, ϕ, the inverse elasticity of abatement, θ, and the standard deviation of the carbon-intensity distribution, σ_z.

Figures D.4-D.6 report the results for parameter values 100 percent above and below the baseline calibration. Overall, the main message remains unchanged: the Mirrleesian planner consistently recovers a large share of the first-best gains and substantially more than the naïve planner.

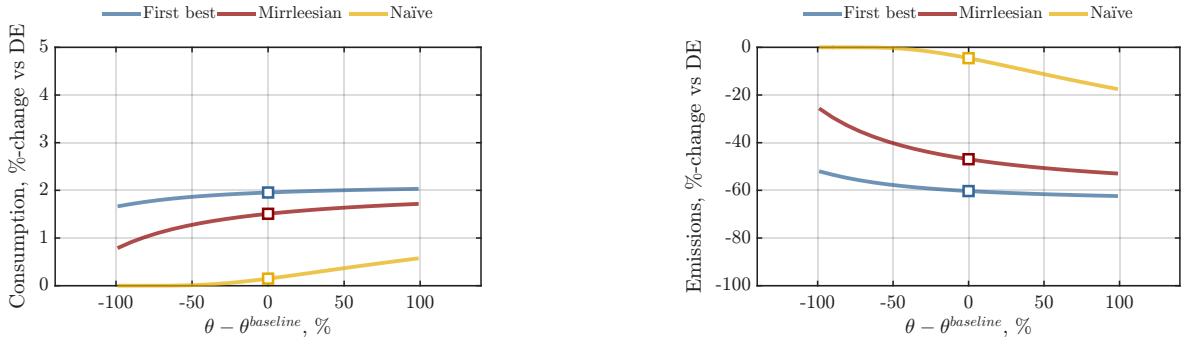
Figure D.4: Comparative static I: environmental externality size



(a) Consumption **(b)** Emissions

Note: The figure plots consumption and emissions under the full-information, Mirrleesian, and naïve planner allocations, expressed as percentage deviations from the decentralized equilibrium, for different values of the environmental externality ϕ. Square symbols denote the baseline steady states.

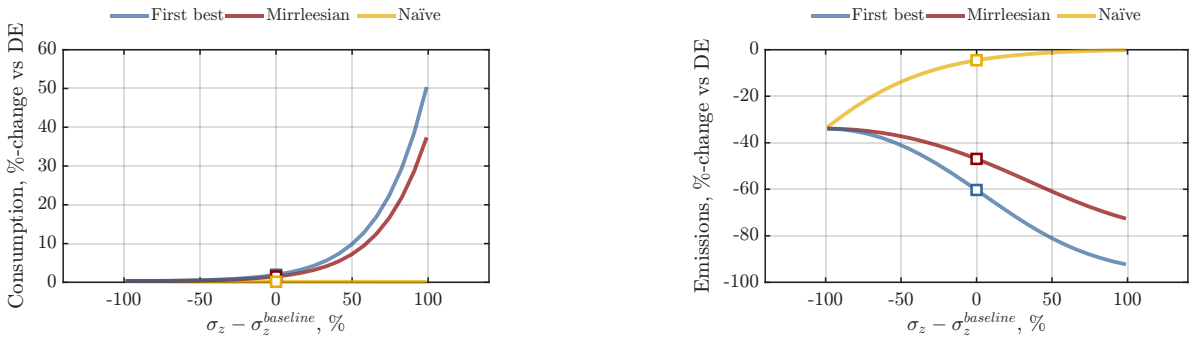
Figure D.5: Comparative static II: abatement elasticity size



(a) Consumption **(b)** Emissions

Note: The figure plots consumption and emissions under the full-information, Mirrleesian, and naïve planner allocations, expressed as percentage deviations from the decentralized equilibrium, for different values of the abatement elasticity θ. Square symbols denote the baseline steady states.

Figure D.6: Comparative static III: carbon intensity dispersion size



(a) Consumption

(b) Emissions

Note: The figure plots consumption and emissions under the full-information, Mirrleesian, and naïve planner allocations, expressed as percentage deviations from the decentralized equilibrium, for different values of the carbon intensity dispersion σ_z. Square symbols denote the baseline steady states.