Taxing Carbon, Not Competition: Optimizing Market and Environment in a World of Imperfections

Yağmur Menzilcioğlu*

October 30, 2026†

Abstract

It is well known that a pollution tax on firms with market power should be below marginal damage. How far below depends on markups and on how much damage a unit of output carries. I measure it for five US carbon-intensive sectors in a dynamic general equilibrium climate model with oligopoly. The optimal tax is 33 to 75 percent of the Pigouvian tax and varies little within a sector, so a sector-wide tax is nearly first-best. Ignoring market power costs 0.11 percent of consumption. Market power lowers these sectors' emissions by about 40 percent relative to perfect competition.

JEL Codes: D43, E62, H21, H23, L13, O44, Q43, Q54, Q58.

Keywords: Carbon taxes, Market power, Oligopoly, Climate change, Externalities, Optimal taxation, Dynamic integrated assessment modeling, Pigouvian tax, Energy and the macroeconomy, Environment and growth.

*Nicholas School of the Environment, Duke University, email: ym53@duke.edu.

†I am deeply grateful for the guidance and support of my advisor, Toshihiko Mukoyama. I also thank Arik Levinson, Mark Huggett, Jan Eeckhout, Per Krusell, John Hassler, Basile Grassi, Vincent Sterk, Pablo Ottonello, Mark Bils, Tiago Cavalcanti, Joseph Hanna, and Tsolmon Otgon for their valuable feedback and suggestions. This research was supported by a Duke SCALES Fellowship and, at an earlier stage, by Georgetown University's Summer Dissertation Fellowship. All errors are my own.

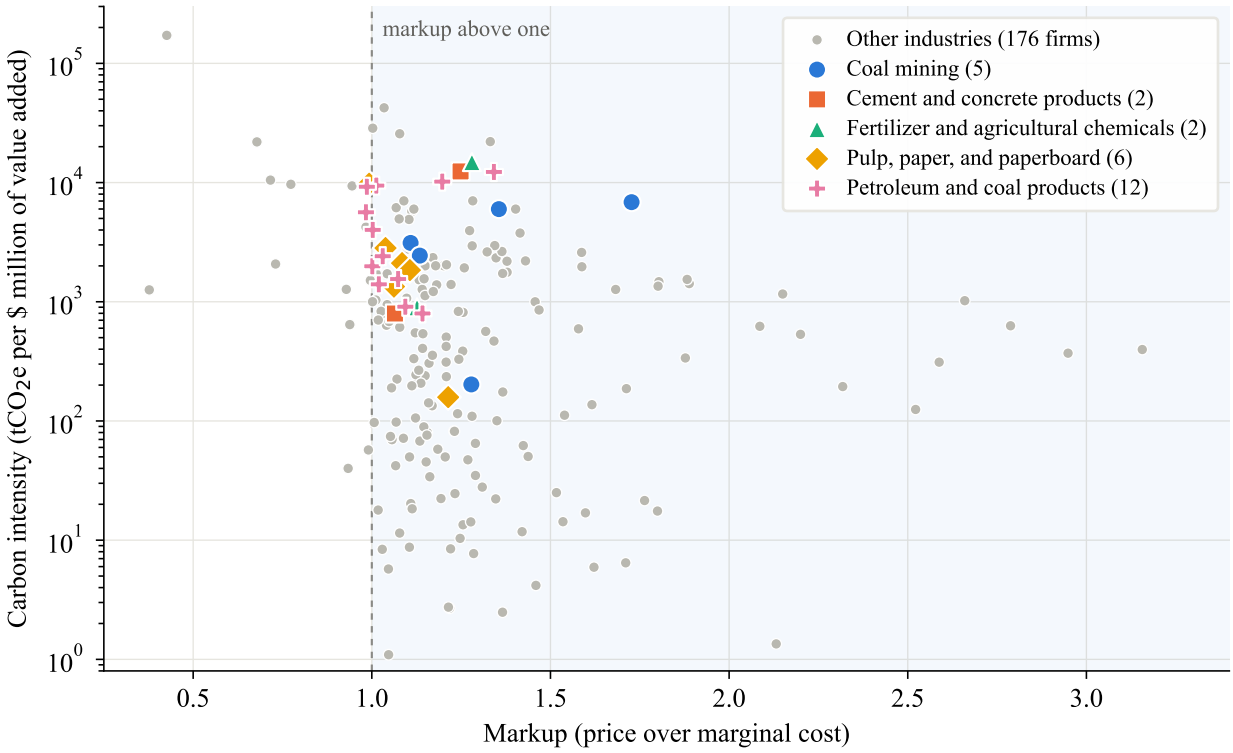
1 Introduction

When a production externality, such as pollution, is the only distortion, the Pigouvian tax, equal to marginal external damage, is the first-best policy. When polluting firms also have market power, it is not: these firms already produce too little, and a tax equal to marginal damage reduces their output by more than is socially optimal. The first-best policy then requires two instruments, a tax on pollution and a subsidy on output. In practice the environmental regulator does not control both instruments, because competition policy belongs to another authority. Mansur (2013) notes that in the United States the Environmental Protection Agency and the Department of Justice have different mandates. How different is the optimal carbon tax for an imperfectly competitive industry from the tax for a competitive one, does it have to differ across firms, and how much is lost if market power is ignored? The answers depend on market structure: on the elasticity of demand that firms face and on how concentrated the industry is, which in turn reflects how much firms differ in productivity.

Market power is a relevant concern for climate policy because many carbon-intensive industries are concentrated and their firms price above marginal cost. Figure 1 plots the carbon intensity of US publicly traded firms that report emissions to the Environmental Protection Agency's Greenhouse Gas Reporting Program (2024) against their markups in 2017. Carbon intensity is the firm's greenhouse gas emissions, in metric tons of CO_2 equivalent, per million dollars of value added, measured as revenue minus the cost of goods sold. Markups are estimated from S&P Global Market Intelligence (2023) data with the production function approach, following Traina (2018). Of the 203 firms, 92 percent have a markup above one, and so do 86 percent of the firms with above-median carbon intensity. The median markup is 1.17, against 1.09 for all Compustat firms in that year, and the sales-weighted average is 1.18 in both groups, so firms that emit have at least as much market power as other publicly traded firms. Markups are not higher at more carbon-intensive firms; the correlation between the two is weakly negative. For the design of a carbon tax, what matters is not this correlation but the level of markups among firms that emit: the correction to the Pigouvian tax that I derive in Section 2 is the firm's price divided by its elasticity of demand per unit of output, and it does not depend on the firm's carbon intensity. The sales-weighted average markup of carbon-intensive industries rose from about 1.15 in 2010 to about 1.25 in 2021 to 2023, and their economic profit rates, measured as in Eeckhout (2025), are positive throughout. Appendix A reports these series and describes the data.

That market power lowers the optimal pollution tax has been known since Buchanan (1969) and Barnett (1980), and that the efficient tax differs across firms with their market power since Lee (1975). Oates and Strassmann (1984) argued early that the adjustment matters little in practice. More recently, studies of individual industries have measured it and found that it can matter a great deal: Fowlie, Reguant, and Ryan (2016) for the US cement industry and Cardoso (2025) for US domestic aviation find that a carbon price equal to the social cost of carbon can lower welfare. These studies model one industry in partial equilibrium, with rich industry detail. The macroeconomic literature on optimal carbon taxation, in contrast, typically treats carbon-intensive

Figure 1: Markups and carbon intensity of US publicly traded firms that report greenhouse gas emissions, 2017.



Note: Each point is a firm. Firms in the five sectors of the model are highlighted; the number of firms is in parentheses. The shaded area marks markups above one. Sources: markups are calculated from S&P Global Market Intelligence (2023); carbon intensities from US Environmental Protection Agency (2024) and S&P Global Market Intelligence (2023).

industries as perfectly competitive, as in Nordhaus (1993), Golosov et al. (2014), and Barrage and Nordhaus (2024), or adds distortions other than market power in those industries: distortionary taxes in Barrage (2020), an innovation externality in Acemoglu et al. (2012), where the optimal policy combines a carbon tax with subsidies to clean research, factor misallocation in Kim (2025), and private information about emissions in Chiavari and Kohlhas (2026). An early exception is Hassler, Krusell, and Olovsson (2010), who derive the optimal tax on a monopoly supplier of oil in a stylized model and show that it can be zero. Annicchiarico, Correani, and Di Dio (2018) study how an emission cap changes markups and entry in a model with oligopolistic firms, without climate damages or an optimal tax. I bring the two together by introducing oligopoly in carbon-intensive sectors into a dynamic general equilibrium climate model. This gives up the industry detail and holds the number of firms fixed. In exchange, it treats several sectors in one framework under a common carbon price, computes climate damages and the path of the social cost of carbon consistently with the allocation each policy produces, and measures welfare in units of consumption across sectors and policies over a long horizon. The correction itself is static: in each period it is the firm's price divided by its elasticity of demand, and its ratio to the Pigouvian tax rises by at most

four percentage points over two centuries. The dynamic model supplies what a static one takes as given: the path of the social cost of carbon under each policy, and a welfare measure that counts damages over a long horizon.

The direction of the adjustment is not in question. What is not known is how large it is across carbon-intensive industries, whether it has to be tailored to individual firms, and how costly it is to get it wrong. This paper measures these for five US sectors: coal, cement, fertilizer, paper, and petroleum. To do so, I derive the optimal output tax for each firm of an oligopoly in the dynamic general equilibrium climate model of Golosov et al. (2014). It equals the Pigouvian tax minus an output subsidy equal to the firm's price divided by its elasticity of demand, the counterpart of the formula of Barnett (1980) for a monopolist. A firm with market power already restricts its output, so a smaller tax is needed to bring its output to the first-best level. With the social cost of carbon at \$50 per ton of CO_2 in 2017, the optimal tax for the median firm ranges from 33 percent of the Pigouvian tax in fertilizer to 75 percent in petroleum, or from \$17 to \$37 per ton. The correction is larger in sectors with higher markups. The ranking follows markups rather than carbon intensity, because in all five sectors marginal damage accounts for about nine tenths of the first-best price of output, a consequence of the five sectors standing in for all emissions. The correction varies little across firms within a sector, so a tax that is uniform within each sector achieves nearly all of the welfare gain of the firm-specific tax. Ignoring market power altogether and levying the Pigouvian tax costs 0.11 percent of consumption, about 18 billion dollars per year. The correction is sensitive to how and when markups are measured: with gross output rather than value added markups, the optimal tax is 77 to 94 percent of the Pigouvian tax, and with the markups of 2022 rather than 2017 it is lower in coal, cement, and fertilizer. These magnitudes depend on the five sectors standing in for all emissions. If instead each firm's tax reflected only the damage from its own emissions, as in studies of a single industry, marginal damage would be about a tenth of the first-best price rather than nine tenths. The optimal output tax would then be negative in every sector, and a tax equal to the social cost of carbon would lower welfare relative to no policy, as Fowlie, Reguant, and Ryan (2016) and Cardoso (2025) find for cement and aviation. The cost of ignoring market power is about the same in both cases, 0.11 to 0.12 percent of consumption (Appendix C.14).

The model extends the dynamic general equilibrium climate model of Golosov et al. (2014) by replacing its competitive energy sectors with carbon-intensive intermediate input sectors that are imperfectly competitive. In each sector, a fixed number of firms produce differentiated inputs and compete in quantities, as in Atkeson and Burstein (2008), so that a firm's markup rises with its market share. Firms produce with labor, differ in productivity, and emit carbon in proportion to their output. The carbon cycle and the damages from climate change are as in Golosov et al. (2014), which keeps the social cost of carbon a constant share of output.

The optimal output tax derived in this model is the Pigouvian tax minus an output subsidy equal to the firm's price divided by its elasticity of demand. Correcting two distortions in general requires two instruments, a tax on emissions and a subsidy on output. Here a single tax attains the first-best allocation because each producer has one decision margin: it chooses its output, and

its emissions follow from its output, so the two instruments act on the same margin and can be combined. I state the policy as a tax on output. It can equivalently be expressed as a firm-specific tax per ton of carbon, and I report results in those units. An emissions price combined with a rebate per unit of output, as in the output-based allocation of emission permits studied by Fowlie, Reguant, and Ryan (2016), has the same structure: the emissions price is the Pigouvian component and the rebate is the output subsidy. With a rebate that differs across firms it implements the optimal output tax, and with a single rebate in each sector it implements the sector-wide tax evaluated in Section 5. Gersbach and Requate (2004) show that for symmetric firms competing in quantities, an emission tax partly refunded in proportion to market shares attains the first-best allocation when marginal damage exceeds the market power distortion. The result relies on a fixed number of firms and on output being the only margin of abatement; with entry or with an abatement technology, a second instrument is needed.

I calibrate the model to five US sectors in 2017. In each sector, the sizes of firms reproduce the revenue shares of the 4, 8, 20, and 50 largest firms and the Herfindahl-Hirschman Index reported in the Economic Census, and the elasticity of substitution between firms is set so that the model's sales-weighted average markup equals the sales-weighted average of the markups estimated from Compustat for firms in the same sector. Because intermediate producers in the model use only labor, I convert the estimated gross output markups to a value added basis. Each sector's share of output and its emissions match the input-output accounts and the Greenhouse Gas Reporting Program, the five sectors' emissions are scaled to world emissions, and climate damages are calibrated to a social cost of carbon of \$50 per ton of CO_2 in 2017. Without policy, long-run warming reaches about 3.9°C by 2100, just below the laissez-faire path of Golosov et al. (2014). Market power already restrains emissions: with the same firms pricing at marginal cost, emissions would be about 60 percent higher and warming in 2100 about 1.0°C higher. Asker et al. (2024) estimate the same effect in the oil market, where the market power of OPEC lowered emissions by 67 gigatons of CO_2 between 1970 and 2021. Under the optimal policy, warming stays near 1.6°C. This path is cooler than in models in which abatement is costly: the five sectors account for about 1 percent of output but stand in for all emissions, so reducing their output lowers emissions at little cost in output.

Finally, I compare policies by their welfare loss relative to the first best. When producers are oligopolistic, no policy loses 24 percent of consumption. The Pigouvian tax, which ignores market power, loses 0.11 percent. This is the relevant mistake in practice: Mansur (2013) notes that the Environmental Protection Agency's planning model assumes a competitive electricity industry. A sector-wide tax, set as if the firms of a sector were identical, loses 0.001 percent and requires no information on individual firms. When producers are in fact perfectly competitive, the tax designed for oligopolists loses 0.16 percent, more than ignoring market power when it exists. When the five sectors stand in for all emissions, a carbon tax should therefore be adjusted for market power sector by sector, firm-specific taxes are not needed, and the welfare at stake in the adjustment is modest relative to the gain from pricing carbon at all; when each firm's tax reflects only the damage from its own emissions, the adjustment matters more, because the optimal tax becomes negative and a

tax equal to the social cost of carbon lowers welfare.

Beyond the papers discussed above, this paper relates to three sets of work. On the theory of pollution taxes under oligopoly, Levin (1985) shows that when firms differ in costs and in pollution per unit of output, a uniform tax per unit of output, and even a tax proportional to each firm's pollution, can increase total pollution for some shapes of demand, because the tax shifts output toward the more polluting firms although industry output falls; Simpson (1995) shows that a related reallocation can make the optimal uniform tax in a duopoly exceed marginal damage. Ebert (1991) derives the optimal emission tax for a symmetric Cournot oligopoly, and Long and Soubeyran (2005) show that when firms differ in costs the optimal emission tax differs across firms and favors the more efficient ones. Shaffer (1995) shows that a single tax on output attains the first-best allocation for identical oligopolists when the external cost depends on output. Katsoulacos and Xepapadeas (1995) and Lee (1999) show that when the number of firms responds to the tax, the optimal tax can be above or below marginal damage, and Requate (2006) surveys this literature. Among studies of individual industries, Ryan (2012) finds that environmental regulation raised the cost of entry in the US cement industry and lowered product market surplus, Fowlie (2009) and Mansur (2013) study how imperfect competition in electricity markets changes emissions leakage and the choice between taxes and tradable permits, Cullen and Reynolds (2023) how entry, exit, and startup costs shape investment in generation, and Preonas (2024) and Ganapati, Shapiro, and Walker (2020) estimate that carbon and energy costs are passed through incompletely where sellers have market power. In macroeconomic models with heterogeneous firms, Klenow, Pasten, and Ruane (2025) find that in Chile a carbon tax would reallocate inputs toward firms with higher markups and raise allocative efficiency, and Jo and Karydas (2024) find that the exit of firms that cannot easily substitute away from dirty energy lowers the carbon tax needed to reach carbon neutrality. In my model carbon intensity is common to the firms of a sector, which is the case in which Levin (1985) shows that a uniform tax reduces pollution unambiguously, and the number of firms is fixed, so these two channels are absent. The reallocation of output across firms emphasized by Levin (1985) and Simpson (1995) is small at the observed degree of concentration.

The rest of the paper is organized as follows. Section 2 presents the model, reviews the constant social cost of carbon formula of Golosov et al. (2014), and derives the optimal output tax for oligopolistic carbon-intensive producers. Section 3 describes the solution method and the calibration to five US sectors. Section 4 reports the optimal taxes and the paths of emissions, temperature, and output with and without policy. Section 5 evaluates the welfare cost of alternative policies, including a uniform Pigouvian tax and a sector-wide tax, under oligopoly and under perfect competition. Section 6 concludes.

2 Model

Consider a version of the multi-sector neoclassical growth model outlined in Golosov et al. (2014). Time is discrete and the time horizon is infinite. There is a representative household with

a utility function:

$$\sum_{t=0}^{\infty} \beta^t U(C_t),$$

The social planner's problem is to maximize the representative household's discounted lifetime utility subject to technology, feasibility, and carbon cycle constraints of the economy.

There are multiple production sectors in the economy: a final goods sector, and M intermediate goods sectors. The final goods sector produces the consumption and investment good Y_t, and each intermediate goods sector $m \in \{1, ..., M\}$ produces a composite input $X_{m,t}$ for use in the final goods sector.

The feasibility constraint in the final goods sector, denoted as sector 0, is:

$$C_t + K_{t+1} = Y_t + (1 - \delta)K_t, \tag{1}$$

where K_t is capital stock and $\delta \in (0, 1)$ is the depreciation rate. The following aggregate production function describes output in the final goods sector:

$$Y_t = F_{0,t}\left(K_t, N_{0,t}, \{X_{m,t}\}_{m=1}^{M}, S_t\right), \tag{2}$$

where K_t and $N_{0,t}$ are capital and labor used in sector 0 at time t (capital is used only in final goods production, so it carries no sector index), $X_{m,t}$ is the intermediate composite input of sector $m \in \{1, ..., M\}$ used in final-goods production at time t, and S_t is the climate variable which affects output. The production function $F_{0,t}$ depends on K_t, $N_{0,t}$, $\{X_{m,t}\}_{m=1}^{M}$, and S_t. This model views climate as sufficiently well represented by one variable, S_t, which is the amount of carbon in the atmosphere. In studies from natural sciences, for example, Schneider (1989), National Research Council (2010), and IPCC (2021), the current atmospheric carbon concentration is considered to be approximating the current climate well. Assume that S_t affects final goods production only. The amount of carbon in the atmosphere, S_t, is measured as atmospheric carbon concentration above pre-industrial times in billions of tons of carbon (GtC) units.

The intermediate input sectors each produce a composite input $X_{m,t}$ for use in the final goods sector, and there are a finite number of sectors M. These intermediate production factors could be carbon-intensive industries such as coal ($m = 1$), cement ($m = 2$), fertilizer ($m = 3$), paper ($m = 4$), and petroleum ($m = 5$). I assume that each intermediate input composite m is produced by J_m number imperfectly substitutable products. Each intermediate product $j = 1, \ldots, J_m$ in sector m has its own technology $F_{m,j,t}$ to produce the intermediate good $X_{m,j,t}$ at time t by employing variable labor input $N_{m,j,t}^{var}$ to produce output, and a fixed cost $\Omega_{m,j,t}$, paid in units of labor, incurred regardless of the level of output $X_{m,j,t}$ produced. I proceed to assume that production remains linear in labor input $N_{m,j,t}^{var}$:

$$X_{m,j,t} = F_{m,j,t}(N_{m,j,t}^{var}) \geq 0, \tag{3}$$

with total labor demanded by firm j in sector m is:

$$N_{m,j,t} = N^{var}_{m,j,t} + \Omega_{m,j,t}. \tag{4}$$

In each period, production factors are mobile across sectors. The final goods sector is perfectly competitive, whereas studying different intermediate goods market structures is this paper's objective.

I assume that intermediate production factors emit carbon to the atmosphere. Let $E_{m,j,t}$ be the amount of carbon emitted by firm j in sector m at time t. I generalize Golosov et al.'s (2014) modeling assumption by letting firm-level emissions $E_{m,j,t}$ be a general function of the intermediate input $X_{m,j,t}$ produced by firm j in sector m at time t, instead of assuming that the intermediate inputs are in one-to-one correspondence with its carbon emissions. The firm-specific emission function is generalized to allow for a nonlinearities and firm-specific carbon intensities:

$$E_{m,j,t} = g_{m,j}(X_{m,j,t}), \tag{5}$$

where $g_{m,j}(X_{m,j,t})$ is a general function that captures the relationship between the intermediate input $X_{m,j,t}$ and the emissions $E_{m,j,t}$ for firm j in sector m at time t.

At the sector level, total emissions $E_{m,t}$ are defined as the sum of emissions from all firms in sector m at time t. Thus, for each sector m, emissions are given by:

$$\begin{aligned} E_{m,t} &= \sum_{j=1}^{J_m} E_{m,j,t}, \\ &= \sum_{j=1}^{J_m} g_{m,j}(X_{m,j,t}). \end{aligned} \tag{6}$$

Finally, aggregate emissions at the economy-wide level, denoted E_t, are defined as the sum of sector-level emissions across all sectors $m = 1, \ldots, M$ at time t. Thus, aggregate emissions are given by:

$$\begin{aligned} E_t &= \sum_{m=1}^{M} E_{m,t}, \\ &= \sum_{m=1}^{M} \sum_{j=1}^{J_m} g_{m,j}(X_{m,j,t}). \end{aligned} \tag{7}$$

The climate variable S_t is affected by the aggregate carbon emissions, E_t, and the carbon cycle. To describe the evolution of the climate, let $\tilde{S}_t$ be a function that maps a history of human-made (anthropogenic) emissions into the current atmospheric concentration level, S_t. History is defined

to start at the time of industrialization, denoted as T periods before period 0:

$$S_t = \tilde{S}_t \left(E_{-T}, E_{-T+1}, ..., E_t \right). \tag{8}$$

2.1 Additional Assumptions

In this section, I summarize the three key assumptions outlined in Golosov et al. (2014) that simplify the planning problem analyzed in the next section. Then, I introduce two additional assumptions that simplify the carbon tax formulation in the presence of imperfect competition in the intermediate sectors. Golosov et al. (2014) discuss the plausibility of their three assumptions. First, households have a logarithmic utility function:

Assumption 1 $U(C) = \log(C)$.

Second, they assume that production damages are multiplicative:

Assumption 2 *The production technology can be represented as:*

$$F_{0,t}\left(K_t, N_{0,t}, \{X_{m,t}\}_{m=1}^{M}, S_t\right) = [1 - D(S_t)]\tilde{F}_{0,t}\left(K_t, N_{0,t}, \{X_{m,t}\}_{m=1}^{M}\right), \tag{9}$$

where $1 - D(S_t) = \exp[-\gamma_t(S_t - \bar{S})]$ *with* $\bar{S}$ *being the pre-industrial atmospheric* CO_2 *concentration, and* γ_t *is the marginal damage measured as a share of GDP per marginal unit of carbon in the atmosphere assumed to be constant, and before damages production function* $\tilde{F}_{0,t}$ *is strictly increasing and strictly concave in all its inputs* $K_t, N_{0,t}$, *and* $\{X_{m,t}\}_{m=1}^{M}$, *and it exhibits constant returns to scale.*

Third, they assume a simplified carbon cycle:

Assumption 3 *The function* $\tilde{S}_t$ *is linear with the following depreciation structure:*

$$S_t - \bar{S} = \sum_{s=0}^{t+T} (1 - d_s) E_{t-s},$$

where $d_s \in [0, 1]$ *for all* s *and* T *denotes the number of periods since industrialization and* E_{t-s} *follows equations (6) and (7).*

The fraction $1 - d_s$ represents the amount of carbon left in the atmosphere s periods in the future. A three-parameter family determines the depreciation structure, where (i) ψ_L is the share of carbon emitted into the atmosphere that stays in it forever; (ii) $1 - \psi_0$ is the share of the remaining emissions exiting the atmosphere immediately, and (iii) ψ is the geometric decay rate of the remaining share of emissions. This parameterization corresponds to a linear system with the depreciation rate at horizon s given by:

$$1 - d_s = \psi_L + (1 - \psi_L)\psi_0(1 - \psi)^s. \tag{10}$$

Golosov et al. (2014) show that this depreciation structure is consistent with the existence of two "virtual carbon stocks" S_1 (the part that remains in the atmosphere forever) and S_2 (the part that depreciates at rate ψ), with $S_{1,t} = S_{1,t-1} + \psi_L E_t$ and $S_{2,t} = (1-\psi)S_{2,t-1} + \psi_0(1-\psi_L)E_t$, and $S_t = S_{1,t} + S_{2,t}$. Golosov et al. (2014) use these three assumptions to obtain a closed-form solution to the optimal aggregate social cost of carbon to output ratio. This expression simplifies the computation of the model. I add two assumptions on the intermediate goods sectors.

Assumption 4 *The composite of sector m is a constant elasticity of substitution (CES) aggregate of the J_m intermediate goods in the sector:*

$$X_{m,t} = \left[\sum_{j=1}^{J_m} \kappa_{m,j} X_{m,j,t}^{\frac{\eta_m - 1}{\eta_m}} \right]^{\frac{\eta_m}{\eta_m - 1}}, \tag{11}$$

where $\eta_m > 1$ is the elasticity of substitution and $\kappa_{m,j} > 0$ are share parameters with $\sum_{j=1}^{J_m} \kappa_{m,j} = 1$.

Assumption 5 *In each sector m, the J_m producers compete in quantities. Each producer chooses its output taking as given the outputs of the other producers in its sector, final output Y_t, and the wage w_t, and it internalizes its effect on the sector composite $X_{m,t}$ and the sector price $p_{m,t}$.*

Assumption 5 follows Atkeson and Burstein (2008).

2.2 The Planning Problem

The social planner's problem is to maximize the representative household's discounted lifetime utility subject to technology, feasibility, and carbon cycle constraints of the economy. The social planner's problem is:

$$\max_{\left\{C_t, K_{t+1}, N_{0,t}, \left\{X_{m,t}, \{N_{m,j,t}, E_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}, E_t, S_t, Y_t\right\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U(C_t)$$

subject to (1), (2), (3), (4) (5), (6), (7), (8), (11), and $N_t = N_0 + \sum_{m=1}^{M} \sum_{j=1}^{J_m} N_{m,j,t}$, where N_t is the inelastic and exogenous labor supply at period t (normalized to unity). Importantly, the social planner accounts for the climate externality in the planning problem.

2.3 The Social Cost of Carbon

Following Hassler, Krusell, and Smith (2016), I define the social cost of carbon (SCC) as the marginal externality damage of carbon emissions, keeping constant behavior in the given allocation. If one further assumes that the saving rate is constant (which does not tend to vary so much over time in the data), then the SCC expression simplifies significantly.

Assumption 6 *The planner's solution yields a constant saving rate $s \in (0,1)$ over time.*

Proposition 1 of Golosov et al. (2014) shows under Assumptions 1-3 and 6, the SCC simplifies be a constant fraction of aggregate output:

$$\frac{\text{SCC}_{m,j,t}}{Y_t} = \bar{\gamma}_t \left(\frac{\psi_L}{1-\beta} + \frac{(1-\psi_L)\psi_0}{1-(1-\psi)\beta} \right). \tag{12}$$

I provide the full derivation of equation (12) in Appendix B.1.

The optimal SCC (OSCC), denoted by $\Lambda_{m,j,t}$, is the marginal external damage of a unit of carbon emissions from firm j in sector m evaluated at the optimal allocation and is equal to the marginal external damage from a unit of carbon emissions. Let $\hat{\Lambda}_{m,j,t}$ denote the optimal SCC to output ratio, or $\hat{\Lambda}_{m,j,t} \equiv \Lambda_{m,j,t}/Y_t$. Finally, since aggregate emissions are the sum of firm-level emissions, $\Lambda_{m,j,t}$ is equal across all firms and sectors.

2.4 Equilibrium

The government levies a per-unit tax $\tau_{m,j,t}$ on the output of producer j in sector m and rebates the proceeds to the consumer as a lump-sum transfer T_t. A negative tax is a subsidy, financed by a negative transfer. All trade takes place at time 0. Let q_t be the time-0 price of the final good delivered at time t, and let the rental rate r_t, the wage w_t, and the intermediate goods prices $p_{m,j,t}$ be expressed in units of the time-t final good.

Definition 1 *Given taxes* $\{\tau_{m,j,t}\}$ *and initial capital* K_0*, an oligopoly equilibrium is an allocation* $\left\{C_t, K_{t+1}, Y_t, N_{0,t}, \left\{X_{m,t}, \{X_{m,j,t}, N_{m,j,t}, N^{var}_{m,j,t}, E_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}, E_t, S_t\right\}_{t=0}^{\infty}$*, prices* $\left\{q_t, r_t, w_t, \left\{\{p_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}\right\}_{t=0}^{\infty}$*, transfers* $\{T_t\}_{t=0}^{\infty}$*, and profits* Π *such that:*

(i) Given prices, transfers, and profits, $\{C_t, K_{t+1}\}_{t=0}^{\infty}$ *solves the consumer's problem:*

$$\max_{\{C_t, K_{t+1}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t U(C_t) \quad \textit{subject to} \quad \sum_{t=0}^{\infty} q_t(C_t + K_{t+1}) = \sum_{t=0}^{\infty} q_t \left[(1 + r_t - \delta)K_t + w_t N_t + T_t\right] + \Pi.$$

(ii) Given prices and the climate variable $\{S_t\}$*,* $\{Y_t, K_t, N_{0,t}, \{X_{m,j,t}\}\}_{t=0}^{\infty}$ *solves the final goods producer's problem:*

$$\Pi_0 \equiv \max \sum_{t=0}^{\infty} q_t \left[Y_t - r_t K_t - w_t N_{0,t} - \sum_{m=1}^{M} \sum_{j=1}^{J_m} p_{m,j,t} X_{m,j,t} \right]$$

subject to (2) and (11).

(iii) For every producer j *in sector* m*, given the tax, the wage, final output, and the outputs of the other producers in its sector,* $\{X_{m,j,t}, N_{m,j,t}, N^{var}_{m,j,t}\}_{t=0}^{\infty}$ *solves:*

$$\Pi_{m,j} \equiv \max \sum_{t=0}^{\infty} q_t \left[(p_{m,j,t} - \tau_{m,j,t})X_{m,j,t} - w_t N_{m,j,t}\right]$$

subject to (3), (4), and the final goods producer's inverse demand for its good, $p_{m,j,t} = p_{m,j,t}(X_{m,j,t})$*, implied by (ii).*

(iv) Profits are $\Pi = \Pi_0 + \sum_{m=1}^{M}\sum_{j=1}^{J_m}\Pi_{m,j}$*, and the government budget is balanced:* $\sum_{t=0}^{\infty} q_t T_t = \sum_{t=0}^{\infty} q_t \sum_{m=1}^{M}\sum_{j=1}^{J_m}\tau_{m,j,t}X_{m,j,t}$*.*

(v) The final goods and labor markets clear: $C_t + K_{t+1} = Y_t + (1-\delta)K_t$ *and* $N_t = N_{0,t} + \sum_{m=1}^{M}\sum_{j=1}^{J_m} N_{m,j,t}$*; the capital market clears because the capital stock* K_t *of (i) is the capital input of (ii).*

(vi) Emissions and the climate variable satisfy (5), (6), (7), and (8).

In this equilibrium no agent internalizes the effect of emissions on the climate: the final goods producer takes S_t as given in (ii), and intermediate producers do not account for their emissions in (iii). The final goods producer's first-order condition for good j in sector m gives the inverse demand in (iii):

$$\frac{p_{m,j,t}}{p_{m,t}} = \kappa_{m,j} X_{m,j,t}^{-\frac{1}{\eta_m}} X_{m,t}^{\frac{1}{\eta_m}}, \tag{13}$$

where $p_{m,t} \equiv \partial Y_t / \partial X_{m,t}$ is the price of the sector m composite, which satisfies $p_{m,t}X_{m,t} = \sum_{j=1}^{J_m} p_{m,j,t}X_{m,j,t}$ and

$$p_{m,t} = \left[\sum_{j=1}^{J_m} \kappa_{m,j}^{\eta_m} p_{m,j,t}^{1-\eta_m}\right]^{\frac{1}{1-\eta_m}}.$$

Under Assumption 5, the producer's first-order condition in (iii) in the absence of taxes is a markup over marginal cost:

$$p_{m,j,t} = \frac{\sigma_m(s_{m,j,t})}{\sigma_m(s_{m,j,t}) - 1} \frac{w_t}{F'_{m,j,t}(N^{var}_{m,j,t})},$$

where $w_t / F'_{m,j,t}(N^{var}_{m,j,t})$ is the marginal cost of producer j and $\sigma_m(s_{m,j,t})$ is the price elasticity of demand for its good, a function of its sales share $s_{m,j,t} = p_{m,j,t}X_{m,j,t} / \sum_{i=1}^{J_m} p_{m,i,t}X_{m,i,t}$ as in Atkeson and Burstein (2008):

$$\sigma_m(s_{m,j,t}) = \left[\frac{1}{\eta_m}(1 - s_{m,j,t}) + \frac{X_{m,t}/p_{m,t}}{|\partial X_{m,t}/\partial p_{m,t}|} s_{m,j,t}\right]^{-1}. \tag{14}$$

The second term in brackets contains the inverse of the price elasticity of demand for the sector m composite. Appendix B derives equation (14).

Golosov et al. (2014) assume that all production sectors are perfectly competitive. When the climate externality is the only distortion, the optimal policy is to tax carbon emissions at the marginal externality damage at the optimal output level, i.e., the OSCC, Λ_t, known as the Pigouvian tax.

However, as argued in the introduction, imperfect competition is a more realistic market structure assumption for carbon-intensive production sectors. Under imperfectly competitive carbon-intensive intermediate production sectors, there are two distortions in each producer's decision:

a climate externality and a market distortion. In general, correcting two distortions requires two policy instruments, a tax on emissions and a subsidy on output, which is known as the Tinbergen rule, named after Tinbergen (1952). In this model, each producer makes a single choice, its output, and its emissions are a function of that output (equation (5)). The two instruments therefore act on the same margin and can be combined into a single firm-specific output tax. This fails if producers have a second margin, for example if they can reduce emissions per unit of output or enter and exit the market.

I state the optimal policy as a tax on output rather than on emissions.[1] Because each producer's emissions are a function of its output, a firm-specific output tax can equivalently be expressed as a firm-specific tax per ton of carbon, by dividing it by the producer's marginal emissions. I report the optimal tax in those units in Section 4. With emissions per unit of output fixed, a tax on output and a tax on carbon are therefore the same instrument; the policies compared in Section 5 differ only in whether the rate per ton is uniform, sector-specific, or firm-specific. The output tax is the more general statement because its second component, the output subsidy, does not depend on emissions: a producer with market power and no emissions would receive a subsidy, which a tax on emissions cannot deliver. What a carbon tax cannot do, if its rate per ton must be the same for all producers, is vary with each producer's markup. I evaluate such policies in Section 5. That the efficient pollution tax differs across producers with their market power, and is lower for producers facing less elastic demand, was first shown by Lee (1975) in a partial equilibrium setting.

2.5 Optimal Output Tax

I use the following terminology. The first-best allocation is the solution to the planning problem in Section 2.2. A policy is first-best if the equilibrium it induces coincides with that allocation. The Pigouvian tax, equal to the marginal external damage, is first-best when the climate externality is the only distortion, but not when producers have market power or there is any other distortion. The optimal output tax derived below is first-best when producers have market power. By a constrained policy tool I mean that the regulator of the climate externality takes as given a preexisting distortion, market power, that is not and will not be corrected by competition policy, so its policy tool must account for both distortions.

Unless environmental and competition regulators operate harmoniously, a regulator is restricted with a constrained policy tool to improve aggregate welfare. As Buchanan (1969) pointed out, a polluting firm with market power is subject to two distortions that move its output in opposite directions: it does not bear the cost of its pollution, which makes its output too high, and it restricts output to raise its price, which makes its output too low. A Pigouvian tax corrects the first distortion and worsens the second, so its net effect on welfare depends on their relative size and

[1] The allocation is the same whether the tax is levied on the intermediate producer's output or on the final goods producer's purchases of the input, since in either case the producer's revenue is the same function of its output. This is the standard result that the incidence of a tax does not depend on which side of the market pays it (Gruber (2011), Fullerton and Metcalf (2002)), which Weyl and Fabinger (2013) show also holds under monopoly and symmetric imperfect competition.

can be negative under monopoly. In a partial equilibrium calculation, Oates and Strassmann (1984) estimate that the gain from reduced pollution is about ten times the loss from reduced output. I quantify the two effects in general equilibrium in Section 5. The optimal output tax will combine the policies that address each distortion individually: the Pigouvian tax and the output subsidy. Proposition 1 states this result.

The optimal output tax is not the solution to a separate maximization problem. It is the tax under which the first-best allocation is an equilibrium allocation, so I construct it from the solution to the planning problem. Let a superscript $*$ denote the solution to the planning problem in Section 2.2. Define the prices that support this allocation as the marginal products of the final goods technology evaluated at it:

$$p^*_{m,j,t} \equiv \frac{\partial Y_t}{\partial X_{m,t}} \frac{\partial X_{m,t}}{\partial X_{m,j,t}}, \quad w^*_t \equiv \frac{\partial Y_t}{\partial N_{0,t}}, \quad r^*_t \equiv \frac{\partial Y_t}{\partial K_t}, \quad \text{and} \quad q^*_t \equiv \beta^t \frac{U'(C^*_t)}{U'(C^*_0)}, \tag{15}$$

where all derivatives are evaluated at the planner's allocation. Let $s^*_{m,j,t} = p^*_{m,j,t} X^*_{m,j,t} / \sum_{i=1}^{J_m} p^*_{m,i,t} X^*_{m,i,t}$ be producer j's sales share at these prices and quantities, and let $\sigma^*_{m,j,t} \equiv \sigma_m(s^*_{m,j,t})$ be the price elasticity of demand in equation (14) evaluated at the planner's allocation. The optimal output tax on intermediate producer j in sector m is:

$$\tau^*_{m,j,t} \equiv g'_{m,j}(X^*_{m,j,t}) \Lambda_{m,j,t} - \frac{p^*_{m,j,t}}{\sigma^*_{m,j,t}}, \tag{16}$$

where $g'_{m,j}(X^*_{m,j,t})$ is the producer's marginal emissions, $dE_{m,j,t}/dX_{m,j,t}$, and $\Lambda_{m,j,t}$ is the optimal SCC, which is by definition evaluated at the planner's allocation. Every term on the right-hand side is a function of the planner's allocation alone. The tax is therefore a sequence of numbers, not a schedule that varies with the producer's own choices, and in condition (iii) of Definition 1 the producer takes it as given. The first term is the Pigouvian tax and the second is the output subsidy. Since the subsidy is positive, the optimal output tax is always smaller than the Pigouvian tax and can even be negative in the presence of market power.

Proposition 1 *Let $\{\tau^*_{m,j,t}\}$ be given by equation (16), and suppose that the tax proceeds and firm profits are rebated lump-sum to the representative consumer, that every producer's output is positive in the solution to the social planner's problem, and that each producer's profit is concave in its own output. (i) If taxes are $\{\tau^*_{m,j,t}\}$, then the solution to the social planner's problem, the prices in equation (15), and the transfers and profits they imply form an oligopoly equilibrium. (ii) Conversely, if the solution to the social planner's problem is an oligopoly equilibrium allocation under taxes $\{\tau_{m,j,t}\}$, then $\tau_{m,j,t} = \tau^*_{m,j,t}$ for all m, j, and t.*

Proof.

Under Assumptions 1, 2, 3, and 6, the planner's first-order condition with respect to the variable

labor input of producer j in sector m is:

$$\frac{dX_{m,j,t}}{dN^{var}_{m,j,t}}\left[\frac{\partial Y_t}{\partial X_{m,t}}\frac{\partial X_{m,t}}{\partial X_{m,j,t}} - \frac{dE_{m,j,t}}{dX_{m,j,t}}\Lambda_{m,j,t}\right] = \frac{\partial Y_t}{\partial N_{0,t}}. \tag{17}$$

By the definitions in equation (15), at the planner's allocation this condition is $F'_{m,j,t}\left[p^*_{m,j,t} - g'_{m,j}(X^*_{m,j,t})\Lambda_{m,j,t}\right] = w^*_t$.

Under Assumptions 4 and 5, in an equilibrium with taxes $\{\tau_{m,j,t}\}$ the oligopolistic producer's first-order condition with respect to its variable labor input, given the outputs of the other producers in its sector, is:

$$\frac{dX_{m,j,t}}{dN^{var}_{m,j,t}}\left[p_{m,j,t}\frac{\sigma_m(s_{m,j,t})-1}{\sigma_m(s_{m,j,t})} - \tau_{m,j,t}\right] = w_t. \tag{18}$$

(i) Evaluate equation (18) at the planner's allocation, the prices in equation (15), and $\tau_{m,j,t} = \tau^*_{m,j,t}$. The term in brackets becomes $p^*_{m,j,t} - p^*_{m,j,t}/\sigma^*_{m,j,t} - \tau^*_{m,j,t} = p^*_{m,j,t} - g'_{m,j}(X^*_{m,j,t})\Lambda_{m,j,t}$, so the producer's first-order condition coincides with the planner's condition (17) and holds. Since the producer's profit is concave in its own output, $X^*_{m,j,t}$ is producer j's best response to the other producers' outputs $\{X^*_{m,i,t}\}_{i\neq j}$. The remaining equilibrium conditions also hold. The final goods producer's first-order conditions hold by the definition of the prices in equation (15). The consumer's Euler equation at prices q^*_t and r^*_t is the planner's Euler equation, and the consumer's budget constraint holds because profits and tax proceeds are rebated and the planner's allocation is feasible. Markets clear because the planner's allocation satisfies the resource constraints.

(ii) Suppose the planner's allocation is an equilibrium allocation under taxes $\{\tau_{m,j,t}\}$. The final goods producer's first-order conditions imply that the equilibrium prices are the marginal products in equation (15). Since every producer's output is positive, equation (18) holds at these prices and quantities. Subtracting the planner's condition (17) gives $\tau_{m,j,t} = g'_{m,j}(X^*_{m,j,t})\Lambda_{m,j,t} - p^*_{m,j,t}/\sigma^*_{m,j,t} = \tau^*_{m,j,t}$. □

Lemma 1 in Appendix B.2 shows that the two conditions of Proposition 1 hold under the functional forms of Section 3.

Two special cases show the role of each term. When producers take prices as given, the second term is absent, and the optimal tax is the Pigouvian tax $g'_{m,j}(X^*_{m,j,t})\Lambda_{m,j,t}$, as in Golosov et al. (2014). When there is no carbon externality, the first term is absent, and the optimal policy is an output subsidy equal to $p^*_{m,j,t}/\sigma^*_{m,j,t}$, the producer's price divided by its elasticity of demand at the planner's allocation.

The result that a single firm-specific tax attains the first best relies on each producer having one decision margin. If producers could also reduce emissions per unit of output, an output tax would give them no incentive to do so, and two instruments would be needed: a tax on emissions and a subsidy on output (Barnett 1980; Baumol and Oates 1988), or on environmental innovation (Katsoulacos and Xepapadeas 1996; Acemoglu et al. 2012). Similarly, with entry and exit the tax also affects the number of producers: the optimal tax may then be below or above the marginal

damage (Lee 1999), and adding a second instrument such as a license fee raises welfare relative to an emission tax alone (Katsoulacos and Xepapadeas 1995). Adding capital to the intermediate technology would not change the result, as long as capital is rented each period: a producer's marginal cost would then be its unit cost of labor and capital, and equation (16) involves only its price, its elasticity of demand, and marginal damage. The saving rate would no longer be exactly constant, because the return to capital in the intermediate sectors is net of marginal damage; it would lie between $\alpha\beta$ and $\beta(\alpha + \sum_m a_m \nu_m)$, where a_m is the capital share in sector m. Since the five sectors account for about 1 percent of output, the social cost of carbon would remain within about 1 percent of a constant share of output (Appendix B.6). The form of competition matters less: under price competition with differentiated products each producer still has one margin, and equation (16) holds with the corresponding elasticity of demand, which is larger than under quantity competition, so the optimal tax is closer to the Pigouvian tax. At the calibrated concentration the difference is small: the optimal tax for the median firm is within half a percentage point of the Pigouvian-tax share under quantity competition in every sector (Appendix C.10). Appendix B.7 discusses three further cases: product markets narrower than the sector, market power of the buyers of carbon-intensive goods, and existing regulation.

3 Quantitative Analysis

I completely characterize the generalized multi-sector model used for the quantitative analysis in Golosov et al. (2014) described above. Assume that Assumptions 1, 2, 3, 4, 5, and 6 hold. I present quantitative model results for a baseline calibration. I explain the calibration process of the new components of my extension of the Golosov et al.'s (2014) model to match the structure of major carbon-intensive intermediate production sectors in the US in 2017.

Assume that there are M intermediate production sectors each with differing degrees of carbon intensity. Further, suppose that within each sector there are J_m firms that produce imperfectly substitutable types of the intermediate production factor, which is a reasonable assumption given these intermediate sectors face varying qualities of goods and high transportation costs. Within sector m, type j firm produces its output using the following constant returns to scale technology:

$$X_{m,j,t} = \chi_{m,j,t} N^{var}_{m,j,t}, \tag{19}$$

where $\chi_{m,j,t}$ is the exogenous labor productivity of the producer j in sector m that grows at a constant rate $g_{\chi_{m,j}}$ and $N^{var}_{m,j,t}$ is the variable labor input used in j's production. In addition to the variable labor input, each producer j in sector m faces a fixed exogenous production cost, denoted in units of labor, $\Omega_{m,j,t}$, that it has to pay in order to produce, regardless the level of output it produces. The fixed costs represent the overhead that must be paid regardless of the output level, e.g., maintenance of physical plants, regulatory compliance, and upfront capital recovery. In oligopolistic sectors, these costs can explain limited firm entry. Thus, the total labor demand of

producer j in sector m is given by:

$$N_{m,j,t} = N^{var}_{m,j,t} + \Omega_{m,j,t}. \tag{20}$$

Deviating further from Golosov et al. (2014), I assume that carbon emissions from each intermediate input production are a general, and potentially nonlinear, function of total production of the intermediate input $X_{m,j,t}$ and the carbon intensity of the production process $\theta_{m,j,t}$. The carbon emissions from the production of intermediate input j in sector m are given by:

$$E_{m,j,t} = \theta_{m,j,t} X^{\zeta_{m,j}}_{m,j,t}, \tag{21}$$

where $\zeta_{m,j}$ is the degree of nonlinearity in the carbon emissions function. This functional generalization allows for the possibility that the carbon emissions from production are not linear in output, which is a reasonable assumption given that the carbon emissions from production can be affected by the technology used and the scale of production. Moreover, it harbors the Golosov et al. (2014) model as a special case, where both $\theta_{m,j,t}$ and $\zeta_{m,j}$ are equal to unity for all m and j. Additionally, if $\theta_{m,j,t} = 0$, then the production process is a zero carbon emissions process, which is the case for renewable energy sources. Finally, I assume that the carbon intensity of the production process $\theta_{m,j,t}$ is time-varying, to reflect the fact that the carbon intensity of production processes can change over time due to technological progress or changes in the energy mix used in production.

The sector composite $X_{m,t}$ that enters the final goods production is a CES aggregate of the different types of products in sector m as outlined in Assumption 4, equation (11). The final goods production function is of Cobb-Douglas specification:

$$Y_t = \exp\{-\gamma_t(S_t - \bar{S})\} A_{0,t} K^{\alpha}_t N^{1-\alpha-\sum_{m=1}^{M}\nu_m}_{0,t} \prod_{m=1}^{M} X^{\nu_m}_{m,t},$$

where $A_{0,t}$ is the exogenous total factor productivity (TFP) in the final goods sector that grows at a constant rate g_0. Labor employed by all producers has to sum up to the exogenous labor supply level N_t. Finally, I assume that there is full capital depreciation, given that each period is a decade.

The Cobb-Douglas production function and complete capital depreciation assumptions deliver a constant saving rate under optimality conditions, which was the assumption behind the OSCC expression derived in equation (12). Next, I characterize the solutions to the planning problem and three types of market equilibria.

With emissions possibly nonlinear in output and heterogeneous firms within each sector, neither the optimal allocation nor the equilibria have a closed form, unlike in Golosov et al. (2014). I solve the planning problem, the competitive equilibrium, and the oligopoly equilibrium with the same fixed-point iteration on the allocation of labor: given a guess for the labor used by the final goods sector and by each intermediate producer, the producers' first-order conditions give their outputs and an updated allocation of labor, and I iterate until the guessed and updated allocations coincide. Appendix B.8 states the first-order conditions of the three problems, the algorithm, and its starting

Table 1: Calibrated parameters by NAICS sector

Parameter	Coal (2121)	Cement (3273)	Fertilizer (3253)	Paper (3221)	Petroleum (3241)
ν_m	0.0010	0.0013	0.0007	0.0013	0.0056
J_m	93	200	118	85	113
$\bar{\mu}_m^{GO}$	1.33	1.08	1.19	1.16	1.08
$\bar{\mu}_m^{VA}$	1.79	1.38	2.81	1.73	1.38
η_m	2.42	3.71	1.59	2.63	4.39
$\bar{\chi}_m$	3.97×10^5	1.49×10^5	4.12×10^5	2.66×10^5	1.03×10^6
$\sigma_{\log \chi_m}$	0.31	0.08	0.41	0.29	0.20
$g_{\bar{\chi}_m}$	0.09	0.06	0.77	0.84	1.07
θ_m	5.43×10^{-4}	7.07×10^{-4}	7.67×10^{-4}	4.24×10^{-4}	5.14×10^{-4}

Note: The calculations and data sources used to obtain the sectoral parameters are described in detail in the present sections. Each row reports a sector-specific calibrated parameter: The first row presents the sectoral output shares ν_m; the second row presents the number of firms J_m in the model; the third and fourth rows present the average sectoral markups on a gross output basis, $\bar{\mu}_m^{GO}$, and on a value added basis, $\bar{\mu}_m^{VA}$, respectively; the fifth row presents the sectoral elasticity of substitution parameters η_m; the sixth row presents the average sectoral labor productivity $\bar{\chi}_m$, in dollars of value added per worker per year; the seventh row presents the standard deviation of the logged sectoral labor productivity $\sigma_{\log \chi_m}$; the eighth row presents the annual growth rate of sectoral labor productivity $g_{\bar{\chi}_m}$; and the ninth row presents the average sectoral carbon intensity θ_m, in tons of carbon per dollar of value added.

values, step size, and stopping rule.

3.1 Calibration

Given the model structure and the parametric assumptions presented, next I describe the calibration of the structural model parameters. One period in the model corresponds to a decade. I calibrate the model to match the US economy in 2017, which is the earliest year for which I have the sectoral concentration data.

I model the five four-digit NAICS sectors with the highest carbon emissions per dollar of gross output in 2017 among those with a gross output share above 0.1 percent: Coal mining (2121), Cement and concrete product manufacturing (3273), Fertilizer and other agricultural chemical manufacturing (3253), Pulp, paper, and paperboard mills (3221), and Petroleum and coal products manufacturing (3241). I exclude utilities (NAICS 22), the most carbon-intensive sector, because regulators set their prices, so the markup condition of the model does not describe them.

Table 1 reports the parameters of the intermediate producers in each of the five sectors, which the following subsections calibrate.[2] Columns are ordered by 2017 emissions per dollar of gross output, the criterion used to select the sectors. The model's carbon intensity θ_m (last row) is instead measured per dollar of value added, consistent with ν_m and $\bar{\chi}_m$, and statements about carbon intensity below refer to θ_m.

The five sectors account for 6.2 percent of US CO_2 emissions on average over 2010–2023, and the

[2]The next five sectors by carbon emissions per dollar of gross output are Basic chemical manufacturing (3251), Iron and steel mills and ferroalloy manufacturing (3311), Grain and oilseed milling (3112), Resin, synthetic rubber, and artificial fibers manufacturing (3252), and Metal ore mining (2122).

United States for 14.6 percent of world emissions over 2011–2023 (Appendix C.7). I next describe the construction of the dataset containing the sectoral economic statistics I use and how I obtain the sectoral parameters used in the calibration of the model.

3.1.1 Sectoral Output Shares

I compute the sectoral output shares ν_m from the 2017 Input-Output Use Table of the Bureau of Economic Analysis (BEA), mapped from BEA industries to 2017 NAICS codes with the concordance in BEA (2024a). For each four-digit NAICS sector, gross output GO_m is the total commodity output ("T001") of the commodities the sector produces, and value added VA_m is industry value added ("VAPRO") summed over the sector's industries; GDP is total value added across all industries. Because the intermediate technology is linear in labor, with no materials input, a sector's output in the model corresponds to its value added. I therefore set $\nu_m = \mathrm{VA}_m/\mathrm{GDP}$, and use gross output only for the share of intermediate inputs in revenue, $s_{M,m} = 1 - \mathrm{VA}_m/\mathrm{GO}_m$, which converts markups in Section 3.1.3.

The first row of Table 1 reports ν_m. Among the five sectors, the more carbon-intensive ones tend to have smaller output shares (correlation -0.37): fertilizer is the most carbon-intensive and has the smallest share, while petroleum has below-median carbon intensity and by far the largest share.

Because the weights $\kappa_{m,j}$ sum to one, the composite X_m has the scale of a single firm's output, so it is not comparable to the sum of firm outputs $\sum_j X_{m,j}$ reported as gross output. Its scale does not affect the results: X_m enters only through $X_m^{\nu_m}$ in final goods production, so a different normalization would multiply output by a constant that the calibration of $A_{0,0}$ to GDP absorbs, and firms' revenue shares depend only on relative outputs. Appendix C.2 compares sectoral employment in the oligopoly equilibrium with the data.

3.1.2 Market Structure

The number of firms in each sector and their revenue shares are calibrated to the concentration statistics of the 2017 Economic Census (US Census Bureau 2023), at the national level for each four-digit NAICS sector. For each sector, the Census reports the number of firms, the share of sector revenue of the 4, 8, 20, and 50 largest firms, and the Herfindahl-Hirschman Index (HHI) of the 50 largest firms, which is the sum of their squared revenue shares in percent. The four largest firms account for 34 percent of revenue in coal, 13 percent in cement, 31 percent in fertilizer, 45 percent in paper, and 42 percent in petroleum, and the HHI ranges from 74 in cement to 661 in paper. Among the five sectors, the more carbon-intensive ones are less concentrated: the correlation between the HHI and carbon intensity is -0.77.

The Census reports the combined revenue share of the 4, 8, 20, and 50 largest firms and the HHI of the 50 largest, but not the shares of individual firms. I assign a revenue share to each of the 50 largest firms by assuming that shares decline smoothly with rank, and choose the curve so

that it reproduces the four concentration ratios and the HHI.[3] The remaining firms form a fringe of equal-sized firms that account for the rest of sector revenue, from 5 percent in petroleum to 55 percent in cement.[4] The second row of Table 1 reports the resulting number of firms J_m. The largest firm's revenue share is 12 percent in coal, 4 percent in cement, 9 percent in fertilizer, 14 percent in paper, and 13 percent in petroleum. Markets are defined nationally and at the four-digit level, so concentration is understated where markets are regional: the cement sector (3273) includes concrete products and has 4,182 firms, whereas regional cement markets are served by a few plants (Fowlie, Reguant, and Ryan 2016).

3.1.3 Average Sectoral Markups

The average markup $\bar{\mu}_m$ measures market power in sector m. I compute it as the sales-weighted average of the markups of publicly traded firms in the sector, estimated from Compustat (S&P Global Market Intelligence 2023) with the method described in Appendix A. These firms own facilities that emit 39 percent of the five sectors' emissions. Cement is the exception: of the three firms in its sample, only Monarch Cement makes cement, and it accounts for 1 percent of the sector's emissions, while most of the sales weight is on Forterra, a maker of concrete pipe and precast products. Two thirds of cement emissions come from plants owned by subsidiaries of firms listed abroad, whose markups are not in Compustat; the US Geological Survey similarly reports that 79 percent of US cement capacity was foreign-owned in 2017 (US Geological Survey 2020). Appendix C.12 reports this coverage by sector and repeats the analysis with the markups of the listed firms that own each sector's emitting facilities; because the value added conversion is not defined for some of these markups, that calibration uses gross output markups; in cement they imply a markup of 1.34 instead of 1.08 and an optimal tax of 74 percent of the Pigouvian tax instead of 92 percent, and the conclusions of Section 5 are unchanged. The Compustat markups are gross output markups, because the production function approach treats materials as a flexible input. The model's intermediate producers use only labor, so the markup that governs their pricing is a value added markup, which I obtain from the gross output markup as $\bar{\mu}_m^{VA} = \bar{\mu}_m^{GO}(1 - s_{M,m})/(1 - \bar{\mu}_m^{GO} s_{M,m})$, where $s_{M,m}$ is the share of intermediate inputs in the sector's revenue (Basu 2019). The conversion is defined only when $\bar{\mu}_m^{GO} s_{M,m} < 1$ and becomes very sensitive to the markup as this product approaches one. It is well defined for the 2017 sector averages, but not for some individual firms or for the 2022 markups and the markups of the emitting firms, so comparisons with individual firms (Appendix C.1) and those two robustness calibrations use gross output markups. In the model, a firm's markup depends on its revenue share, $\mu_{m,j} = \sigma_m(s_{m,j})/(\sigma_m(s_{m,j}) - 1)$, with $\sigma_m(\cdot)$ given by equation (14). I choose η_m so that the sales-weighted average markup in the equilibrium without

[3] The log of a firm's revenue share is piecewise linear in the log of its rank, with knots at ranks 1, 4, 8, 20, and 50. The five knot values are chosen so that the shares match the four concentration ratios and the HHI exactly.

[4] I cap the number of firms at 200 per sector to limit computation time. The cap binds only in cement, where each fringe firm then has 0.37 percent of sector revenue, more than the 50th largest firm (0.19 percent). The 50 largest firms in the model therefore account for 46.5 percent of revenue instead of 44.8 percent, and their HHI is 74.5 instead of 73.5; the 4-, 8-, and 20-firm shares are matched exactly.

policy in 2017 equals the target:

$$\sum_{j=1}^{J_m} s_{m,j}\,\mu_{m,j}(\eta_m) = \bar{\mu}_m^{VA}, \tag{22}$$

where $s_{m,j}$ are the revenue shares of Section 3.1.2. The resulting markups range from 1.72 for the median firm to 1.92 for the largest firm in coal, 1.38 to 1.43 in cement, 2.74 to 2.97 in fertilizer, 1.63 to 1.88 in paper, and 1.30 to 1.49 in petroleum. Appendix C.1 compares them with the markups of individual Compustat firms, which are not targeted.

The average sectoral markups $\bar{\mu}_m$ are presented in the third and fourth rows of Table 1, and η_m is presented in the fifth row. For a given degree of concentration, a lower average markup implies a higher elasticity of substitution η_m. Concentration also matters: petroleum has the same average markup as cement (1.38) but a higher η_m (4.39 against 3.71), because more of its markup comes from the market shares of its largest firms. The sectors with the highest average sectoral markups among the top ten carbon-intensive sectors are petrochemical manufacturing (3251), fertilizer manufacturing (3253), and iron and steel mills (3311). Coal mining (2121), which ranks highest on the gross output markup, falls to fourth once the value added conversion is applied: its materials share (0.43) is well below the ten-sector average (0.69), so the conversion raises its markup by less. There is no clear correlation between the average sectoral markups and the carbon intensity of the sector, but in every carbon-intensive sector the average markup is above one, so prices exceed marginal cost. I present the history of average sectoral markups $\bar{\mu}_m$ values used in the calibration of η_m in Appendix C.3. Appendix C.9 reports the results with gross output markups, which imply higher elasticities (η_m from 4.7 to 19, and at the upper bound of 20 in petroleum).

3.1.4 Firm Productivities

I use the average sectoral labor productivity in 2017 as the calibration target for the level of labor productivity $\chi_{m,j,t}$ in each intermediate sector $m = 1, \ldots, M$. I calculate it using the sectoral value added in the IO Table dataset and the sectoral total employment in the US Census Bureau (2023) dataset for 2017. For each four-digit NAICS sector, value added is the variable "VAPRO" in the BEA (2024a) dataset, which I merge with total employment from the US Census Bureau (2023) dataset. The base year average sectoral labor productivity $\bar{\chi}_m$, presented in the sixth row of Table 1, is the sector's total value added divided by its total employment in 2017.

Within each sector, I set firm productivities so that the revenue shares of Section 3.1.2 are the equilibrium shares without policy in 2017. A firm's revenue share depends on its CES weight and on its price, $s_{m,j} = \kappa_{m,j}^{\eta_m}(p_{m,j}/p_m)^{1-\eta_m}$ with $p_{m,j} = \mu_{m,j}w/\chi_{m,j}$, so revenue shares identify only the combination $\kappa_{m,j}^{\eta_m}\chi_{m,j}^{\eta_m-1}$. I assume that CES weights are proportional to productivity, $\kappa_{m,j} = \chi_{m,j}/\sum_k \chi_{m,k}$, so that more productive firms also produce the varieties buyers value more; the weights are fixed at their 2017 values, and because all firms in a sector share the same productivity growth rate, they would not change if recomputed in later periods. Productivities are then given

by

$$\chi_{m,j} \propto \left(s_{m,j}\, \mu_{m,j}^{\eta_m - 1} \right)^{\frac{1}{2\eta_m - 1}},$$

where the markup $\mu_{m,j} = \sigma_m(s_{m,j})/(\sigma_m(s_{m,j}) - 1)$ depends only on the firm's share (Appendix C.4), and I scale them so that the mean of $\log \chi_{m,j}$ equals $\log \bar{\chi}_m$.[5] The seventh row of Table 1 reports the resulting standard deviation of log productivity, $\sigma_{\log \chi_m}$, which is implied by the Census concentration data. It is smaller than the dispersion of value added per employee across Compustat firms in the same sectors (0.60, 0.18, 0.41, 0.30, and 0.57), which reflects differences in prices and markups as well as in physical productivity. Average sectoral labor productivity $\bar{\chi}_m$ displays no clear correlation with the carbon intensity of the sector (corr $= -0.27$): petroleum refining combines the highest labor productivity with below-median carbon intensity.

3.1.5 Average Sectoral Productivity Growth Rate

Labor productivity in each sector grows at a constant annual rate $g_{\bar{\chi}_m}$, which I set to the average annual growth of labor productivity from 2000 to 2023 in the Integrated Industry-Level Production Accounts (ILPA) of the BEA and the Bureau of Labor Statistics (BEA 2024b). The ILPA are published for three-digit industries, so each sector takes the rate of the industry that contains it: Mining, except oil and gas (212) for coal, Nonmetallic mineral products (327) for cement, Chemical products (325) for fertilizer, Paper products (322) for paper, and Petroleum and coal products (324) for petroleum. The eighth row of Table 1 reports the rates, in percent per year, and Appendix C.5 shows the annual series. Productivity growth is weakly negatively correlated with carbon intensity (correlation -0.32).

3.1.6 Carbon Intensity

I assume that carbon intensity is constant over time and common to the firms in a sector, $\theta_{m,j,t} = \theta_m$, and that emissions are proportional to output, $\zeta_m = 1$ (Section 3.1.7). I set θ_m to the sector's emissions per dollar of value added in 2017. Emissions are the facility-level emissions reported to the EPA's Greenhouse Gas Reporting Program (GHGRP) (US Environmental Protection Agency 2024), aggregated to four-digit 2017 NAICS sectors with the concordance of US Census Bureau (2017), and value added is from the BEA Input-Output table (Section 3.1). The GHGRP covers 129 four-digit sectors from 2010 to 2023 and reports emissions of all greenhouse gases in metric tons of CO_2 equivalent, which I convert to metric tons of carbon by multiplying by 12/44, the ratio of the atomic mass of carbon to the molecular mass of CO_2. The world emissions target below and the carbon cycle (Appendix C.6) are measured the same way, in CO_2 equivalent, so the three are consistent. The cost is that a ton of methane or nitrous oxide in CO_2 equivalent is treated as a ton of CO_2 in the carbon cycle, although methane leaves the atmosphere within about

[5] The assumption matters for how output, and therefore emissions, is divided among firms with given revenue shares, because a firm's output is proportional to $s_{m,j}\chi_{m,j}/\mu_{m,j}$. The largest coal firm, with 11.5 percent of revenue, emits 15.4 percent of the sector's emissions under this assumption; with equal CES weights it would emit 24.0 percent, and with equal productivities 10.7 percent.

Table 2: Baseline calibration of aggregate parameters

Parameter	Description	Value	Source
α	Output share of capital	0.3	Golosov et al. (2014)
β	Discount factor	0.985^{10}	Golosov et al. (2014)
δ	Capital depreciation rate	1.0	Golosov et al. (2014)
g_0	Annual TFP growth rate	0.0073	World Bank (2021) & WDI (2024)
N	Exogenous labor supply (millions)	142.5	OEWS (2017)
s_D	World-to-US GDP ratio	6.59	WDI (2024)
$\bar{\gamma}$	Climate change damage elasticity	5.6×10^{-5}	Calibrated to \$50/tCO$_2$ SCC, IWG 2021
ψ_L	Permanent emission share	0.2	Golosov et al. (2014)
ψ_0	Decadal retained emission share	0.1988	NOAA (2025a) & OWID (2023)
ψ	Decadal geometric emission decay rate	0.0445	NOAA (2025a) & OWID (2023)
ϕ_E	Emissions scaling constant	133.24	Internally calibrated (see text)
$\bar{S}$	Pre-industrial carbon concentration	596.4	NOAA (2025a)
$S_{1,-1}$	Permanent carbon stock entering 2017	174.876	OWID (2023)
$S_{2,-1}$	Transitory carbon stock entering 2017	95.122	NOAA (2025a) & OWID (2023)
ζ_m	Output elasticity of emissions	1.0	Assumption (see text)
$\Omega_{m,j}$	Fixed cost share	0	Assumption (see text)
K_0	Initial capital stock (\$2021 trillions)	47.6	Internally calibrated (see text) & WDI (2024)
$A_{0,0}$	Final goods sector initial TFP	3.2455×10^4	Internally calibrated (see text)

a decade. This matters mainly for coal mining, where CO_2 is 1 percent of reported emissions and the rest is methane; it is 78 to 87 percent in fertilizer and paper and above 99 percent in cement and petroleum. Appendix C.13 repeats the analysis with CO_2 emissions only. The ninth row of Table 1 reports θ_m in tons of carbon per dollar of value added. Because the model's sectoral output differs from observed value added, I scale each θ_m by a sector-specific factor so that the sector's emissions in the no policy oligopoly equilibrium equal its observed 2017 emissions; Table 1 reports the unscaled values and Appendix C.2, Table 6 the scale factors.

3.1.7 Other Structural Parameters

I summarize the calibration strategy for the remaining structural parameters of my model $\left\{\alpha, \beta, \delta, N, g_0, \bar{\gamma}, \psi_0, \psi_L, \psi, \left\{\zeta_m, \{\Omega_{m,j}\}_{j=1}^{J_m}\right\}_{m=1}^{M}\right\}$, the initial conditions $\{K_0, A_{0,0}, \bar{S}, S_{1,-1}, S_{2,-1}\}$, and other calibration targets in Table 2. The paragraphs that follow explain the choices that do not follow Golosov et al. (2014).

I follow Golosov et al. (2014) for α, β, and δ. Unlike them, I set the annual growth rate of total factor productivity in the final goods sector, g_0, to 0.73 percent, the GDP-weighted average of country-level TFP growth from 2007 to 2017. Country TFP growth is "Total factor productivity (TFP) in log difference, percent" from the Cross-Country Database of Productivity (World Bank 2021), and the weights are real GDP data, "GDP, PPP (constant 2021 international \$)," from the World Development Indicators (WDI 2024). I set labor supply N to total US employment in 2017, 142.5 million, from the Occupational Employment and Wage Statistics (OEWS 2017), so that the intermediate sectors, capital, and $A_{0,0}$ are all on the US scale, and I assume no population growth.

I keep the functional form of damages in Golosov et al. (2014) but not their value of $\bar{\gamma}$, 2.379×10^{-5}, which implies an output loss of only 1.4 percent at 3°C. I instead calibrate $\bar{\gamma}$ so that the

social cost of carbon in 2017 is \$50 per ton of CO_2, close to the central estimate of the Interagency Working Group (IWG 2021). Because output is on the US scale while emissions and the carbon cycle are global, the planner's marginal damage is scaled by the ratio of world to US GDP in 2017, $s_D = 6.59$ (WDI 2024), the standard gap between the global and the domestic social cost of carbon. This gives $\bar{\gamma} = 5.6 \times 10^{-5}$, 2.37 times their value, and an output loss of 3.3 percent at 3°C, close to the 3.1 percent of DICE-2023 (Barrage and Nordhaus 2024) and below the 7 to 8 percent that Howard and Sterner (2017) estimate for non-catastrophic damages, or 9 to 10 percent including catastrophic risks. The parameters and initial stocks of the carbon cycle also differ from theirs; Appendix C.6 describes their calibration.

In the baseline, emissions are linear in output, $\zeta_m = 1$, and there are no fixed costs, $\Omega_{m,j} = 0$, although the model and the solution algorithm allow for both. Fixed costs do not enter the firms' markup conditions, and the data leave no room for them: model employment already exceeds Census employment in every sector (Table 6).

The initial capital stock K_0, used in production in 2017, is 47.6 trillion dollars: the saving out of the preceding decade's output at the model's saving rate $\alpha\beta$, with that output equal to ten times US GDP in 2007, 18.47 trillion dollars (WDI 2024).

The five sectors are US sectors, whereas the carbon cycle is global. I bridge the two with a constant ϕ_E that scales modeled emissions to world emissions. Following Barrage (2020), I calibrate ϕ_E and $A_{0,0}$ to two 2017 targets: ϕ_E so that emissions in the first period of the oligopoly equilibrium without policy equal world greenhouse gas emissions in CO_2 equivalent, about 14.2 GtC per year (OWID 2023), the series to which the carbon cycle is fitted, and $A_{0,0}$ so that output equals US GDP. Neither parameter enters a first-order condition, so I solve the equilibrium once with both set to one and then set each to match its target; because emissions are linear in ϕ_E and output is linear in $A_{0,0}$ given the allocation, both targets are matched exactly. Unlike in Barrage's two-sector model, the allocation has no closed form here, so it comes from the numerical solution rather than from a formula. Since the five sectors' modeled emissions equal their observed 2017 level of 0.106 GtC, $\phi_E = 133$ is the ratio of world emissions to the five sectors' emissions.

Holding ϕ_E fixed over time implies that emissions outside the five sectors move in proportion to theirs. This abstracts from carbon leakage, the relocation of emissions to regions without climate policy (Hoel 1996), and means that the optimal policy in the model should be read as one applied to all emissions. The five sectors' share of US emissions has been stable since 2010, but the US share of world emissions has declined (Appendix C.7), so a ton of the five sectors' emissions stands for more world emissions over time than the constant ϕ_E implies. Appendix C.14 shows how the results depend on ϕ_E: a higher value raises the share of marginal damage in the price of the sectors' output and moves the optimal tax closer to the Pigouvian tax, while the welfare cost of ignoring market power does not depend on it.

4 Quantitative Results

I quantitatively evaluate the model by simulating the optimal and equilibrium paths of allocations and prices under the baseline calibration. The baseline results show that market power lowers emissions and warming relative to perfect competition, the optimal tax is 33 to 75 percent of the Pigouvian tax and differs across sectors with their markups, and it differs little across the firms of a sector.

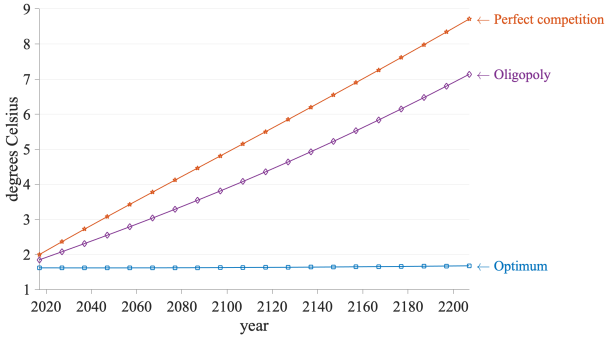
Figure 2 summarizes key simulation outcomes. Panel (a) shows global mean surface temperature change relative to pre-industrial levels. The oligopoly equilibrium path without policy is anchored to observed 2017 emissions and is the counterpart to Golosov et al.'s (2014) laissez-faire scenario. The temperature path starts at 1.8°C above pre-industrial levels in 2017 and reaches about 3.9°C by 2100 and 7.1°C by 2207.[6] It is close to their laissez-faire path, which reaches 4.0°C by 2100 and 9.6°C by 2200: slightly above it until the 2080s and below it after. The perfect competition path is the counterfactual of removing the five sectors' markups. It raises emissions by about 60 percent and warming to about 4.9°C by 2100 and 8.7°C by 2207, above their laissez-faire path until the 2150s and below it after. The gap between the two paths measures the emission restraint that market power currently provides: warming is lower under oligopoly by 1.0°C in 2100 and 1.6°C in 2207. The oligopoly path is within the range of 3.75°C to 4.25°C in 2100 in which the baseline projections of the integrated assessment models compared by Gillingham et al. (2018) are bunched at a common climate sensitivity, and the competition path is above it. Under the optimal policy, temperature stays near 1.6°C over the whole horizon (1.61°C in 2017, 1.63°C in 2117, and 1.67°C in 2207), well below the 2.5°C by 2100 in their optimal path, reflecting the higher social cost of carbon to which I calibrate damages.[7] Because the five sectors account for about 1 percent of output but stand in for all emissions, reducing their output is a cheap way to cut emissions: the optimal path corresponds to an immediate reduction in emissions of about 92 percent at a cost of 1.6 percent of output.

Panel (b) shows the ratio of optimal to equilibrium aggregate output, net of climate damages. The ratio starts below one: in 2017 optimal output is about 1.5 percent lower than equilibrium output under either market structure, because the optimal policy reduces carbon-intensive production. It rises above one as avoided climate damages accumulate, by 2057 relative to perfect competition and by 2087 relative to oligopoly. By 2207, optimal output exceeds equilibrium output by about 13 percent under oligopoly and 26 percent under perfect competition. The gap is smaller under

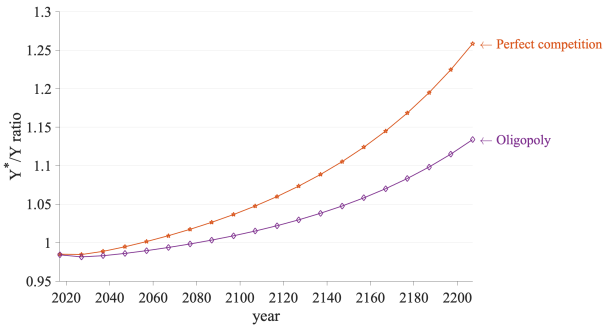
[6] As in Golosov et al. (2014), temperature is the long-run warming implied by the atmospheric carbon concentration, without a lag adjustment, and each period's value reflects emissions through the end of that decade. It is therefore higher than observed warming at the same date: the 2017 concentration alone implies 1.6°C, compared with observed warming of 0.9°C to 1.1°C across estimates reported by NOAA (2025b), C3S (2026), and NASA (2025).

[7] Golosov et al. (2014) report an optimal tax of $56.9 per ton of carbon, or about $15.5 per ton of CO_2, for 2010. I calibrate to $50 per ton of CO_2 in 2017, or $183 per ton of carbon, which is about 3.2 times their value. In both models the social cost of carbon is the product of world output, the damage elasticity, and the discounted share of an emitted ton that remains in the atmosphere. World output is about twice as large in my calibration ($140 trillion per year in 2017, in 2021 dollars at purchasing power parity, compared with their $70 trillion for 2010), and my damage elasticity is 2.4 times theirs. These are partly offset by my estimated carbon cycle, in which the discounted share of a ton remaining in the atmosphere is about two-thirds of theirs.

Figure 2: Model simulations



(a) Long-run warming implied by atmospheric carbon concentration, relative to pre-industrial levels



(b) Ratio of optimal to equilibrium output

oligopoly because market power already restrains carbon-intensive production, so less warming and less damage accumulate without policy.

4.1 Market Structure and Optimal Policy

Table 3 measures the two distortions in dollars at the equilibrium without policy. The markup wedge of a firm is its price minus its marginal cost: the output subsidy of equation (16), evaluated at the equilibrium without policy instead of at the first-best allocation. Per ton of CO_2, it equals the sector's value added per ton times one minus the sales-weighted average of its firms' inverse markups, so it is largest where both are high. This wedge is between \$118 in cement and \$298 in paper, or 2.4 to 6.0 times the social cost of carbon. In the baseline calibration, in which each ton emitted by the five sectors stands for 133 tons of world emissions, the wedge is between 2 and 4.5 percent of the marginal damage. The wedge is not the optimal subsidy. Under the optimal policy the price of carbon-intensive output rises, and the subsidy, which is proportional to the price, rises with it: at the first-best allocation it is 28 to 66 percent of the Pigouvian tax.

Figure 3 shows the optimal output tax relative to the Pigouvian tax for the median firm of each sector. The ratio is below one in every sector and period, highest in cement and petroleum, the two sectors with the lowest markups, and lowest in fertilizer, the sector with the highest. In 2017 it is 75 percent in petroleum, 71 percent in cement, 57 percent in paper, 55 percent in coal, and 33 percent in fertilizer; with gross output markups it ranges from 77 to 94 percent (Appendix C.9).[8] The ratios rise slowly over time, by at most four percentage points by 2207, because labor productivity growth lowers private marginal cost relative to the social cost of carbon; the subsidy is proportional to the price, so it falls relative to the Pigouvian tax. The increase is negligible in cement and coal, where labor productivity is nearly flat, and largest in fertilizer, paper, and petroleum, where it grows fastest. I show the optimal output tax levels, measured in 2021 US dollars per unit of output,

[8] Coal mining's emissions are mostly methane. Counting CO_2 only, its optimal output tax is a subsidy (Appendix C.13).

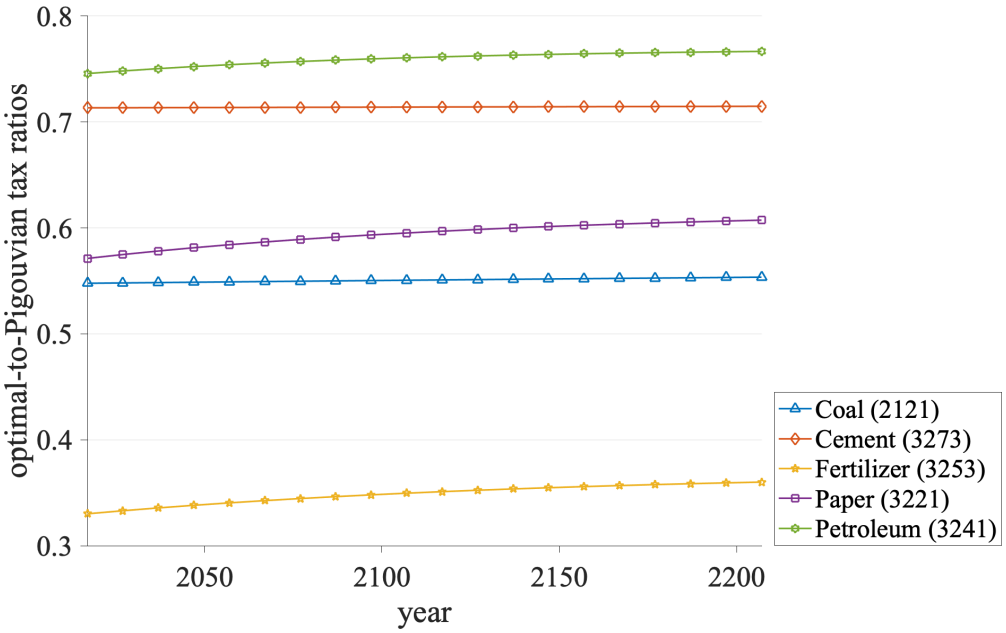
Table 3: Markup Wedge per Ton of Carbon Dioxide at the Equilibrium without Policy, 2017

	Coal	Cement	Fertilizer	Paper	Petroleum
Sales-weighted markup	1.79	1.38	2.81	1.73	1.38
Value added per ton of CO_2 ($)	553	425	392	709	585
Markup wedge per ton of CO_2 ($)	243	118	252	298	161
Markup wedge relative to the social cost of carbon					
Own emissions, times	4.9	2.4	5.0	6.0	3.2
Scaled emissions ($\phi_E = 133$), percent	3.6	1.8	3.8	4.5	2.4
At the first-best allocation					
Output subsidy / Pigouvian tax, percent	45	29	66	44	28

Note: The markup wedge is price minus marginal cost, which equals the firm's price divided by its elasticity of demand, summed over the firms of a sector and divided by the sector's emissions. It is computed at the oligopoly equilibrium without policy, which reproduces the 2017 sales-weighted markups, value added, and emissions of each sector. Dollar values are in 2021 US dollars. The social cost of carbon is \$50 per ton of CO_2. With own emissions, the wedge is divided by the sector's own emissions; with scaled emissions, each ton stands for $\phi_E = 133$ tons of world emissions, as in the baseline calibration. The last row reports the optimal output subsidy as a share of the Pigouvian tax at the first-best allocation, weighted by firms' emissions under perfect competition.

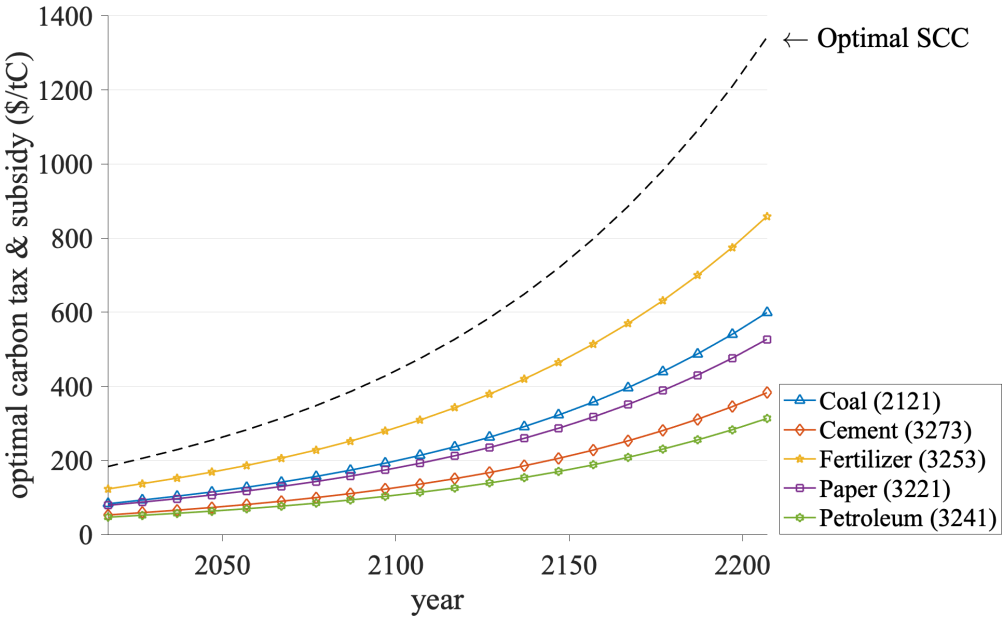
where a unit is the quantity worth one dollar of value added in 2017, in Appendix C.8, Figure 15.

Figure 3: Ratio of the optimal output tax to the Pigouvian tax for the median productivity firm in five carbon-intensive sectors.



Because emissions are proportional to output ($\zeta_m = 1$), the output tax can be expressed per ton of emissions by dividing it by the firm's emissions per unit of output. In 2017, with the social cost of carbon at \$50 per ton of CO_2 (\$183 per ton of carbon), the optimal tax for the median firm is \$37 per ton of CO_2 in petroleum, \$36 in cement, \$29 in paper, \$27 in coal, and \$17 in fertilizer. Since the optimal carbon taxes are a combination of the Pigouvian tax, that is equal to the marginal external damages from the carbon emissions, and the output subsidy, that is equal to the marginal

Figure 4: Optimal carbon tax decomposition into Pigouvian tax and output subsidy for the median productivity firm in five carbon-intensive sectors.



output distortion, I decompose the optimal carbon taxes into these two components. The marginal external damages from a unit of carbon emitted is uniform across all sectors and firms, and is equal to the OSCC.

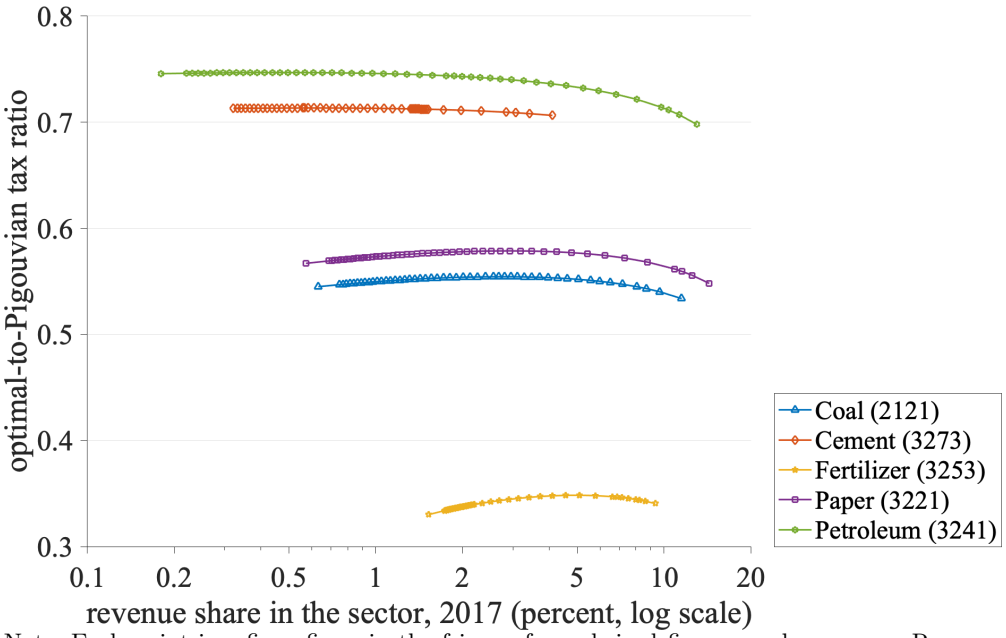
Figure 4 decomposes the optimal tax of the median firm into the Pigouvian tax, the social cost of carbon (dashed line), and the output subsidy of each sector. The optimal tax is positive in every sector and period, so in this calibration the climate externality exceeds the output distortion throughout.

The subsidy is largest in fertilizer and coal, the two sectors with the highest markups, and smallest in cement and petroleum, the two with the lowest. Because marginal damage is close to nine tenths of the first-best price in every sector (Appendix C.14), differences in the subsidy relative to the Pigouvian tax reflect differences in markups, and the same ordering appears in Figure 3.

The optimal tax also differs across firms within a sector, because a firm's markup rises with its market share. With the concentration observed in the Census data, these differences are small. In 2017 the optimal tax for the firm at the 99th percentile of the productivity distribution is 54 percent of the Pigouvian tax in coal, 71 percent in cement, 34 percent in fertilizer, 56 percent in paper, and 71 percent in petroleum, within four percentage points of the median firm's. The optimal tax is positive for every firm in every period. Figure 16 in Appendix C.8 shows the optimal output subsidy by firm productivity percentile in each sector over time.

Two opposing forces explain why the differences are small. The optimal subsidy equals the firm's price divided by its elasticity of demand. A more productive firm has a lower marginal cost and therefore a lower price, which lowers its subsidy and raises its tax. It also has a larger market share and therefore a less elastic demand, which raises its subsidy and lowers its tax. For most

Figure 5: Ratio of the optimal output tax to the Pigouvian tax for every firm, by revenue share, 2017



Note: Each point is a firm; firms in the fringe of equal-sized firms are shown once. Revenue shares are those of the equilibrium without policy, which reproduce the concentration statistics of the 2017 Economic Census; a firm's share rises with its productivity. The model does not generate the dispersion of markups observed across Compustat firms of similar size (Appendix C.1).

firms, market shares are small and the two forces nearly offset: from the 10th to the 90th percentile the optimal tax varies by at most two percentage points in every sector. For the largest firms the second force dominates, most visibly in paper and petroleum, where the largest firms hold 14 and 13 percent of sector revenue. The optimal tax relative to the Pigouvian tax is therefore single-peaked in productivity: flat among the smallest firms, slightly higher in the middle of the distribution, and lowest for the most productive firm in every sector except fertilizer (Figure 5). In fertilizer, where goods are least substitutable ($\eta_m = 1.59$), the first force is stronger, and the smallest firms have the lowest optimal tax relative to the Pigouvian tax.

It follows that a tax that is uniform within a sector is close to the optimal tax of every firm in it. Section 5 shows that such a tax achieves nearly all of the welfare gain of the firm-specific tax. Differences across firms are larger when goods within a sector are closer substitutes: with gross output markups, the optimal tax for the firm at the 99th percentile is 83 percent of the Pigouvian tax in paper and 84 percent in petroleum, against 91 and 94 percent for the median firm (Appendix C.9).

5 Policy Experiments

I now evaluate the welfare consequences of alternative policies, focusing on cases where carbon-intensive sectors are either oligopolistic or perfectly competitive. As established in Section 2.5, the optimal output tax under oligopoly is firm-specific and depends on a firm's carbon intensity, its price, and its elasticity of demand. In Golosov et al. (2014) and Barrage (2020), in contrast, emissions come from a competitive energy sector with a representative producer, so the carbon tax is a single rate per ton and does not vary across firms. I consider two policy settings.

In the first, the true market structure is Cournot oligopoly with heterogeneous firms. I compare three scenarios: (i) an economy-wide uniform tax per ton, equal to marginal external damages (the Pigouvian tax); (ii) a sector-wide tax, equal to the optimal output tax in an economy where every firm in a sector has the sector's average productivity; and (iii) no policy intervention, serving as the benchmark decentralized equilibrium. To derive the carbon tax from an output tax, I divide by each firm's carbon intensity. The resulting tax per ton of emissions is constant across firms in scenario (i), but sector-specific in scenario (ii). The first-best benchmark is the planner's allocation, which the firm-specific optimal output tax of Section 2.5 decentralizes under oligopoly and the Pigouvian tax decentralizes under perfect competition. Policies (i) and (ii) are suboptimal under oligopoly, because policy (i) ignores markups and policy (ii) ignores the differences in markups across firms within a sector.

In the second setting, I assume the true market structure is perfect competition. I compare: (i) the firm-specific optimal output tax designed for oligopolistic sectors with heterogeneous firms from Section 2.5, and (ii) the sector-wide tax described above, which thus appears in both settings. I also report welfare under no intervention. For each case, I compute the welfare loss relative to the planner's first-best allocation, as a consumption-equivalent loss and as an equivalent variation in billions of 2021 dollars per year (Table 4).

Welfare is computed over 40 model periods, twice the horizon shown in Section 4. The planner prices carbon at its infinite-horizon social cost, which includes the damage from the share of emissions that remains in the atmosphere permanently, so a welfare measure computed over a short horizon omits part of that damage and is not maximized by the planner's allocation. Over 40 periods it is: taxes 10 percent above and below the Pigouvian tax both lower welfare under perfect competition. Appendix C.10 shows that the losses under the policy scenarios change little beyond 40 periods, whereas the losses without policy keep rising with the horizon, because warming continues for as long as there is no policy.

The top panel of Table 4 assumes that the carbon-intensive sectors are Cournot oligopolies with heterogeneous firms. Without intervention, the oligopoly equilibrium results in a 24.5 percent consumption-equivalent welfare loss (equivalent to 3.9 trillion dollars per year). Imposing an economy-wide uniform tax, which is first-best under perfect competition, reduces this loss to 0.112 percent (18 billion dollars per year). The sector-wide tax reduces it further to 0.001 percent (0.1 billion dollars per year). The standard result is that all users of carbon should face the same carbon tax, because the damage from a ton of carbon does not depend on where it is emitted, unless an-

Table 4: Welfare Losses Associated with Policy Experiments

Market Structure	Policy Scenario	CE Loss (%)	EV ($2021 bil. per year)
Cournot Oligopoly	No Policy Intervention	24.485	3,864
	Uniform tax per ton (Pigouvian)	0.112	17.7
	Sector-wide tax per ton	0.001	0.1
	Firm-specific tax per ton (optimal under oligopoly)	0.000	0.0
Perfect Competition	No Policy Intervention	37.746	5,957
	Firm-specific tax per ton (designed for oligopoly)	0.160	25.2
	Sector-wide tax per ton	0.151	23.9
	Uniform tax per ton (Pigouvian)	0.000	0.0

Note: The first column reports the true market structure of the carbon-intensive sectors. The second column reports the policy scenario. The uniform tax per ton (Pigouvian) equals the social cost of carbon and is the optimal policy under perfect competition; the firm-specific tax per ton is the optimal policy under oligopoly. The third column reports the permanent percentage reduction in consumption that would make a representative agent indifferent between the policy scenario and the first-best allocation, i.e., the consumption-equivalent (CE) welfare loss. The fourth column reports the equivalent variation (EV) of the welfare loss in billions of constant 2021 US dollars per year, calculated by applying the percentage loss to annual aggregate consumption in the first-best allocation in 2017 ($15.8 trillion). Welfare is computed over 40 model periods (400 years), with climate damages scaled by the world-to-US GDP ratio s_D, the same scaling the planner uses to price carbon. The first-best allocation maximizes this welfare measure, so every loss is nonnegative.

other distortion cannot be corrected directly as in Hoel (1996), where trade policy is constrained. Market power is such a distortion. The uniform tax captures 99.5 percent of the gain from the first-best policy under oligopoly. The remaining loss is the cost of ignoring market power, and the sector-wide tax recovers 99 percent of it. The sector-wide tax is 34 to 74 percent of the Pigouvian tax across the five sectors and is close to the first-best tax of every firm except the largest, because the optimal tax varies little across firms within a sector (Section 4). Firm-specific taxes therefore add almost nothing to a tax that differs only by sector.

These magnitudes depend on how much climate damage a unit of the sectors' output carries. In the baseline the five sectors stand in for all emissions, and marginal damage is about nine tenths of the first-best price of their output. When each firm is instead taxed on its own emissions only, the cost of ignoring market power is about the same, at 0.12 percent of consumption, but the gain from pricing carbon in these sectors falls to 0.06 percent: the optimal output tax is negative in all five sectors, and the Pigouvian tax lowers welfare relative to no policy, consistent with the findings of Fowlie, Reguant, and Ryan (2016) and Cardoso (2025) (Appendix C.14).

The bottom panel considers the case where carbon-intensive sectors are perfectly competitive. Here, the unregulated equilibrium produces a 37.7 percent welfare loss (6.0 trillion dollars per year). Imposing the firm-specific output tax that is first-best for oligopolistic firms leads to a 0.160 percent loss (25 billion dollars per year), more than the loss from ignoring market power in the top panel. The sector-wide tax loses 0.151 percent (24 billion dollars per year). Both taxes are too low for competitive firms: weighted by emissions, the firm-specific tax is 55 percent of the Pigouvian tax in coal, 71 percent in cement, 34 percent in fertilizer, 56 percent in paper, and 72 percent in petroleum, and the sector-wide tax is within about three percentage points of these values. The two taxes cost about the same because no firm dominates its sector under perfect competition: the largest firm accounts for between 6 percent of sector emissions in cement and 22 percent in paper.

Pricing carbon at 34 to 74 percent of its social cost thus still captures more than 99 percent of the gain from the first-best policy.

Together, these results highlight that pricing carbon close to its social cost delivers nearly all of the welfare gain under either market structure. Both possible mistakes are small: ignoring market power when it exists costs 0.112 percent of consumption, and correcting for it when firms are in fact competitive costs 0.151 to 0.160 percent. Across the alternative calibrations in Appendix C.10, which mistake is more costly varies: in the baseline the second costs about 1.4 times the first, while over 20 periods the first costs about twice the second, and neither is more than about one percent of the loss without policy. When firms are oligopolistic, the sector-wide tax is nearly first-best, so the correction for market power can be made at the sector level. A regulator who is unsure of the market structure can hedge: half of the sector-level correction loses 0.043 percent under oligopoly and 0.022 percent under perfect competition, less than either full policy loses when it is the wrong one.

These conclusions are robust in seven respects, documented in Appendix C.10. First, the losses under the policy scenarios change little when welfare is computed over 60 periods. Second, with elasticities calibrated to gross output markups the optimal tax is closer to the Pigouvian tax, and the uniform tax under oligopoly loses 0.014 percent. Third, the results do not depend on measuring markups precisely: if the regulator's estimate of every firm's markup is off by 10 percentage points, either tax loses less than 0.01 percent. Measurement error that differs across firms is more costly for the firm-specific tax, which loses 0.016 percent when firm-level markups are measured with noise of 10 percentage points, more than ten times the loss of the sector-wide tax; this is a further reason to prefer the sector-level correction. Fourth, the form of competition matters little: applying the tax designed for quantity competition when firms compete in prices loses 0.001 percent. Fifth, the conclusions do not depend on the base year (Appendix C.11). When the concentration of 2022 is used in place of 2017, the tax ratios change by at most about one percentage point and no welfare loss under a policy changes by more than 0.01 percentage points. When markups are also set to their 2022 values, which in coal and fertilizer were the highest since 2010, the value added conversion is not defined in fertilizer, so this calibration uses gross output markups and is compared with the gross output calibration. The optimal tax then falls from 77 to 58 percent of the Pigouvian tax in coal, from 92 to 78 percent in cement, and from 86 to 61 percent in fertilizer, and the uniform tax under oligopoly loses 0.040 percent of consumption instead of 0.014 percent; the sector-wide tax loses 0.007 percent. Sixth, the conclusions do not depend on which firms' markups are used (Appendix C.12). When the markup targets are those of the listed firms that own the sectors' emitting facilities, which are also used on a gross output basis, the optimal tax falls from 92 to 74 percent of the Pigouvian tax in cement and changes by five percentage points or less in the other sectors, the uniform tax under oligopoly loses 0.023 percent instead of 0.014 percent, and the sector-wide tax loses 0.005 percent. Seventh, with CO_2 only in place of all greenhouse gases (Appendix C.13), the optimal tax in coal mining, whose emissions are mostly methane, becomes a subsidy, while the tax ratios of the other four sectors change by at most two percentage points.

Ignoring market power still costs 0.11 percent of consumption, but the sector-wide tax then loses 0.059 percent under oligopoly, so it recovers about half of that cost rather than nearly all of it.

The conclusions do change in two cases. The first is the one described above: when marginal damage is a small share of the price of a sector's output, the Pigouvian tax can lower welfare and the sector-wide tax recovers little or none of the gain (Appendix C.14). The second is market definition: if product markets are narrower than the four-digit sectors, markups differ more across the firms of a sector and the sector-wide tax loses more relative to the firm-specific tax. Appendix B.7 discusses this case, together with market power of the buyers of carbon-intensive goods and existing regulation.

6 Conclusion

This paper asks how a carbon tax should be adjusted when carbon-intensive industries are imperfectly competitive and the regulator cannot correct market power directly. In a dynamic general equilibrium climate model with oligopolistic carbon-intensive sectors, the optimal output tax on each firm is the Pigouvian tax minus an output subsidy equal to the firm's price divided by its elasticity of demand. It attains the first-best allocation with a single instrument, because each firm's emissions follow from its output.

I quantify the adjustment for five US sectors, with firm sizes matched to the concentration data of the Economic Census and markups matched to estimates from Compustat. The optimal tax for the median firm is 33 to 75 percent of the Pigouvian tax, or \$17 to \$37 per ton of CO_2 when the social cost of carbon is \$50. The adjustment is larger in sectors with higher markups. It varies little across the firms of a sector, so a tax that is uniform within each sector achieves nearly all of the welfare gain of the firm-specific tax. Market power itself lowers the emissions of these sectors by about 40 percent relative to perfect competition.

The welfare at stake in the adjustment is modest. Ignoring market power and levying the Pigouvian tax costs 0.11 percent of consumption, about 18 billion dollars per year, and captures 99.5 percent of the gain from the first-best policy; correcting for market power when firms are in fact competitive costs somewhat more, 0.15 to 0.16 percent. For policy, as long as the taxed sectors stand in for all emissions, this means that a carbon price should be lower where markups are high, that the adjustment can be made by sector, for example through output-based rebates, without information on individual firms, and that pricing carbon matters far more than getting the adjustment right. Firm-specific taxes add almost nothing, and they are more exposed to errors in measured markups than a sector-wide tax.

These results hold with market structure given. The number of firms is fixed, there is no trade, and emissions can be reduced only by reducing output. Studies of individual industries find that a carbon price equal to the social cost of carbon can lower welfare. Fowlie, Reguant, and Ryan (2016) find this for cement even with production capacity held fixed, and larger losses once firms exit and imports rise. Cardoso (2025) finds it for aviation with fleets and networks held fixed, because the

surplus lost through markups per ton abated exceeds the social cost of carbon. My results differ because in my calibration the five sectors stand in for all emissions, so marginal damage is a large share of the price of their output and the markup distortion is small relative to the externality. When firms are taxed on their own emissions only, the model reproduces the finding of these studies: the optimal output tax is negative and the Pigouvian tax lowers welfare relative to no policy.

Five further limitations matter for the magnitudes. The five sectors account for about 1 percent of output but stand in for all emissions, so reducing their output is a cheap way to cut emissions, and the optimal path is cooler than in models in which abatement is costly. Intermediate producers use only labor, so the model overstates employment in these sectors. The size of the adjustment depends on how markups are measured: it is smaller with gross output markups, 6 to 23 percent of the Pigouvian tax instead of 25 to 67 percent, and, on a gross output basis, larger with the markups of 2022 or with those of the listed firms that own the emitting facilities; the welfare cost of ignoring market power stays small in every case. Emissions are measured as all greenhouse gases in CO_2 equivalent; counting CO_2 only leaves the other sectors unchanged but turns the optimal tax on coal mining, whose emissions are mostly methane, into a subsidy (Appendix C.13). Because the carbon cycle treats a ton of methane in CO_2 equivalent as a ton of long-lived CO_2, a carbon cycle with separate gases, or a separate social cost of methane, would price coal mining's emissions more accurately. Finally, the model prices carbon only: local air pollutants emitted with it would add to marginal damage and leave the output subsidy unchanged, so the optimal tax would be a larger share of total marginal damage but would remain below it. Entry and exit, an abatement technology, and differences in carbon intensity across the firms of a sector, which Kim (2025) and Klenow, Pasten, and Ruane (2025) find to be lower at more productive firms, are natural extensions.

References

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2012). "The Environment and Directed Technical Change." *American Economic Review* 102.1, pp. 131–166.

Annicchiarico, Barbara, Luca Correani, and Fabio Di Dio (2018). "Environmental policy and endogenous market structure". *Resource and Energy Economics* 52, pp. 186–215.

Asker, John, Allan Collard-Wexler, Charlotte De Canniere, Jan De Loecker, and Christopher R. Knittel (Nov. 2024). *Two Wrongs Can Sometimes Make a Right: The Environmental Benefits of Market Power in Oil*. Working Paper 33115. National Bureau of Economic Research. http://www.nber.org/papers/w33115.

Atkeson, Andrew and Ariel Burstein (2008). "Pricing-to-Market, Trade Costs, and International Relative Prices." *American Economic Review* 98.5, pp. 1998–2031.

Barnett, A. H. (1980). "The Pigouvian Tax Rule Under Monopoly." *American Economic Review* 70.5, pp. 1037–1041.

Barrage, Lint (2020). "Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy". *The Review of Economic Studies* 87.1, pp. 1–39.

Barrage, Lint and William Nordhaus (2024). "Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model." *Proceedings of the National Academy of Sciences* 121.13.

Basu, Susanto (2019). "Are Price-Cost Markups Rising in the United States? A Discussion of the Evidence". *Journal of Economic Perspectives* 33.3, pp. 3–22.

Baumol, William J. and Wallace E. Oates (1988). *The Theory of Environmental Policy.* 2nd ed. Cambridge University Press.

Buchanan, James M. (1969). "External Diseconomies, Corrective Taxes, and Market Structure." *American Economic Review* 59.1, pp. 174–177.

Cardoso, Diego S. (2025). *Second-Best Carbon Taxation in a Differentiated Oligopoly.* Working Paper 7435943. SSRN. %7Bhttps://doi.org/10.2139/ssrn.7435943%7D.

Chiavari, Andrea and Alexandre Kohlhas (2026). *Mirrleesian Carbon Taxation.* Working Paper. University of Oxford. %7Bhttp://perseus.iies.su.se/~akohl/mct_web.pdf%7D.

Copernicus Climate Change Service (2026). *Climate Indicators: Temperature.* Tech. rep. Accessed: 2026-10-08. Copernicus Climate Change Service (C3S), European Centre for Medium-Range Weather Forecasts. https://climate.copernicus.eu/climate-indicators/temperature.

Cullen, Joseph A. and Stanley S. Reynolds (2023). "Market Dynamics and Investment in the Electricity Sector". *International Journal of Industrial Organization* 89, p. 102954.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (2020). "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics* 135.2, pp. 561–644.

Ebert, Udo (1991). "Pigouvian Tax and Market Structure: The Case of Oligopoly and Different Abatement Technologies". *FinanzArchiv / Public Finance Analysis* 49.2, pp. 154–166.

Eeckhout, Jan (Feb. 2025). "Presidential Address 2024: The Value and Profits of Firms." *Journal of the European Economic Association* 23.1, pp. 52–118.

Forbes (Aug. 2015). *U.S. Coal Company Alpha Natural Resources Files For Bankruptcy.* Forbes. Accessed: 2026-10-26. https://www.forbes.com/sites/nathanvardi/2015/08/03/u-s-coal-company-alpha-natural-resources-files-for-bankruptcy/.

Forster, P., T. Storelvmo, K. Armour, W. Collins, J.-L. Dufresne, D. Frame, D.J. Lunt, T. Mauritsen, M.D. Palmer, M. Watanabe, M. Wild, and H. Zhang (2021). "The Earth's Energy Budget, Climate Feedbacks, and Climate Sensitivity". *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change.* Ed. by V. Masson-Delmotte, P. Zhai, A. Pirani, S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press. Chap. 7, pp. 923–1054.

Fowlie, Meredith, Mar Reguant, and Stephen P Ryan (2016). "Market-Based Emissions Regulation and Industry Dynamics." *Journal of Political Economy* 124.1, pp. 249–302.

Fowlie, Meredith L (July 2009). "Incomplete Environmental Regulation, Imperfect Competition, and Emissions Leakage." *American Economic Journal: Economic Policy* 1.2, pp. 72–112.

Fullerton, Don and Gilbert E. Metcalf (2002). "Chapter 26 Tax Incidence." *Handbook of Public Economics*. Vol. 4, pp. 1787–1872.

Ganapati, Sharat, Joseph S. Shapiro, and Reed Walker (Apr. 2020). "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes." *American Economic Journal: Applied Economics* 12.2, pp. 303–342.

Gersbach, Hans and Till Requate (Mar. 2004). "Emission Taxes and Optimal Refunding Schemes." *Journal of Public Economics* 88.3–4, pp. 713–725.

Gillingham, Kenneth, William Nordhaus, David Anthoff, Geoffrey Blanford, Valentina Bosetti, Peter Christensen, Haewon McJeon, and John Reilly (2018). "Modeling Uncertainty in Integrated Assessment of Climate Change: A Multimodel Comparison". *Journal of the Association of Environmental and Resource Economists* 5.4, pp. 791–826.

Golosov, Mikhail, John Hassler, Per Krusell, and Aleh Tsyvinski (2014). "Optimal Taxes on Fossil Fuel in General Equilibrium." *Econometrica* 82.1, pp. 41–88.

Gruber, Jonathan (2011). *Public Finance and Public Policy.* 3rd ed. New York: Worth Publishers.

Hassler, J., P. Krusell, and A.A. Smith (2016). "Chapter 24 - Environmental Macroeconomics". *Handbook of Macroeconomics*. Ed. by John B. Taylor and Harald Uhlig. Vol. 2. Elsevier, pp. 1893–2008.

Hassler, John, Per Krusell, and Conny Olovsson (May 2010). "Oil Monopoly and the Climate." *American Economic Review* 100.2, pp. 460–464.

Hoel, Michael (1996). "Should a Carbon Tax Be Differentiated Across Sectors?": *Journal of Public Economics* 59.1, pp. 17–32.

Howard, Peter H. and Thomas Sterner (Sept. 2017). "Few and Not So Far Between: A Meta-analysis of Climate Damage Estimates." *Environmental and Resource Economics* 68.1, pp. 197–225.

Interagency Working Group on Social Cost of Greenhouse Gases (2021). *Technical Support Document: Social Cost of Carbon, Methane, and Nitrous Oxide Interim Estimates under Executive Order 13990.* Tech. rep. Accessed: 2026-09-14. United States Government. https://www.regulations.gov/document/EPA-HQ-OPPT-2021-0057-0097.

IPCC (2021). "Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change." Ed. by V. Masson-Delmotte, Zhai P., Pirani A., S.L. Connors, C. Péan, S. Berger, N. Caud, Y. Chen, L. Goldfarb, M.I. Gomis, M. Huang, K. Leitzell, E. Lonnoy, J.B.R. Matthews, T.K. Maycock, T. Waterfield, O. Yelekçi, R. Yu, and B. Zhou. Cambridge, United Kingdom and New York, NY, USA: Cambridge University Press.

Jo, Ara and Christos Karydas (2024). *Firm Heterogeneity, Industry Dynamics and Climate Policy.* Working Paper 4325565. SSRN. %7Bhttps://doi.org/10.2139/ssrn.4325565%7D.

Katsoulacos, Yannis and Anastasios Xepapadeas (1995). "Environmental Policy under Oligopoly with Endogenous Market Structure." *The Scandinavian Journal of Economics* 97.3, pp. 411–420.

— (1996). "Environmental Innovation, Spillovers and Optimal Policy Rules." *Environmental Policy and Market Structure*. Ed. by Carlo Carraro, Yannis Katsoulacos, and Anastasios Xepapadeas. Kluwer Academic Publishers, pp. 143–150.

Kim, Seho (2025). *Optimal Carbon Taxes and Misallocation across Heterogeneous Firms.* Working Paper 4844030. SSRN. %7Bhttps://doi.org/10.2139/ssrn.4844030%7D.

Klenow, Peter J, Ernesto Pasten, and Cian Ruane (2025). *Carbon Taxes and Misallocation in Chile.* Working Paper. Stanford University. http://klenow.com/Chile_Carbon_Taxes.pdf.

Lan, Xin and Ralph Keeling (2025a). *Trends in Atmospheric CO2, CH4, N2O, SF6.* Accessed: 2025-06-01. https://gml.noaa.gov/ccgg/trends/data.html.

Lee, Dwight R. (1975). "Efficiency of Pollution Taxation and Market Structure." *Journal of Environmental Economics and Management* 2, pp. 69–72.

Lee, Sang-Ho (1999). "Optimal Taxation for Polluting Oligopolists with Endogenous Market Structure." *Journal of Regulatory Economics* 15.3, pp. 293–308.

Levin, Dan (1985). "Taxation Within Cournot Oligopoly." *Journal of Public Economics* 27.3, pp. 281–290.

Lindsey, Rebecca and LuAnn Dahlman (2025b). *Climate Change: Global Temperature.* Tech. rep. Accessed: 2026-10-08. National Oceanic and Atmospheric Administration (NOAA) Climate.gov. https://www.climate.gov/news-features/understanding-climate/climate-change-global-temperature.

Long, Ngo Van and Antoine Soubeyran (Feb. 2005). "Selective Penalization of Polluters: An Inf-Convolution Approach". *Economic Theory* 25.2, pp. 421–454.

Mansur, Erin T. (Aug. 2013). "Prices versus Quantities: Environmental Regulation and Imperfect Competition." *Journal of Regulatory Economics* 44.1, pp. 80–102.

NASA (2025). *Global Temperature - Earth Indicator.* Tech. rep. Accessed: 2026-10-08. National Aeronautics and Space Administration (NASA). https://science.nasa.gov/earth/explore/earth-indicators/global-temperature/.

National Research Council (2010). *Advancing the Science of Climate Change*. Washington, DC: The National Academies Press.

NOAA Global Monitoring Laboratory (2024). *Annual Greenhouse Gas Index (AGGI).* Dataset. Accessed: 2025-06-01. National Oceanic and Atmospheric Administration (NOAA) Global Monitoring Laboratory. https://gml.noaa.gov/aggi/aggi.html.

Nordhaus, William D. (1993). "Optimal Greenhouse-Gas Reductions and Tax Policy in the "DICE" Model." *American Economic Review* 83.2, pp. 313–317.

Oates, Wallace E. and Diana L. Strassmann (June 1984). "Effluent Fees and Market Structure." *Journal of Public Economics* 24.1, pp. 29–46.

Preonas, Louis (July 2024). "Market Power in Coal Shipping and Implications for U.S. Climate Policy." *Review of Economic Studies* 91.4, pp. 2508–2537.

Requate, Till (2006). "Environmental Policy under Imperfect Competition." *The International Yearbook of Environmental and Resource Economics 2006/2007: A Survey of Current Issues*. Ed. by Tom Tietenberg and Henk Folmer. Cheltenham: Edward Elgar, pp. 120–207.

Reuters (Apr. 2016). *Leading Global Coal Miner Peabody Files for Bankruptcy*. Reuters. Accessed: 2026-10-26. https://www.reuters.com/article/business/leading-global-coal-miner-peabody-files-for-bankruptcy-idUSKCN0XA0NX/.

Ritchie, Hannah, Pablo Rosado, and Max Roser (2023). *CO2 and Greenhouse Gas Emissions*. Dataset. https://ourworldindata.org/co2-and-greenhouse-gas-emissions.

Ryan, Stephen P. (2012). "The Costs of Environmental Regulation in a Concentrated Industry." *Econometrica* 80.3, pp. 1019–1061.

S&P Global Market Intelligence (2023). *Compustat Daily Updates - Fundamentals Annual*. Dataset. Accessed: 2025-05-04. https://wrds-www.wharton.upenn.edu/.

Schneider, Stephen H. (1989). "The Greenhouse Effect: Science and Policy." *Science* 243.4892, p. 771.

Shaffer, Sherrill (1995). "Optimal Linear Taxation of Polluting Oligopolists." *Journal of Regulatory Economics* 7.1, pp. 85–100.

Simpson, R. David (1995). "Optimal Pollution Taxation in a Cournot Duopoly." *Environmental and Resource Economics* 6.4, pp. 359–369.

The Wall Street Journal (Oct. 2016). *Arch Coal Emerges From Chapter 11*. The Wall Street Journal. Accessed: 2026-10-26. https://www.wsj.com/articles/arch-coal-emerges-from-chapter-11-1475684184.

Tinbergen, Jan (Jan. 1952). *On the Theory of Economic Policy*. North-Holland Publishing Company. http://hdl.handle.net/1765/15884.

Traina, James (2018). *Is Aggregate Market Power Increasing? Production Trends using Financial Statements*. Tech. rep. SSRN. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3120849.

US Bureau of Economic Analysis (2024a). *Input-Output Accounts Data, Use Tables*. Dataset. Accessed: 2024-04-10. https://www.bea.gov/industry/input-output-accounts-data.

— (2024b). *Integrated Industry-Level Production Accounts (KLEMS)*. Dataset. Accessed: 2024-02-24. https://www.bea.gov/data/special-topics/integrated-industry-level-production-account-klems.

US Bureau of Labor Statistics (2017). *Occupational Employment and Wage Statistics (OEWS) Tables*. Dataset. Accessed: 2025-06-24. US Department of Labor. https://www.bls.gov/oes/.

US Census Bureau (2017). *North American Industry Classification System (NAICS) Concordance Tables*. Dataset. Accessed: 2025-06-10. https://www.census.gov/naics/.

— (2023). *Selected Sectors: Concentration of Largest Firms for the US: 2017*. Dataset. Accessed: 2024-02-14. https://data.census.gov/table/ECNSIZE2017.EC1700SIZECONCEN?q=Concentration%20Ratio.

US Energy Information Administration (2017). *Coal Production Declines in 2016, with Average Coal Prices Below Their 2015 Level.* Tech. rep. Accessed: 2025-06-24. US Department of Energy. https://www.eia.gov/todayinenergy/detail.php?id=29472.

US Environmental Protection Agency (2024). *Greenhouse Gas Reporting Program (GHGRP).* Dataset. Accessed: 2024-02-24. https://www.epa.gov/ghgreporting/data-sets.

US Geological Survey (2016). *Cement Statistics and Information.* Tech. rep. Accessed: 2025-06-24. https://www.usgs.gov/centers/national-minerals-information-center/cement-statistics-and-information.

— (2020). "Cement". *2017 Minerals Yearbook, Volume I: Metals and Minerals.* Accessed: 2026-10-26. Reston, VA: US Geological Survey. https://www.usgs.gov/media/files/cement-2017-xlsx.

Weyl, E. Glen and Michal Fabinger (2013). "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition". *Journal of Political Economy* 121.3, pp. 528–583.

World Bank (2021). *Global Productivity: A Cross-Country Database of Productivity.* Ed. by Alistair Dieppe. Accessed: 2025-06-30. https://data360.worldbank.org/en/dataset/WB_ASPD.

— (2024). *World Development Indicators (WDI).* Dataset. Accessed: 2025-06-01. World Bank. https://databank.worldbank.org/.

Online Appendix

Taxing Carbon, Not Competition:
Optimizing Market and Environment in a World of Imperfections

Appendix A Data

For the motivational evidence presented in Section 1 and data used in the model calibration in Section 3, I use two primary data sources: (i) US Environmental Protection Agency's (2024) Greenhouse Gas Reporting Program (GHGRP) data, for annual facility-level GHG emissions, and (ii) the S&P Global Market Intelligence's (2023) Compustat for financial statements of private sector. GHGRP data is available from 2010 to 2023, and has detailed information on the emissions of individual facilities in the US. Compustat data is available from 1950 to 2024, and contains information on the financial statements of publicly traded private sector firms in the United States with public equity or debt.

GHGRP reports emissions data in units of metric tons of carbon dioxide equivalent (CO_2e) for each facility that is subject to the reporting requirements of the program. Each facility has a unique facility identification number that is used to track emissions over time even if the facility's name or address changes. The evidence on firms that report emissions in Section 1 and in Appendix C.3 uses emissions from all of these facility types, aggregated to parent companies; the calibration uses the direct emitters, the facility type under which the five modeled sectors report. GHGRP requires reporting of GHG emissions and other relevant information from large emissions sources and suppliers of certain products in the US. GHGRP divides emissions based on the type of facility, such as direct emitter, onshore oil and gas production, gathering and boosting, transmission pipelines, local distribution companies (LDC) direct emissions, sulfur hexafluoride (SF6) from electrical equipment, and suppliers of products such as coal, oil, and natural gas.

Moreover, EPA also publishes data on the parent company ownership structure of facilities. I designate emissions attributable to each parent company by assigning a fraction of emissions of each facility based on the parent company's ownership share. Matching these facility-level emissions to parent company's that own them allows me to aggregate facility-level emissions to the parent company-level.

On the other hand, Compustat data contain firm-level balance sheet and income statement data. I use Compustat to construct three variables of interest: firm-level value added, economic profit rates, and markups. I restrict the sample to domestic firms, which have standard industry format in USD with Foreign Incorporation Codes (FIC) in the USA. I keep utilities (NAICS codes starting with 22) in the sample; their prices are regulated, but they include no firm in the five sectors of the model, and the markups of those sectors are the same with or without them. I also exclude firms with negative or missing assets, sales, cost of goods sold, operating expenses, or gross plants, property, and equipment (PPE). Finally, I also exclude observations in which acquisitions are larger than 5 percent of the value of total assets.

I also download price deflators from NIPA Table 1.1.9 to convert nominal Compustat variables to real variables. I use the GDP deflator from line 1 to deflate all nominal variables, except for deflating capital expenditures, which I deflate using the nonresidential fixed investment deflator from line 9.

I construct firm-level measure of capital using the perpetual inventory method. First, I initialize the capital stock using the first PPE observations for each firm. Then iterate forward on capital using the accumulation equation $K_{i,t} = K_{i,t-1} + I_{i,t} - \delta K_{i,t-1}$, where I compute net investment using changes to net PPE. I assume the depreciation rate δ is 0.07. Finally, to obtain a measure of the real capital stock, I deflate the net investment by the investment goods deflator. Moving forward, I use only real variables when I construct firm-level measures.

First, I construct a measure of firm-level value added by subtracting the cost of goods sold (COGS) from revenue (REVT) for each observation. Next, I calculate firm-level economic profit rates by following Eeckhout (2025). For profits (earnings) of the firm, I use the accounting profits $\pi_{i,t}$, also known as "net income":

$$\pi_{i,t} = \text{REVT}_{i,t} - \text{COGS}_{i,t} - \text{XSGA}_{i,t} - \text{XINT}_{i,t} - \text{DP}_{i,t} + \text{NOPI}_{i,t} + \text{SPI}_{i,t} - \text{TXPD}_{i,t},$$

where REVT is the revenue, COGS is the cost of goods sold, XSGA is the selling, general and administrative expenses, XINT is the interest expense, DP is the sum of depreciation tangible fixed assets and amortization of intangibles, NOPI is the net operating income before interest and taxes, SPI is the special items, and TXPD is the tax liability paid by the firm. Since some shareholders maintain some assets in the firm, Eeckhout (2025) argues that the accounting profits do not adequately measure the opportunity cost of those funds the shareholders expect to obtain, and which should be deduced to obtain the economic profits. He constructs economic profits as the difference between the accounting profits and the opportunity cost of capital, which is the product of the firm's shareholder equity and the cost of debt. I define a firm's shareholder equity E as equal to the bookvalue of the firm, which is equal to its retained earnings RE plus the value of its common stock CEQ, minus its goodwill GDWL,

$$E_{i,t} = \text{RE}_{i,t} + \text{CEQ}_{i,t} - \text{GDWL}_{i,t}.$$

I calculate the firm-level cost of debt as the ratio of the firm's interest expense (XINT) to the firm's total debt, equivalent to Total Liabilities (LT) in Compustat:

$$\text{cost of debt}_{i,t} = \frac{\text{XINT}_{i,t}}{\text{LT}_{i,t}}.$$

After I construct the firm-level shareholder equity and cost of debt, I can calculate its economic profits as:

$$\pi^e_{i,t} = \pi_{i,t} - \text{cost of debt}_{i,t} \cdot E_{i,t-1}.$$

A firm's economic profit rate is calculated as profits as a share of the firm's revenue: $\pi^e_{i,t}/\text{REVT}_{i,t}$.

I merge firm-level economic profit rates constructed from Compustat with the firm-level GHG emissions data constructed from the GHGRP data by matching the parent company names and years across the two datasets before I proceed to construct the markup measure.

Finally, I construct a measure of firm-level markups by following the production function approach framework proposed by De Loecker, Eeckhout, and Unger (2020). De Loecker, Eeckhout, and Unger (2020) derive a firm-level markup, $\mu_{i,t}$, formula as:

$$\mu_{i,t} = \theta^V_{i,t} \frac{P^Q_{i,t} Q_{i,t}}{P^V_{i,t} V_{i,t}}, \tag{23}$$

where $\theta^V_{i,t}$ is the output elasticity of a variable input, $P^Q_{i,t} Q_{i,t}$ is the output (sales), and $P^V_{i,t} V_{i,t}$ is the variable input cost (operating expenses). Following Traina (2018), I use operating expenses, OPEX, as a measure of variable input costs, which is the sum of COGS and XSGA. Whereas De Loecker, Eeckhout, and Unger (2020) use only COGS as a measure of variable input costs, but Traina (2018) argues that OPEX is a better measure of variable input costs because it includes all the costs associated with producing and selling a product, including marketing and management. While, sales and input costs are measurable, the output elasticity is estimated in three stages. First, I estimate industry-level Cobb-Douglas production functions with variable inputs and capital using the following equation:

$$q_{i,t} = \theta_v v_{i,t} + \theta_k k_{i,t-1} + \omega_{i,t} + \varepsilon_{i,t}, \tag{24}$$

where $q_{i,t}$ is the log of sales, $v_{i,t}$ is the log of variable inputs, $k_{i,t-1}$ is the log of capital stock, $\omega_{i,t}$ is the log of productivity, and $\varepsilon_{i,t}$ is the error term. The productivity process is assumed to follow an AR(1) process:

$$\omega_{i,t} = \rho \omega_{i,t-1} + \zeta_{i,t},$$

where $\zeta_{i,t}$ is the innovation to productivity.

In the first stage, I remove the idiosyncratic measurement error from the production process by calculating predicted sales from the following sales-weighted regression with firm and year fixed effects:

$$q_{i,t} = \theta_v v_{i,t} + \theta_k k_{i,t-1} + \xi_i + \xi_t + \epsilon_{i,t}.$$

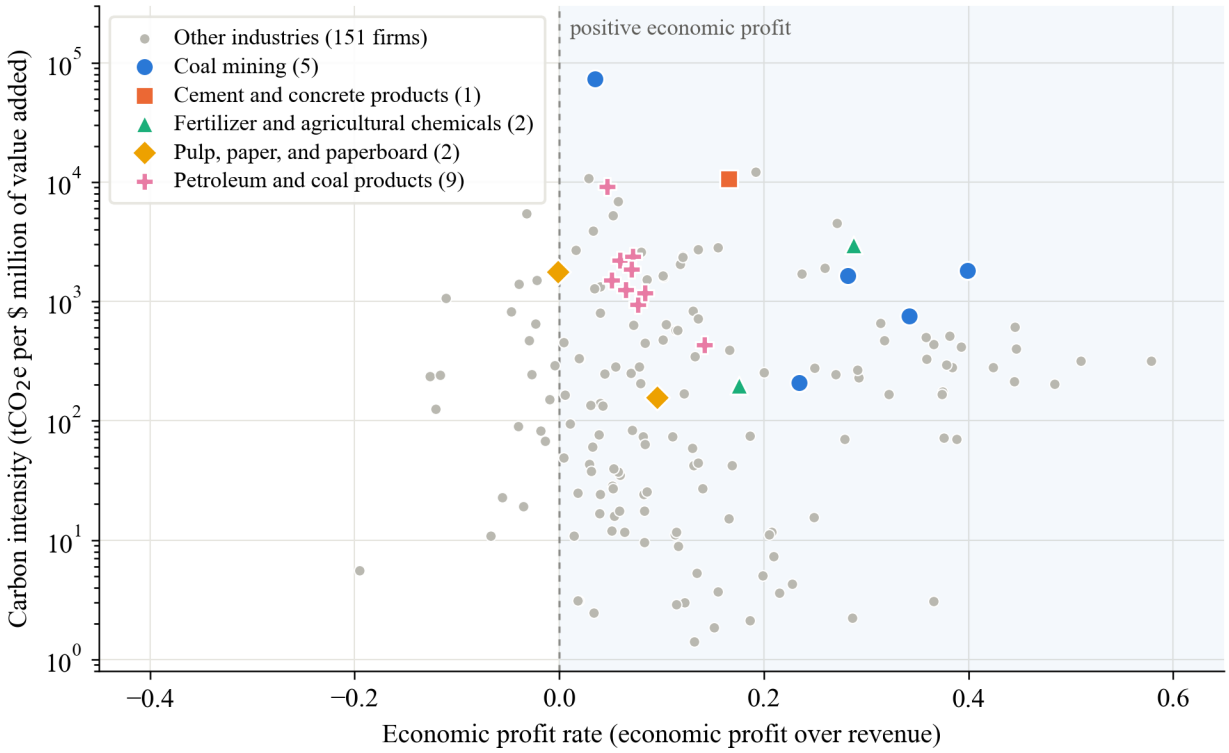
In the second stage, I use the predicted sales from the first stage to derive the implied productivity process from equation (24). I regress this productivity process on its lagged value to recover the persistence parameter ρ and to obtain innovation to productivity $\zeta_{i,t}$.

In the third stage, I use a Generalized Method of Moments (GMM) estimator to estimate the output elasticities, $\theta \equiv [\theta_v, \theta_k]$, by assuming the following moment conditions hold:

$$\mathbb{E}\left[\hat{\zeta}_{i,t}(\theta) \begin{pmatrix} v_{i,t} \\ k_{i,t-1} \end{pmatrix}\right] = 0.$$

The GMM estimator of θ_v is used to construct the firm-level markups following equation (23).

Figure 6: Economic profit rates and carbon intensity of US publicly traded firms that report greenhouse gas emissions, 2022.



Sources: Economic profit rates are calculated from S&P Global Market Intelligence (2023); carbon intensities from US Environmental Protection Agency (2024) and S&P Global Market Intelligence (2023).

Additional evidence on market power. Figure 6 plots the carbon intensity of US publicly traded firms with reported emissions against their economic profit rates in 2022; most firms have positive economic profit rates. Figures 7 and 8 show the sales-weighted average markup and economic profit rate of these firms from 2010 to 2023. The average markup rose from about 1.15 in 2010 to about 1.25 in 2021 to 2023. The average economic profit rate is positive in every year, between 2 percent in 2020 and 11 percent in 2021 and 2022.

Appendix B Theoretical Derivations & Proofs

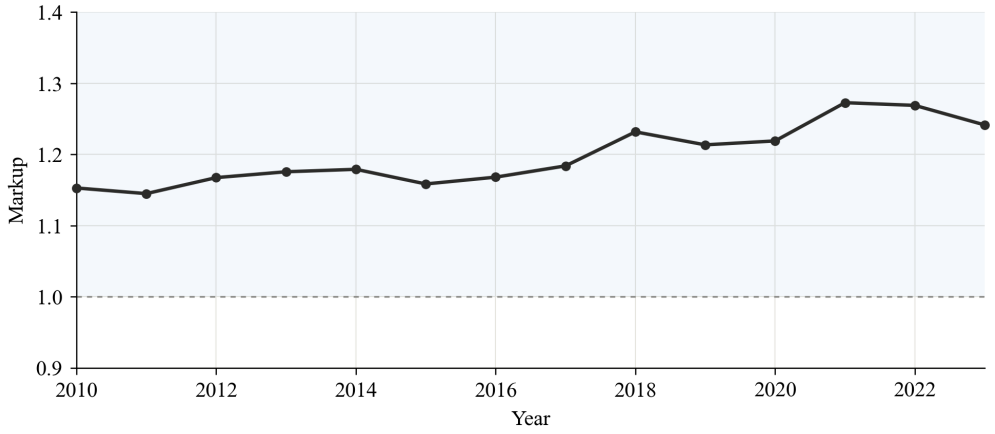
B.1 Derivation of the Constant Social Cost of Carbon

The social cost of carbon emitted at time t by firm j in sector m, in consumption units at this point is given by:

$$\text{SCC}_{m,j,t} = \sum_{s=0}^{\infty} \beta^s \frac{U'(C_{t+s})}{U'(C_t)} \frac{\partial Y_{t+s}}{\partial S_{t+s}} \frac{\partial S_{t+s}}{\partial E_t} \frac{\partial E_t}{\partial E_{m,t}} \frac{\partial E_{m,t}}{\partial E_{m,j,t}}.$$

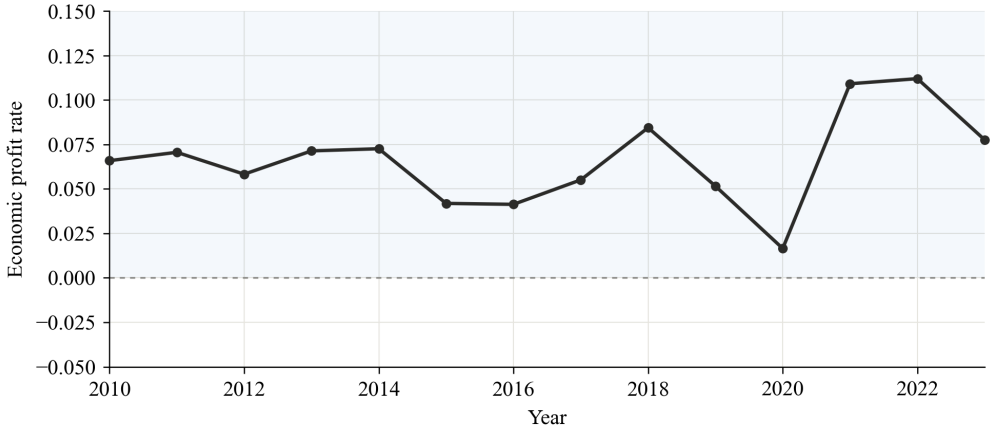
The $\text{SCC}_{m,j,t}$ captures the externality of carbon emission from production of producer j in sector m. It depends on structural parameters in complicated ways. When Assumptions 1, 2, and 3 are

Figure 7: Sales-weighted average markups of US firms that report greenhouse gas emissions, 2010 to 2023.



Note: The shaded area indicates markups greater than one.

Figure 8: Sales-weighted average economic profit rates of US firms that report greenhouse gas emissions, 2010 to 2023.



Note: The shaded area indicates positive economic profit rates.

satisfied, the expression for the SCC simplifies dramatically:

$$\begin{aligned}\text{SCC}_{m,j,t} &= \frac{\partial E_t}{\partial E_{m,t}} \frac{\partial E_{m,t}}{\partial E_{m,j,t}} \sum_{s=0}^{\infty} \beta^s C_t \frac{Y_{t+s}}{C_{t+s}} \gamma_{t+s}(1-d_s), \\ &= \sum_{s=0}^{\infty} \beta^s C_t \frac{Y_{t+s}}{C_{t+s}} \gamma_{t+s}(1-d_s),\end{aligned}$$

which expresses the costs of the externality only in terms of exogenous parameters and the endogenous aggregate saving rate, along with the contemporaneous level of aggregate consumption. If furthermore Assumption 6 holds, then the SCC expression simplifies significantly.

The baseline model parameterization implies that as long as the saving rates are constant across

time in both the planner's solution and the representative consumer's solution, the saving rates are equal across these solutions. As long as the saving rate is constant and equal across the planner's solution and the representative consumer's solution, both the optimal and decentralized allocations will yield comparable dynamic structures. Assumption 6 follows Golosov et al. (2014) and simplifies the optimal SCC expression to:

$$\text{SCC}_{m,j,t} = Y_t \left[\sum_{s=0}^{\infty} \beta^s \gamma_{t+s} (1 - d_s) \right],$$

which expresses the firm-specific SCC as a proportion of GDP and a function of exogenous parameters. This result is identical to Proposition 1 of Golosov et al. (2014). A final expression for their SCC is obtained by assuming that the time path for the damage parameter is constant, i.e., $\gamma_{t+j} = \bar{\gamma}_t$ for all j and d_j defined as in equation (10):

$$\frac{\text{SCC}_{m,j,t}}{Y_t} = \bar{\gamma}_t \left(\frac{\psi_L}{1-\beta} + \frac{(1-\psi_L)\psi_0}{1-(1-\psi)\beta} \right). \tag{25}$$

B.2 Conditions of Proposition 1

Lemma 1 *Suppose that Assumptions 1–5 hold, that the final goods technology is Cobb-Douglas,* $\tilde{F}_{0,t} = A_{0,t} K_t^{\alpha} N_{0,t}^{1-\alpha-\sum_{m=1}^{M} \nu_m} \prod_{m=1}^{M} X_{m,t}^{\nu_m}$*, that intermediate production is linear,* $X_{m,j,t} = \chi_{m,j,t} N_{m,j,t}^{var}$*, that emissions are* $E_{m,j,t} = \theta_{m,j,t} X_{m,j,t}^{\zeta_{m,j}}$ *with* $\zeta_{m,j} \geq 1$*, that* $J_m \geq 2$ *for every sector* m*, and that fixed costs satisfy* $\sum_{m=1}^{M} \sum_{j=1}^{J_m} \Omega_{m,j,t} < N_t$*. Then: (a) every producer's output is positive in the solution to the social planner's problem; and (b) given positive outputs of the other producers in its sector, final output, the wage, and the tax, each producer's profit is strictly concave in its own output, so the first-order condition (18) is necessary and sufficient for its best response.*

Proof.

(a) Positive consumption in every period is feasible, so with logarithmic utility the planner's solution has $C_t > 0$ and therefore $Y_t > 0$, which under the Cobb-Douglas technology requires $N_{0,t} > 0$ and $X_{m,t} > 0$ for every m. Suppose $X_{m,j,t} = 0$ for some producer j, and move $\varepsilon > 0$ units of labor from the final goods sector to producer j. Let $\rho_m \equiv (\eta_m - 1)/\eta_m \in (0,1)$. The sector composite satisfies $X_{m,t}^{\rho_m} = \sum_{i=1}^{J_m} \kappa_{m,i} X_{m,i,t}^{\rho_m}$, so it rises by an amount of order ε^{ρ_m}, and so does final output in period t. The lower labor input in final goods reduces output in period t by an amount of order ε. Emissions rise by $\theta_{m,j,t}(\chi_{m,j,t}\varepsilon)^{\zeta_{m,j}}$, which is at most of order ε since $\zeta_{m,j} \geq 1$, and this reduces output in period t and in every later period by amounts of order ε whose discounted sum is finite. Since $\rho_m < 1$, the gain exceeds the losses for ε small enough, which contradicts optimality.

(b) With the Cobb-Douglas technology, the final goods producer's expenditure on sector m is $\nu_m Y_t$. Using the inverse demand (13), producer j's revenue is

$$p_{m,j,t} X_{m,j,t} = \nu_m Y_t \frac{a}{a+b}, \quad \text{where} \quad a \equiv \kappa_{m,j} X_{m,j,t}^{\rho_m} \quad \text{and} \quad b \equiv \sum_{i \neq j} \kappa_{m,i} X_{m,i,t}^{\rho_m} > 0.$$

The ratio $a/(a+b)$ is increasing and strictly concave in a, and a is increasing and concave in $X_{m,j,t}$, so revenue is strictly concave in $X_{m,j,t}$. The producer's cost, $(\tau_{m,j,t} + w_t/\chi_{m,j,t})X_{m,j,t} + w_t\Omega_{m,j,t}$, is linear in $X_{m,j,t}$, so profit is strictly concave. Marginal revenue is unbounded as $X_{m,j,t}$ goes to zero, so a best response is positive and satisfies the first-order condition, and by strict concavity any output that satisfies the first-order condition is the best response. □

B.3 The General Model: Intermediate Sector m Producer j's Problem

I derive the first-order condition of sector m producer j's problem with respect to the choice of variable labor input $N^{var}_{m,j,t}$. First, rewrite the per-period profit maximization problem for sector m producer j by substituting its price function from equation (13):

$$\max_{N_{m,j,t},N^{var}_{m,j,t}} \kappa_{m,j} p_{m,t} X_{m,j,t}^{1-\frac{1}{\eta_m}} X_{m,t}^{\frac{1}{\eta_m}} - w_t N_{m,j,t} \quad \text{s.t. } p_{m,t} = \nu_m \frac{Y_t}{X_{m,t}},$$
$$N_{m,j,t} = N^{var}_{m,j,t} + \Omega_{m,j,t}, \text{ and}$$
$$X_{m,j,t} = F_{m,j,t}(N^{var}_{m,j,t}).$$

As outlined in Assumption 5, the sector m producer j internalizes its influence over the intermediate sector m composite and price but not the final output and the general price level. Thus, the sector m producer j's first-order condition with respect to the choice of variable labor input $N^{var}_{m,j,t}$ is:

$$\max_{N_{m,j,t},N^{var}_{m,j,t}} \kappa_{m,j} p_{m,t} X_{m,j,t}^{1-\frac{1}{\eta_m}} X_{m,t}^{\frac{1}{\eta_m}} - w_t N_{m,j,t} \quad \text{s.t. } p_{m,t} = \nu_m \frac{Y_t}{X_{m,t}},$$
$$N_{m,j,t} = N^{var}_{m,j,t} + \Omega_{m,j,t}, \text{ and}$$
$$X_{m,j,t} = F_{m,j,t}(N^{var}_{m,j,t}).$$

As outlined in Assumption 5, the sector m producer j internalizes its influence over the intermediate sector m composite and price but not the final output and the general price level. Thus, the sector m producer j's first-order condition with respect to the choice of variable labor input $N^{var}_{m,j,t}$ is:

$$\kappa_{m,j} \frac{dX_{m,j,t}}{dN^{\text{var}}_{m,j,t}} \left[\frac{\partial p_{m,t}}{\partial X_{m,t}} \frac{\partial X_{m,t}}{\partial X_{m,j,t}} X_{m,j,t}^{1-\frac{1}{\eta_m}} X_{m,t}^{\frac{1}{\eta_m}} + \frac{1}{\eta_m} p_{m,t} X_{m,t}^{\frac{1}{\eta_m}-1} \frac{\partial X_{m,t}}{\partial X_{m,j,t}} X_{m,j,t}^{1-\frac{1}{\eta_m}} \right.$$
$$\left. + \left(\frac{\eta_m - 1}{\eta_m}\right) p_{m,t} X_{m,t}^{\frac{1}{\eta_m}} X_{m,j,t}^{-\frac{1}{\eta_m}} \right] = w_t,$$

$$\kappa_{m,j} p_{m,t} X_{m,t}^{\frac{1}{\eta_m}} X_{m,j,t}^{-\frac{1}{\eta_m}} \cdot \frac{dX_{m,j,t}}{dN^{\text{var}}_{m,j,t}} \left[\frac{\partial p_{m,t}}{\partial X_{m,t}} \frac{p_{m,j,t} X_{m,j,t}}{p^2_{m,t}} + \frac{1}{\eta_m} \frac{p_{m,j,t} X_{m,j,t}}{p_{m,t} X_{m,t}} + \left(\frac{\eta_m - 1}{\eta_m}\right) \right] = w_t,$$

$$p_{m,j,t}F'_{m,j,t}(N^{\text{var}}_{m,j,t})\left[\frac{\partial p_{m,t}}{\partial X_{m,t}}\frac{s_{m,j,t}X_{m,t}}{p_{m,t}}+\frac{1}{\eta_m}s_{m,j,t}+\left(\frac{\eta_m-1}{\eta_m}\right)\right]=w_t,$$

$$\left[\frac{\partial p_{m,t}}{\partial X_{m,t}}\frac{X_{m,t}}{p_{m,t}}s_{m,j,t}+\frac{1}{\eta_m}s_{m,j,t}+\left(\frac{\eta_m-1}{\eta_m}\right)\right]^{-1}=\frac{p_{m,j,t}}{w_t/F'_{m,j,t}(N^{\text{var}}_{m,j,t})}.$$

The right-hand side of the last equation line is price over marginal cost for sector m producer j, which is the producer j's markup. The markup expression can be rewritten as $\frac{\sigma(s_{m,j,t})-1}{\sigma(s_{m,j,t})}$, where $\sigma(s_{m,j,t})$ is the price elasticity of sector m producer j's demand. Thus,

$$\frac{\sigma(s_{m,j,t})}{\sigma(s_{m,j,t})-1}=\left[\frac{\partial p_{m,t}}{\partial X_{m,t}}\frac{X_{m,t}}{p_{m,t}}s_{m,j,t}+\frac{1}{\eta_m}s_{m,j,t}+\left(\frac{\eta_m-1}{\eta_m}\right)\right]^{-1},$$

which implies that:

$$\sigma(s_{m,j,t})=\left[\frac{1}{\eta_m}(1-s_{m,j,t})+\frac{X_{m,t}/p_{m,t}}{|\partial X_{m,t}/\partial p_{m,t}|}s_{m,j,t}\right]^{-1},$$

where $\frac{X_{m,t}/p_{m,t}}{|\partial X_{m,t}/\partial p_{m,t}|}$ is the absolute value of the inverse price elasticity of sector m composite demand.

B.4 Planning Problem

Building upon the multi-sector neoclassical growth model with multiple intermediate production sectors, I consider the case that the accumulated emissions from carbon-intensive intermediate input production result in external damages, which affect the final good production possibilities. The social planner will choose allocations taking this negative externality into account when maximizing aggregate welfare subject to technology, feasibility, and carbon cycle conditions:

$$\max_{\left\{C_t,K_{t+1},N_t,N_{0,t},\left\{X_{m,t},E_{m,t},\left\{N_{m,j,t},N^{\text{var}}_{m,j,t},X_{m,j,t},E^f_{m,j,t}\right\}_{j=1}^{J_m}\right\}_{m=1}^{M},E^f_t,S_t,Y_t\right\}_{t=0}^{\infty}}\sum_{t=0}^{\infty}\beta^t\log(C_t)$$

$$\begin{aligned}
\text{subject to } \forall t,\quad C_t &= Y_t+(1-\delta)K_t-K_{t+1},\\
Y_t &= \exp[-\gamma_t(S_t-\bar{S})]A_{0,t}K_t^{\alpha}(N_{0,t})^{1-\alpha-\sum_{m=1}^{M}\nu_m}\prod_{m=1}^{M}X_{m,t}^{\nu_m},\\
X_{m,t} &= \left[\sum_{j=1}^{J_m}\kappa_{m,j}X_{m,j,t}^{\frac{\eta_m-1}{\eta_m}}\right]^{\frac{\eta_m}{\eta_m-1}},\quad \forall m,\\
X_{m,j,t} &= \chi_{m,j,t}N^{\text{var}}_{m,j,t},\quad \forall m,j,\\
N_{m,j,t} &= N^{\text{var}}_{m,j,t}+\Omega_{m,j,t},\quad \forall m,j,\\
E_{m,j,t} &= \theta_{m,j,t}X_{m,j,t}^{\zeta_{m,j}},\quad \forall m,j,
\end{aligned}$$

$$
\begin{aligned}
E_{m,t} &= \sum_{j=1}^{J_m} E_{m,j,t}, \quad \forall m, \\
E_t &= \sum_{m=1}^{M} E_{m,t}, \\
A_{0,t+1} &= (1+g_A)A_{0,t}, \\
\chi_{m,j,t+1} &= (1+g_{\chi_{m,j}})\chi_{m,j,t}, \quad \forall m,j, \\
C_t, K_{t+1}, X_t, X_{j,t}, X_{m,j,t} &\geq 0, \quad \forall m,j, \\
N_t &= N_{0,t} + \sum_{m=1}^{M}\sum_{j=1}^{J_m} N_{m,j,t}, \\
S_t &= \sum_{s=0}^{t+T}(1-d_s)E_{t-s},
\end{aligned}
$$

Definition. The optimal allocations $\left\{C_t, K_{t+1}, N_{0,t}, \left\{X_{m,t}, E_{m,t}, \{N_{m,j,t}, N^{\text{var}}_{m,j,t}, X_{m,j,t}, E_{m,j,t}, \text{SCC}_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}, E_t, S_t, S_{1,t}, S_{2,t}, Y_t\right\}_{t=0}^{\infty}$ that solve the planning problem satisfy the following equations for every period t:

$$
\frac{1}{C_t} = \beta \frac{1}{C_{t+1}} \left(\alpha \frac{Y_{t+1}}{K_{t+1}}\right), \tag{26}
$$

$$
\chi_{m,j,t}\left[\frac{\nu_m \kappa_{m,j}}{X_{m,t}^{\frac{\eta_m-1}{\eta_m}} X_{m,j,t}^{\frac{1}{\eta_m}}} - \theta_{m,j,t}\zeta_{m,j}\frac{\text{SCC}_{m,j,t}}{Y_t}X_{m,j,t}^{\zeta_{m,j}-1}\right] = \frac{1-\alpha-\sum_{m=1}^{M}\nu_m}{N_{0,t}}, \tag{27}
$$

$$
C_t = Y_t + K_{t+1}, \tag{28}
$$

$$
Y_t = \exp[-\gamma_t(S_t - \bar{S})]A_{0,t}K_t^{\alpha}(N_{0,t})^{1-\alpha-\sum_{m=1}^{M}\nu_m}\prod_{m=1}^{M}X_{m,t}^{\nu_m},
$$

$$
X_{m,j,t} = \chi_{m,j,t}N^{var}_{m,j,t}, \ \forall m,j
$$

$$
N_{m,j,t} = N^{var}_{m,j,t} + \Omega_{m,j,t}, \ \forall m,j
$$

$$
X_{m,t} = \left[\sum_{j=1}^{J_m}\kappa_{m,j}X_{m,j,t}^{\frac{\eta_m-1}{\eta_m}}\right]^{\frac{\eta_m}{\eta_m-1}}, \ \forall m
$$

$$
N_t = N_{0,t} + \sum_{m=1}^{M}\sum_{j=1}^{J_m}N_{m,j,t},
$$

$$
E_{m,j,t} = \theta_{m,j,t}X_{m,j,t}^{\zeta_{m,j}}, \ \forall m,j
$$

$$E_{m,t} = \sum_{j=1}^{J_m} E_{m,j,t}, \ \forall m$$

$$E_t = \sum_{m=1}^{M} E_{m,t},$$

$$S_t = S_{1,t} + S_{2,t},$$

$$S_{1,t} = S_{1,t-1} + \psi_L E_t,$$

$$S_{2,t} = (1-\psi) S_{2,t-1} + \psi_0 (1-\psi_L) E_t,$$

$$\text{SCC}_{m,j,t} = \frac{\partial E_t}{\partial E_{m,t}} \frac{\partial E_{m,t}}{\partial E_{m,j,t}} \sum_{s=0}^{\infty} \beta^s C_t \frac{Y_{t+s}}{C_{t+s}} \gamma_{t+s} (1-d_s), \tag{29}$$

Conditions (26) and (28) are satisfied if and only if the saving rate is constant at $\alpha\beta$. Then by Proposition 1 from Golosov et al. (2014), the optimal SCC_t expression in equation (29) simplifies to $\Lambda_{m,j,t}$. Moreover, the ratio $\text{SCC}_{m,j,t}/Y_t$ that appears in equation (27) is constant at level $\hat{\Lambda}_{m,j,t}$ as shown in equation (12).

B.5 Decentralized Equilibrium

B.5.1 Competitive Equilibrium

Definition. In a competitive equilibrium, given the exogenous paths of aggregate labor supply $\{N_t\}_{t=0}^{\infty}$, taxes $\left\{\left\{\{\tau_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}\right\}_{t=0}^{\infty}$, and fixed costs $\left\{\left\{\{\Omega_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}\right\}_{t=0}^{\infty}$, the allocations $\left\{C_t, K_{t+1}, Y_t, N_{0,t}, \left\{X_{m,t}, E_{m,t}, \{N_{m,j,t}, N^{\text{var}}_{m,j,t}, X_{m,j,t}, E_{m,j,t}, \text{SCC}_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}, E_t, S_t, S_{1,t}, S_{2,t}\right\}_{t=0}^{\infty}$, prices $\left\{q_t, w_t, r_t, \left\{p_{m,t}, \{p_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}\right\}_{t=0}^{\infty}$, transfers $\{T_t\}_{t=0}^{\infty}$, and profits $\{\Pi_0, \Pi_{m,j}, \Pi\}$ that satisfy the set of conditions for every period t:

$$\chi_{m,j,t} \left[\frac{\nu_m \kappa_{m,j}}{X_{m,t}^{\frac{\eta_m - 1}{\eta_m}} X_{m,j,t}^{\frac{1}{\eta_m}}} - \frac{\tau_{m,j,t}}{Y_t} \right] = \frac{1 - \alpha - \sum_{m=1}^{M} \nu_m}{N_{0,t}}, \ \forall m, j,$$

$$X_{m,t} = \left[\sum_{j=1}^{J_m} \kappa_{m,j} X_{m,j,t}^{\frac{\eta_m - 1}{\eta_m}} \right]^{\frac{\eta_m}{\eta_m - 1}}, \ \forall m,$$

$$p_{m,j,t} = \kappa_{m,j} p_{m,t} X_{m,t}^{\frac{1}{\eta_m}} X_{m,j,t}^{-\frac{1}{\eta_m}}, \ \forall m, j,$$

$$\frac{1}{C_t} = \beta \frac{1}{C_{t+1}} r_{t+1},$$

$$C_t = Y_t + K_{t+1},$$

$$Y_t = \exp[-\gamma_t(S_t - \bar{S})]A_{0,t}K_t^{\alpha}(N_{0,t})^{1-\alpha-\sum_{m=1}^{M}\nu_m}\prod_{m=1}^{M}X_{m,t}^{\nu_m},$$

$$X_{m,j,t} = \chi_{m,j,t}N_{m,j,t}^{var},\ \forall m, j,$$

$$N_{m,j,t} = N_{m,j,t}^{var} + \Omega_{m,j,t},\ \forall m, j,$$

$$E_{m,j,t} = \theta_{m,j,t}X_{m,j,t}^{\zeta_{m,j}},\ \forall m, j,$$

$$N_t = N_{0,t} + \sum_{m=1}^{M}\sum_{j=1}^{J_m}N_{m,j,t},$$

$$E_{m,t} = \sum_{j=1}^{J_m}E_{m,j,t},\ \forall m,$$

$$E_t = \sum_{m=1}^{M}E_{m,t},$$

$$S_t = S_{1,t} + S_{2,t},$$

$$S_{1,t} = S_{1,t-1} + \psi_L E_t,$$

$$S_{2,t} = (1-\psi)S_{2,t-1} + \psi_0(1-\psi_L)E_t,$$

$$q_t = \beta^t\left(C_t/C_0\right)^{-1},$$

$$r_t = \alpha Y_t/K_t,$$

$$w_t = \left(1 - \alpha - \sum_{m=1}^{M}\nu_m\right)Y_t/N_{0,t},$$

$$p_{m,t} = \nu_m Y_t/X_{m,t},$$

$$\sum_{t=0}^{\infty}q_t T_t = \sum_{t=0}^{\infty}q_t\sum_{m=1}^{M}\sum_{j=1}^{J_m}\tau_{m,j,t}X_{m,j,t},$$

$$\Pi_0 = \sum_{t=0}^{\infty}q_t\left[Y_t - r_t K_t - w_t N_{0,t} - \sum_{m=1}^{M}p_{m,t}X_{m,t}\right],$$

$$\Pi_{m,j} = \sum_{t=0}^{\infty}q_t\left[(p_{m,j,t} - \tau_{m,j,t})X_{m,j,t} - w_t N_{m,j,t}\right],$$

$$\Pi = \Pi_0 + \sum_{m=1}^{M}\sum_{j=1}^{J_m}\Pi_{m,j}.$$

B.5.2 Equilibrium with Oligopolistic Intermediate Sectors

Definition. In an equilibrium with oligopolistically competitive intermediate good producers, given the exogenous paths of labor supply $\{N_t\}_{t=0}^{\infty}$, taxes $\left\{\left\{\{\tau_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}\right\}_{t=0}^{\infty}$, and fixed costs $\left\{\left\{\{\Omega_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}\right\}_{t=0}^{\infty}$ the allocations $\left\{C_t, K_{t+1}, Y_t, N_{0,t}, \left\{X_{m,t}, E_{m,t}, \{N_{m,j,t}, N_{m,j,t}^{\text{var}}, X_{m,j,t}, E_{m,j,t}, s_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}, E_t, S_t, S_{1,t}, S_{2,t}\right\}_{t=0}^{\infty}$, prices $\left\{q_t, w_t, r_t, \left\{p_{m,t}, \{p_{m,j,t}\}_{j=1}^{J_m}\right\}_{m=1}^{M}\right\}_{t=0}^{\infty}$, transfers $\{T_t\}_{t=0}^{\infty}$, and profits $\{\Pi_0, \Pi_{m,j}, \Pi\}$ that satisfy the set of conditions for every period t:

$$\chi_{m,j,t}\left[\frac{\nu_m \kappa_{m,j}}{X_{m,t}^{\frac{\eta_m-1}{\eta_m}} X_{m,j,t}^{\frac{1}{\eta_m}}}\frac{\sigma(s_{m,j,t})-1}{\sigma(s_{m,j,t})} - \frac{\tau_{m,j,t}}{Y_t}\right] = \frac{1-\alpha-\sum_{m=1}^{M}\nu_m}{N_{0,t}}, \forall m, j,$$

$$X_{m,t} = \left[\sum_{j=1}^{J_m} \kappa_{m,j} X_{m,j,t}^{\frac{\eta_m-1}{\eta_m}}\right]^{\frac{\eta_m}{\eta_m-1}}, \forall m,$$

$$p_{m,j,t} = \kappa_{m,j} p_{m,t} X_{m,t}^{\frac{1}{\eta_m}} X_{m,j,t}^{-\frac{1}{\eta_m}}, \forall m, j,$$

$$s_{m,j,t} = \frac{p_{m,j,t} X_{m,j,t}}{p_{m,t} X_{m,t}}, \forall m, j,$$

$$\frac{1}{C_t} = \beta \frac{1}{C_{t+1}} r_{t+1},$$

$$C_t = Y_t + K_{t+1},$$

$$Y_t = \exp[-\gamma_t(S_t - \bar{S})] A_{0,t} K_t^{\alpha} (N_{0,t})^{1-\alpha-\sum_{m=1}^{M}\nu_m} \prod_{m=1}^{M} X_{m,t}^{\nu_m},$$

$$X_{m,j,t} = \chi_{m,j,t} N_{m,j,t}^{var}, \forall m, j,$$

$$N_{m,j,t} = N_{m,j,t}^{var} + \Omega_{m,j,t}, \forall m, j,$$

$$E_{m,j,t} = \theta_{m,j,t} X_{m,j,t}^{\zeta_{m,j}}, \forall m, j,$$

$$N_t = N_{0,t} + \sum_{m=1}^{M}\sum_{j=1}^{J_m} N_{m,j,t},$$

$$E_{m,t} = \sum_{j=1}^{J_m} E_{m,j,t}, \forall m,$$

$$E_t = \sum_{m=1}^{M} E_{m,t},$$

$$S_t = S_{1,t} + S_{2,t},$$

$$S_{1,t} = S_{1,t-1} + \psi_L E_t,$$

$$S_{2,t} = (1-\psi)S_{2,t-1} + \psi_0(1-\psi_L)E_t,$$

$$q_t = \beta^t \left(C_t/C_0\right)^{-1},$$

$$r_t = \alpha Y_t / K_t,$$

$$w_t = \left(1 - \alpha - \sum_{m=1}^{M} \nu_m\right) Y_t / N_{0,t},$$

$$p_{m,t} = \nu_m Y_t / X_{m,t}, \; \forall m,$$

$$\sum_{t=0}^{\infty} q_t T_t = \sum_{t=0}^{\infty} q_t \sum_{m=1}^{M} \sum_{j=1}^{J_m} \tau_{m,j,t} X_{m,j,t},$$

$$\Pi_0 = \sum_{t=0}^{\infty} q_t \left[Y_t - r_t K_t - w_t N_{0,t} - \sum_{m=1}^{N} p_{m,t} X_{m,t} \right],$$

$$\Pi_{m,j} = \sum_{t=0}^{\infty} q_t \left[(p_{m,j,t} - \tau_{m,j,t}) X_{m,j,t} - w_t N_{m,j,t} \right], \; \forall m,$$

$$\Pi = \Pi_0 + \sum_{m=1}^{M} \sum_{j=1}^{J_m} \Pi_{m,j}.$$

B.6 Capital in the Intermediate Sectors

Suppose intermediate producer j in sector m produces $X_{m,j,t} = \chi_{m,j,t} K_{m,j,t}^{a_m} N_{m,j,t}^{1-a_m}$ with capital rented each period, so that $K_t = K_{0,t} + \sum_m \sum_j K_{m,j,t}$, and keep all other assumptions of Section 2. The planner equates the marginal product of capital in final goods production, $R_t \equiv \alpha Y_t / K_{0,t}$, to its marginal value in every intermediate use, net of marginal damage:

$$\left(\frac{\partial Y_t}{\partial X_{m,j,t}} - g'_{m,j}(X_{m,j,t}) \Lambda_{m,j,t} \right) a_m \frac{X_{m,j,t}}{K_{m,j,t}} = R_t.$$

Multiplying by $K_{m,j,t}$, summing over producers and sectors, and using $\sum_j \frac{\partial Y_t}{\partial X_{m,j,t}} X_{m,j,t} = \nu_m Y_t$ gives $R_t K_t = \kappa_t Y_t$ with

$$\kappa_t \equiv \alpha + \sum_{m=1}^{M} a_m \left(\nu_m - \frac{\sum_j g'_{m,j}(X_{m,j,t}) \Lambda_{m,j,t} X_{m,j,t}}{Y_t} \right).$$

The term in parentheses is the cost of production in sector m as a share of output; it is nonnegative and at most ν_m, so $\alpha \leq \kappa_t \leq \alpha + \sum_m a_m \nu_m$. With logarithmic utility and full depreciation, the Euler equation in terms of the saving rate $s_t \equiv K_{t+1}/Y_t$ is

$$\frac{s_t}{1 - s_t} = \frac{\beta \kappa_{t+1}}{1 - s_{t+1}},$$

so the saving rate is constant if and only if κ_t is. Solving forward, $\frac{1}{1-s_t} = \sum_{i=0}^{\infty} \beta^i \prod_{k=1}^{i} \kappa_{t+k}$, and therefore

$$\alpha\beta \;\leq\; s_t \;\leq\; \beta\Big(\alpha + \sum_{m=1}^{M} a_m \nu_m\Big).$$

Without capital in the intermediate sectors, $a_m = 0$ and $s_t = \alpha\beta$. In the calibration $\sum_m \nu_m = 0.0099$ and $\beta = 0.86$, so for any capital shares the saving rate lies between 0.258 and 0.266, and the ratio of the social cost of carbon to output, which depends on the saving rate only through the terms $(1 - s_t)/(1 - s_{t+j})$, is within 1.2 percent of the constant in equation (12).

B.7 Scope: Product Markets, Vertical Relations, and Existing Regulation

Product markets within a sector. Suppose the goods of sector m are grouped into product markets, with elasticity of substitution η_m^w between goods within the same market and $\eta_m^a < \eta_m^w$ across markets (superscripts w and a for within and across). Equation (16) holds as stated, with the elasticity of demand under quantity competition given by

$$\frac{1}{\sigma_{m,j,t}} = \frac{1}{\eta_m^w}\,(1 - s_{m,j,t}) + \frac{1}{\eta_m^a}\,s_{m,j,t}\,(1 - S_{m,j,t}) + s_{m,j,t} S_{m,j,t},$$

where $s_{m,j,t}$ is the producer's sales share in its product market and $S_{m,j,t}$ is the share of that market in the sector. A producer's share of its product market is larger than its share of the sector, so markups differ more across producers than in the baseline. The average correction is still determined by the average markup, which the calibration matches, but the sector-wide tax then loses more relative to the firm-specific tax. The conclusion that the correction can be made at the sector level therefore depends on markets being as unconcentrated as the national four-digit data indicate; it is least secure in cement, where markets are regional.

Vertical relations. In the model, carbon-intensive goods are sold to a competitive final goods sector. If the buyers also have market power and resell at a markup, output is restricted twice. Let $M_{m,j,t} \geq 1$ be the product of the gross markups along the chain and $P_{m,j,t}$ the price paid by final users. The tax on the carbon-intensive producer that attains the first-best allocation is then

$$\tau^*_{m,j,t} = g'_{m,j}(X^*_{m,j,t})\Lambda_{m,j,t} - P^*_{m,j,t}\,\frac{M_{m,j,t} - 1}{M_{m,j,t}},$$

which is equation (16) with the cumulative markup in place of the producer's own. The calibration uses producers' own markups and therefore understates the correction in sectors whose buyers have market power. Because each producer has a single margin, the allocation does not depend on whether the tax is levied on the producer or on its buyers.

Existing regulation. Equation (16) gives the total net tax per unit of output, so existing carbon prices and other taxes on these goods count toward it, and the optimal additional tax is correspondingly lower. A binding standard on emissions per unit of output lowers marginal emissions, and with them the Pigouvian term, and leaves the output subsidy unchanged, so the optimal tax

is a smaller share of the Pigouvian tax than in the baseline. Under a binding cap on the sector's emissions, a tax changes the price of permits and not emissions, and the correction for market power has to be made through the allocation of permits, for example in proportion to output.

B.8 Solution Algorithm

B.8.1 The Planning Problem

Given the full capital depreciation assumption and the constant SCC-to-output ratio result, I obtain a numerical solution to the planner's problem defined in Appendix B.4, pinned down by the first-order condition with respect to the variable labor input used in sector m intermediate input j production, $N^{var}_{m,j,t}$, as formulated in equation (17). Unlike Golosov et al.'s (2014) solution method, there is no closed-form solution for the optimal paths when I generalize the carbon emissions function to be nonlinear in output. Given the functional assumptions made, the optimal paths for each sector $m = 1, \ldots, M$ intermediate input $j = 1, \ldots, J_m$ production simultaneously satisfy:

$$\chi_{m,j,t}\left[\frac{\nu_m \kappa_{m,j}}{X_{m,t}^{\frac{\eta_m - 1}{\eta_m}} X_{m,j,t}^{\frac{1}{\eta_m}}} - \theta_{m,j,t}\zeta_{m,j}\frac{\text{SCC}_{m,j,t}}{Y_t}X_{m,j,t}^{\zeta_{m,j}-1}\right] = \frac{1-\alpha-\sum_{m=1}^{M}\nu_m}{N_{0,t}}, \tag{30}$$

where $\text{SCC}_{m,j,t}$ is the social cost of carbon from producer j in sector m, and $\hat{\Lambda}_{m,j,t} \equiv \text{SCC}_{m,j,t}/Y_t$ is a constant function of the exogenous parameters, as shown in equation (12). The optimal paths of each intermediate sector $m = 1, \ldots, M$ good $j = 1, \ldots, J_m$ production satisfies:

$$X_{m,j,t} = \xi^*_{m,j,t} X_{m,t}^{-(\eta_m - 1)}, \tag{31}$$

where $\xi^*_{m,j,t}$ for the efficient allocations are given by:

$$\xi^*_{m,j,t} \equiv \left(\frac{\nu_m \kappa_{m,j}\chi_{m,j,t}N_{0,t}}{1-\alpha-\sum_{m=1}^{M}\nu_m + \chi_{m,j,t}\hat{\Lambda}_{m,j,t}N_{0,t}\theta_{m,j,t}\zeta_{m,j}X_{m,j,t}^{\zeta_{m,j}-1}}\right)^{\eta_m},$$

where $\theta_{m,j,t}\zeta_m X_{m,j,t}^{\zeta_{m,j}-1}$ is the marginal carbon emissions from production of intermediate input j in sector m at time t. Even though $\xi^*_{m,j,t}$ is a function of the endogenous path of firm-specific output $X_{m,j,t}$, I obtain a numerical solution for the optimal paths of intermediate input production $X_{m,j,t}$, as follows.

First, I guess the paths of total labor demand $N_{m,j,t}$ of each producer j in sector m and the labor demand $N_{0,t}$ of the final goods sector. Specifically, I guess a constant path of final goods sector labor demand $N_{0,t}$ at 70 percent of the exogenous aggregate labor supply level $\bar{N}$, and a constant path of total labor demand $N_{m,j,t}$ of each producer j in sector m at $(1-0.7)\cdot\bar{N}/(M\cdot\max\{J_1,\ldots,J_M\})$. Then, I calculate the firm-specific $\xi^*_{m,j,t}$ as a function of the guessed $X_{m,j,t}$ for each producer j in sector m implied by the guessed values of $N_{m,j,t}$ and $N_{0,t}$, and the exogenous parameters at time t.

Next, I obtain sector-level output $X_{m,t}$ by aggregating the firm-specific outputs $X_{m,j,t}$ implied

by equation (31) across all producers j in sector m and period t by using the firm-specific $\xi^*_{m,j,t}$ obtained in the last step. Using the sector-level output $X_{m,t}$, I obtain updated firm-specific outputs $X_{m,j,t}$ from equation (31), the variable labor inputs $N^{var}_{m,j,t}$ they require, and total labor inputs $N_{m,j,t}$ by adding the fixed costs $\Omega_{m,j,t}$. The updated labor input of the final goods sector $N_{0,t}$ is the remaining labor supply.

Once I obtain the updated paths of labor inputs for every producer-sector-time combination, I compare the updated paths of labor inputs $N_{m,j,t}$ and $N_{0,t}$ with the guessed paths. If the paths are not sufficiently close to each other, I set the new guess to a convex combination that puts weight 0.1 on the updated paths and 0.9 on the guessed paths, and repeat the process until convergence.

B.8.2 Decentralized Equilibrium

In a decentralized equilibrium, the intermediate producer j in sector m chooses its path of variable labor input $N^{var}_{m,j,t}$ to maximize its profits, given the path of taxes $\{\tau_{m,j,t}\}_{t=0}^{\infty}$, which are set exogenously by the government. Notably, the intermediate producers do not internalize their influences over the final goods production possibilities through the effect of their accumulated emissions on climate change. First, I describe the competitive equilibrium, as implemented in Golosov et al. (2014), and then consider the case of equilibrium with oligopolistic intermediate sectors.

The Competitive Equilibrium.

In a competitive equilibrium, each producer is a price taker. Specifically, intermediate good producers do not internalize the pricing power they have through their control over their own and sectoral outputs. Appendix B.5.1 gives the complete competitive equilibrium definition.

Producer j in sector m chooses its variable labor input $N^{var}_{m,j,t}$ to satisfy

$$\chi_{m,j,t}\left[\frac{\nu_m \kappa_{m,j}}{X_{m,t}^{\frac{\eta_m-1}{\eta_m}} X_{m,j,t}^{\frac{1}{\eta_m}}} - \hat{\tau}_{m,j,t}\right] = \frac{1-\alpha-\sum_{m=1}^{M}\nu_m}{N_{0,t}}, \tag{32}$$

where $\hat{\tau}_{m,j,t} \equiv \tau_{m,j,t}/Y_t$ is the per-unit output tax on the producer relative to aggregate output, and $\hat{\tau}_{m,j,t} = 0$ in the unregulated equilibrium. This condition differs from equation (30) in that the tax replaces the social cost of the producer's emissions. It implies

$$X_{m,j,t} = \xi^{CE}_{m,j,t} X_{m,t}^{-(\eta_m-1)}, \qquad \xi^{CE}_{m,j,t} \equiv \left(\frac{\nu_m \kappa_{m,j} \chi_{m,j,t} N_{0,t}}{1-\alpha-\sum_{m=1}^{M}\nu_m + \chi_{m,j,t} N_{0,t} \hat{\tau}_{m,j,t}}\right)^{\eta_m}.$$

Unlike $\xi^*_{m,j,t}$, $\xi^{CE}_{m,j,t}$ does not depend on the producer's output, so given $N_{0,t}$ sectoral output has the closed form $X_{m,t}^{\eta_m-1} = \sum_{j=1}^{J_m} \kappa_{m,j} (\xi^{CE}_{m,j,t})^{(\eta_m-1)/\eta_m}$. Because $N_{0,t}$ is endogenous, I solve for the equilibrium with the algorithm of Appendix B.8.1, with $\xi^{CE}_{m,j,t}$ in place of $\xi^*_{m,j,t}$.

The Oligopoly Equilibrium.

In an oligopoly equilibrium, each intermediate producer j in sector m internalizes its influence over the sectoral composite and composite price but not the final output and the general price level.

The definition of the complete oligopoly equilibrium is available in Appendix B.5.2.

Each firm takes into account that its output moves its own price $p_{m,j,t}$, given by equation (13), and the sectoral price $p_{m,t} = \nu_m Y_t / X_{m,t}$, through its effect on $X_{m,t}$. Its first-order condition with respect to variable labor input is

$$\chi_{m,j,t}\left[\frac{\nu_m \kappa_{m,j}}{X_{m,t}^{\frac{\eta_m-1}{\eta_m}} X_{m,j,t}^{\frac{1}{\eta_m}}}\frac{\sigma_m(s_{m,j,t})-1}{\sigma_m(s_{m,j,t})} - \hat{\tau}_{m,j,t}\right] = \frac{1-\alpha-\sum_{m=1}^{M}\nu_m}{N_{0,t}}, \tag{33}$$

where $\sigma_m(s_{m,j,t})$ is the price elasticity of demand in equation (14). It differs from equation (32) by the inverse markup $(\sigma_m(s_{m,j,t})-1)/\sigma_m(s_{m,j,t})$, which depends on the firm's sales share and hence on its output. The condition implies

$$X_{m,j,t} = \xi^O_{m,j,t} X_{m,t}^{-(\eta_m-1)}, \qquad \xi^O_{m,j,t} \equiv \left(\frac{\nu_m \kappa_{m,j} \chi_{m,j,t} N_{0,t}}{1-\alpha-\sum_{m=1}^{M}\nu_m + \chi_{m,j,t} N_{0,t} \hat{\tau}_{m,j,t}} \frac{\sigma_m(s_{m,j,t})-1}{\sigma_m(s_{m,j,t})}\right)^{\eta_m}.$$

I solve for the equilibrium with the algorithm of Appendix B.8.1, with $\xi^O_{m,j,t}$ in place of $\xi^*_{m,j,t}$. At each iteration, the sales shares are computed from the guessed outputs as $s_{m,j,t} = \kappa_{m,j}(X_{m,j,t}/X_{m,t})^{(\eta_m-1)/\eta_m}$, which follows from equation (13).

Numerical settings.

The three problems, and the variants with exogenous taxes used in Section 5, use the same starting values and stopping rule. The initial guess assigns 70 percent of labor to final goods production and divides the rest equally across producers. Each new guess puts weight 0.1 on the updated allocation and 0.9 on the previous guess. Iteration stops when the sum of absolute differences between the guessed and updated labor shares falls below 10^{-4}, both for the final goods sector over all periods and for the intermediate producers over all periods, sectors, and producers. The cap is 1,000 iterations. In the main runs the algorithm converges in at most 115 iterations.

Appendix C Calibration Details

C.1 Model Fit

The solved model reproduces every calibration target: each sector's share of value added, its average markup, the revenue shares of its 4, 8, 20, and 50 largest firms and their HHI (up to the cap on the number of firms in cement), and its 2017 emissions, along with aggregate output, world emissions, and the social cost of carbon. Table 5 reports the targeted moments and a set of moments that are not targeted.

The distribution of markups across firms is not targeted. In the model a firm's markup rises with its revenue share, so the largest firm has the highest markup in each sector. The model's markup for the largest firm is close to that of the largest S&P Global Market Intelligence (2023) firm in petroleum and 0.02 to 0.09 above it in the other four sectors, both measured as gross output

Table 5: Model fit, 2017

	Coal	Cement	Fertilizer	Paper	Petroleum
A. Targeted (model / data)					
Emissions (MtC per year)	10.4 / 10.4	18.0 / 18.0	10.2 / 10.2	10.8 / 10.8	56.9 / 56.9
Revenue share, 4 largest firms	0.34 / 0.34	0.13 / 0.13	0.31 / 0.31	0.45 / 0.45	0.42 / 0.42
Revenue share, 8 largest firms	0.53 / 0.53	0.19 / 0.19	0.52 / 0.52	0.64 / 0.64	0.63 / 0.63
Revenue share, 20 largest firms	0.77 / 0.77	0.32 / 0.32	0.69 / 0.69	0.83 / 0.83	0.86 / 0.86
Revenue share, 50 largest firms	0.92 / 0.92	0.47 / 0.45	0.82 / 0.82	0.94 / 0.94	0.95 / 0.95
HHI, 50 largest firms	455 / 455	75 / 74	395 / 395	661 / 661	618 / 618
Sales-weighted markup	1.79 / 1.79	1.38 / 1.38	2.81 / 2.81	1.73 / 1.73	1.38 / 1.38
B. Not targeted (model / data)					
Employment (thousands)	112 / 48	192 / 170	49 / 32	154 / 96	836 / 108
Wage bill / revenue	0.56 / 0.21	0.72 / 0.36	0.36 / 0.17	0.58 / 0.30	0.72 / 0.10
Markup of the largest firm (gross output)	1.37 / 1.28	1.09 / 1.06	1.20 / 1.13	1.18 / 1.16	1.09 / 1.09
Std. dev. of log markups (gross output)	0.01 / 0.13	0.00 / 0.08	0.00 / 0.16	0.01 / 0.05	0.00 / 0.09

Note: Each cell reports the value in the solved oligopoly equilibrium without policy in 2017 and its counterpart in the data. Aggregate targets, also reproduced, are US output ($21.6 trillion per year), world greenhouse gas emissions (14.2 GtC per year in CO_2 equivalent), and the social cost of carbon ($50 per ton of CO_2). Emissions are from the US Environmental Protection Agency (2024); revenue shares, the HHI, and employment are from the 2017 Economic Census; the data counterpart of the wage bill share is payroll over value added. Markups of individual firms are compared on a gross output basis, because the value added conversion is unstable for firms whose markup times the sector's intermediate share is close to one: the data are gross output markups of S&P Global Market Intelligence (2023) firms in the same four-digit sector, and the model's markups are converted to gross output by inverting the relation in Section 3.1.3. The largest firm is the largest by sales, and the standard deviation is across firms with a gross output markup of at least one. The cement sector has three such firms.

markups. The model does not reproduce the dispersion of markups across firms, which is about ten to ninety times larger in the data. That dispersion is unrelated to firm size: in a regression of the log gross output markup on log sales with sector-by-year effects, for Compustat firms in the five sectors from 2010 to 2023, the coefficient is -0.001 with a standard error of 0.006. A model in which markups depend on market shares cannot generate it. To the extent that it reflects true differences in market power across firms of similar size, the firm-specific optimal tax varies more across firms than the model implies, and a sector-wide tax loses more.

C.2 Employment and the Scaling of Carbon Intensity

Table 6 reports two implications of abstracting from capital in the intermediate sectors. First, the model allocates more workers to these sectors than the US Census Bureau (2023) reports. When labor is the only input, the wage bill is a much larger share of sector revenue in the model (0.56, 0.72, 0.36, 0.58, and 0.72 for the five sectors) than payroll is of value added in the 2017 Economic Census (0.21, 0.36, 0.17, 0.30, and 0.10). The gap is largest in petroleum, where payroll is a tenth of value added. In the model a sector's employment depends only on its output share and its markups, not on the level of labor productivity, so this gap cannot be closed without adding capital to the intermediate technology. The five sectors employ 0.9 percent of the labor force in the model, against 0.3 percent in the data, so labor in final goods production is understated by about 0.6 percent. The optimal tax as a share of the Pigouvian tax is not affected by the composition of costs: by equation

Table 6: Carbon intensity scale factors and model employment, 2017

Sector	θ_m scale factor	Census employment	Model workers	Ratio
Coal (2121)	0.258	48,466	112,362	2.32
Cement (3273)	0.825	170,336	191,907	1.13
Fertilizer (3253)	0.269	32,377	49,488	1.53
Paper (3221)	0.353	95,798	154,134	1.61
Petroleum (3241)	0.086	107,931	836,119	7.75

Note: The scale factor is the ratio of observed to modeled sector emissions in 2017 before scaling. "Census employment" is total employment by four-digit NAICS sector from the 2017 Economic Census. "Model workers" is labor allocated to the sector in the solved 2017 oligopoly equilibrium, $\pi_{m,0}N$.

(36) it depends on the elasticity of demand and on marginal damage as a share of the price, and the latter is determined by each sector's revenue, markups, and emissions, which the calibration matches. For the same reason the level of labor productivity $\bar{\chi}_m$, which in the data reflects capital per worker as well as efficiency, only sets the units of sector output. Its growth rate does matter for the paths of output and emissions, and the measured growth of labor productivity includes capital deepening. Second, and for the same reason, modeled output measured in units of 2017 value added exceeds observed value added. I therefore scale each sector's carbon intensity θ_m by the factor in the first column, so that modeled emissions equal observed emissions in every sector in 2017. Table 5 reports the targeted moments and a set of moments that are not targeted.

C.3 Sectoral Markups & Economic Profit Rates

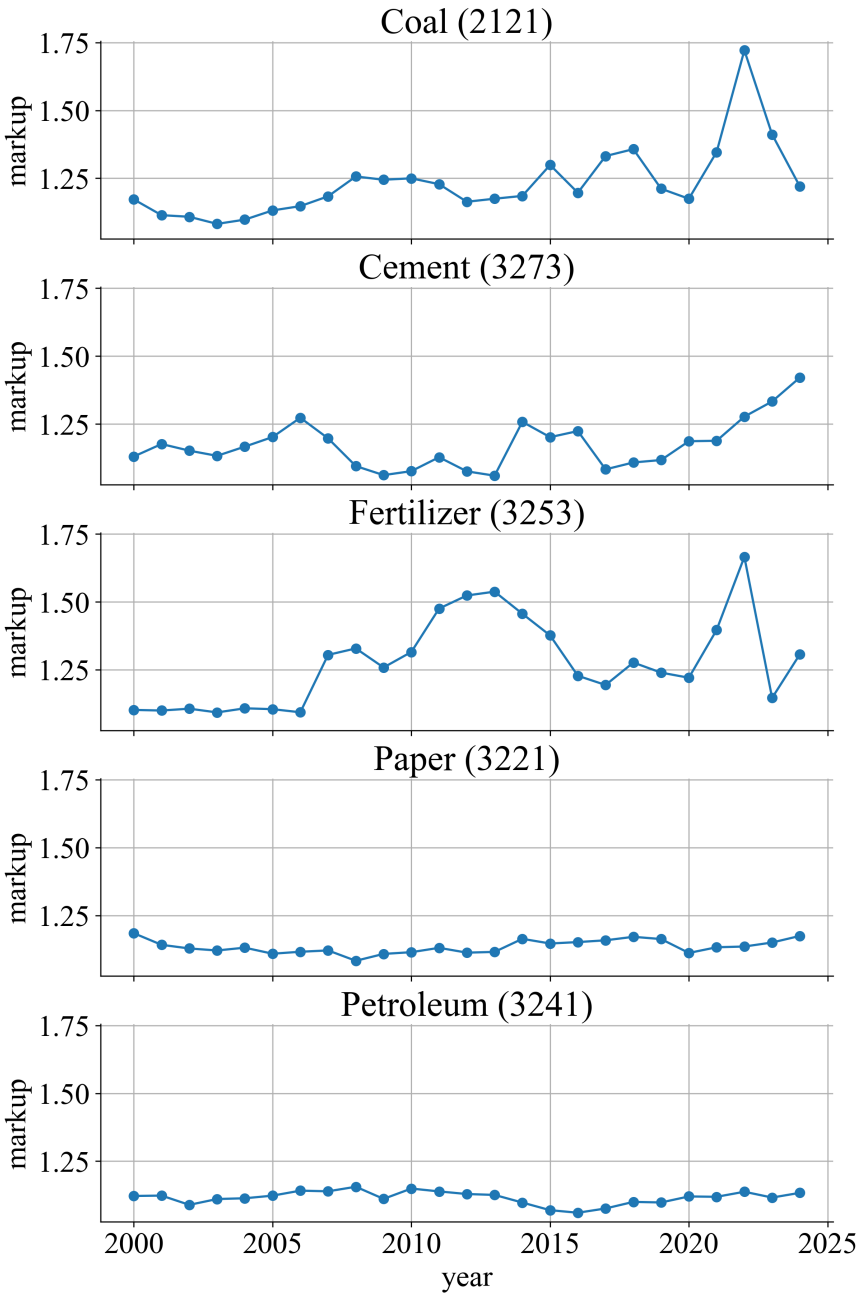
I estimate the sectoral markups and calculate economic profit rates as described in Appendix A. I present the sales-weighted average markups for the top five most carbon-intensive sectors in Figure 9. I use the 2017 average sectoral markups as the base year for the calibration of the model. The average markups for the top five most carbon-intensive sectors consistently exceed one, indicating that these sectors have pricing power and can set prices above their marginal costs.

Moreover, although not directly used in calibration, I also present the sales-weighted average economic profit rates for the top five most carbon-intensive sectors in Figure 10.

Although majority of the sectors historially have positive economic profit rates, the sectoral economic profit rates for some sectors are negative in some years. Notably, the coal mining sector (2121) had negative economic profit rates from 2011 to 2015, and most sharply in 2015, as coal demand from the electric power sector fell with the shale-gas-driven decline in natural gas prices and export demand for metallurgical coal weakened; Alpha Natural Resources, Arch Coal, and Peabody Energy filed for bankruptcy between August 2015 and April 2016 (Forbes 2015; The Wall Street Journal 2016; Reuters 2016), and coal production in 2016 was the lowest since 1978 (US Energy Information Administration 2017). The cement sector (3273) had negative economic profit rates from 2008 to 2013, after the financial crisis slowed housing and infrastructure construction, and again in 2016 (US Geological Survey 2016).

Except the few outliers, the average economic profit rates for the top five most carbon-intensive

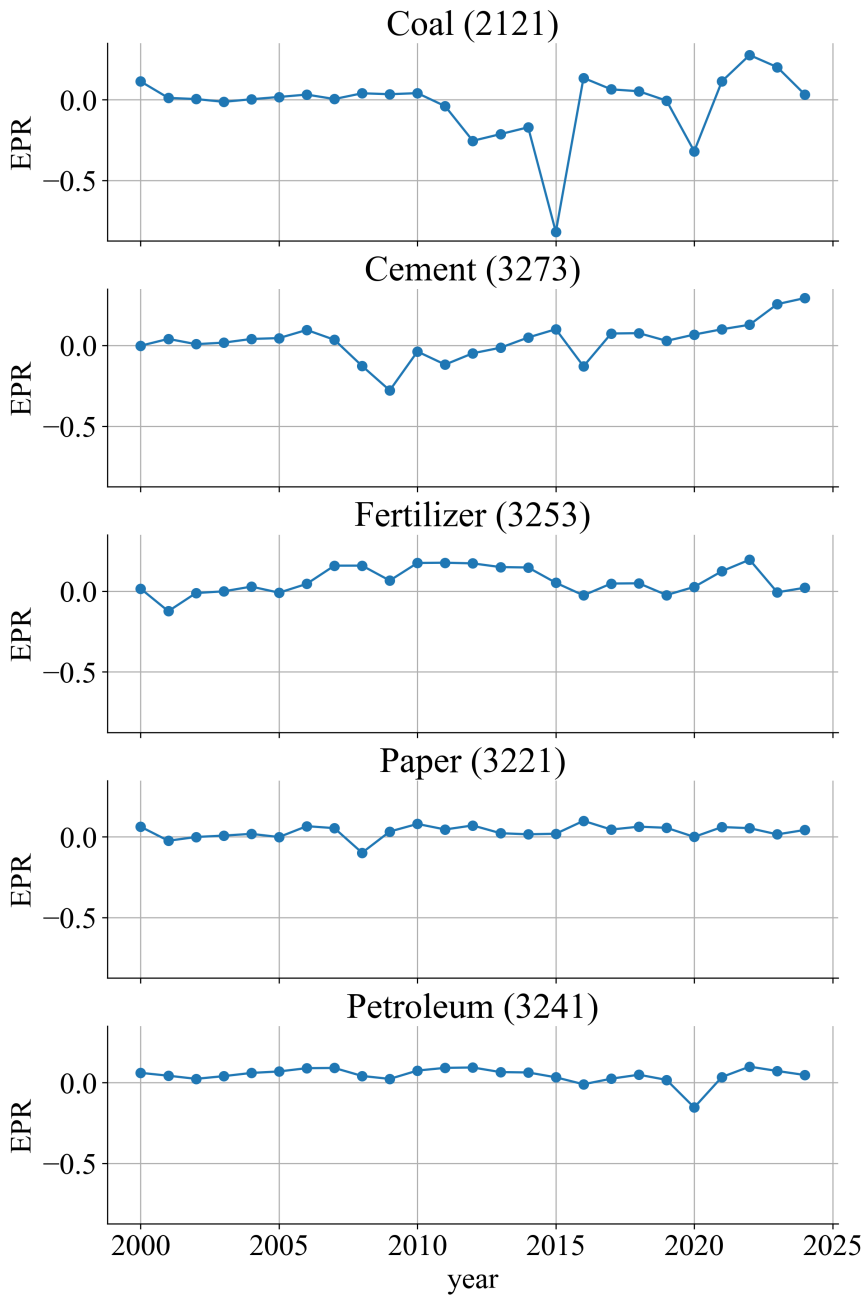
Figure 9: History of sectoral sales-weighted average markups of the top five most carbon-intensive sectors



Note: Sales-weighted average gross output markups of Compustat firms with a markup of at least one, by four-digit NAICS sector. Sectors are ordered by their 2017 carbon emissions per dollar of gross output.

sectors are positive, indicating that these sectors have been able to generate excess returns above their cost of capital and lack of competition in these sectors product markets.

Figure 10: History of sectoral sales-weighted average economic profit rates of the top five most carbon-intensive sectors



C.4 Firm Productivities Consistent with Observed Revenue Shares

Fix a sector m and drop the subscripts m and t. Without policy, producer j prices at $p_j = \mu_j w/\chi_j$, with $\mu_j = \sigma(s_j)/(\sigma(s_j)-1)$ and $\sigma(\cdot)$ given by equation (14). Because the composite enters final goods production with a Cobb–Douglas share, the sectoral demand elasticity is one, so μ_j depends only on s_j and η.

Lemma 2 *Let $s_1, \ldots, s_J$ be positive shares that sum to one, let $\eta > 1$, and let $\kappa_j = \chi_j / \sum_k \chi_k$. If, for all producers i and j,*

$$\frac{\chi_j}{\chi_i} = \left(\frac{s_j \, \mu_j^{\eta-1}}{s_i \, \mu_i^{\eta-1}} \right)^{\frac{1}{2\eta-1}}, \qquad \mu_j = \frac{\sigma(s_j)}{\sigma(s_j) - 1}, \tag{34}$$

then the shares $s_1, \ldots, s_J$, with markups $\mu_1, \ldots, \mu_J$, are an equilibrium of the sector without policy. Conversely, in any equilibrium with shares s_j, productivities satisfy (34).

Proof. From the demand (13), $X_j = \kappa_j^\eta (p_j/p_m)^{-\eta} X_m$, so $s_j = p_j X_j / (p_m X_m) = \kappa_j^\eta (p_j/p_m)^{1-\eta}$. Substituting $p_j = \mu_j w / \chi_j$ and $\kappa_j \propto \chi_j$, the terms common to the sector cancel from the ratio of two shares:

$$\frac{s_j}{s_i} = \left(\frac{\chi_j}{\chi_i} \right)^{2\eta-1} \left(\frac{\mu_j}{\mu_i} \right)^{1-\eta}.$$

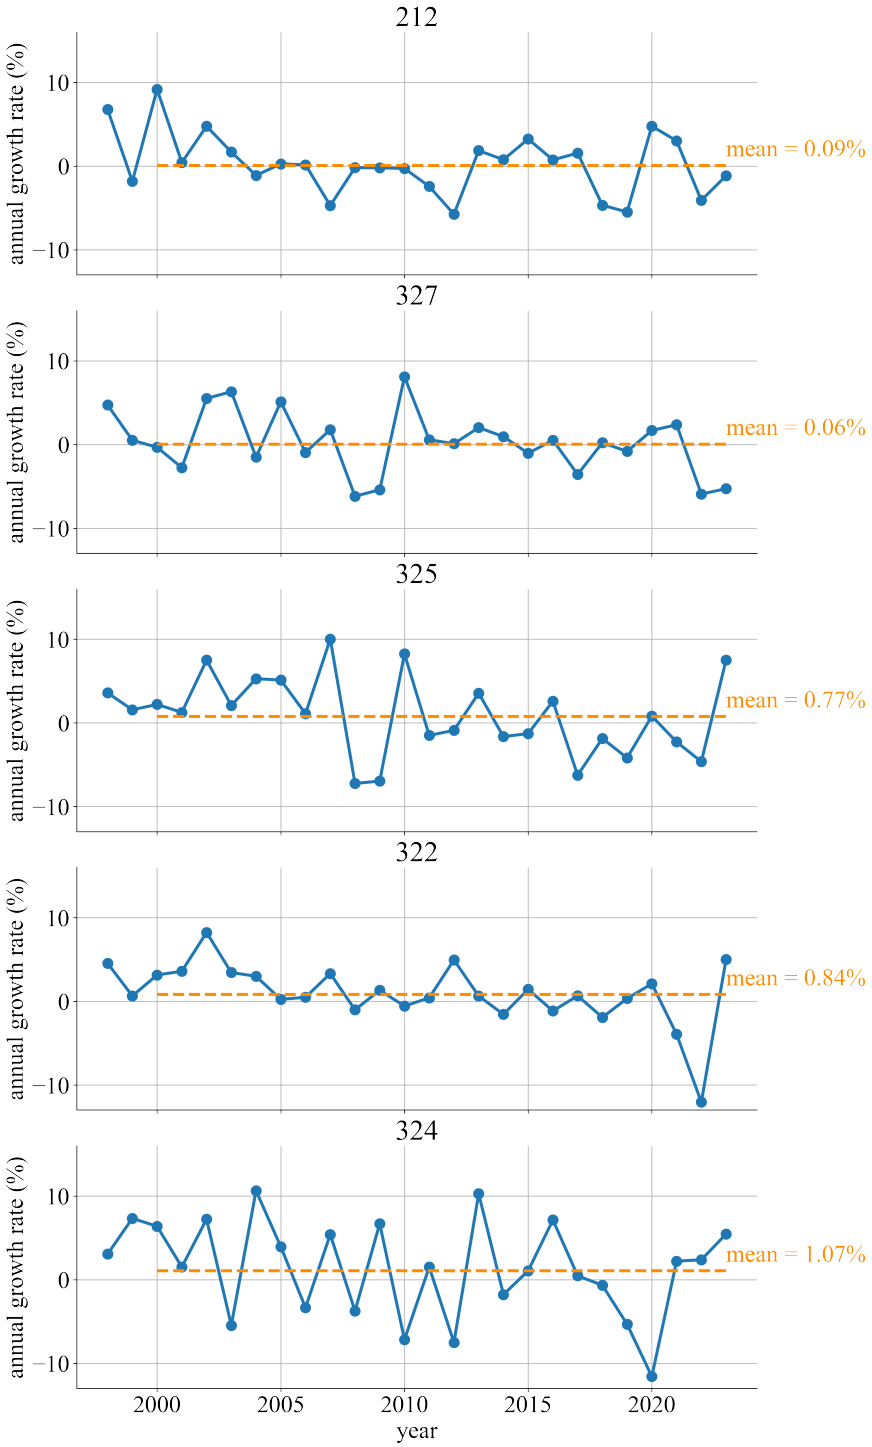
If (34) holds, then when each producer charges $\mu_j w / \chi_j$ the resulting shares are proportional to s_j, and since both sum to one they equal s_j; the markups μ_j are then the ones that correspond to these shares, so every producer's pricing condition holds. Conversely, in an equilibrium with shares s_j the display holds for every pair, and solving it for χ_j / χ_i, which is possible because $2\eta - 1 > 0$, gives (34). □

The lemma does not establish that this equilibrium is unique. In the calibrated model, the solution algorithm of Appendix B.8, started from equal labor allocations, converges to it: the solved equilibrium reproduces the Census concentration ratios and HHI (Appendix C.1).

C.5 Sectoral Labor Productivity Growth

The blue dotted line in Figure 11 represents the annual labor productivity growth rate for each sector, while the orange dashed line represents the average annual sectoral labor productivity growth rate for years 2000-2023, which is used for the calibration of the model. Additionally the average annual sectoral labor productivity growth rates are printed in orange on the right side of each sector's subplot.

Figure 11: Historical sectoral labor productivity growth rates of the the three-digit Production Account (BEA) sectors that contain the five most carbon-intensive sectors



Note: The mapping between the four-digit NAICS codes of the top five most carbon-intensive sectors and the three-digit BEA codes and their industry descriptions is as follows: 1. 2121: 212, Coal extraction; 2. 3273: 327, Nonmetallic mineral products; 3. 3253: 325, Chemical products; 4. 3221: 3221, Paper products; 5. 3241: 3241, Petroleum and coal products.

C.6 Carbon Cycle and Climate Module

Given my model's different base year, I adjust the parameters of the carbon cycle and climate module to match the observed atmospheric carbon concentration and annual carbon emissions data from 1960 until 2020. First, I adjust the initial pre-industrial carbon concentration $\bar{S}$ to match the NOAA (2024) estimate of 280 parts per million before the start of the industrial revolution and convert it to gigatons of carbon (GtC) using the conversion factor of 2.13 GtC per ppm, which results in $\bar{S} = 596.4$ GtC. Second, I keep the share of permanent emissions ψ_L at 0.2 as reported by IPCC (2021), which is the same as in Golosov et al. (2014). However, I deviate from Golosov et al. (2014) when I estimate the decadal retained emission share ψ_0 and the decadal geometric emission decay rate ψ to match the path of atmospheric carbon concentration implied by linear equation system for the carbon cycle restated below in equation (35) and the path of observed atmospheric carbon concentration and the annual carbon emissions given the initial pre-industrial carbon concentration $\bar{S}$, using the following linear system of equations:

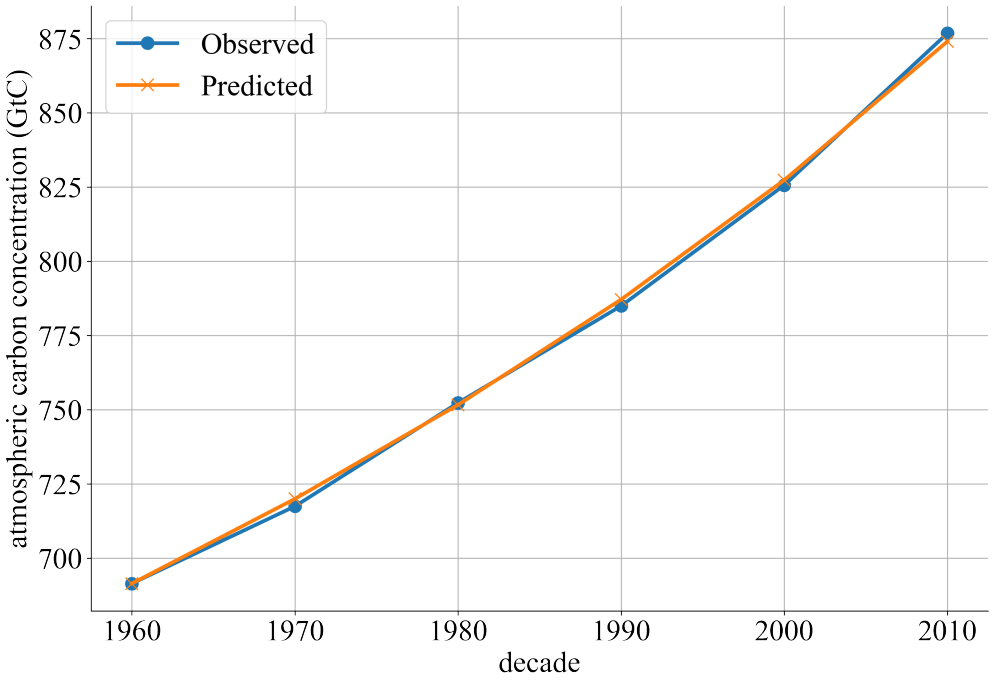
$$\begin{aligned} S_{1,t} &= S_{1,t-1} + \psi_L E_t, \\ S_{2,t} &= (1-\psi) S_{2,t-1} + \psi_0 (1-\psi_L) E_t, \\ S_t &= \bar{S} + S_{1,t} + S_{2,t}. \end{aligned} \tag{35}$$

I obtain the historical atmospheric carbon concentration data from the NOAA (2025a) dataset, which provides the annual atmospheric carbon concentration data from 1960 until 2020 in units of parts per million (ppm). I convert the atmospheric carbon concentration data from ppm to GtC using the conversion factor of 2.13 GtC per ppm. I obtain the annual carbon emissions data from the OWID (2023) dataset, which provides the annual carbon emissions data for all countries in the world from 1850 until 2020 in units of metric tons of CO_2 equivalent emissions. I convert the annual carbon emissions data from metric tons of CO_2 equivalent emissions to GtC by multiplying the reported emissions by $12/44/10^{-9}$.

I estimate the decadal retained emission share ψ_0 to be 0.1988 and the decadal geometric emission decay rate ψ to be 0.0445, which are the values that match the observed atmospheric carbon concentration from 1960 until 2020 given the initial pre-industrial carbon concentration $\bar{S}$ and the annual carbon emissions data from the OWID (2023) dataset. I calibrate the initial permanent carbon stock $S_{1,-1}$ by summing up annual emissions from 1850 to 2017 and multiplying this sum by ψ_L, which equals 174.876 GtC. I calculate the initial transitory carbon stock $S_{2,-1}$ by subtracting the initial permanent carbon stock $S_{1,-1}$ from the atmospheric carbon concentration in 2017, which is 866.399 GtC as reported by the NOAA (2025a) dataset, resulting in the initial transitory carbon stock to be 95.122 GtC. Figure 12 shows the calibration of the carbon cycle and climate module, where the blue line represents the observed atmospheric carbon concentration from 1960 until 2020, and the orange line represents the model-implied atmospheric carbon concentration given the calibrated parameters from 1960 until 2020.

I deviate from Golosov et al. (2014) when I calibrate the carbon cycle because my model has

Figure 12: Calibration of the carbon cycle and climate module



a different base year, and I aim to match the observed atmospheric carbon concentration and annual carbon emissions data rather than using reports from the IPCC. As Figure 12 shows, the model-implied atmospheric carbon concentration closely follows the observed atmospheric carbon concentration from 1960 until 2020, indicating that the calibrated parameters of the carbon cycle and climate module are consistent with the observed data.

C.7 Carbon Shares

Figure 13 shows the historical global share of US carbon emissions from 1960 until 2020, which is calculated as the ratio of US carbon emissions to the total global carbon emissions using data from OWID (2023). The red dashed horizontal line represents the average of this ratio from 2010 until 2023, which is 14.6 percent.

Figure 14 shows the share of US CO_2 emissions from the five sectors, the sectors' CO_2 emissions reported to the US Environmental Protection Agency (2024) divided by total US CO_2 emissions from OWID (2023). These five most carbon-intensive sectors, ranked by carbon emissions per dollar of gross output, are coal mining (NAICS 2121), cement manufacturing (NAICS 3273), fertilizer manufacturing (NAICS 3253), pulp and paper mills (NAICS 3221), and petroleum refineries (NAICS 3241). The red dashed horizontal line represents the average of this ratio from 2010 until 2023, which is 6.2 percent.

Over time, the share of global emissions from the US has been declining, reflecting the global shift of carbon emissions from developed to developing countries. However, the share of US emissions from these top five most carbon-intensive sectors in the US have been relatively stable over time

Figure 13: Historical global share of US carbon emissions

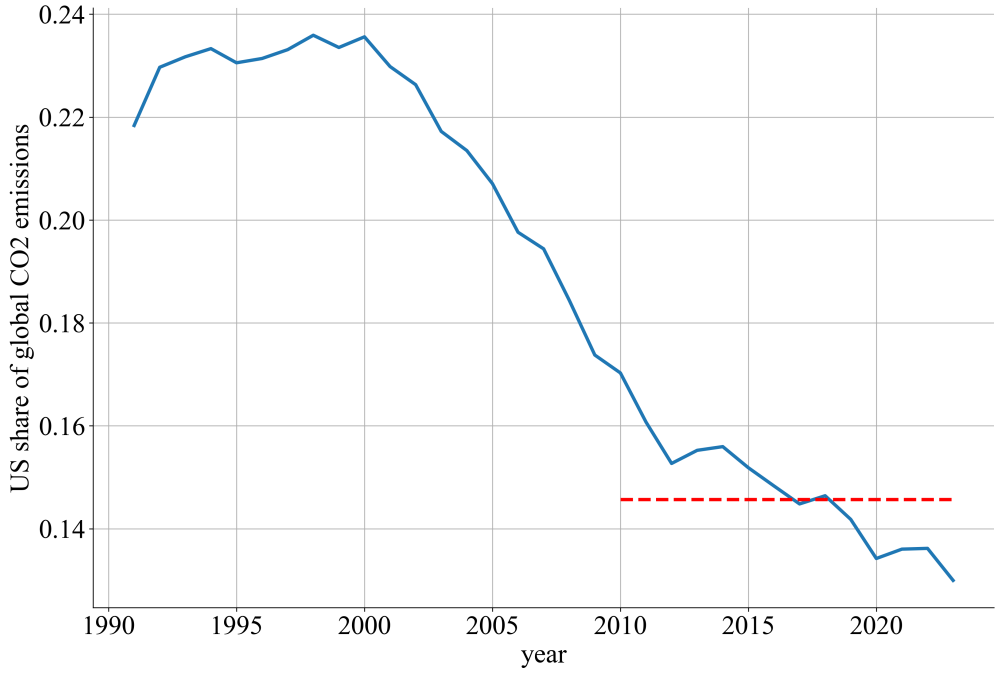
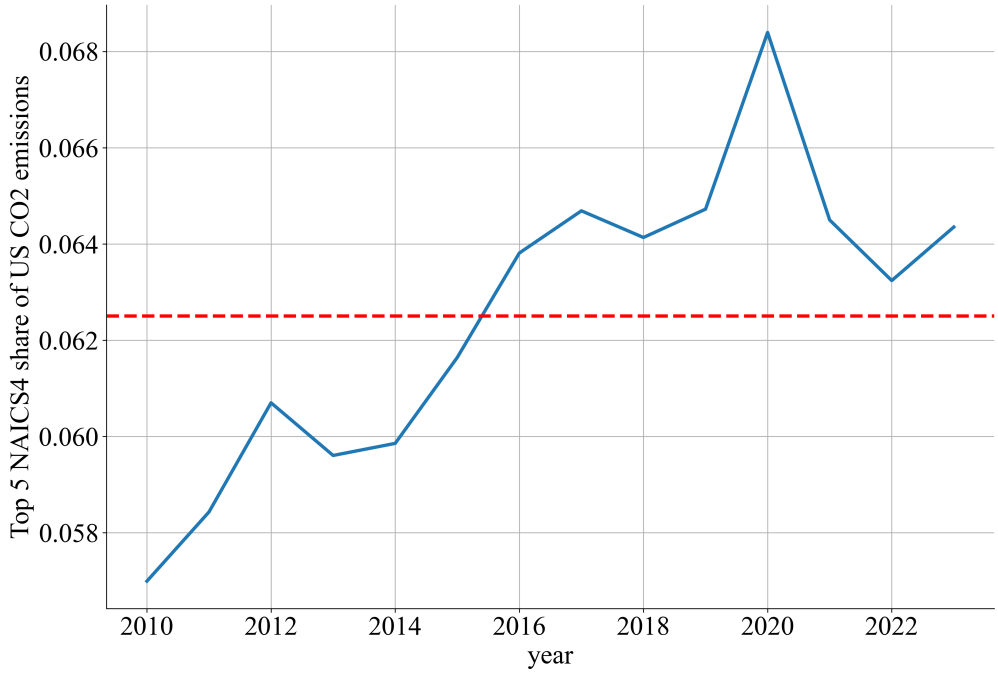


Figure 14: Historical US carbon emissions share of the top five most carbon-intensive sectors in 2017



with a slight upward trend, reflecting their continued importance in the US economy and their high carbon intensity. The two averages imply that the five sectors account for about 0.9 percent of world CO_2 emissions, a ratio of about 110, close to the scaling constant of 109 in the calibration with

CO_2 only (Appendix C.13); the baseline value of 133 instead relates the five sectors' greenhouse gas emissions to world greenhouse gas emissions, both in CO_2 equivalent.

C.8 Additional Results

Figure 15 displays optimal output taxes for the median productivity firm in five carbon-intensive sectors under perfect competition and oligopoly. Under perfect competition, the optimal tax equals the Pigouvian tax (marginal external damages per unit of output). Under oligopoly, optimal taxes are generally lower, reflecting firms' market power. Optimal policy varies by sector. Output taxes are measured in 2021 US dollars per dollar of value added, where one equals a 100 percent tax rate.

The optimal output taxes under perfect competition only reflect the climate externality correction. Thus, the optimal output taxes under perfect competition decrease as the carbon intensity of the sector decreases, reflecting the greater emissions and damages from a unit of value added in more carbon-intensive sectors. On the other hand, the optimal output taxes under oligopoly reflect both the climate externality correction and the output distortion correction. Thus there is no clear relationship between the optimal output taxes under oligopoly and the carbon intensity of the sector. The optimal output taxes under oligopoly are lower than those under perfect competition, reflecting the market power of firms. However, the difference between the oligopoly and perfect competition output taxes depends on the scale of the output distortion, which is not necessarily related to the carbon intensity of the sector.

Figure 16 shows the optimal output subsidy per ton of carbon for firms at the 10th, 25th, 50th, 75th, 90th, and 99th percentiles of the productivity distribution in each of the five sectors. The black dashed line shows the OSCC. The optimal firm-specific carbon tax equals the OSCC minus the firm-specific output subsidy, which is the vertical distance between the dashed line and the firm's subsidy. The subsidy stays below the OSCC for every firm in every period, so the optimal carbon tax is always positive.

C.9 Robustness to the Markup Concept — Gross Output Elasticities

In the baseline markup calibration, I convert the sales-weighted gross output markup to a value added markup, and calibrate η_m to the latter. In this section, I instead choose η_m so that the solved model's sales-weighted average markup equals the gross output markup. The revenue shares of firms, and so the concentration statistics, are the same as in the baseline; firm productivities are recomputed so that these shares remain the equilibrium shares, and θ_m, ϕ_E, $A_{0,0}$, and $\bar{\gamma}$ are recalibrated to match the same targets as in Section 3.1. In petroleum the gross output markup (1.08) is below what the observed concentration implies for any elasticity in the range the solution method can handle; I set $\eta_m = 20$ there, which gives an average markup of 1.12.

Table 7 reports η_m under the two markup concepts, and the ratio of the optimal output tax to the Pigouvian tax for the median productivity firm in 2017 under each. The optimal tax remains below the Pigouvian tax in every sector, but the gross output calibration produces optimal taxes that are closer to it, 6 to 23 percent below the Pigouvian tax, compared to 25 to 67 percent in

Table 7: Optimal-to-Pigouvian tax ratio under alternative markup concepts

	η_m		Optimal / Pigouvian tax	
Sector	Value added	Gross output	Value added	Gross output
Coal (2121)	2.42	4.72	0.55	0.77
Cement (3273)	3.71	14.74	0.71	0.92
Fertilizer (3253)	1.59	7.90	0.33	0.86
Paper (3221)	2.63	13.69	0.57	0.91
Petroleum (3241)	4.39	20.00	0.75	0.94

Note: Median-productivity firm in each sector, 2017. "Value added" is the baseline calibration. "Gross output" calibrates η_m to the gross output markup, with the same revenue shares; in petroleum η_m is set to 20.

the baseline. The oligopoly path without policy, which is anchored to observed emissions, and the optimal path are unchanged. The perfect competition counterfactual is closer to oligopoly: emissions are 1.2 times those under oligopoly instead of 1.6 times, and long-run warming in 2207 is 7.7°C instead of 8.7°C.

C.10 Welfare Losses: Horizon, Markup Concept, and Form of Competition

Table 8 repeats the policy experiments of Section 5 for alternative scoring horizons and for elasticities calibrated to gross output markups (Appendix C.9).

Scoring horizon.

With 20 periods the planner's allocation does not maximize the welfare measure: under perfect competition a tax 10 percent below the Pigouvian tax scores 0.01 percent higher. For the same reason, at that horizon taxes below the Pigouvian tax appear nearly costless under perfect competition, and the comparison of the two mistakes is reversed. With 40 and 60 periods, taxes 10 percent below and above the Pigouvian tax both score lower, and the losses under the policy scenarios change by less than 0.01 percentage points between the two horizons. The losses without policy keep rising with the horizon, because warming continues for as long as there is no policy.

Markup concept. With gross output markups the optimal tax is closer to the Pigouvian tax, so ignoring market power under oligopoly costs less, 0.014 percent. The sector-wide tax loses 0.005 percent under oligopoly and 0.007 percent under perfect competition. The firm-specific tax applied to competitive firms loses 0.044 percent, about six times the loss of the sector-wide tax. With higher elasticities a firm's markup depends mostly on its market share, so the corrections of the largest firms stand out: weighted by emissions under perfect competition, the firm-specific tax is 83 percent of the Pigouvian tax in paper and 85 percent in petroleum, against 91 and 94 percent for the sector-wide tax. All of these losses are small relative to the loss without policy.

Mismeasured markups. The optimal output tax is the Pigouvian tax minus the firm's price times its Lerner index, $(\mu_{m,j}-1)/\mu_{m,j}$. Table 9 reports the welfare loss when the regulator sets this tax with a mismeasured markup, holding the Pigouvian component and prices at their first-best values, and firms are oligopolistic. In panel A every firm's markup is over- or understated by the same number of percentage points. In panel B each firm's markup is measured with independent

Table 8: Welfare losses (CE, %): horizon and markup concept

		Value added markups			Gross output markups
Market Structure	Policy Scenario	20 periods	40 periods	60 periods	40 periods
Cournot Oligopoly	No Policy Intervention	16.207	24.485	27.646	24.586
	Uniform tax per ton	0.160	0.112	0.108	0.014
	Sector-wide tax per ton	0.003	0.001	0.001	0.005
Perfect Competition	No Policy Intervention	26.614	37.746	41.661	28.807
	Firm-specific tax per ton	0.080	0.160	0.167	0.044
	Sector-wide tax per ton	0.075	0.151	0.158	0.007

Table 9: Welfare losses (CE, %) when the regulator mismeasures markups; oligopoly

Error	Firm-specific tax	Sector-wide tax
A. Common error in all firms' markups (percentage points)		
A. Common error in all firms' markups (percentage points)		
−40	0.051	0.053
−20	0.016	0.020
−10	0.004	0.007
0	0.000	0.001
+10	0.002	0.002
+20	0.010	0.009
+40	0.039	0.038
B. Independent error in each firm's markup (standard deviation, percentage points)		
5	0.003	0.001
10	0.011	0.001
20	0.043	0.001

Note: Welfare is computed over 40 periods as in Table 4. Markups are bounded below by one. Panel B reports the mean over five draws of the errors.

noise; the sector-wide tax is not affected, on the assumption that the errors average out across a sector's firms. An error of 10 percentage points common to all firms costs less than 0.01 percent of consumption under either tax, and an error of 40 points costs less than ignoring market power altogether (0.112 percent). Firm-level noise is more costly for the firm-specific tax: with a standard deviation of 10 percentage points it loses 0.016 percent, against 0.001 percent for the sector-wide tax, and with 20 points it loses 0.058 percent, about half the cost of ignoring market power.

Mode of competition. I repeat the experiments assuming that firms compete in prices. The optimal output tax has the same form, with the Bertrand elasticity $\eta_m(1 - s_{m,j}) + s_{m,j}$ in place of the Cournot elasticity. At the calibrated concentration the two are close: the optimal tax for the median firm is 55, 72, 33, 57, and 75 percent of the Pigouvian tax under Bertrand competition, within half a percentage point of the Cournot values. Applying the tax designed for Cournot competition when firms compete in prices loses 0.001 percent of consumption, and the reverse mistake loses 0.001 percent.

Table 10: Market structure and markups, 2017 and 2022

		Coal	Cement	Fertilizer	Paper	Petroleum
Number of firms	2017	334	4,182	730	157	941
	2022	218	3,879	766	156	790
Revenue share of 4 largest firms (%)	2017	34.2	12.8	31.1	45.2	42.0
	2022	39.7	16.4	38.9	44.6	49.5
HHI of 50 largest firms	2017	455	74	395	661	618
	2022	523	101	531	665	764
Markup, gross output	2017	1.33	1.08	1.19	1.16	1.08
	2022	1.72	1.28	1.67	1.14	1.14
Markup, value added	2017	1.79	1.38	2.81	1.73	1.38
	2022	3.85	4.62	—	1.58	1.90
η_m: baseline (2017)		2.42	3.71	1.59	2.63	4.39
η_m: 2022 concentration, 2017 markups		2.45	3.73	1.60	2.64	4.66
η_m: 2022 concentration and markups (gross output)		2.59	4.82	2.75	18.93	20.00

Note: Number of firms, revenue shares, and HHI are from the 2017 and 2022 Economic Census, national, four-digit NAICS. Markups are sales-weighted averages over Compustat firms in the sector (in 2022: 6, 2, 8, 4, and 12 firms). Value added markups use the 2017 materials shares in both years. The Census flags 40–50% of cement receipts as imputed in 2022.

C.11 Base Year: Market Structure and Markups in 2022

The model is calibrated to 2017. I repeat the calibration of market structure with data for 2022, the only other year for which it can be done for all five sectors: the Economic Census did not publish concentration statistics for mining before 2017, and the facility-level emissions data begin in 2010. Table 10 compares the two years. Between 2017 and 2022 concentration rose in coal, cement, fertilizer, and petroleum, and was unchanged in paper. Sales-weighted markups were close to their 2017 values in paper and higher in the other four sectors. In coal and fertilizer the 2022 markups are the highest of the 2010–2023 period and fall back in 2023, to 1.41 and 1.15 on a gross output basis. In cement the sample has two to five firms in each year, and the 2017 markup (1.08 on a gross output basis) is lower than in any year from 2014 to 2023.

I consider two alternative calibrations. In the first, firms' revenue shares are fitted to the 2022 concentration statistics, and the elasticities η_m are chosen as in Section 3 so that the sales-weighted average markup in the solved equilibrium equals its 2017 target. This isolates the change in market structure. In the second, the markup targets are those of 2022 as well. In both, sectoral output shares, carbon intensities, average labor productivity, and the materials shares used to convert gross output markups into value added markups are held at their 2017 values, and the macroeconomic parameters are recalibrated to the same targets as in the baseline.

Table 11 reports the results. With the market structure of 2022 and the markups of 2017, the optimal tax relative to the Pigouvian tax changes by at most about one percentage point in every sector, and no welfare loss under a policy scenario changes by more than 0.01 percentage points. The increase in concentration between 2017 and 2022 is too small to matter: the HHI remains below 800 in every sector, and the elasticities that reproduce a given average markup barely change.

With the markups of 2022, the correction for market power is larger in the four sectors whose

Table 11: Welfare losses (CE, percent) and optimal taxes: base year

Market Structure	Policy Scenario	Baseline (2017)	2022 concentration, 2017 markups	Gross output markups (2017)	2022 concentration and markups (gross output)
Cournot Oligopoly	No Policy Intervention	24.485	24.492	24.586	24.557
	Uniform tax per ton	0.112	0.109	0.014	0.040
	Sector-wide tax per ton	0.001	0.001	0.005	0.007
Perfect Competition	No Policy Intervention	37.746	38.009	28.807	31.026
	Firm-specific tax per ton	0.160	0.155	0.044	0.100
	Sector-wide tax per ton	0.151	0.144	0.007	0.036
Optimal tax relative to the Pigouvian tax, 2017, weighted by emissions					
	Coal	0.55	0.55	0.74	0.57
	Cement	0.71	0.71	0.92	0.78
	Fertilizer	0.34	0.34	0.84	0.60
	Paper	0.56	0.56	0.83	0.76
	Petroleum	0.72	0.72	0.85	0.82

Note: Welfare losses are consumption-equivalent, relative to the first-best allocation, computed over 40 periods as in Table 4. The lower panel reports the firm-specific optimal output tax of oligopolistic firms divided by the Pigouvian tax, weighted by firms' emissions under perfect competition.

markups rose. With these markups the value added conversion is not defined in fertilizer, where the 2022 markup times the intermediate share exceeds one, so this calibration uses gross output markups and is compared with the gross output calibration of Appendix C.9. Weighted by emissions, the optimal tax falls from 74 to 57 percent of the Pigouvian tax in coal, from 92 to 78 percent in cement, from 84 to 60 percent in fertilizer, from 83 to 76 percent in paper, and from 85 to 82 percent in petroleum. Ignoring market power then costs 0.040 percent of consumption instead of 0.014 percent, and applying the tax designed for oligopolistic firms to perfectly competitive firms costs 0.100 percent instead of 0.044 percent. Both are well below one percent of the loss without policy, and the sector-wide tax loses 0.007 percent under oligopoly, against 0.005 percent in the gross output calibration. The loss without policy under perfect competition rises from 29 to 31 percent, because higher markups restrain more emissions.

The size of the correction for market power therefore depends on the year in which markups are measured, as it depends on the markup concept (Appendix C.9). The conclusions of Section 5 do not: the welfare cost of ignoring market power is small relative to the gain from pricing carbon, the correction can be made at the sector level, and both possible mistakes, ignoring market power when it exists and correcting for it when firms are competitive, stay small.

C.12 Markup Sample and Emitting Firms

The markup target of each sector is the sales-weighted average markup of the publicly listed US firms classified in that sector. Table 12 reports how much of each sector's emissions in the Greenhouse Gas Reporting Program these firms account for in 2017, splitting each facility's emissions across its parent companies by ownership share. The firms in the markup sample account for 39 percent of the five sectors' emissions: 47 percent in coal, 62 percent in fertilizer, 28 percent in paper, 48 percent in petroleum, and 1 percent in cement, where Monarch Cement is the only cement producer among the three sample firms and two thirds of emissions come from subsidiaries of firms listed abroad. Listed firms of any kind account for 73 percent of emissions and privately

Table 12: Share of sector emissions by type of parent company, 2017 (%)

	Coal	Cement	Fertilizer	Paper	Petroleum	All five
In the sector's markup sample	47	1	62	28	48	39
US-listed, not in the sample	16	8	5	27	18	15
Listed or incorporated abroad	0	66	10	0	12	18
Privately held or state-owned	22	8	10	27	7	11
Not classified	15	18	13	19	16	16

Note: Emissions are those of facilities reporting to the US Environmental Protection Agency (2024) with a primary NAICS code in the sector, divided among their parent companies by ownership share, as reported in the same source. Parents within the markup sample are identified for all facilities. The remaining parents are classified, from company filings, until the classified parents cover 80 percent of each sector's 2017 emissions; the rest are not classified. Firms whose shares traded over the counter are counted as listed in the United States.

held and state-owned firms for 11 percent, about a quarter in coal and paper; the remaining 16 percent is owned by parents that are not classified. The calibration also excludes firms with an estimated markup below one, which in 2017 removes two refiners and one paper producer that together account for 16 and 5 percent of their sectors' emissions; without this exclusion the value added markup target is 1.29 instead of 1.38 in petroleum and 1.68 instead of 1.73 in paper, and changes by 0.03 or less in the other sectors.

To assess what this implies for the results, I recalibrate the elasticities η_m to an alternative target: the average markup of the listed US firms that own each sector's emitting facilities, weighted by their emissions in the sector and including firms with markups below one. These are 8, 3, 4, 11, and 11 firms, which account for 47, 9, 67, 41, and 64 percent of sector emissions. Their gross output markups are higher than the baseline in coal (1.41 against 1.33), cement (1.34 against 1.08), and fertilizer (1.24 against 1.19), and lower in paper (1.13 against 1.16) and petroleum (1.04 against 1.08). Converted to value added, the cement markup would be 10.9, because the markup times the intermediate share is close to one, so this calibration uses gross output markups and is compared with the gross output calibration of Appendix C.9. The cement figure rests on three firms. Table 13 reports the results. The optimal tax for the median firm falls from 92 to 74 percent of the Pigouvian tax in cement and by four to five percentage points in coal and fertilizer, and changes by three points or less in paper and petroleum. Ignoring market power costs 0.023 percent of consumption instead of 0.014 percent, the sector-wide tax loses 0.005 percent under oligopoly, and applying the tax designed for oligopolistic firms to perfectly competitive firms costs 0.060 percent instead of 0.044 percent. As with the base year, the size of the correction depends on whose markups are measured, and the conclusions of Section 5 do not.

C.13 Carbon Dioxide Only

The baseline measures each sector's emissions as all greenhouse gases reported to the US Environmental Protection Agency (2024), in CO_2 equivalent. The world emissions target is measured the same way, and the carbon cycle is fitted to world greenhouse gas emissions in CO_2 equivalent (Appendix C.6). Methane, however, leaves the atmosphere within about a decade, with a perturba-

Table 13: Welfare losses (CE, %) and optimal taxes: markups of emitting firms

Market Structure	Policy Scenario	Baseline	Gross output markups	Emitters' markups (gross output)
Cournot Oligopoly	No policy intervention	24.485	24.586	24.570
	Uniform tax per ton (Pigouvian)	0.112	0.014	0.023
	Sector-wide tax per ton	0.001	0.005	0.005
Perfect Competition	No policy intervention	37.746	28.807	29.429
	Firm-specific tax per ton (designed for oligopoly)	0.160	0.044	0.060
	Sector-wide tax per ton	0.151	0.007	0.016
Optimal tax relative to the Pigouvian tax, 2017, weighted by emissions				
	Coal	0.55	0.74	0.70
	Cement	0.71	0.92	0.74
	Fertilizer	0.34	0.84	0.80
	Paper	0.56	0.83	0.82
	Petroleum	0.72	0.85	0.85

tion lifetime of 11.8 years (Forster et al. 2021, Table 7.15), so the baseline overstates the long-run damage from sectors whose emissions are largely methane. In 2017, CO_2 accounts for 1 percent of the reported emissions of coal mining, most of the rest being methane from underground mines, 78 percent in fertilizer, where most of the rest is nitrous oxide, 87 percent in paper, and more than 99 percent in cement and petroleum. I repeat the calibration with CO_2 emissions only. Everything else is as in the baseline: the carbon intensities are rescaled to match each sector's 2017 CO_2 emissions, the world emissions target becomes world CO_2 emissions, about 10 GtC per year, so that the scaling constant ϕ_E changes from 133 to 109, and the damage parameter is recalibrated to the same social cost of carbon. The carbon cycle is left as in the baseline, fitted to emissions in CO_2 equivalent, so this calibration is not fully consistent; it serves to show which results depend on counting methane.

Table 14 reports the results. The optimal tax relative to the Pigouvian tax changes by less than one percentage point in cement and petroleum and falls by about two percentage points in fertilizer and paper, where CO_2 is 78 and 87 percent of emissions. Marginal damage per unit of output rises in cement and petroleum, but it is already close to the whole first-best price, so by equation (36) the ratio barely moves. In coal mining, emissions of CO_2 alone are so small that marginal damage is a small share of the first-best price, and the optimal output tax under oligopoly is negative: a subsidy of 3.6 times the Pigouvian tax, as in the case of own emissions in Appendix C.14. Ignoring market power costs 0.113 percent of consumption, about as in the baseline. The sector-wide tax loses 0.059 percent under oligopoly instead of 0.001 percent, and applying the taxes designed for oligopolistic firms to perfectly competitive firms costs 0.163 percent with firm-specific taxes and 0.375 percent with sector-wide taxes. The cost of ignoring market power is therefore unchanged, but the sector-wide tax recovers only about half of it, and the optimal tax in coal mining depends on whether its methane emissions are counted.

C.14 Emissions Scaling

In the baseline calibration each ton of carbon emitted by the five sectors stands for $\phi_E = 133$ tons of world emissions, so that the model matches world emissions in 2017 (Section 3.1). Marginal damage is then a large share of the first-best price of carbon-intensive output, between 0.89 and

Table 14: Welfare losses (CE, %) and optimal taxes: CO_2 only

Market Structure	Policy Scenario	Baseline (all gases)	CO_2 only
Cournot Oligopoly	No policy intervention	24.485	18.442
	Uniform tax per ton (Pigouvian)	0.112	0.113
	Sector-wide tax per ton	0.001	0.059
Perfect Competition	No policy intervention	37.746	28.613
	Firm-specific tax per ton (designed for oligopoly)	0.160	0.163
	Sector-wide tax per ton	0.151	0.375
Optimal tax relative to the Pigouvian tax, 2017, weighted by emissions			
	Coal	0.55	−1.69
	Cement	0.71	0.71
	Fertilizer	0.34	0.34
	Paper	0.56	0.55
	Petroleum	0.72	0.71

Note: Welfare losses are consumption-equivalent, relative to the first-best allocation, computed over 40 periods as in Table 4. In the second column, sector emissions are the CO_2 emissions reported to the GHGRP; all other parameters are calibrated as in the baseline. The lower panel reports the firm-specific optimal output tax of oligopolistic firms divided by the Pigouvian tax, weighted by firms' emissions under perfect competition; a negative value is a subsidy.

0.95 for the median firm. Dividing equation (16) by the Pigouvian tax gives

$$\frac{\tau^*_{m,j,t}}{\frac{dE_{m,j,t}}{dX_{m,j,t}}\Lambda_{m,j,t}} = 1 - \frac{1}{\sigma_m(s_{m,j,t})} \Big/ \frac{\frac{dE_{m,j,t}}{dX_{m,j,t}}\Lambda_{m,j,t}}{p_{m,j,t}}, \tag{36}$$

so the optimal tax as a share of the Pigouvian tax depends on this damage share as well as on the elasticity of demand. Table 15 repeats the oligopoly experiments of Section 5 for lower values of ϕ_E, holding all other parameters at their baseline values. With $\phi_E = 1$ firms are taxed on their own emissions only. Emissions from outside the five sectors are then taken as given and not modeled in this exercise.

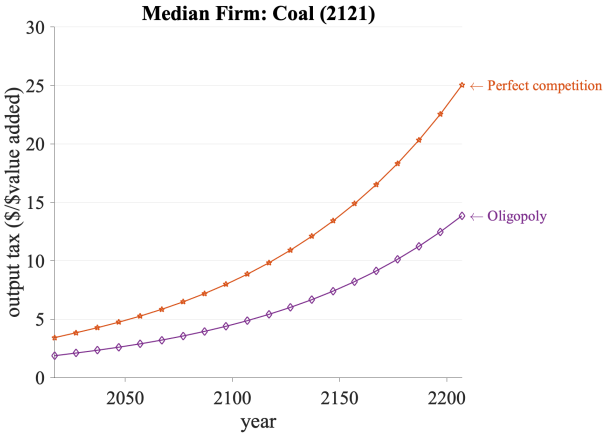
The cost of ignoring market power does not depend on the scaling: the Pigouvian tax loses 0.112 to 0.122 percent of consumption for every value of ϕ_E. What changes is the gain from pricing carbon, which falls from 24.5 percent of consumption to 0.06 percent as ϕ_E falls from 133 to 1. With $\phi_E = 10$ the optimal tax for the median firm ranges from −15 percent of the Pigouvian tax in fertilizer to 55 percent in cement, and the Pigouvian tax still captures 88 percent of the gain from the first-best policy. With $\phi_E = 1$ the output subsidy exceeds the Pigouvian tax in every sector, so the optimal output tax is negative, the first-best allocation has 36 percent more emissions than the equilibrium without policy, and the Pigouvian tax loses about twice as much as no policy. The sector-wide tax recovers more than 80 percent of the cost of ignoring market power for ϕ_E of 10 or more. With $\phi_E = 1$ it is a large output subsidy in every sector, and the solution algorithm does not return a valid allocation for it. For values of ϕ_E above the baseline, the damage share approaches one and, by equation (36), the optimal tax as a share of the Pigouvian tax approaches $1 - 1/\sigma_m(s_{m,j,t})$, which is at most about four percentage points above its baseline value in every sector.

Table 15: Emissions scaling: welfare losses and optimal taxes under oligopoly

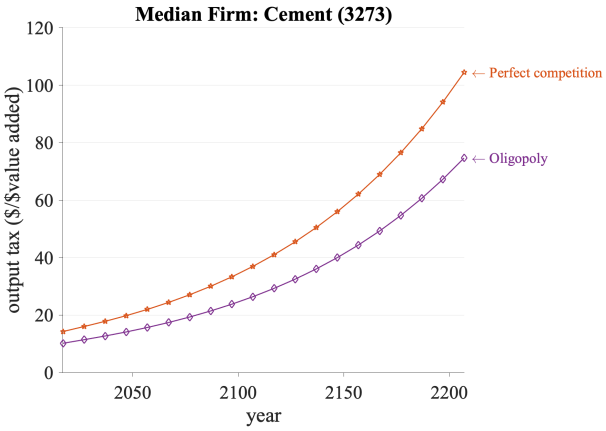
ϕ_E	133	30	10	1
A. Welfare loss relative to the first best (CE, %)				
No policy intervention	24.485	4.646	0.987	0.057
Uniform tax per ton (Pigouvian)	0.112	0.113	0.116	0.122
Sector-wide tax per ton	0.001	0.002	0.018	—
First-best emissions / no-policy emissions, 2017	0.076	0.281	0.605	1.359
B. Optimal tax / Pigouvian tax, median firm, 2017				
Coal	0.55	0.42	0.10	−4.34
Cement	0.71	0.67	0.55	−1.09
Fertilizer	0.33	0.20	−0.15	−4.89
Paper	0.57	0.41	0.01	−5.49
Petroleum	0.75	0.66	0.45	−2.49
C. Marginal damage / first-best price, median firm, 2017				
Coal	0.92	0.72	0.46	0.08
Cement	0.95	0.82	0.60	0.13
Fertilizer	0.94	0.78	0.55	0.11
Paper	0.89	0.65	0.39	0.06
Petroleum	0.90	0.68	0.41	0.07

Note: $\phi_E = 133$ is the baseline; with $\phi_E = 1$ each firm is taxed on its own emissions only. A dash marks a case in which the solution algorithm does not return a valid allocation: with $\phi_E = 1$ the sector-wide tax is a large output subsidy in every sector, and the iteration stalls with some firms at zero output, which suggests that no equilibrium with all firms producing exists under that tax.

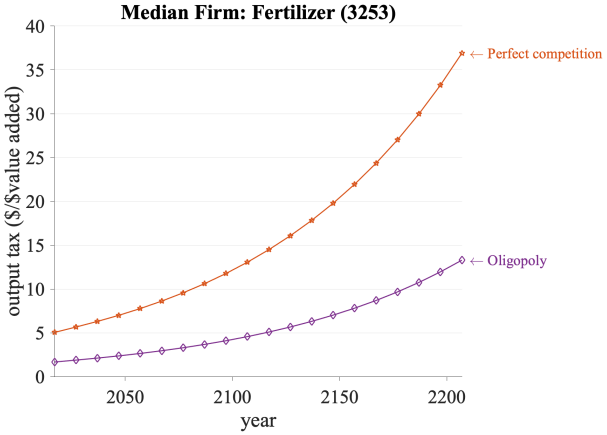
Figure 15: Output taxes under perfect competition vs. oligopoly for the median productivity firm in five carbon-intensive NAICS sectors.



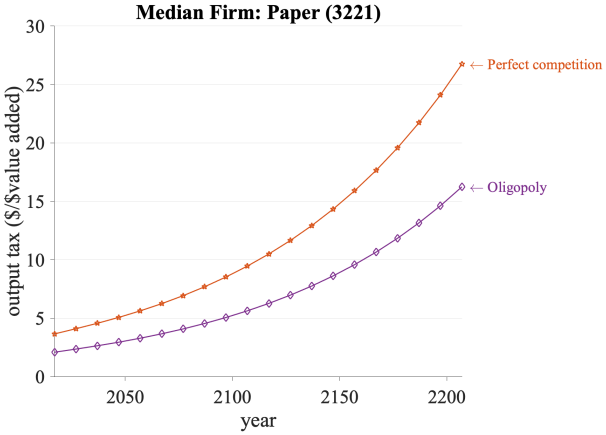
(a) Coal mining



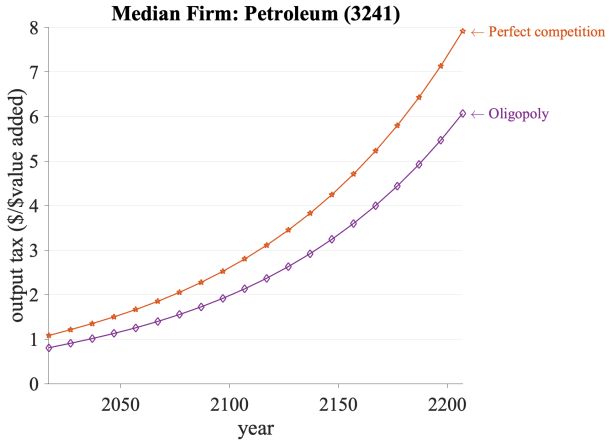
(b) Cement manufacturing



(c) Fertilizer manufacturing

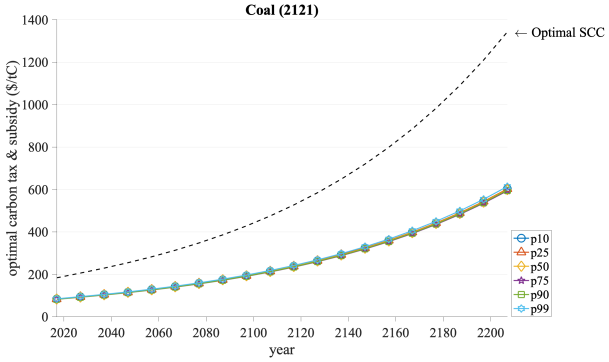


(d) Paper manufacturing

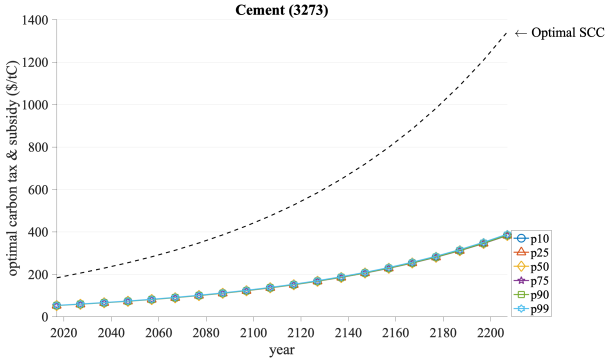


(e) Petroleum and coal products manufacturing

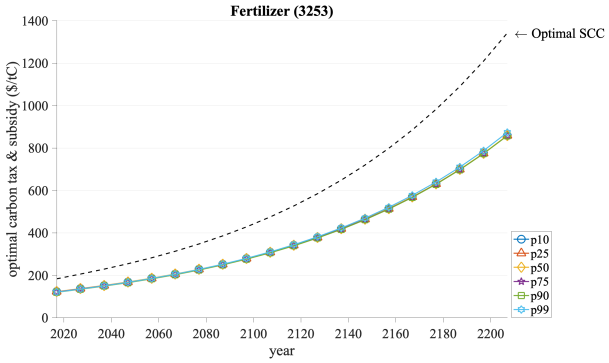
Figure 16: Optimal output subsidy per ton of carbon by firm productivity percentile, and the optimal social cost of carbon, in five carbon-intensive sectors.



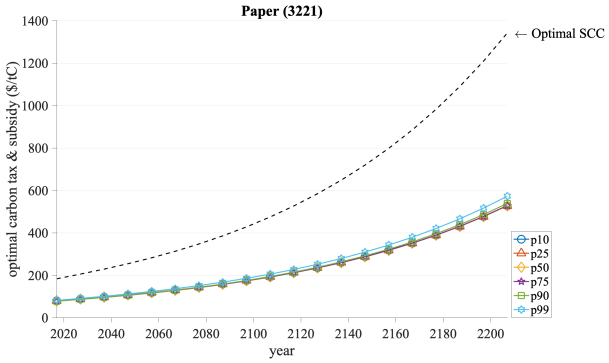
(a) Coal mining



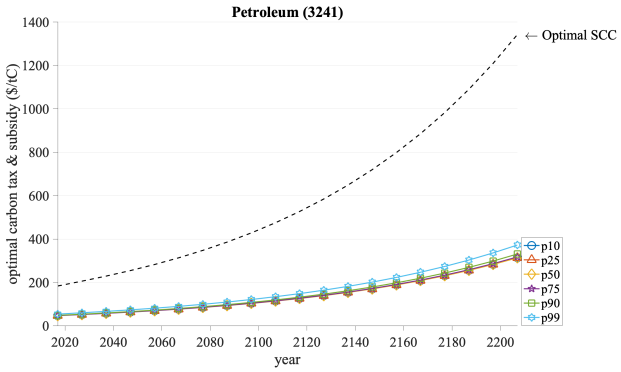
(b) Cement manufacturing



(c) Fertilizer manufacturing



(d) Paper manufacturing



(e) Petroleum and coal products manufacturing