

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Essays in Environmental Macroeconomics

Permalink

<https://escholarship.org/uc/item/2t99f1qw>

ISBN

9798186386070

Author

Gao, Yang

Publication Date

2026-06-09

Peer reviewed|Thesis/dissertation

University of California
Santa Barbara

Essays in Environmental Macroeconomics

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Economics

by

Yang Gao

Committee in charge:

Professor Peter Rupert, Chair
Professor Yueyuan Ma
Professor Gregory Casey, Williams College

June 2026

The Dissertation of Yang Gao is approved.

Professor Yueyuan Ma

Professor Gregory Casey, Williams College

Professor Peter Rupert, Committee Chair

April 2026

Essays in Environmental Macroeconomics

Copyright © 2026

by

Yang Gao

To my loving family.

Acknowledgements

I would like to express my deepest gratitude to my exceptional advisors, Peter Rupert, Yueyuan Ma, and Gregory Casey, for their unwavering support, guidance, and mentorship throughout my PhD journey. I am especially grateful to my main advisor, Peter Rupert, whose invaluable guidance and steadfast support have shaped the way I think about research, develop meaningful research questions, and communicate ideas effectively. I am also sincerely thankful to Yueyuan Ma and Gregory Casey for their insightful feedback, encouragement, and guidance. I feel particularly fortunate to have worked at the intersection of Yueyuan's and Greg's research fields. Since joining UCSB, Yueyuan has been a constant source of support and has generously shared her expertise with me. I am equally grateful to Greg, who offered timely guidance when he was a visiting scholar at UCSB, at a moment when I had only recently developed an interest in environmental macroeconomics and was still seeking direction in the field. As both a coauthor and mentor, Greg played an important role in encouraging me to continue pursuing this area of research. Together, the three of you have profoundly shaped my development as an economist. Your distinct perspectives, insights, and comments have strengthened my work in countless ways. Beyond academics, I am deeply grateful for your kindness, understanding, and encouragement, which helped me navigate not only the challenges of research but also the many ups and downs of these years.

Many wonderful collaborators and mentors have also shaped my journey. I have been fortunate to work with Peter Kruse-Andersen, an excellent coauthor from whom I have learned a great deal. I am also deeply grateful to Tamma Carleton, Ishan Nath, and Levi Crews for their mentorship during my time as a research assistant. I also thank my undergraduate research mentor, Chris Meissner, who deepened my understanding of economics and encouraged me to pursue a PhD in the United States.

I would like to thank the many faculty members with whom I have had valuable discussions about my research and career. In particular, I am grateful to Kyle Meng, Antony Millner, Don Fullerton, Olivier Deschênes, Christopher Costello, Javier Birchenall, Julia Lowell, Lint Barrage, Kieran Walsh, Simon Margolin, Laura Murphy, Peter Kuhn, Shuo Zhang, Andrew Liu, Lizi Yu, Xintong Yang, Han Xu, Shaoshuang Yang, and Yunyu Shu.

The Environmental Economics and Macroeconomics reading groups at UCSB provided an engaging, supportive, and intellectually stimulating environment throughout my PhD. I am especially grateful to Alexander Abajian, Caroline Zhang, and Kite Liu for their enduring friendship and support during this journey. I also deeply appreciate my colleagues and friends, including Ariel Gao, Qian Li, Jimena Rico-Straffon, Flavio Malagutti, Camilo Abbate, Roberto Amaral Santos, Woongchan Jeon, Zebang Xu, Yuanzhe Liu, Ruth Morales, Sandy Sum, Liona Liu, Wenbei Zhang, Yaoyi Bai, Lei Yue, Jianyu Xu, Dan Qiao, Chong Liu, Xin Jiang, Ming Yin, Hong Wang, Yifan Qiao, Ruoyun Tan, Qiqi Li, Guyue Huang, Lianke Qin, Yuran Sun, Yepeng Liu, Liangyu Li, Jiaxuan Cai, Jialu Ran, and Pengyue Zhu, for their encouragement, friendship, and companionship along the way.

This thesis is lovingly dedicated to my husband, Xuandong Zhao, a brilliant computer scientist and my greatest source of love, strength, and encouragement. His unwavering support, patience, and belief in me have carried me through every step of this journey. In moments of doubt and difficulty, his companionship has brought me comfort, joy, and the courage to keep going. Watching his own dedication and perseverance through his PhD and postdoctoral journey has inspired me more than words can express. His love has made this path lighter, richer, and more meaningful, and I am deeply grateful to have shared this chapter of my life with him. During my PhD, our son Henry came into our lives and filled our family with a new kind of joy, hope, and love. His presence has brought warmth and light to each day, reminding me of what truly matters and giving this

milestone an even deeper meaning. I am especially grateful to my loving Mom and Dad, whose unconditional love, quiet sacrifices, and unwavering faith in me have supported me at every stage of my life. The education and values they instilled in me from a young age laid the foundation for this achievement. I am also deeply grateful to my husband's family for their kindness, warmth, and encouragement throughout this journey. This achievement reflects not only my own efforts, but also the love, strength, and sacrifices of the family who have supported me so generously. To all of you, I offer my deepest love and gratitude.

Curriculum Vitæ

Yang Gao

Education

2026	Ph.D. in Economics, University of California, Santa Barbara
2019	B.A. in Economics, Renmin University of China
2017–2018	Undergraduate Exchange Student, University of California, Davis

Research Fields

Macroeconomics, Environmental and Energy Economics

Publications

Casey, Gregory, **Yang Gao**, and Peter K. Kruse-Andersen. 2026. “Some Macroeconomic Evidence on Substitution between Clean and Dirty Energy.” *AEA Papers and Proceedings* 116: 522–526.

“Energy Efficiency Dynamics and Climate Policy,” with Gregory Casey and Peter K. Kruse-Andersen. Forthcoming in the *Journal of the Association of Environmental and Resource Economists*.

Working Papers

“Beyond Emissions: The Implications of Capital Misallocation on Optimal Carbon Tax,” Solo Authored.

Work in Progress

“Second-Best Climate Policy with Realistic Transition Dynamics,” with Gregory Casey and Peter K. Kruse-Andersen.

“The Effects of Price Deregulation in Ghana’s Oil Sector,” with Alexander Abajian and Flavio Malagutti.

Honors and Awards

2024	Dissertation Summer Fellowship, UCSB
2024	Economics Graduate Research Quarter Fellowship, UCSB
2022	Best Second-Year Research Paper Award, UCSB
2022	Gretler Foundation Fellowship, UCSB
2020	Distinction in Ph.D. Preliminary Qualifying Examination (Macroeconomics), UCSB
2018	National Scholarship (Top 0.2% in China Nationwide), RUC

2018	Academic Perfection Recognition, UCD
2017	Chancellor’s Exchange Student Scholarship, RUC
2016–2017	Academic Scholarship, RUC

Conference Presentations

2026	West Coast Micro-to-Macro Development Workshop (poster)
2025	AERE Summer Conference
2024	Midwest Macroeconomics Meeting
2024	The Occasional Workshop in Environmental and Resource Economics

Research Positions

2021–2022	<i>Research Assistant</i> for Tamma Carleton, UCSB
2018	<i>Research Assistant</i> for Christopher M. Meissner, UCD

Professional Service

Referee: Journal of Environmental Economics and Management;
European Economic Review

Teaching Experience

2022–2023	<i>Head Teaching Assistant</i> , UCSB ECON 101: Intermediate Macroeconomic Theory (2022–2023)
2019–2026	<i>Teaching Assistant</i> , UCSB ECON 135: Monetary Economics (2026); ECON 181: International Finance (2024, 2026); ECON 180: International Trade (2021–2023); ECON 101: Intermediate Macroeconomic Theory (2021, 2025); ECON 204C: Ph.D. Macroeconomic Theory III (2021); ECON 2: Principles of Economics–Macro (2020); PSTAT 109: Statistics for Economics (2020); ECON 10A: Intermediate Microeconomic Theory (2019)

Permissions and Attributions

This dissertation contains material that is coauthored with Gregory Casey and Peter K. Kruse-Andersen.

Chapters 2 and 3 are based on joint work with these coauthors. The ideas, analysis, and writing were developed collaboratively. I contributed to the formulation of the research questions, the analysis, and the writing of these chapters.

Chapter 1 has not been previously published. Chapter 2 has been accepted for publication in the *Journal of the Association of Environmental and Resource Economists*. Chapter 3 is based on the published article “Some Macroeconomic Evidence on Substitution between Clean and Dirty Energy,” published in *AEA Papers and Proceedings*, 116, 522–526, 2026.

Abstract

Essays in Environmental Macroeconomics

by

Yang Gao

This dissertation investigates the macroeconomic and welfare implications of climate change and climate policies, with particular emphasis on firm dynamics and market frictions. The three chapters approach these questions using a combination of quantitative macroeconomic modeling, theoretical analysis, and econometric evidence from micro- and macro-level data.

The first chapter studies optimal carbon taxation when capital is misallocated across firms. Using data on U.S. public firms, I document that firm-level revenue-based emissions intensity is negatively correlated with the return to capital, indicating that capital is misallocated toward dirty firms. Embedding this pattern into a misallocation framework, I demonstrate a reallocation mechanism: a carbon tax can reduce resource misallocation and boost aggregate productivity by reallocating capital from dirty to clean firms. I develop a dynamic, general-equilibrium climate-economy model with heterogeneous firms and tangibility-dependent financial frictions: dirty firms hold more tangible assets and face looser collateral constraints, attracting excess capital. Calibrated to firm-level balance sheet and emissions data, the model implies that a carbon tax raises aggregate productivity over a wide range of tax rates. The allocative efficiency gains generate welfare improvements beyond those from internalizing emissions, increasing the optimal carbon tax by 75% relative to a benchmark without pre-existing misallocation. Consistent with the same mechanism, removing financial frictions to equalize returns further reduces economy-wide emissions intensity.

The second chapter, coauthored with Gregory Casey and Peter K. Kruse-Andersen, studies the effectiveness of climate change mitigation policies in reducing carbon emissions, focusing on the channel of economy-wide energy efficiency. Using U.S. data, we estimate an impulse response function (IRF) that characterizes how energy efficiency responds to energy price shocks. Standard climate-economy models cannot replicate the slow transition dynamics observed in the data. We build a tractable model that nests the existing literature and can closely replicate the empirical IRF. The slow dynamics in our model imply that carbon taxes reduce short-run and cumulative emissions less than predicted by standard models. The slow dynamics also imply that higher carbon taxes and greater reductions in output are needed to achieve a given environmental target.

The third chapter, coauthored with Gregory Casey and Peter K. Kruse-Andersen, further investigates the aggregate elasticity of substitution between clean and dirty energy using time-series data from the United States. To match the data, a standard macro climate-economy model requires an elasticity that is close to one in the long run but well below one in the short run. Impulse response functions show that the elasticity converges slowly to its long-run value after a shock to dirty energy prices. Our findings suggest that carbon taxes reduce emissions more slowly than standard models predict.

JEL Codes: E22, E24, G32, H23, O33, O44, O47, Q43, Q54, Q58

Contents

Curriculum Vitae	viii
Permissions and Attributions	x
Abstract	xi
1 Beyond Emissions: The Implications of Capital Misallocation on Optimal Carbon Tax	1
1.1 Introduction	1
1.2 Related Literature	6
1.3 Empirical Evidence	9
1.4 Simple Model	18
1.5 Quantitative Model	32
1.6 Model Calibration	43
1.7 Quantitative Results	46
1.8 Conclusion	58
1.9 Acknowledgments	59
2 Energy Efficiency Dynamics and Climate Policy	61
2.1 Introduction	61
2.2 Related Literature	64
2.3 Motivating Empirical Evidence	66
2.4 Model	69
2.5 Calibration	78
2.6 Policy Analyses	89
2.7 Conclusion	100
2.8 Acknowledgments	101
3 Some Macroeconomic Evidence on Substitution between Clean and Dirty Energy	102
3.1 Introduction	102
3.2 Theoretical Framework	104

3.3	Descriptive Evidence	105
3.4	Impulse Response Functions	106
3.5	Implications for Climate Policy	107
3.6	Conclusion	109
3.7	Acknowledgments	110
A	Appendix for Chapter 1	111
B	Appendix for Chapter 2	124

Chapter 1

Beyond Emissions: The Implications of Capital Misallocation on Optimal Carbon Tax

1.1 Introduction

Climate change is a pressing global challenge that demands urgent and well-designed policy interventions. Among these, carbon taxes play a central role. They not only internalize the externality from carbon emissions but also reallocate production factors from high-emission-intensity (“dirty”) to low-emission-intensity (“clean”) firms. In a frictionless economy with efficient resource allocation, such reallocation would exacerbate allocative inefficiency. However, in the presence of pre-existing distortions that cause resource misallocation, it may alter *allocative efficiency* in a distinct way.

Using firm-level CO₂ emissions linked to balance sheet data for U.S. publicly listed firms, I document a significant negative relationship between firms’ emissions intensity and their marginal revenue product of capital (MRPK), indicating that capital is systematically

misallocated toward dirty firms. To shed light on the source of this misallocation, I further show that dirty firms hold more tangible assets and exhibit higher leverage. Building on this evidence, I propose a novel mechanism of *tangibility-dependent financial frictions* that endogenously explains the observed dirty-biased capital misallocation: because tangible assets ease collateral constraints, dirty firms with more tangible assets face looser borrowing limits and attract excess capital. This paper quantifies how carbon taxes interact with these distortions in a dynamic general equilibrium climate-economy framework—a channel that has been largely overlooked in the literature. In this setting, a carbon tax can not only internalize the climate externality but also enhance aggregate productivity by reallocating capital more efficiently from dirty to clean firms, yielding a potential “*double dividend*”. This paper further examines how this reallocation mechanism affects the second-best optimal carbon tax in markets with financial frictions that generate capital misallocation.

To provide analytical insights, I first build a simple, static model featuring firm-specific capital distortions à la Hsieh and Klenow (2009), heterogeneous emission intensities, and a negative externality from emissions as a by-product of production. The economy exhibits two market imperfections: climate externality and capital misallocation. I allow for a negative covariance between firms’ capital distortions and emission intensities, reflecting the empirical pattern that clean firms tend to face higher returns to capital. I prove that a carbon tax can generate aggregate TFP gains up to a theoretical threshold. Intuitively, by imposing higher levies on dirty firms, a carbon tax reallocates capital toward clean firms, reducing MRPK dispersion and mitigating capital misallocation, thereby raising TFP. Consequently, the carbon tax enhances welfare by improving allocative efficiency, beyond the standard benefit of internalizing emissions. The optimal carbon tax under pre-existing capital misallocation thus exceeds the Pigouvian benchmark that would apply

in the absence of such misallocation¹. Beyond carbon taxes, I also prove that eliminating capital misallocation alone reduces the economy’s aggregate emissions intensity by shifting resources from dirty to clean firms. Therefore, alleviating misallocation not only boosts TFP but also delivers direct environmental benefits.

To quantify policy impacts, I develop a quantitative, dynamic climate-economy model with tangibility-dependent financial frictions that replicates firm-level empirical patterns. The model features heterogeneous firms evolving over time in a general equilibrium setting. Consistent with the empirical evidence, firms differ along two key dimensions. First, they vary in the share of tangible capital in their asset structure, which determines their emissions intensity. Second, they face financial frictions that link access to credit and borrowing capacity to tangible capital: because tangible assets are more pledgeable, clean firms with more intangible assets encounter tighter collateral constraints and underinvest relative to their productivity, whereas dirty firms with more tangible assets are largely unconstrained and tend to overinvest. Cumulative emissions from firms’ production impose a linear utility cost on both firms and workers, capturing the negative climate externality in the model. Relative to the simple model, the quantitative model endogenizes capital distortions through the bindingness of firms’ borrowing constraints. As a result, firm-level capital distortions now adjust endogenously to carbon taxes, providing a more realistic depiction of firm dynamics and factor reallocation in response to climate policy. Furthermore, the quantitative model incorporates a variable aggregate capital supply, enabling capital accumulation to respond endogenously to policy changes.

The quantitative model is calibrated to balance sheet and emissions data for U.S. public firms. I evaluate the laissez-faire equilibrium outcomes in the baseline model with

¹In both the simple and quantitative models, the climate externality is modeled as a linear additive utility cost, independent of capital misallocation. Hence, the allocative efficiency, i.e. aggregate TFP, is unrelated to climate damages. Throughout the paper, the “efficient allocation” refers to the state in which marginal products are equalized and aggregate productivity attains its theoretical maximum.

tangibility-dependent financial frictions and compare them with those in a benchmark model without financial frictions, as commonly used in the macro-climate economy literature. The baseline model successfully reproduces several firm-level moments observed in the data that are not specifically targeted during model calibration. In particular, it endogenously generates a negative correlation between emissions intensity and MRPK; higher capital, output, and leverage among dirty firms; and higher cash holdings among clean firms. These results validate my financial frictions mechanism in explaining the empirically observed misallocation patterns. In contrast, the benchmark model without financial frictions fails to replicate these empirical moments.

Using the calibrated model, I examine counterfactual economies in the steady state after carbon taxes are introduced. I compare aggregate, distributional, and welfare outcomes under varying carbon tax rates in two settings: the baseline economy featuring tangibility-dependent financial frictions and the benchmark economy without such frictions. The reallocation mechanisms derived from the simple model carry over to the quantitative analysis. In the baseline model, carbon taxes enhance aggregate productivity relative to the laissez-faire equilibrium across a broad range of tax rates, with the largest TFP gain occurring at \$275/tCO₂. The effect is driven by the reallocation of capital from dirty to clean firms, which alleviates capital misallocation by narrowing the gap in MRPK between these groups. In contrast, in the benchmark model, higher carbon taxes monotonically reduce TFP by amplifying the dispersion in marginal products and moving the economy further away from the efficient allocation.

Consequently, a carbon tax in the baseline economy can deliver a “double dividend”: it simultaneously mitigates the climate externality and improves allocative efficiency, whereas in the benchmark economy it serves only the traditional Pigouvian purpose. In the baseline model, the steady-state optimal carbon tax is \$210/tCO₂, about 75 percent higher than the optimal tax \$120/tCO₂ in the benchmark. Hence, financial frictions drive

up the optimal carbon tax due to additional TFP gains from the tax, beyond its effect on reducing emissions: higher TFP raises aggregate capital demand by increasing the marginal return to investment, which partially offsets the decline in aggregate capital supply caused by a higher carbon tax relative to the benchmark. In the baseline economy, the optimal carbon tax (\$210/tCO₂) remains below the level that would maximize TFP gains (\$275/tCO₂), because even though higher TFP could be achieved, aggregate capital and output are largely depressed by a higher carbon tax; thus, excessively high taxes can harm welfare in the dynamic model. In terms of welfare gains, the steady-state consumption at the optimal tax in the baseline model rises by 0.32 percent relative to laissez-faire, compared to only 0.14 percent under the optimal carbon tax in the benchmark economy. Moreover, financial frictions substantially widen the range of carbon tax rates that improve welfare relative to laissez-faire.

Financial frictions alone also have important environmental implications. As an alternative policy counterfactual to carbon taxes, removing these frictions generates sizable aggregate and environmental effects. Eliminating financial frictions raises TFP by 2.3 percent and aggregate output by 6.2 percent. In the dynamic model, these productivity gains increase the equilibrium interest rate and induce additional capital accumulation, further amplifying the output response. Aggregate emissions rise by 4.8 percent, primarily due to the scale effect of higher output, while the economy-wide emissions intensity declines by 1.4 percent, reflecting a larger share of clean production as capital reallocates toward clean firms. Overall, the welfare gain from removing financial frictions is an order of magnitude larger than that achieved under the optimal carbon tax in the presence of financial frictions.

1.2 Related Literature

A closely related concurrent study by Kim (2023) documents a similar negative correlation between firms' emissions intensity and marginal products. He attributes this pattern to the assumption that clean firms have higher firm-level TFP, since more productive firms exhibit higher equilibrium marginal products in standard models with financial frictions. In contrast, this paper shows that higher-TFP firms are not necessarily cleaner. Drawing on richer empirical evidence and insights from the literature, I propose a distinct mechanism: dirty firms hold more tangible assets and face looser collateral constraints than clean firms, giving rise to heterogeneous financial frictions that systematically favor dirty firms. This mechanism endogenizes aggregate TFP as a function of tangibility-dependent financial frictions and yields sharper predictions for the effects of carbon taxes on firm dynamics and allocative efficiency.

My paper closely relates to several other studies on the allocative efficiency impacts of environmental policies in the presence of correlated distortions. Klenow et al. (2024) document that revenue distortions are correlated with fossil fuel intensity among Chilean firms and find that carbon taxes can enhance allocative efficiency by reallocating resources from dirty, low-markup firms to clean, high-markup firms. My work differs in several aspects. I focus on the implications of capital distortions from heterogeneous financial frictions among U.S. firms. My dynamic climate-economy model with richer firm dynamics allows me to quantify policy impacts and to explore both positive and normative questions on the optimal carbon tax in the presence of capital misallocation, in contrast to their static framework where misallocation arises from exogenous wedges. Kabir and Tan (2025) show that green financing can increase emissions if it reduces capital misallocation, where they present a sufficient statistics approach to evaluate the circumstances under which such unintended consequences can occur and show that this is a valid concern for

over 53.9% of industries in the United States. Focusing on water pollution rather than carbon emissions, Qi et al. (2021) document that misallocation in China favors smaller, more polluting firms, and examine the effects of either removing such misallocation or implementing water pollution regulations through a quantification exercise.

Second, my paper contributes to an extensive body of work on second-best carbon pricing in the presence of pre-existing market distortions. Much of this literature employs a representative-agent framework in computable general equilibrium or macroeconomic models and emphasizes the interactions between carbon taxation, distortionary fiscal instruments, and revenue recycling (e.g., Goulder, 1995; Bovenberg and Goulder, 1996; Parry et al., 1999; Fullerton and Kim, 2008; Barrage, 2020). In contrast, I study the optimal carbon tax in a macro-climate-economy model with capital misallocation across heterogeneous firms. My analysis reveals the welfare benefit of carbon taxes through their positive effects on allocative efficiency.

Third, my paper builds on the extensive macroeconomic literature on resource misallocation and aggregate productivity losses, which can arise from exogenous wedges in static settings (Hsieh and Klenow, 2009) or from endogenous distortions due to financial frictions in dynamic models (Midrigan and Xu, 2014; Moll, 2014; Buera and Moll, 2015; Gopinath et al., 2017). More recently, Liu and Xu (2024) examine climate-change-induced factor misallocation and its effects on TFP. My research goes beyond the productivity implications of capital misallocation to highlight its environmental and climate policy implications. I show that removing capital misallocation in favor of dirty firms can reduce economy-wide emissions intensity alongside the gains in TFP, and that accounting for capital misallocation also has important implications for the optimal carbon tax.

Fourth, my work relates to the literature on firms' asset structure and financial constraints. A large body of research emphasizes that asset tangibility plays a central role in determining firms' access to credit (e.g., Almeida and Campello, 2007; Caggese and

Perez-Orive, 2022; Falato et al., 2022), and that firms investing in intangible capital are more likely to face financial constraints (e.g., Hall, 2002; Brown et al., 2009; Dell’Ariccia et al., 2021). I contribute to this literature by examining the environmental bias of financial frictions that depend on firms’ asset tangibility and pledgeability.

Lastly, my paper contributes to the growing literature on climate finance. Iovino et al. (2023) document the U.S. corporate tax effectively subsidizes carbon emissions, as dirty firms benefit disproportionately from debt tax shields enabled by their higher leverage and asset tangibility. Papoutsis et al. (2022) show that the European Central Bank’s market-neutral corporate bond purchases are unintentionally tilted toward brown sectors, reflecting their larger bond issuance and higher tangible capital. Lanteri and Rampini (2025) find that financial constraints in the commercial shipping industry hinder smaller firms from adopting expensive cleaner technologies, leading to investment in cheaper, older, and dirtier capital. Relative to these studies, I focus on U.S. public firms that differ in their emissions intensity and examine how financial frictions distort capital allocation in favor of dirtier firms due to differences in their asset tangibility. Capelle et al. (2023) highlight the importance of adoption of newer capital-embedded technologies in lowering the costs of mitigation policies by noticing the firm-level heterogeneity in environmental performance.

The structure of the paper is as follows. Section 1.3 describes the data, documents the empirical evidence on capital misallocation favoring dirty firms, and provides supporting evidence for the underlying mechanism. Section 1.4 presents a simple, static model and derives analytical results. Section 1.5 develops a quantitative, dynamic climate-economy model in general equilibrium with tangibility-dependent financial frictions. Section 1.6 outlines the model calibration. Section 1.7 validates the model and analyzes policy counterfactuals, including the implications for the optimal carbon tax. Section 1.8 concludes.

1.3 Empirical Evidence

This section documents a robust empirical relationship between firms’ emissions intensity, return on capital, and asset tangibility. The results highlight how capital is misallocated in favor of high-emissions-intensity “dirty” firms, potentially driven by financial frictions related to asset tangibility.

1.3.1 Data Description

My empirical analysis focuses on U.S. public firms that generate carbon emissions. The data are collected from two data sources: firm-level balance sheet data from Compustat and firm-level carbon emissions from S&P Trucost.

Firm Balance Sheet Variables. Firm-level balance sheet data for all U.S.-headquartered firms are obtained from Compustat over the period 2002-2021.² Firms’ output is measured by sales (Compustat item SALE). Tangible capital is constructed from the tangible fixed assets, specifically property, plant, and equipment (PPE; Compustat items PPEGT and PPENT), using the perpetual inventory method. Intangible capital is taken from intangible fixed assets (Compustat item INTAN). The total capital stock is the sum of tangible and intangible capital. Total assets are obtained from Compustat item AT and the number of employees from Compustat item EMP. Debt is defined as the sum of short- and long-term debt (Compustat items DLC and DLTT). Firm size is measured by total assets. Firm age is defined as the number of years since the founding year, based on information provided by Jay Ritter’s website; if the founding year is unavailable, the first year the firm appears in Compustat is used as a proxy. For regression analysis, industry-by-year fixed

²In the model calibration in Section 1.6, I also extend the Compustat firm sample, including those with at least one carbon emissions observation, back to 1980. This longer time series is essential for estimating productivity parameters, allowing for more consistent and unbiased results. Further details will be provided in that section.

effects are further included, with industries classified by two-digit SIC codes. Nominal values are deflated by GDP deflators, and tangible assets by investment goods price deflators, both sourced from Federal Reserve Economic Data. The sample is restricted to exclude firms in the financial (SIC codes 6000-6799) and utility (SIC codes 4900-4999) sectors. For the empirical analysis, firm-years with missing values for key variables are excluded.

In the following discussion, the marginal revenue product of capital (MRPK) for firms is calculated as the ratio of revenue output to total capital; the marginal revenue product of labor (MRPL) as the ratio of revenue output to the number of employees. The approximation of marginal products by average products follows directly from the assumption of Cobb-Douglas production function for firms³. Firms' asset tangibility is defined as the proportion of tangible assets to total assets, and leverage is defined as the ratio of total debt to total assets.

Firm Carbon Emissions. Data on firm-level carbon emissions, spanning 2002-2021, are obtained from the S&P Trucost database. The firm-level greenhouse gas emissions, quantified in CO₂ equivalent (CO₂e), are subsequently merged with firm-level balance sheet data, restricting the sample to Compustat firms for which emissions data are available. Trucost quantifies firm-level scope 1, 2, and often scope 3 greenhouse gas emissions. Being the data provider with the broadest coverage for firm-level direct CO₂ emissions (Busch et al., 2022; Iovino et al., 2023), Trucost collects emissions data from company reports and public disclosures. It ensures reliable data via robust quality control and standardization, broadening coverage by estimating non-disclosed emissions with sector-specific models.

³See equations (1.9) and (1.52). Also for empirical analysis, the omission of the constant capital share α and the constant price markup η does not affect the regression coefficients. Potential variations in factor shares across industries are addressed by including industry-by-year fixed effects.

In the baseline analysis, I use firm-level emissions defined as the sum of scope 1 emissions, which measure direct emissions from company-owned or controlled resources, and scope 2 emissions, which capture indirect emissions from purchased electricity or energy. I chose to use both scope 1 and 2 emissions as many of the firms' emissions come not only from their own plants but also from their reliance on external energy sources, such as electricity, to run daily operations, which is not necessarily captured in scope 1. This approach provides a more comprehensive view of the total emissions associated with the company's activities and energy consumption. Firms' emissions intensity is measured as the ratio of emissions to sales, expressed in metric tons of CO₂e per thousand dollars of sales. The baseline sample for the empirical analysis covers 2,150 U.S. non-financial public firms, with a total of 15,830 firm-year observations. Table A.1 presents the summary statistics of the constructed firm sample.

1.3.2 Cross-sectional Link between Emissions Intensity and MRPK

In the absence of distortions, capital should be allocated toward more productive firms until the marginal revenue product of capital (MRPK) is equalized across firms. Capital misallocation arises when market distortions or policy interventions introduce wedges between factor prices and marginal products in firms' capital decisions, generating dispersion in MRPK across firms. Firms facing greater capital distortions exhibit higher MRPK and are unable to achieve an allocatively efficient level of capital, while others with lower MRPK may hold more capital than is efficient.

This section examines how MRPK varies with emissions intensity. To this end, I estimate the following ordinary least squares (OLS) regression using the constructed

firm-level dataset:

$$Y_{i,t} = \beta_0 + \beta_1 \log(\text{Emissions Intensity}_{i,t}) + X_{i,t} + \gamma_{s,t} + \epsilon_{i,t} \quad (1.1)$$

In this baseline specification, $\log(\text{Emissions Intensity}_{i,t})$ represents the natural logarithm of emissions intensity of firm i in year t . I apply the log of emissions intensity to address the concern that the distribution of emissions intensity across firms is right-skewed. The main dependent variable, $Y_{i,t}$, is the natural logarithm of firms' marginal revenue product of capital, $\log(\text{MRPK})$. $X_{i,t}$ are firm-level controls including firm size, measured by $\log(\text{total assets})$, and firm age, measured by $\log(\text{age})$. These controls are included in a subset of the regression specifications. The term, $\gamma_{s,t}$, captures 2-digit SIC industry-by-year fixed effects, controlling for variations within industries over time. Standard errors, $\epsilon_{i,t}$, are clustered at the firm-year level.

I show that capital is inefficiently over-allocated to dirty firms and under-allocated to clean firms, which is illustrated by the baseline results reported in Table 1.1. Columns (1) and (2) of Table 1.1 reveal a negative correlation between firms' emissions intensity and MRPK that is statistically significant at the 0.1% confidence level. This relationship remains robust even after accounting for firm-level controls. Figure 1.1 further illustrates this strong negative relationship with a binned firm plot.

Robustness. To strengthen the findings on the negative relationship between firms' emissions profiles and their MRPK, I conduct several robustness checks. First, I apply the scope 1 direct emissions alone, rather than the sum of scope 1 and 2 emissions, to calculate firms' emissions intensity. Columns (1) and (2) of Appendix Table A.2 indicate that this negative relationship remains statistically significant at the 0.1% level. Second, I replace sales with gross value added—calculated as sales minus the cost of goods sold—to

	Dep. Var.: $\log(MRPK)$		Dep. Var.: $Asset\ Tangibility$		Dep. Var.: $Leverage$	
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(Emissions\ Intensity)$	-0.196*** (0.020)	-0.198*** (0.021)	0.074*** (0.004)	0.074*** (0.004)	0.015** (0.004)	0.015** (0.004)
Firm Controls	N	Y	N	Y	N	Y
Industry $\times$ Year FE	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.315	0.341	0.651	0.651	0.113	0.134
Observations	15,830	15,830	15,830	15,830	15,830	15,830

Standard errors two-way clustered at the firm-year level are reported in parentheses.

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 1.1: Firm-level results between emissions intensity and balance sheet outcomes

construct firms' MRPK. Columns (3) and (4) of Appendix Table A.2 confirm that the coefficient remains negative and statistically significant at the 0.1% level. Third, I use the level instead of the log transformation of emissions intensity as an alternative independent variable in specification (1.1). As shown in columns (5) and (6) of Appendix Table A.2, the coefficient remains negative and statistically significant at the 0.1% level. Lastly, I assess the relationship between firms' marginal revenue product of labor (MRPL) and emissions intensity. As detailed in Appendix Table A.3, the regression coefficients for $\log(MRPL)$ vary in sign and largely insignificant across specifications, suggesting the correlation between MRPL and emissions intensity is neither consistent nor robust. This indicates that, although there are notable associations with firms' emissions profiles in terms of capital allocation, similar patterns in labor allocation do not consistently emerge from the data.




Figure 1.1: Firm-level emissions intensity and marginal revenue product of capital

Notes: the figure shows a binned scatter plot with 50 bins of firms (solid blue dots) and a linear regression fit (red dashed line) adjusted for industry-year fixed effects. Marginal revenue product of capital (MRPK) is measured by the ratio of output to total capital; emissions intensity is measured in metric tons of CO₂ per thousand dollar sales.

1.3.3 The Mechanism behind the MRPK-Emissions Intensity Link

In this section, I explain the mechanism through which dirty firms experience lower capital distortions, as evidenced by their lower MRPK. Previous studies have shown that dirty firms or sectors tend to own more tangible assets (Papoutsis et al., 2022; Iovino et al., 2023). Tangible assets—such as buildings, machinery, and equipment—can more easily be used as collateral. I therefore propose that the greater tangibility of assets in dirty firms enhances their ability to obtain external financing, thereby reducing distortions in their access to capital and resulting in a lower MRPK.

In columns (3) and (4) of Table 1.1, I highlight a significant positive relationship between a firm's logarithm of emissions intensity and its asset tangibility using regression specification (1.1). Asset tangibility is measured by tangible capital scaled by total assets where tangible capital is measured by property, plant, and equipment (PPE), and

the coefficient for this relationship is statistically significant at the 0.1% level. Panel (a) of Figure 1.2 further illustrates this strong relationship, suggesting that firms with substantial emissions tend to have a higher proportion of tangible capital in their asset structure due to the fact that carbon emissions are often linked to the use of tangible capital in the production process.

Tangible capital plays a crucial role in external finance, as it is more readily used as collateral compared to intangible capital. This means that firms with a higher ratio of tangible assets can secure more financing since tangible assets are viewed by lenders as a more reliable source of collateral (see e.g. Almeida and Campello, 2007; Caggese and Perez-Orive, 2022; Falato et al., 2022). In contrast, firms with fewer tangible assets tend to face tighter borrowing constraints due to the lower collateral value of their assets. As such, the tangibility of firms' assets serves as a proxy for their pledgeability, which consequently determines the firms' borrowing capacity. In this context, I further demonstrate that dirty firms are able to sustain higher levels of debt due to their increased tangibility and pledgeability of assets. As illustrated in columns (5) and (6) of Table 1.1, there is a significant positive relationship between firms' emissions intensity and leverage (defined as the debt-to-assets ratio), which holds at the 1% significance level using regression specification (1.1). Panel (b) in Figure 1.2 further depicts this positive correlation. This distinction in firms' share of tangible assets and its impact on their financing capacity has significant implications for the potential financial constraints faced by dirtier versus cleaner firms.

As a result, the observed negative relationship between emissions intensity and MRPK is consistent with the mechanism of heterogeneous financial constraints faced by firms with different emissions characteristics. The financial frictions channel has been crucial in explaining the variation in marginal products and the capital misallocation in the literature (see e.g. Moll, 2014; Midrigan and Xu, 2014; Buera and Moll, 2015; Gopinath et al., 2017).

Dirty firms, which typically possess more tangible assets, face fewer collateral constraints and can operate closer to their desired levels of capital, contributing to their lower MRPK. In contrast, clean firms with fewer tangible assets encounter tighter collateral constraints, hindering their ability to secure sufficient financing and thereby preventing them from reaching their desired capital levels, resulting in a higher MRPK. Complementing this mechanism, Appendix Table A.4 provides robust evidence that clean firms tend to hold more cash in their assets. This behavior is consistent with findings from Cunha and Pollet (2020), who document that financially constrained firms often accumulate more cash for precautionary motives due to insufficient external finance. Therefore, financial frictions based on asset tangibility play a pivotal role in shaping the financial landscape for both high- and low-emissions firms, influencing their access to capital and ultimately affecting the allocative efficiency of capital across these firms.




Figure 1.2: Firm-level emissions intensity, asset tangibility and leverage ratio

Notes: the figure shows a binned scatter plot with 50 bins of firms (solid blue dots) and a linear regression fit (red dashed line) adjusted for industry-year fixed effects. Asset tangibility is measured by the ratio of property, plant and equipment (PPE) to total assets; leverage by the ratio of debt to total assets; emissions intensity by metric tons of CO₂ equivalent per thousand dollar sales.

Tangibility vs. Productivity. Kim (2023) identifies a similar negative correlation

	Dep. Var.: $\log(Emissions\ Intensity)$			
	Tangibility Only (1)	Productivity Only (2)	Both (3)	Both + Firm Controls (4)
<i>Asset Tangibility</i>	2.271*** (0.175)		2.271*** (0.174)	2.256*** (0.171)
$\log(Productivity)$		-0.012 (0.015)	-0.011 (0.013)	-0.021 (0.014)
Firm Controls	N	N	N	Y
Industry×Year FE	Y	Y	Y	Y
Adjusted R^2	0.648	0.577	0.649	0.649
Observations	15,830	15,830	15,830	15,830

Standard errors two-way clustered at the firm-year level are reported in parentheses.

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table 1.2: Explaining firm-level emissions intensity: asset tangibility vs. productivity

between firms' emissions intensity and their marginal products, where he attributes the observed MRPK-emissions intensity relationship to the higher physical productivity, i.e. firm-level TFP, of clean firms. However, my proposed mechanism provides a different explanation, emphasizing the role of asset tangibility in explaining the MRPK-emissions intensity connection, independent of firm-level productivity differences.

To empirically test and compare these mechanisms, I examine whether firm-level productivity or asset tangibility more robustly explains emissions intensity using the following regression:

$$\log(Emissions\ Intensity_{i,t}) = \beta_0 + \beta_1 Tangibility_{i,t} + \beta_2 \log(Productivity_{i,t}) + \gamma_{s,t} + X_{i,t} + \epsilon_{i,t} \quad (1.2)$$

where the new variable $\log(Productivity_{i,t})$ is the logarithm of estimated firm-level productivity⁴.

The results, presented in Table 1.2, consistently show that asset tangibility has a

⁴The firm-level TFP is not directly observable from the balance sheet data. Its estimation will be discussed in detail in Section 1.6.

strong, positive, and statistically significant relationship with emissions intensity, with or without the productivity variable. Conversely, the coefficients on firms' productivity are insignificant across specifications. Notably, conditional on including the other explanatory variable, the inclusion of tangibility markedly increases the regression's explanatory power (R^2), while adding productivity barely changes the value. These findings indicate no strong connection between higher firm TFP and cleaner operations.

Further analysis, presented in Appendix Table A.5, demonstrates that incorporating asset tangibility as an explanatory variable substantially reduces the explanatory power of emissions intensity on MRPK, accounting for 50% of the observed MRPK-emissions intensity relationship. In contrast, adding productivity has a smaller and even negligible impact, suggesting that the potential link between emissions intensity and the marginal product of capital is not mediated through firms' productivity features. This evidence underscores the pivotal role of the firms' differential asset structure in explaining the MRPK-emissions intensity connection, supporting the tangibility-based financial frictions channel proposed in this paper⁵.

1.4 Simple Model

This section develops a simple model that provides analytical results on how carbon taxes affect allocative efficiency in the presence of empirically observed capital misallocation biased toward dirty firms, and how addressing such misallocation influences environmental outcomes. These analytical insights are further illustrated through numerical simulations. The simple model builds on Hsieh and Klenow (2009) but incorporates firm-specific emissions intensities that are negatively correlated with capital distortions.

⁵In addition, Kim (2023) suggests that labor is also misallocated toward dirty firms. However, my empirical findings show no systematic correlation between firms' MRPL and emission intensity.

1.4.1 Model Setup

The model is static. The economy is populated by N monopolistically competitive firms producing differentiated intermediate goods. The final output (Y) is a Dixit-Stiglitz CES bundle of intermediate goods:

$$Y = \left(\sum_{i=1}^N y_i^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad (1.3)$$

where y_i denotes the real output of intermediate producer i , and σ denotes the elasticity of substitution between varieties. Each intermediate firm produces output with a Cobb-Douglas production function:

$$y_i = z_i k_i^\alpha l_i^{1-\alpha} \quad (1.4)$$

where z_i is the physical productivity, k_i is capital, l_i is labor, and α is the capital share of output. Let p_i denote the price of variety i in the units of final good. The profit-maximization of the final producer suggests that firms face an isoelastic demand function for their output:

$$y_i = p_i^{-\sigma} Y \quad (1.5)$$

where I normalize the price of final output to 1. Firms generate carbon emissions, e_i , as a by-product from their production, which is given by:

$$e_i = \omega_i \text{rev}_i \quad (1.6)$$

where ω_i is the firm-specific constant emissions share in terms of revenue (Sogalla et al., 2024), and $\text{rev}_i := p_i y_i$ is firms' total revenue output. Thus ω_i describes firms' revenue-

based emissions intensity, a measure of how much carbon emissions a firm emits for each unit of revenue it generates.

Firms choose their price, capital, and labor to maximize their profits taking (distortion-inclusive) factor prices as given:

$$\max_{p_i, k_i, l_i} (1 - \tau_e \omega_i) p_i y_i - (1 + \tau_{ki})(r + \delta) k_i - w l_i \quad (1.7)$$

where w denotes the wage, r denotes the real interest rate, δ denotes the depreciation rate, τ_e denotes a unit tax on firms' carbon emissions, and τ_{ki} denotes a firm-specific wedge that distorts capital relative to labor, which I refer to as the firm-specific capital distortion henceforth. Because firms pay a total of $\tau_e e_i$ in carbon taxes, equation (1.6) implies this levy is equivalent to an ad-valorem tax at rate $\tau_e \omega_i$, which appears in the profit function as an output wedge that distorts firms' choices and, consequently, resource allocation across firms. In this section, I treat τ_{ki} and its potential relationship with ω_i as exogenous. They are later endogenized in the model of Section 1.5.

The optimality condition of price gives:

$$p_i = \eta \left(\frac{r + \delta}{\alpha} \right)^\alpha \left(\frac{w}{1 - \alpha} \right)^{1 - \alpha} \frac{(1 + \tau_{ki})^\alpha}{z_i (1 - \tau_e \omega_i)} \quad (1.8)$$

where $\eta = \frac{\sigma}{\sigma - 1}$ denotes firms' constant markup of price over marginal cost. Firms' optimal decisions with respect to capital and labor yield:

$$\text{MRPK}_i := \frac{\alpha p_i y_i}{\eta k_i} = \frac{(1 + \tau_{ki})(r + \delta)}{1 - \tau_e \omega_i} \quad (1.9)$$

$$\text{MRPL}_i := \frac{(1 - \alpha) p_i y_i}{\eta l_i} = \frac{w}{1 - \tau_e \omega_i} \quad (1.10)$$

where equation (1.9) suggests that firms' marginal revenue product of capital (MRPK)

is set equal to the cost of capital times the composite distortion $\frac{1+\tau_{ki}}{1-\tau_e\omega_i}$, and equation (1.10) states that their marginal revenue product of labor (MRPL) equals wage rate times the distortion $\frac{1}{1-\tau_e\omega_i}$. With the Cobb-Douglas production function, the marginal revenue product of each factor is proportional to the average revenue product (or revenue-based productivity) of the factor.

In the simple static model, I assume a fixed aggregate factor supply. In the dynamic model of Section 1.5, I will relax the assumption of fixed capital supply. The capital and labor market clearing conditions are given by:

$$\bar{K} = \sum_{i=1}^N k_i \quad (1.11)$$

$$\bar{L} = \sum_{i=1}^N l_i \quad (1.12)$$

To close the model economy, I assume that the revenues from carbon taxes ($\tau_e e_i$) and distortions on capital ($\tau_{ki}(r+\delta)k_i$) from all firms are lump-sum rebated to a representative household who simply consumes the final output.

I now have all the necessary components to define a competitive equilibrium of the simple model, which is given in Definition 1.1.

Definition 1.1. A *competitive equilibrium* is a set of prices (r, w, p_i) and allocations of (k_i, l_i, y_i, e_i, Y) , such that given carbon tax (τ_e), exogenous firm productivity (z_i), emissions intensity (ω_i) and capital distortions (τ_{ki}): (1) intermediate and final-output firms maximize profits, i.e. (1.3), (1.4), (1.6), (1.8), (1.9), and (1.10) hold, (2) factor markets clear, i.e. (1.11) and (1.12) hold.

1.4.2 Allocative Efficiency in the Model

In the laissez-faire (i.e. $\tau_e = 0$) competitive equilibrium of the simple model outlined in Section 1.4.1, firms disregard the climate damages arising from the carbon emissions, which justifies the use of a carbon tax to correct for the negative externality. In addition, capital misallocation arises from the firm-specific τ_{ki} and marginal revenue product of capital, which are negatively correlated with the firm-specific emissions intensity as suggested by the empirical evidence in Section 1.3. Therefore, the model essentially features two market imperfections, the climate externality and allocative inefficiency. Two key policy questions thus relate to: 1) what are the impacts of carbon taxes on resource (mis)allocation and aggregate TFP provided the pre-existing capital misallocation systematically correlated with firm-level emissions feature, and 2) what are the implications of this reallocation channel on the optimal carbon tax that maximizes social welfare? The optimal carbon tax in my context therefore differs substantially from a conventional Pigouvian setting, as it represents a second-best solution that balances the policy's interaction with two market distortions.

Carbon Tax, Reallocation, and TFP. To begin with, I formally define the aggregate TFP of the model economy in *Lemma 1.2*.

Lemma 1.2. *The aggregate production function can be written as the product of an effective aggregate TFP and the aggregate capital-labor composite.*

$$Y = \text{TFP} \times \bar{K}^\alpha \bar{L}^{1-\alpha} \quad (1.13)$$

The effective aggregate TFP can be expressed as a function of the firm-level productivity

(z_i) , emissions intensity (ω_i) , capital distortion (τ_{ki}) , and carbon tax (τ_e) ,

$$\text{TFP} = \frac{\left(\sum_{i=1}^N \left(\frac{(1-\tau_e\omega_i)z_i}{(1+\tau_{ki})^\alpha} \right)^{\sigma-1} \right)^{\frac{\sigma}{\sigma-1}}}{\left(\sum_{i=1}^N \frac{(1-\tau_e\omega_i)^\sigma z_i^{\sigma-1}}{(1+\tau_{ki})^{1+\alpha(\sigma-1)}} \right)^\alpha \left(\sum_{i=1}^N \frac{(1-\tau_e\omega_i)^\sigma z_i^{\sigma-1}}{(1+\tau_{ki})^{\alpha(\sigma-1)}} \right)^{1-\alpha}} \quad (1.14)$$

Proof: I provide a detailed proof in Appendix Section A.

Next, I define resource misallocation, which lowers the allocative efficiency and aggregate total factor productivity (TFP). Adopting the terminology used in Foster et al. (2008), Hsieh and Klenow (2009), and Gopinath et al. (2017), I define firm-level revenue-based total factor productivity (TFPR) as the product of price p_i times physical productivity z_i :

$$\begin{aligned} \text{TFPR}_i &:= p_i z_i \\ &= \eta \left(\frac{\text{MRPK}_i}{\alpha} \right)^\alpha \left(\frac{\text{MRPL}_i}{1-\alpha} \right)^{1-\alpha} \\ &\propto \frac{(1+\tau_{ki})^\alpha}{1-\tau_e\omega_i} \end{aligned} \quad (1.15)$$

Equation (1.15) shows that the firm-level TFPR is related to both factors' marginal revenue product. Firms with higher capital distortions τ_{ki} , or higher tax distortions $\tau_e\omega_i$, face higher marginal revenue products and a higher (more distorted) TFPR. In the model, aggregate TFP achieves its most efficient level when resources are efficiently allocated across firms. In the efficient allocation, more factors are allocated to firms with higher productivity z_i , until there is no dispersion of the MRPK, MRPL, and TFPR across firms, i.e. they face the same or zero distortions on output and capital⁶. However, idiosyncratic distortions lead to a dispersion in marginal revenue products. This dispersion, termed

⁶In this paper, allocative efficiency refers to the state where aggregate TFP is maximized that is independent of climate externality. Consequently, this TFP-maximizing allocation in general does not coincide with the welfare-maximizing allocation once climate externality damage is considered.

resource misallocation, in turn lowers aggregate TFP relative to its efficient level. We measure this misallocation by defining the *misallocation of capital, labor, and resources* (MoK, MoL, and MoR) as the dispersion of the natural logarithm of firm-level MRPK, MRPL, and TFPR, respectively:

$$\text{MoK} := \text{Var}(\log \text{MRPK}_i) \quad (1.16)$$

$$\text{MoL} := \text{Var}(\log \text{MRPL}_i) \quad (1.17)$$

$$\text{MoR} := \text{Var}(\log \text{TFPR}_i) \quad (1.18)$$

MoR (misallocation of resources) is a key determinant of aggregate TFP. A higher MoR signifies worse allocative efficiency, and consequently, lower aggregate TFP. A higher value of MoK (misallocation of capital) or MoL (misallocation of labor), implying a worse capital or labor allocation, can both contribute to higher MoR⁷. To analyze the impact of carbon taxes (τ_e) on factor and resource misallocation, I present the log-approximated expressions for MoK, MoL, and MoR below⁸:

$$\text{MoK} = \text{Var}(\tau_{ki}) + \text{Var}(\omega_i)\tau_e^2 + 2 \text{Cov}(\tau_{ki}, \omega_i)\tau_e \quad (1.19)$$

$$\text{MoL} = \text{Var}(\omega_i)\tau_e^2 \quad (1.20)$$

$$\text{MoR} = \alpha^2 \text{Var}(\tau_{ki}) + \text{Var}(\omega_i)\tau_e^2 + 2\alpha \text{Cov}(\tau_{ki}, e_i)\tau_e. \quad (1.21)$$

In Section 1.3, I establish a negative relationship between MRPK and emissions intensity (ω_i) in the data. Equation (1.9) informs that in the laissez-faire market outcome ($\tau_e = 0$), firm-level capital wedges τ_{ki} directly account for deviations of MRPK from its efficient level. Therefore the negative relationship between MRPK and ω_i translates to a negative

⁷ $\text{MoR} = \alpha^2 \text{MoK} + (1 - \alpha)^2 \text{MoL} + 2\alpha(1 - \alpha) \text{Cov}(\log \text{MRPK}_i, \log \text{MRPL}_i)$

⁸These expressions are derived under the assumption that capital and output wedges, τ_{ki} and $\tau_e\omega_i$, are sufficiently close to zero.

$\text{Cov}(\tau_{ki}, \omega_i)$.

To understand the underlying mechanisms through which carbon taxes affect aggregate TFP, I first examine their impact on factor and resource misallocation. Equation (1.20) demonstrates that a higher τ_e unambiguously increases MoL, as it introduces a higher dispersion in MRPL driven by the heterogeneous emissions intensities. In contrast, given a negative $\text{Cov}(\tau_{ki}, \omega_i) < 0$, a higher τ_e induces two countervailing effects on MoK in equation (1.19): it increases MoK through the $\text{Var}(\omega_i)$ term (reflecting new output distortions from varying emissions intensities that amplify the MRPK dispersion), but decreases MoK through the $\text{Cov}(\tau_{ki}, \omega_i)$ term (reflecting a corrective reallocation effect that reduces the MRPK dispersion). Consequently, the net impact of τ_e on MoK is ambiguous, depending on the relative strength of these two forces.

As a result, a higher τ_e also has an ambiguous effect on the overall misallocation of resources (MoR) in equation (1.21), the intuition behind which is twofold. On one hand, τ_e creates output distortion, particularly when firms have heterogeneous ω_i , which tends to increase the dispersion of marginal products, thereby exacerbating misallocation of both factors. On the other hand, τ_e directly incentivizes capital reallocation away from the dirty firms by increasing the cost of operations by more for this group. As cleaner firms typically demonstrate a higher MRPK before the tax is applied, this shift actively reduces the dispersion of MRPK, especially between dirty and clean firms, leading to improved capital allocation. The ultimate impact depends on which of these two counteracting forces dominates at different tax levels. I now formally state the impact of a carbon tax on resource misallocation (MoR) and TFP in *Proposition 1.3*. It establishes a non-monotonic effect of τ_e on TFP conditional on the distributions of τ_{ki} and ω_i , and the existence of a certain threshold value of τ_e where TFP performance is improved the most.

Proposition 1.3. *Assume 1) $\{\log z_i, \log(1 + \tau_{ki}), \log(1 - \tau_e \omega_i)\}$ follow a joint normal*

distribution, and 2) $\text{Cov}(\tau_{ki}, \omega_i) < 0$. Then there exists a threshold value

$$\bar{\tau}_e = \frac{-\alpha \text{Cov}(\tau_{ki}, \omega_i)}{\text{Var}(\omega_i)} > 0, \quad (1.22)$$

such that aggregate TFP (MoR) is strictly increasing (decreasing) in τ_e for $0 < \tau_e < \bar{\tau}_e$, and strictly decreasing (increasing) in τ_e for $\tau_e > \bar{\tau}_e$.

Proof: The detailed proof is provided in Appendix Section A.

Figure 1.3 illustrates the reallocation mechanism using simulations from the simple model. I consider three settings that vary the extent to which misallocation is tilted toward dirty firms: $\text{Cov}(\tau_{ki}, \omega_i) \in \{0, -0.025, -0.05\}$. $\text{Cov}(\tau_{ki}, \omega_i) = 0$ refers to the case without prior capital misallocation, and $\text{Cov}(\tau_{ki}, \omega_i) < 0$ refers to the case of capital misallocation in favor of dirty firms, which is observed from the data. $\text{Cov}(\tau_{ki}, \omega_i) = -0.025$ and -0.05 illustrate weaker and stronger capital misallocation toward dirty firms, respectively. Panel 1.3(a) shows that, when misallocation is tilted toward dirty firms, a higher τ_e reduces MRPK dispersion and alleviates capital misallocation by reallocating capital from dirty to clean firms; the correction is stronger when the initial misallocation is larger. Panel 1.3(b) shows that τ_e always increases MRPL dispersion, thereby exacerbating labor misallocation. These opposing effects on capital and labor allocation determine how τ_e influences the overall misallocation of resources and aggregate TFP. Panels 1.3(c) and (d) replicate *Proposition 1.3*: with capital misallocation tilted toward dirty firms, τ_e improves allocative efficiency up to a threshold at which TFP is maximized and resource misallocation is minimized. Intuitively, at modest tax rates the reduction in capital misallocation dominates, raising TFP; beyond a critical threshold, the increase in labor misallocation prevails. Accordingly, there exists a tax interval $\tau_e \in [0, \bar{\tau}_e]$ within which τ_e enhances TFP performance in the model economy, and the TFP-maximizing upper bound $\bar{\tau}_e$ increases with the extent to which capital misallocation is tilted toward dirty

firms. By contrast, when there is no pre-existing capital misallocation ($\text{Cov}(\tau_{ki}, \omega_i) = 0$), a carbon tax distorts the allocation of both factors and unambiguously lowers TFP. These results underscore an allocative efficiency channel that much of the existing literature often abstracts from when assessing the impacts of carbon taxes.


Optimal Carbon Tax. I now study the welfare implications of the allocative efficiency channel on the optimal carbon tax design. To keep the focus on tax-induced reallocation across firms, the household side is designed to be simple. A representative household supplies fixed endowments of capital $\bar{K}$ and labor $\bar{L}$ at competitive factor prices, and consumes the final good. The household also endures a utility cost from carbon emissions. Preferences are given by

$$U = \log Y - mE \quad (1.23)$$

with aggregate emissions $E = \sum_{i=1}^N e_i$ and marginal external cost of emissions $m > 0$. Because climate damages enter additively, they do not affect firms' static input choices; the tax influences allocations solely through firms' optimal decisions.

Let τ_e^* denote the welfare-maximizing carbon tax subject to the competitive equilibrium in Definition 1.1. In an environment with two market imperfections—climate externality and capital misallocation— τ_e^* is the second-best (i.e., constrained-optimal) carbon tax, because the constrained planner cannot perfectly undo the misallocation distortion. In the absence of capital misallocation, the optimal carbon tax serves solely a Pigouvian purpose by internalizing emissions damages. With Assumption 1 in *Proposition 1.3*, the planner's problem is equivalent to

$$\max_{\tau_e} -\frac{\sigma}{2} \text{Var}(\omega_i) \tau_e^2 - \sigma \alpha \text{Cov}(\tau_{ki}, \omega_i) \tau_e - mE(\tau_e). \quad (1.24)$$

Figure 1.3: Impacts of τ_e on allocative efficiency and welfare: an illustrative example

Notes: The simulation features 10,000 firms. I set $\alpha = 0.30$, $\sigma = 3$, $\bar{K} = 2$, $\bar{L} = 1$, and $m = 0.002$. Firm log productivity $\log(z_i)$ is normal with $\mathbb{E}(\log z_i) = 0$ and $\text{Var}(\log z_i) = 0.1$, and is independent of τ_{ki} or ω_i . Capital distortions τ_{ki} and emissions intensity ω_i are drawn from a bivariate normal: $\mathbb{E}(\tau_{ki}) = 0$, $\text{Var}(\tau_{ki}) = 0.05$; $\mathbb{E}(\omega_i) = 0.1$, $\text{Var}(\omega_i) = 0.05$; $\text{Corr}(\tau_{ki}, \omega_i) \in \{0, -0.5, -1\}$, implying $\text{Cov}(\tau_{ki}, \omega_i) \in \{0, -0.025, -0.05\}$. The figure plots misallocation indicators, aggregate TFP, and welfare in consumption-equivalent variation, each normalized to the laissez-faire benchmark ($\tau_e = 0$), across different values of $\text{Cov}(\tau_{ki}, \omega_i)$. More negative covariance reflects stronger misallocation of capital toward dirtier firms.

yielding the first-order condition

$$\sigma \text{Var}(\omega_i) \tau_e^* = -m \left. \frac{\partial E}{\partial \tau_e} \right|_{\tau_e = \tau_e^*} - \sigma \alpha \text{Cov}(\tau_{ki}, \omega_i) \quad (1.25)$$

The planner's first-order condition does not admit a closed-form solution. However, the simplified characterization in (1.25) decomposes the marginal benefits and costs associated with τ_e^* , offering insight into how τ_e^* depends on capital misallocation. The left-hand side captures the marginal misallocation cost of increasing τ_e that lowers TFP. The right-hand side comprises (i) the climate benefit from marginal externality correction, $-m \left. \frac{\partial E}{\partial \tau_e} \right|_{\tau_e = \tau_e^*} > 0$, and (ii) a reallocation benefit that raises TFP, $-\sigma \alpha \text{Cov}(\tau_{ki}, \omega_i)$, which is positive when capital distortions favor dirty firms ($\text{Cov}(\tau_{ki}, \omega_i) < 0$).

Two implications follow immediately. First, when capital is inefficiently tilted toward dirty firms, τ_e^* increases with $|\text{Cov}(\tau_{ki}, \omega_i)|$, holding the marginal climate benefit fixed across settings. Figure 1.3(e) shows that as this tilt strengthens, (i) the welfare peak shifts to the right, that is, the optimal carbon tax rises, and (ii) the peak height increases, so the welfare gain at the optimum over the laissez-faire is larger. This is because a stronger tilt expands the scope for the carbon tax to raise TFP by inducing more efficient reallocation of capital away from dirty firms. The TFP-maximizing tax $\bar{\tau}_e$ in Figure 1.3(d) corresponds to the optimum when climate damage is zero (i.e., $m = 0$)⁹; thus, even absent climate damage, a carbon tax can raise welfare by improving resource allocation. Second, the welfare-maximizing τ_e^* exceeds the TFP-maximizing level $\bar{\tau}_e$ because of the positive climate benefit. Consistent with panels (d) and (e) of Figure 1.3, the optimal carbon tax lies to the right of the TFP peak: internalizing emissions pushes the welfare optimum above the purely reallocation-driven level.

These results underscore the carbon tax's "double dividend" in the presence of capital

⁹In the simple model, total factor supply is fixed; hence, when $m = 0$, $\Delta \text{welfare} = \Delta \log Y = \Delta \log \text{TFP}$.

misallocation: it mitigates the climate externality and improves allocative efficiency, together implying a higher optimal carbon tax in this second-best setting than in a first-best no-misallocation benchmark. Ignoring capital misallocation as a market distortion therefore risks underestimating the optimal carbon tax.

1.4.3 Environmental Implications of Addressing Misallocation

I have shown that a carbon tax can improve both aggregate TFP and reduce climate damages, taking the existing capital wedges and their covariance with emissions intensity as given. In this section, I turn the focus to the potential impact of hypothetical policies designed to eliminate the existing capital wedges τ_{ki} such that all firms effectively face the same marginal product of capital. Capital wedges that were larger in cleaner firms lead to their allocation of capital and production activity being lower than the efficient level. Thus, the existing capital market distortions tend to hinder aggregate output and make it dirtier. Next I examine the effects of removing all capital wedges, i.e. setting $\tau_{ki} = 0$, on aggregate TFP, output, and emissions. Intuitively, the removal of these wedges should lead to the efficient allocation of capital across firms, particularly benefiting cleaner firms that currently face larger capital wedges. The induced reallocation can consequently raise aggregate TFP and output to their most efficient state, while simultaneously reducing the economy-wide emissions intensity as cleaner firms now account for a larger share of the aggregate production. I formalize this environmental impact of addressing capital misallocation in *Proposition 1.4*.

Proposition 1.4. *Assume 1) $\{\log z_i, \log \omega_i, \log(1 + \tau_{ki})\}$ follow a joint normal distribution, 2) $\text{Cov}(\tau_{ki}, \omega_i) < 0$ prior to removing τ_{ki} , and 3) intermediate goods are gross substitutes ($\sigma > 1$). Then eliminating capital wedges by setting $\tau_{ki} = 0$ unambiguously reduces the economy-wide emissions intensity $(\frac{E}{Y})$.*

Proof: The detailed proof is provided in Appendix Section A.




Figure 1.4: Impacts of addressing capital misallocation: an illustrative example

Figure 1.4 illustrates the impact of addressing the misallocation of capital that favors dirty firms with the same simulation setup as in Figure 1.3. A more negative $Cov(\tau_{ki}, \omega_i)$ indicates greater misallocation in favor of dirty firms. As capital misallocation decreases, both aggregate TFP and output Y improve, along with a notable reduction in aggregate emissions E . The reduction in E and the improvement in Y both contribute to a decrease in the aggregate emissions intensity of output, $\frac{E}{Y}$. These results support *Proposition 1.4*, showing that equalizing capital access for both dirty and clean firms can yield climate benefits. While removing existing capital distortions is well known to enhance aggregate productivity, the findings above underscore its significant environmental implications, suggesting that policies promoting equal capital access for firms with different emissions features could also help address climate challenges.

1.5 Quantitative Model

This section develops a quantitative, dynamic climate-economy model that incorporates financial frictions linked to firms' heterogeneous emissions intensities via differences in asset tangibility. The model adopts a heterogeneous-firm structure in the spirit of Gopinath et al. (2017), but extends it to a general equilibrium environment with tangibility-based borrowing constraints. The climate module includes a utility cost from cumulative emissions. Unlike the static model, the quantitative framework features a variable aggregate capital supply and endogenizes capital distortions through tangibility-dependent financial frictions, consistent with the empirical evidence in Section 1.3.3.

1.5.1 Firms

Firms produce differentiated products with a Cobb-Douglas production function $y_{i,t} = z_{i,t} k_{i,t}^\alpha l_{i,t}^{1-\alpha}$, where $z_{i,t}$ is firm i 's exogenous idiosyncratic TFP. Firm i 's total capital is a Leontief composite of both tangible ($k_{i,t}^T$) and intangible capital ($k_{i,t}^I$) stock (Caggese and Perez-Orive, 2022),

$$k_{i,t} = \min\left\{\frac{k_{i,t}^T}{\mu_i}, \frac{k_{i,t}^I}{1 - \mu_i}\right\} \quad (1.26)$$

where μ_i denotes the proportion of tangible capital in firm i 's total capital stock, capturing the firm's asset tangibility. For simplicity, it is assumed to be constant over time. Each firm i produces carbon emissions as a by-product and has a firm-specific emissions intensity per unit of revenue output, ω_i :

$$E_{i,t} = \omega_i \text{rev}_{i,t},$$

where $\text{rev}_{i,t} = p_{i,t} y_{i,t}$ (1.27)

The choice of revenue-based emissions intensity can be justified for the following reasons. First, I observe firms' sales revenue instead of physical output in the firm-level data, which facilitates the calculation of revenue-based emissions intensity. Second, the modeling assumption that firms have a constant emissions share in terms of sales has been applied in a wide range of environmental models incorporating firm heterogeneity in emissions intensity (Sogalla et al., 2024). Importantly, I assume ω_i is related to μ_i in the following functional form:

$$\omega_i = \mu_i^\epsilon \quad (1.28)$$

where $\epsilon > 0$ represents a positive elasticity of emissions intensity with respect to the share of tangible capital, in line with the empirical finding in Section 1.3.3 that firms with a higher share of tangible capital tend to be more emissions intensive. The Leontief structure in (1.26) implies that

$$k_{i,t}^T = \mu_i k_{i,t}, \quad k_{i,t}^I = (1 - \mu_i) k_{i,t}, \quad (1.29)$$

I assume for simplicity that both tangible and intangible capital face the same depreciation rate δ , so that the law of motion for total capital follows:

$$k_{i,t+1} = (1 - \delta) k_{i,t} + x_{i,t} \quad (1.30)$$

where firms' total investment $(x_{i,t})$ effectively comprises both tangible $(x_{i,t}^T)$ and intangible investments $(x_{i,t}^I)$:

$$x_{i,t}^T = \mu_i x_{i,t}, \quad x_{i,t}^I = (1 - \mu_i) x_{i,t} \quad (1.31)$$

It follows from condition (1.29) that the effective collateral value of firm i 's total capital, denoted by θ_i , can be expressed as:

$$\theta_i = \mu_i \theta^T + (1 - \mu_i) \theta^I \quad (1.32)$$

where θ^T and θ^I denote the collateral values of tangible and intangible capital, respectively. The key assumption is that $\theta^T > \theta^I$, reflecting that lenders can seize and recover a larger fraction of tangible assets than intangible ones in the event of default. In other words, tangible capital is more pledgeable as collateral than intangible capital. This assumption is supported by extensive empirical and theoretical literature on asymmetric information in credit markets and the role of collateral in easing borrowing constraints (e.g., Almeida and Campello, 2007; Dell'Ariccia et al., 2021; Caggese and Perez-Orive, 2022; Falato et al., 2022). Consequently, condition (1.32) implies that a higher share of tangible capital μ_i increases the collateral value of a firm's total assets.

A perfectly competitive aggregate producer produces final output (Y_t) as a CES bundle of intermediate goods y_{it} :

$$Y_t = \left(\sum_{i=1}^N y_{i,t}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad (1.33)$$

where σ is the elasticity of substitution between intermediate outputs. Each firm i is a monopolistically competitive producer facing a downward-sloping demand curve,

$$y_{it} = p_{it}^{-\sigma} D_t \quad (1.34)$$

where $p_{i,t}$ is firms' output price and D_t denotes the aggregate demand shifter.

Firm i 's objective is to maximize the expected present discounted lifetime utility from

its consumption $(c_{i,t})$:

$$U_i^f = \max_{\{c_{i,t}, x_{i,t}, p_{i,t}, l_{i,t}, k_{i,t+1}, b_{i,t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\log(c_{i,t}) - m \tilde{E}_t \right] \quad (1.35)$$

by choosing labor $(l_{i,t})$, price, investment, capital, and debt $(b_{i,t+1})$. The entrepreneur in each firm passively faces a linear utility cost from the stock of cumulative emissions $(m \tilde{E}_t)$, which they ignore when making production decisions. $m > 0$ is the marginal damage in each period from the cumulative carbon emissions $(\tilde{E})$ in the environment:

$$\tilde{E}_t = \sum_{\tilde{t}=0}^t \sum_{i=1}^N E_{i,\tilde{t}} \quad (1.36)$$

I follow Casey et al. (2026b) in specifying a linear climate damage component in the utility function. Golosov et al. (2014) introduce a climate economy model where they show that damages from CO₂ are linear in welfare under an exponential damage function on output with atmospheric carbon concentrations, log utility and full capital depreciation. Traeger (2023) establishes similar results in a fully calibrated climate model with partial capital persistence. The intuition is that while warming temperature is strongly concave in CO₂ concentration, damages are convex in temperature, yielding an overall linear welfare effect of carbon emissions. I adopt this linear damage formulation straightforwardly, which offers three advantages for the analysis. First, because I focus on U.S. outcomes, linear damages imply that the welfare cost of U.S. emissions is independent of both historical emissions prior to $t = 0$ and emissions originating abroad. Second, recent evidence indicates that non-market damages represent a substantial share of the social cost of carbon (e.g. Climate Impact Lab, 2022; Rennert et al., 2022; EPA, 2022), and these are readily incorporated into the parameter m . Third, with the additively separable utility cost, climate damages do not affect the dynamics of other macroeconomic variables, which greatly simplifies the

model's computational solution.

Firms hire labor, $l_{i,t}$, at a common wage w_t , and borrow using one-period debt, $b_{i,t+1}$ (> 0 if borrowing), from financial markets at a common interest rate r_t . Firms can finance their capital from external borrowing or internal funds (net worth). They face the budget constraint given by,

$$c_{i,t} + x_{i,t} + (1 + r_t)b_{i,t} + \tau_e E_{i,t} = p_{i,t} z_{i,t} k_{i,t}^\alpha l_{i,t}^{1-\alpha} - w_t l_{i,t} + b_{i,t+1} + \overline{\text{TR}}_t \quad (1.37)$$

where firms pay carbon taxes on their emissions if $\tau_e > 0$ and then each firm receives a common transfer $\overline{\text{TR}}_t$. Firms are also subject to capital accumulation (1.30), investment (1.31), demand function (1.34), and firm-specific borrowing constraint,

$$b_{i,t+1} \leq \frac{\theta_i(1 - \delta)k_{i,t+1}}{1 + r_{t+1}} \quad (1.38)$$

Condition (1.38) implies that the collateral value of a firm's capital (θ_i) positively affects its borrowing limit. The borrowing constraint thus indicates that a firm's maximum attainable debt is limited by the value of its pledgeable capital, which, according to condition (1.32), is determined by the tangibility of the firm's assets. This occurs because, if the borrower defaults, the lender can only seize a fraction θ_i of the firm's asset value. In the model, financial frictions arise from these tangibility-based borrowing constraints.

Now denote firm i 's net worth as $a_{it} = k_{it} - b_{it}$. Then the budget constraint can be rewritten as

$$c_{i,t} + a_{i,t+1} = \underbrace{p_{i,t} z_{i,t} k_{i,t}^\alpha l_{i,t}^{1-\alpha} - w_t l_{i,t} - (r + \delta)k_{it} - \tau_e E_{i,t}}_{\pi_{it}} + (1 + r_t)a_{it} + \overline{\text{TR}}_t \quad (1.39)$$

where the firm effectively maximizes a static profit first and then decides on intertemporal

consumption and savings (net worth a'). Thus in my model, the setup where firms own and accumulate capital is equivalent to the setup with a rental market of capital. And the firm's borrowing constraint can be rewritten as

$$k_{i,t+1} \leq \frac{1 + r_{t+1}}{1 + r_{t+1} - \theta_i(1 - \delta)} a_{i,t+1} \quad (1.40)$$

so that firm's maximum capital will be constrained by its net worth accumulation.

Recursive Formulation. The firm's problem can be rewritten in the recursive form. I rewrite $a = k - b$ as a single state variable defined as firm's net worth. I assume that at the end of each period t , firms' productivity shocks $z_{i,t+1}$ can be observed. Given this information, firms can choose to borrow $b_{i,t+1}$ to finance capital $k_{i,t+1}$ after knowing the next-period productivity $z_{i,t+1}$. This assumption effectively makes capital input a static choice and productivity-contingent¹⁰. A firm of type μ , with state (z, a) at the beginning of each period, faces a value function $V(z, a, \mu)$,

$$V(z, a, \mu) = \max_{c, a'} \log c + \beta \mathbb{E} V(z', a', \mu) \quad (1.41)$$

where $V(z, a, \mu)$ is independent of the climate damage term since firms do not account for the negative externality from firms' emissions in making intertemporal decisions; it chooses consumption (c) and next-period net worth (a') subject to the budget constraint,

$$c + \underbrace{a'}_{k' - b'} = \pi + (1 + r) \underbrace{a}_{k - b} + \overline{\text{TR}} \quad (1.42)$$

where π is the firm's maximized profits within the period and $\overline{\text{TR}}$ is the lump-sum transfer

¹⁰In this sense, I will focus on financial frictions as the only source for capital misallocation.

to firms from the government; and the borrowing constraint,

$$k' \leq \kappa' a', \quad (1.43)$$

$$\text{where } \kappa' = \frac{1 + r'}{1 + r' - \theta(1 - \delta)}, \quad (1.44)$$

$$\theta = \mu\theta^T + (1 - \mu)\theta^I \quad (1.45)$$

The Euler equation over consumption can be expressed as:

$$\frac{1}{c} = \beta \mathbb{E} \left[\frac{1}{c'} (1 + r' + \lambda' \kappa') \right] \quad (1.46)$$

in which the Lagrangian multiplier λ on the borrowing constraint enters the Euler equation. Therefore the expectation of a binding borrowing constraint in the future raises the net worth accumulation.

Within the current period, taking interest rate (r), wage (w), and carbon tax rate (τ_e) as given, a firm with the state vector (z, a, μ) solves a static profit maximization by choosing price (p), capital (k), and labor (l):

$$\pi = \max_{p,k,l} py - \tau_e E - wl - (r + \delta)k \quad (1.47)$$

subject to the production function $y = zk^\alpha l^{1-\alpha}$, the demand function $y = p^{-\sigma} D$ with D denoting the aggregate demand shifter, and the borrowing constraint:

$$k \leq \kappa a, \quad (1.48)$$

$$\text{where } \kappa = \frac{1 + r}{1 + r - \theta}, \quad (1.49)$$

$$\theta = \mu\theta^T + (1 - \mu)\theta^I. \quad (1.50)$$

The first-order conditions for labor, capital and price are given by:

$$(1 - \tau_e \mu^\epsilon)(1 - \alpha) \frac{py}{\eta l} = w \quad (1.51)$$

$$(1 - \tau_e \mu^\epsilon) \alpha \frac{py}{\eta k} = r + \delta + \lambda \quad (1.52)$$

$$p = \eta \text{MC} \quad (1.53)$$

where $\eta = \frac{\sigma}{\sigma-1}$ denotes the firm's constant markup of price over marginal cost. I present the full solution to firm's problem in Appendix A. The optimal capital decision indicates that in the laissez-faire setting (i.e. $\tau_e = 0$), all firms would face the same MRPK $:= \frac{\alpha py}{\eta k} = r + \delta$ if the borrowing constraint does not bind for any firm, that is, the Lagrangian multiplier associated with the borrowing constraint λ equals zero. In the case of binding borrowing constraints, financial frictions lead to the dispersion of MRPK across firms, and the return to capital increases with the firm's borrowing tightness captured by a higher λ . On the other hand, the allocation of labor across firms is undistorted by financial frictions in the laissez-faire.

1.5.2 Workers

There is a mass L of workers. Each worker is hand-to-mouth and simply consumes all the labor income ($w_t L$) in each period. Similar to firms, workers also have log utility over the consumption stream and are subject to a linear welfare cost from cumulative emissions. Their lifetime utility, without any uncertainty, is thus defined as:

$$U^w = \sum_{t=0}^{\infty} \beta^t \left[\log(w_t L) - m \tilde{E}_t \right]. \quad (1.54)$$

1.5.3 Government

The government manages the use of carbon tax revenue collected from firms in each period. The government maintains a balanced budget in each period where it simply exhausts all tax revenues to finance the lump-sum transfers to firms. Regarding the tax-rebate scheme, I assume the total revenue is equally split among the firms,

$$\sum_{i=1}^N \overline{\text{TR}}_t = \tau_e \sum_{i=1}^N E_{i,t} \quad (1.55)$$

such that each firm receives a fixed amount of payment $\overline{\text{TR}}_t$ in its period- t budget constraint. Different approaches to tax rebates among heterogeneous agents can potentially lead to varying welfare outcomes (e.g., Fried et al., 2024). However, as long as the redistribution of tax revenues is not structured to favor dirty firms, the positive impact of carbon taxes on allocative efficiency should always hold.

1.5.4 Recursive Competitive Equilibrium

I can now provide the definition of a general equilibrium, with a focus on the steady state (stationary equilibrium) of the model economy¹¹.

A stationary recursive competitive equilibrium conditional on a carbon tax (τ_e) includes value functions $V(z, a, \mu)$ for firms; policy functions capital $k(z, a, \mu)$, labor $l(z, a, \mu)$, price $p(z, a, \mu)$, consumption $c(z, a, \mu)$, and net worth $a'(z, a, \mu)$ for firms; a stationary probability distribution over firms $\Gamma(z, a, \mu)$; constant factor prices (w, r) ; and constant aggregate quantities including capital, output, investment, consumption, flow emissions and tax revenues, such that:

- Taking all prices, taxes, transfers and emissions as given, the policy functions and

¹¹This paper currently focuses on the steady state of the economy. Solving for the entire transition path from the pre-tax to the post-tax stationary equilibrium will be addressed in future work.

value functions solve the maximization problems for firms.

- All markets clear, including
 - Goods market clears: $Y = C + I = \int_Z \int_A \int_\mu p(z, a, \mu) y(z, a, \mu) d\Gamma(z, a, \mu) = D$, where C is the sum of firms' and households' consumption.
 - Labor market clears: $L = \int_Z \int_A \int_\mu l(z, a, \mu) d\Gamma(z, a, \mu)$
 - Asset market clears: $\int_Z \int_A \int_\mu a'(z, a, \mu) d\Gamma(z, a, \mu) = \int_Z \int_A \int_\mu k'(z, a, \mu) d\Gamma(z, a, \mu)$
- Aggregate flow emissions satisfies $E = \int_Z \int_A \int_\mu e(z, a, \mu) d\Gamma(z, a, \mu)$.
- Government maintains a balanced budget: $\bar{T} = \tau_e \int_Z \int_A \int_\mu e(z, a, \mu) d\Gamma(z, a, \mu)$
- The distribution of firms $\Gamma(z, a, \mu)$ is stationary and consistent with the distribution of μ , evolution of z , and decision rules a' :

$$\Gamma(z', a', \mu) = \int_Z \int_A \int_\mu 1 \{ (z', a', \mu) \text{ transition from } (z, a, \mu) \} d\Gamma(z, a, \mu)$$

1.5.5 TFP and Welfare

Measurement of TFP. Because capital choice is determined after observing the productivity shocks, the aggregate production function at each point in time can be expressed as:

Lemma 1.5. *With firms' MRPK defined in equation (1.52), the aggregate production function in the model economy can be written as:*

$$Y = \frac{\left[\int_i \left(\frac{(1-\tau_e \mu_i^\epsilon) z_i}{(\text{mrpk}_i / (r+\delta))^\alpha} \right)^{\sigma-1} di \right]^{\frac{\sigma}{\sigma-1}}}{\left(\int_i \frac{(1-\tau_e \mu_i^\epsilon)^\sigma z_i^{\sigma-1}}{(\text{mrpk}_i / (r+\delta))^{1+\alpha(\sigma-1)}} di \right)^\alpha \left(\int_i \frac{(1-\tau_e \mu_i^\epsilon)^\sigma z_i^{\sigma-1}}{(\text{mrpk}_i / (r+\delta))^{\alpha(\sigma-1)}} di \right)^{1-\alpha}} K^\alpha L^{1-\alpha} \quad (1.56)$$

where $K = \int_i k_i di$ and $L = \int_i l_i di$ are the aggregate capital and labor supply; and the aggregate TFP is measured by

$$\text{TFP} = \frac{\left[\int_i \left(\frac{(1-\tau_e \mu_i^\epsilon) z_i}{(\text{mrpk}_i / (r+\delta))^\alpha} \right)^{\sigma-1} di \right]^{\frac{\sigma}{\sigma-1}}}{\left(\int_i \frac{(1-\tau_e \mu_i^\epsilon)^\sigma z_i^{\sigma-1}}{(\text{mrpk}_i / (r+\delta))^{1+\alpha(\sigma-1)}} di \right)^\alpha \left(\int_i \frac{(1-\tau_e \mu_i^\epsilon)^\sigma z_i^{\sigma-1}}{(\text{mrpk}_i / (r+\delta))^{\alpha(\sigma-1)}} di \right)^{1-\alpha}} \quad (1.57)$$

The derivation of aggregate TFP is similar to that in Lemma 1.2 of the simple model, except that the firms' MRPK is now determined endogenously rather than from exogenous capital wedges. Consequently, TFP endogenously depends on the magnitude of carbon tax, distribution of firm-specific productivity, distribution of capital tangibility, the equilibrium rental rate of capital, the tightness of borrowing constraint for each firm.

Welfare. Since both firms and workers enjoy consumption in the model economy, it's crucial to properly define the welfare function in the heterogeneous agent economy in order to analyze the optimal climate policy. Incorporating the climate damage, I compute measures of firms' welfare, workers' welfare and aggregate welfare at any time t .

Lemma 1.6. *The lifetime welfare of firms (W_t^f) and workers (W_t^w), and the aggregate social welfare ($\bar{W}_t$) at time t takes the form:*

$$W_t^f = \int V_t(z, a, \mu) d\Gamma_t(z, a, \mu) - \sum_{s=t}^{\infty} \beta^{s-t} \frac{m}{1-\beta} E_s \quad (1.58)$$

$$W_t^w = \sum_{s=t}^{\infty} \beta^{s-t} \log(w_s L) - \sum_{s=t}^{\infty} \beta^{s-t} \frac{m}{1-\beta} E_s \quad (1.59)$$

$$\bar{W}_t = \Theta W_t^f + (1-\Theta) W_t^w, \quad (1.60)$$

where Θ is the Pareto weight of firms.

Lemma 1.6 informs that the firms' welfare depends on two terms, an aggregation of the

value function over the probability distribution of firms, as well as the lifetime utility cost from cumulative emissions. The workers' welfare depends on both the present discounted value of wage payments and the lifetime utility cost from emissions similar to the firms. Since the climate change damage arises from cumulative emissions in the environment, in each period, each additional unit of emissions leads to a marginal lifetime utility cost of $\frac{m}{1-\beta}$. The aggregate welfare is a weighted sum of firms' and workers' welfare, with the Pareto weight of firms (Θ) describing the relative importance assigned to the firms' well-being. The straightforward computation of welfare functions allows me to study how variations in carbon taxes τ_e differentially affect the welfare. In the later welfare analysis, I present the welfare outcomes in terms of permanent consumption equivalent variation (CEV), the derivation procedure of which is detailed in Appendix A.

I assume that a benevolent *social planner*'s objective is to maximize the aggregate welfare of the society, $\bar{W}_t$, by varying the carbon tax τ_e . Since I focus on steady state outcomes, the (Pareto) optimal carbon tax in the model is defined to be the τ_e that maximizes $\bar{W}_t$ in the steady state, i.e. the stationary equilibrium.

1.6 Model Calibration

The model is annual. I calibrate the model to the U.S. using firm-level balance-sheet and emissions data from Compustat and Trucost. Table 1.3 summarizes the parameter values for the baseline model of tangibility-dependent financial frictions.

I choose the conventional values for the set of parameters on preferences and production structure $(\beta, \delta, \alpha, \sigma)$. Following Gopinath et al. (2017), I set the discount factor $\beta = 0.87$, depreciation rate $\delta = 0.06$, and capital share $\alpha = 0.35$. The elasticity of substitution is set at $\sigma = 3$, a standard value for measuring imperfect substitution between intermediate products (e.g. Hsieh and Klenow, 2009; Gopinath et al., 2017). In the model, the two

Table 1.3: Baseline calibration

Description	Parameter	Value	Source
Discount factor	β	0.87	Gopinath et al. (2017)
Depreciation rate	δ	0.06	Gopinath et al. (2017)
Capital share	α	0.35	Gopinath et al. (2017)
Elasticity of substitution	σ	3	Standard
Collateral value: tangible	θ^T	0.9	Falato et al. (2022)
Collateral value: intangible	θ^I	0.0	Falato et al. (2022)
Productivity persistence	ρ_z	0.87	Compustat
Productivity standard deviation	σ_z	0.20	Compustat
Tangibility deciles	$\mu_{1,2,\dots,10}$	see Table 1.4	Compustat
Emissions intensity elasticity w/ tangibility	ϵ	0.31	Compustat & Trucost
Flow climate damage	m	0.034	Calibrated

Table 1.4: Tangibility deciles

μ_1	μ_2	μ_3	μ_4	μ_5	μ_6	μ_7	μ_8	μ_9	μ_{10}
0.02	0.04	0.07	0.09	0.12	0.16	0.22	0.31	0.46	0.74

parameters on the pledgeability of assets, θ^T and θ^I , govern the firm-specific tightness of the borrowing constraint. For my baseline model, I choose $\theta^T = 0.9$ and $\theta^I = 0$ based on the collateral values provided in Falato et al. (2022)¹².

I then directly estimate the parameters governing firm productivity, asset tangibility, and emissions - $(\rho_z, \sigma_z, \mu, \epsilon)$ - using firm-level balance sheet data from Compustat and emissions data from Trucost. The productivity process is estimated via the following AR(1) specification:

$$\log \tilde{z}_{it} = \rho_z \log \tilde{z}_{it-1} + \gamma_i + \gamma_{st} + u_{it} \quad (1.61)$$

¹²The authors estimate that the pledgeability of tangible assets lies between 0.9 and 1, while the pledgeability of intangible asset - except for patents and brands - is effectively zero. To allow for the possibility that certain intangible assets, such as intellectual capital, may have some collateral value, I also test the model with a modestly positive $\theta^I = 0.1$. The resulting simulation shows no quantitative differences from the baseline calibration.

where γ_i denotes firm fixed effects and γ_{st} the 2-digit industry-year fixed effects. The term $\tilde{z}_{it}$ is the estimated firm-level productivity constructed using the GMM method in Wooldridge (2009). To mitigate the downward bias in the estimated autoregressive coefficient that arises in short panels with firm fixed effects and a lagged dependent variable, I estimate regression (1.61) using a firm sample extended back to 1980, restricted to firms with at least one carbon-emissions observation and continuous coverage over the extended period. Details on the construction of extended firm sample and the estimation of $\tilde{z}_{it}$ are provided in Appendix A. The regression results yield a persistence parameter $\rho_z = 0.87$, and a standard deviation of productivity shocks $\sigma_z = 0.20$ from the average cross-sectional standard deviation of residuals u_{it} .

In the model, I choose to model 10 distinct types of the firms' asset tangibility (μ_i) based on the empirical distribution. Table 1.4 displays the results. In the data, the asset tangibility is averaged over time for each individual firm. The 10 unique types of μ_i are then set to the average values of 10 deciles of the asset tangibility distribution from the data. In period 0 of the model, each firm's tangibility is exogenously and randomly selected from these 10 types with equal probability, and remains permanent once chosen. To determine ϵ , the elasticity of firm's emissions intensity to asset tangibility, I run the following panel regression of the logarithm of revenue-based emissions intensity on the logarithm of tangibility,

$$\log \omega_{it} = \epsilon \log \mu_{it} + \gamma_{st} + u_{it}, \quad (1.62)$$

where the 2-digit industry-year fixed effect γ_{st} is included. I obtain $\epsilon = 0.31$ from the regression result.

Lastly, I calibrate the flow damage parameter from emissions, m , using the U.S. data of carbon emissions and GDP, as well as the recent estimate of the social cost of carbon.

I choose the value for $m = 0.034$ such that the optimal carbon tax in the benchmark model with no financial frictions ($\kappa = \infty$) matches the recent social cost of carbon emissions at \$120/tCO₂ from EPA (2022)¹³. Since the carbon taxes (τ_e) in the model are in abstract units, I convert the model units of τ_e to the real-world units in U.S. dollars. This is achieved through a back-of-the-envelope calculation by aligning the proportions of potential tax revenue in the U.S. GDP calculated from the model and the data. Details of this procedure are documented in Appendix A.

1.7 Quantitative Results

In this section, I first evaluate the calibrated model and compare its performance to the benchmark model commonly used in the literature in replicating firm-level patterns observed in the data. I then discuss the results from policy counterfactuals focusing on carbon taxes and the removal of financial frictions.

1.7.1 Model Validation

To evaluate the model's ability to replicate key empirical patterns, I compare the baseline model with tangibility-dependent financial frictions to a benchmark model that abstracts from such frictions, a conventional setup in the optimal climate policy literature. In the baseline model, I set the two parameters governing the financial frictions to $\theta^T = 0.9$ and $\theta^I = 0$, and keep all other parameters at their calibrated values reported in Table 1.3. Therefore, firms' borrowing (collateral) constraints relate to their emissions intensity through variation in asset tangibility, generating heterogeneous financial frictions across firms (with clean firms being more financially constrained). I thus refer to the baseline

¹³This calibration target is chosen as the majority of studies on the social cost of carbon and Pigouvian carbon tax do not consider the effects from financial frictions.

model as “HeFF” (heterogeneous financial frictions). The benchmark model, denoted “NoFF” (no financial frictions), removes financial frictions by setting $\kappa = \infty$ while keeping all other parameters unchanged, rendering the borrowing constraint non-binding so that all firms are financially unconstrained. Both models are evaluated under the laissez-faire scenario ($\tau_e = 0$), reflecting the current absence of a nationwide carbon tax in the United States.

Model Fit. Table 1.5 evaluates the models’ performance against the data by reporting a set of firm-level moments in the cross section that were not targeted during calibration. These untargeted moments capture the cross-sectional patterns of capital misallocation and heterogeneity in firms’ financial conditions. Comparing the HeFF model with the data assesses how well the baseline model reproduces empirical firm-level patterns, while comparing between the HeFF and NoFF models demonstrates how these moments change when heterogeneous financial frictions are omitted during calibration. Overall, the HeFF model matches the data moments more closely than NoFF, indicating that incorporating heterogeneous financial frictions improves the model’s ability to capture key aspects of firms’ behavior in emissions, production and financing conditions. Conversely, the NoFF model fails to replicate these empirical patterns.

First, only the HeFF model generates imperfect (yet positive) correlations between firm-level productivity ($\log z$) and both capital and output ($\log k$ and $\log y$), as well as a positive correlation with the marginal revenue product of capital ($\log \text{MRPK}$) seen in the data. Financial frictions in the HeFF model impede capital from flowing to its most productive uses, so firm productivity does not perfectly predict capital or output. This capital misallocation is reflected in systematically higher MRPK among more productive firms: firms experiencing higher productivity shocks have incentives to build up higher capital but might be constrained to do so due to financial frictions, thus $\log(\text{MRPK})$

Table 1.5: Firm-level cross-sectional moments: model vs. data

Moment	Data	HeFF	NoFF
Corr($\log k$, $\log z$)	0.52	0.87	1.00
Corr($\log y$, $\log z$)	0.90	0.98	1.00
Corr($\log \text{MRPK}$, $\log z$)	0.52	0.53	0.00
Corr($\log k$, $\log \omega$)	0.18	0.09	0.00
Corr($\log y$, $\log \omega$)	0.01	0.04	0.00
Corr($\log \text{MRPK}$, $\log \omega$)	-0.28	-0.14	0.00
Corr(leverage, $\log \omega$)	0.12	0.61	0.00
Corr(cash, $\log \omega$)	-0.27	-0.06	0.00

Notes: Moments are not targeted in the model calibration. Model moments (HeFF and NoFF) are computed from the cross section of simulated firms in the stationary equilibrium. Firm leverage is defined as $\frac{\max\{b,0\}}{k}$, and cash holdings as $\frac{|\min\{b,0\}|}{k}$. Data moments are computed from the cross section of firms in the sample and averaged over the period 2002-2021. The construction of firm-level variables in the data follows the definitions provided in Section 1.3.

will be higher for these firms. In contrast, in the absence of financial frictions, capital is allocated efficiently, implying that all firms face the same MRPK. Consequently, $\log z$ is perfectly correlated with $\log k$ and $\log y$, and there is no dispersion in $\log \text{MRPK}$ in the NoFF model.

Next, I demonstrate that the baseline HeFF model successfully replicates the misallocation patterns across firms with differing emissions intensities observed in the data, which arises from the main novel mechanism introduced in this paper. Empirically, firms' emissions intensity ($\log \omega$) is positively correlated with capital ($\log k$), only weakly correlated with output ($\log y$), and negatively correlated with the marginal revenue product of capital ($\log \text{MRPK}$). The baseline HeFF model reproduces these patterns. In the HeFF model, capital is misallocated toward dirtier firms, which are not necessarily more productive. Because dirty firms can borrow more easily against tangible capital, they tend to face weaker borrowing constraints, accumulate more capital without a commensurate

increase in output, and consequently exhibit lower returns to capital (lower MRPK)¹⁴.

The same mechanism also explains the observed relationship between firms' financial conditions and their emissions intensity. Consistent with the data, the HeFF model generates a positive correlation between leverage (debt-to-assets ratio) and $\log \omega$, and a negative correlation between cash holdings (cash-to-assets ratio) and $\log \omega$. These patterns arise naturally from firms' optimal financing decisions in the baseline model. Dirty firms can sustain higher leverage due to their lower shadow costs of external finance, captured by lower λ in model condition (1.52). In contrast, clean firms are more likely to be constrained in their external finance and thus accumulate more internal savings as a precautionary buffer to relax future borrowing constraints when investment opportunities arise.

The NoFF model, which abstracts from financial frictions, exhibits no systematic relationship between $\log \omega$ and firm-level outcomes, suggesting its inability to capture the key misallocation patterns observed in the data. As supporting evidence, Appendix A plots the model-implied cross-sectional variation in firm outcomes against emissions intensity, illustrating the role of tangibility-dependent financial frictions in driving misallocation toward dirty firms.

¹⁴The negative correlation between $\log \text{MRPK}$ and $\log \omega$ in the HeFF model is somewhat muted relative to the data, suggesting additional factors may contribute to this relationship. For instance, some emissions-intensive firms may be older and politically connected firms with outdated capital that enjoy easier access to financing, or there may exist firm-specific frictions or policy distortions correlated with emissions characteristics. Nonetheless, my results provide a conservative lower bound on the impact of capital misallocation on the optimal carbon tax. With stronger misallocation toward dirtier firms, the allocative efficiency gains from a carbon tax would likely be amplified, implying an even higher optimal carbon tax.


1.7.2 Allocative Efficiency and Macroeconomic Implications of Carbon Tax

Having shown the success of the baseline model in replicating key aspects of firm-level data, I now conduct a series of counterfactual policy experiments by introducing and varying carbon taxes τ_e . The experiments evaluate the allocative efficiency, macroeconomic, and distributional impacts of carbon taxes in both the baseline HeFF model and the benchmark NoFF model. The goal is to assess welfare implications and study the optimal carbon tax in each model scenario. Comparing these optima highlights the quantitative importance of accounting for tangibility-dependent financial frictions when designing optimal climate policy in a second-best world with capital misallocation.

For each policy scenario, I solve for the stationary equilibrium of the model conditional on different values of carbon tax τ_e and compute the corresponding model outcomes. The simulations span $\tau_e \in [\$0, \$440]/\text{tCO}_2$. In the discussion below, I refer to the case where $\tau_e = 0$ as the laissez-faire scenario.

Aggregate Implications of Carbon Tax. Figure 1.5 plots the allocative efficiency outcomes, i.e. TFP performance and capital misallocation, against τ_e for HeFF and NoFF models. Figure 1.5(a) describes the responses in aggregate TFP conditional on different values of τ_e . Specifically, it plots the changes relative to laissez faire in the difference between $\log(\text{TFP})$ and its efficient level, $\log(\text{TFP}^e)$. A higher point along each curve indicates that TFP is higher and closer to the efficient level and that the economy is moving towards the most efficient allocation in each model. For the baseline HeFF model with tangibility-dependent financial frictions, the inverse-U-shaped TFP curve against τ_e shows that intuition from the simple model in Section 1.4 also holds in the quantitative model. Consistent with the finding of *Proposition 1.3*, carbon taxes can improve TFP and

there exists a TFP-maximizing tax rate under financial frictions favoring dirty firms. The maximum TFP gain occurs at \$275/tCO₂, because the tax mitigates misallocation by prompting a reallocation of capital from less financially constrained dirty firms to more constrained clean firms, leading to a more productive use of capital. Strikingly, TFP remains improved over the laissez faire for taxes up to \$550/tCO₂¹⁵.



(a) $\log(\text{TFP}) - \log(\text{TFP}^e)$ (p.p.Δ from LF) (b) S.D. of $\log(\text{MRPK})$ (Δ from LF)

Figure 1.5: Impacts of carbon tax on TFP

Notes: All outcomes are shown relative to laissez-faire (LF) levels in each model.

By contrast, a higher τ_e in the NoFF model unambiguously lowers TFP and amplifies the misallocation, as there is no reallocation benefit when financial frictions are absent, and that a higher carbon tax continues to distort the economy farther away from the efficient allocation. Figure 1.5(b) plots the dispersion of firms' return to capital for different τ_e , measured by changes in the standard deviation of $\log(\text{MRPK})$ relative to laissez-faire. Even though changes in the dispersion of MRPK do not map one-to-one to changes in TFP performance, this dispersion metric is an informative indicator of capital misallocation that contributes to TFP losses. In the HeFF model, the dispersion

¹⁵The simulated results beyond $\tau_e = \$440/\text{tCO}_2$ are not shown in Figure 1.5 due to space limitations.




Figure 1.6: Impacts of carbon tax on macroeconomic variables

Notes: All outcomes are shown relative to laissez-faire (LF) levels in each model.

of MRPK consistently declines as τ_e rises, highlighting that carbon taxes can improve capital allocation when heterogeneous financial frictions are present. In contrast, MRPK dispersion rises with τ_e in the NoFF model, showing that ignoring misallocation from financial frictions can cause carbon taxes to merely worsen capital allocation.

In addition, I plot the responses of macroeconomic variables against τ_e in Figure 1.6. Figures 1.6(a) and (b) plot aggregate output and capital, respectively. In both models,

higher taxes dampen the accumulation of capital and depress output in the dynamic setting. However, the HeFF model exhibits a more muted decline compared with NoFF, because higher TFP gains from carbon taxes boost aggregate capital demand by increasing the marginal return to investment, which partially offsets the decline in aggregate capital supply caused by the taxes. Figures 1.6(c) and (d) present aggregate emissions and emissions intensity of output, both of which decrease as carbon tax rises. Similarly, the HeFF model shows a slightly smaller reduction in emissions outcomes relative to NoFF.

Cross-sectional Implications of Carbon Tax. I next examine how cross-sectional outcomes shape the aggregate results discussed above. Figure 1.7 illustrates how a carbon tax reshapes the allocation of capital across firms by reporting firms' share of total capital and their average $\log(\text{MRPK})$ for each emissions intensity decile. For both models, I compare outcomes under $\tau_e = \$275/\text{tCO}_2$ —the TFP-maximizing rate in the baseline HeFF model—to those under *laissez-faire* ($\tau_e = 0$).

Figures 1.7(a) and (b) show that in the *laissez-faire* equilibrium of the HeFF model, capital is initially concentrated among dirty firms, and the average $\log(\text{MRPK})$ declines with emissions intensity. Implementing the TFP-maximizing carbon tax induces a reallocation of capital from dirty to clean firms: the top four deciles experience reductions in capital shares, while the bottom five gain. Because capital is initially skewed toward dirty firms, this reallocation compresses cross-group disparities in capital holdings, moving the economy closer to the efficient allocation characterized by equalized marginal products of capital. Consequently, the MRPK gap between clean and dirty firms narrows, as the negative slope of MRPK with respect to emissions intensity becomes less pronounced. This compression in cross-group MRPK also explains the decline in overall MRPK dispersion shown in Figure 1.5(b), highlighting that the TFP gains in Figure 1.5(a) primarily arise from a more efficient reallocation of capital from dirty to clean firms.




Figure 1.7: Impacts of carbon tax on reallocation of capital across firms

By contrast, Figures 1.7(c) and (d) show that in the NoFF model, the laissez-faire equilibrium is already close to the efficient allocation: capital shares are relatively uniform across emissions-intensity deciles, and MRPK is essentially flat across firm types. An identical increase in τ_e reallocates capital from dirty to clean firms, making the capital distribution less even across these groups. In this case, the carbon tax amplifies MRPK differences between clean and dirty firms, moving the economy away from the efficient allocation. This pattern is consistent with the rise in overall MRPK dispersion and the

decline in TFP associated with the tax in the NoFF model, shown in Figures 1.5(b) and (a) respectively.

Optimal Carbon Tax. For the welfare analysis, I set the Pareto weight of firms (Θ) in the aggregate welfare function in equation (1.60) to match their consumption share in the laissez-faire equilibrium of the baseline HeFF model, which yields $\Theta = 0.54$ ¹⁶. This calibration implicitly assumes that the relative importance, or bargaining power of firms and workers in aggregate welfare is proportional to their respective consumption shares.




Figure 1.8: Impacts of carbon tax on aggregate welfare (CEV)

Notes: All welfare outcomes are shown as permanent consumption equivalents in percentage relative to laissez faire in each model.

Figure 1.8 plots aggregate steady-state welfare as a function of the carbon tax, measured in terms of the permanent consumption equivalent variation (CEV) relative to laissez-faire

¹⁶The firms' consumption share is approximately the same in the laissez-faire equilibria of both models.

in each model.¹⁷ It also reports the steady-state welfare-maximizing carbon tax in each case. In the benchmark NoFF model, the optimal carbon tax (τ_e^{*NoFF}) equals \$120/tCO₂¹⁸, whereas in the baseline HeFF model, the optimal tax (τ_e^{*HeFF}) equals \$210/tCO₂. Thus, the presence of tangibility-dependent financial frictions raises the optimal carbon tax by about 75% relative to the Pigouvian benchmark in the frictionless economy.

This result highlights that when capital misallocation arises from tangibility-dependent financial frictions, the optimal carbon tax exceeds the Pigouvian level, consistent with the theoretical insight from Section 1.4. It also reveals a “double dividend” effect: a carbon tax not only reduces emissions but also enhances TFP in the presence of financial frictions. However, in the baseline model, the optimal carbon tax (\$210/tCO₂) remains below the level that would maximize TFP gains (\$275/tCO₂), since excessively high taxes depress aggregate capital and output, which tends to lower welfare. Nonetheless, the TFP gains in the HeFF model mitigate the contraction in capital accumulation shown in Figure 1.6(b), resulting in a higher optimal carbon tax than in the frictionless benchmark.

The welfare implications also differ in magnitude between models. In the HeFF model, moving from laissez-faire to τ_e^{*HeFF} increases steady-state welfare by the equivalent of a 0.32% permanent rise in consumption, whereas in the NoFF model, the corresponding welfare gain under τ_e^{*NoFF} is only 0.14%. Furthermore, welfare remains above the laissez-faire level for carbon taxes up to \$367/tCO₂ in the HeFF model and only \$262/tCO₂ in the NoFF model, indicating a substantially broader range of welfare-improving tax rates in the presence of financial frictions. In Appendix Section A, I report welfare outcomes for firms and workers, respectively.

¹⁷Appendix A details the computation of welfare in CEV terms.

¹⁸This value serves as the calibration target for the climate damage parameter m in Section 1.6.

1.7.3 Aggregate and Environmental Implications of Removing Financial Frictions

In this section, I abstract from carbon tax and examine the aggregate implications of removing financial frictions by comparing outcomes in the HeFF and NoFF economies under laissez-faire. Comparing these two model economies provides insight into how hypothetical policies that address existing financial frictions—and the associated capital misallocation—could affect aggregate output, emissions, and welfare. Conceptually, this experiment parallels the counterfactual exercise in Section 1.4.3 using the static model.

Table 1.6 presents aggregate outcomes and relative changes in both models. Aggregate output Y in the NoFF economy exceeds that in HeFF by 6.21 percent, owing to productivity gains associated with the removal of financial frictions. In the HeFF model, the TFP loss equals 2.32 percent, which is expressed as the percentage deviation of TFP from the efficient allocation. Removing financial frictions fully eliminates this allocative inefficiency, thereby achieving the efficient allocation of capital across firms in the NoFF model. Unlike the static model, which holds aggregate capital fixed, the quantitative model features an endogenous aggregate capital supply that responds to changes in TFP. As productivity improves, the equilibrium interest rate rises¹⁹, inducing further capital accumulation. This amplification channel magnifies the impact of the TFP improvement on output, yielding a larger increase in Y than the TFP gain alone would imply.

Regarding the environmental outcomes, aggregate emissions E rise by 4.75 percent following the removal of financial frictions, largely due to higher output. The increase in emissions, however, is smaller than the rise in Y , resulting in a 1.37 percent decline in aggregate emissions intensity (E/Y). This pattern is consistent with *Proposition 1.4*: as financial constraints that disproportionately bind clean firms are relaxed, capital

¹⁹The equilibrium interest rate rises from $r = 0.10$ in HeFF to $r = 0.14$ in NoFF.

reallocates from high- to low-emissions producers, lowering the economy-wide emissions intensity.

Lastly, aggregate welfare increases substantially. The elimination of financial frictions generates a welfare gain equivalent to a 4.5 percent permanent increase in consumption - an order of magnitude larger than the 0.32 percent welfare improvement achieved under the optimal carbon tax when financial frictions are present.

Table 1.6: Aggregate implications of addressing financial frictions

Variable	HeFF	NoFF	% Δ
Y	1.205	1.280	+6.21%
E	0.692	0.725	+4.75%
$\frac{E}{Y}$	0.575	0.567	-1.37%
TFP loss	-2.32%	0	
Welfare (CEV)	0	4.5%	

Notes: Output (Y), emissions (E), emissions intensity ($\frac{E}{Y}$) are reported in model units for each economy, with the last column showing percentage differences between NoFF and HeFF. TFP loss is measured as the percent deviation of aggregate TFP from its efficient level. Welfare is expressed as the permanent consumption-equivalent variation relative to the baseline HeFF model.

1.8 Conclusion

This paper examines how carbon taxes interact with financial frictions that distort capital allocation across firms with differing emissions intensity. Using data on U.S. public firms, I document a robust negative relationship between emissions intensity and the marginal revenue product of capital, indicating that capital is inefficiently allocated toward dirtier firms. Moreover, dirty firms hold more tangible capital and exhibit higher leverage. These empirical patterns point to a misallocation of capital in favor of dirty firms due to tangibility-dependent financial frictions: dirty firms possess more tangible assets and have greater borrowing capacity because of the higher collateral value of their assets, whereas

clean firms, with more intangible assets, face tighter constraints in accessing credit.

Building on this evidence, I develop climate-economy models with heterogeneous firms to analyze how carbon taxes affect allocative efficiency and welfare in an environment with capital misallocation. I provide both analytic insights and quantitative results on the “double dividend” effects of carbon taxes: reducing emissions while improving aggregate productivity. The productivity gains arise from tax-driven capital reallocation from less financially constrained, dirty firms to more constrained, clean firms. The resulting welfare improvements imply that the optimal carbon tax under pre-existing capital misallocation exceeds the Pigouvian tax in a benchmark economy without misallocation. Furthermore, policies that address capital misallocation by removing financial frictions alone can reduce aggregate emissions intensity by shifting resources toward clean production, while simultaneously boosting productivity and output, yielding even higher welfare gains than imposing a carbon tax alone. In conclusion, these analyses highlight the importance of recognizing the sources and biases of existing capital misallocation in designing climate policies that jointly advance environmental and economic objectives.

1.9 Acknowledgments

I am deeply grateful to my advisors, Gregory Casey, Yueyuan Ma, and Peter Rupert, for their guidance and support throughout this project. I also thank Alexander Abajian, Eugenie Dugoua, Maarten De Ridder, Max Farrell, Don Fullerton, Stephie Fried, Hugo Hopenhayn, Derek Lemoine, Simon Margolin, Kyle Meng, Antony Millner, Laura Murphy, Stephen Terry, Christian Traeger, Eugene Tan, Zebang Xu, and Pengyue Zhu for their valuable suggestions and feedback. I am grateful as well for the comments from seminar participants at the AERE Summer Conference, the Midwest Macro Meetings, the Occasional Workshop, the West Coast Micro-to-Macro Development Workshop, and the UCSB

Environmental and Macroeconomics reading groups. Any remaining errors are my own.

Chapter 2

Energy Efficiency Dynamics and Climate Policy

2.1 Introduction

Many national governments are instituting policies meant to reduce carbon emissions. How effective will these *climate change mitigation* policies be at reducing carbon emissions? We address this question in a quantitative, dynamic macroeconomic model that includes energy as an input to production. The key benefit of the model is that it has realistic transition dynamics. Transition dynamics are crucial to climate policy analysis for two reasons. First, climate change is a function of the stock of carbon in the atmosphere, rather than the flow of emissions. So, the environmental impact of a mitigation policy depends on how it affects emissions in the short, medium, and long run. Second, many climate policy goals are framed in terms of medium-run emissions. For example, the United States' Nationally Determined Contribution to the Paris Agreement focuses on targets for flow emissions in 2030. Despite the importance of transition dynamics for climate policy analysis, existing climate-economy models are calibrated to static or steady

state moments (e.g., Golosov et al., 2014; Hassler et al., 2016; Hassler and Krusell, 2018; Barrage and Nordhaus, 2024).

We focus on the dynamics of energy efficiency, an important determinant of carbon emissions (Raupach et al., 2007; Peters et al., 2017; Le Quéré et al., 2019). Put differently, we study substitution between energy and non-energy factors of production. Existing work finds that the short-run elasticity of substitution is near zero, but the long-run elasticity is around one (Atkeson and Kehoe, 1999; Hassler et al., 2021a). Climate-economy models generally assume that the elasticity is one in both the short and long run (e.g., Golosov et al., 2014; Hassler et al., 2016; Barrage, 2020; Cruz and Rossi-Hansberg, 2024). This causes standard models to overestimate the degree to which energy efficiency reacts to policy-induced changes in energy prices in the short and medium run (Casey, 2024).

To generate realistic dynamics in our model, we incorporate the insights of Hassler et al. (2021a) and the modeling framework of Leon-Ledesma and Satchi (2019) into a neoclassical production setting where final goods are produced with capital, labor, and energy. The result is a tractable setting in which the elasticity of substitution between energy and non-energy inputs can vary over time due to endogenous technology choice. We consider a model with two types of technology, both of which increase output. One technology augments energy and improves energy efficiency, and the other technology augments a capital-labor composite and decreases energy efficiency. Firms pay an adjustment cost when they alter their technology mix. The size of the adjustment costs determines how slowly energy efficiency responds to a change in energy prices.

We calibrate the model to the energy efficiency transition dynamics observed in the data. We use aggregate U.S. data on energy use and expenditures from the Energy Information Administration (EIA) to estimate an empirical impulse response function (IRF) that shows how economy-wide energy efficiency responds to changes in the real price of energy. The IRF suggests that energy efficiency responds slowly to changes in

energy prices. We calibrate the adjustment costs and short-run substitution elasticity in the model so that it generates an IRF that is similar to the one estimated in the data. The calibrated model fits the data closely. We also show that the approach used in leading climate-economy models cannot replicate the patterns observed in the data, because it assumes that energy efficiency responds quickly to changes in energy prices.

To analyze the importance of slow transition dynamics, we use the model to study the effectiveness of climate change mitigation policies. We focus on comparing our results to the standard approach used in the existing literature.

Slow transition dynamics significantly hinder the ability of carbon taxes to reduce cumulative carbon emissions. Carbon taxes increase the price of energy and induce substitution away from energy and toward non-energy inputs. In other words, they increase energy efficiency. In our model and the data, this process happens slowly, as firms adjust their technology choices in response to changes in tax-inclusive energy prices. In standard models, however, the adjustment process happens immediately.

We start by examining a constant carbon tax. Our model predicts that a 50 percent ad valorem carbon tax reduces cumulative carbon emissions by 22 percent over the next decade, relative to a no-policy scenario. In the standard model, by contrast, the same tax reduces emissions by 35 percent. By this metric, accounting for the slow transition dynamics in the data reduces the estimated effectiveness of carbon taxes by around one third.

We then examine the least-cost path of carbon taxes that achieves an exogenous policy target for cumulative emissions reductions. Given the slow evolution of energy efficiency, our model requires ad valorem carbon taxes that are roughly 16% (6 p.p.) higher to achieve the target, relative to the prediction of the standard model. The reductions in dirty energy use also occur later in the century. The welfare costs of meeting the target are small in both models, and roughly 19% (0.09 p.p.) higher when accounting for slow

transition dynamics.

2.2 Related Literature

Starting with the DICE model (Nordhaus, 1993), many studies have investigated climate policy using growth models (e.g., Golosov et al., 2014; Hassler et al., 2021b; Barrage and Nordhaus, 2024; Cruz and Rossi-Hansberg, 2024). When energy efficiency is included in these models, it is generally done with a Cobb-Douglas aggregate production function, implying that the short- and long-run elasticities of substitution are both equal to one. This implies that energy efficiency reacts quickly to changes in energy prices. We build on this literature by constructing a model that matches the empirical dynamics of energy efficiency and showing that these dynamics have significant consequences for the effectiveness of mitigation policies.

There is also a long literature documenting empirical facts about energy use (e.g., Van der Werf, 2008; Koetse et al., 2008), including the fact that the elasticity of substitution differs in the short and long run (e.g., Pindyck, 1979; Hassler et al., 2021a). Additionally, there is a long literature presenting models that explain the difference between the short- and long-run elasticities. Older models relied on putty-putty or putty-clay models with adjustment costs and many types of capital (e.g., Pindyck and Rotemberg, 1983; Atkeson and Kehoe, 1999), while newer models stress directed technical change in more aggregated settings (e.g., André and Smulders, 2014; Hassler et al., 2021a). We contribute to this literature by estimating the full transition dynamics that occur as an economy moves from the short run to the long run and by demonstrating the importance of these slow dynamics for the effectiveness of climate change mitigation policies.

Our work also contributes to a small but growing literature on transition dynamics and climate policy. The most closely related papers are those of Hassler et al. (2021a) and

Casey (2024). Hassler et al. (2021a) show that a model with directed technical change can match both the short- and long-run data patterns.¹ Casey (2024) builds on these insights to study climate policy. He uses a model of directed technical change to study the role of energy efficiency transition dynamics in determining the effectiveness of carbon taxes. The model is calibrated to steady state moments only. He utilizes a short-run elasticity of zero and does not report a long-run elasticity, which cannot be found analytically in standard directed technical change models. By utilizing the specific model of directed technical change from Leon-Ledesma and Satchi (2019), we overcome this problem and specify a model that can exactly match both the short- and long-run elasticities. Thus, we are able to study the impacts of energy efficiency dynamics on climate policy in a more empirically-grounded model.

More generally, our work is part of a growing literature on slow dynamics in macro climate-economy models. Like us, Hinkelmann (2024a) builds on this existing directed technical change literature by utilizing the Leon-Ledesma and Satchi (2019) framework. He considers a richer model with several margins of adjustment and similarly shows that transition dynamics matter for the design of carbon taxes that can implement medium-run climate policy goals in the United States. Unlike us, he does not match the transition dynamics to data.^{2,3} Airaud et al. (2025) take a somewhat different approach and embed the framework of Hassler et al. (2021a) in an open-economy business cycle model to study the impact of climate policy on both economic and environmental outcomes. Campiglio et al. (2026) study optimal climate policy in a model with capital adjustment costs. They calibrate their model to the outcomes of granular energy models, rather than directly to the transition dynamics observed in aggregate data. Like us, they find that transition

¹See also André and Smulders (2014) who focus on matching the long-run patterns in a similar model.

²See also Farajpour and Hinkelmann (2021), who confirm the importance of distinguishing between short- and long-run elasticities, but do not include transition dynamics.

³Hinkelmann (2024b) uses related tools to study substitution between electricity and fossil fuels.

dynamics play an important role in climate policy. Our work complements these earlier studies in two ways. First, and most importantly, we match our model to the transition dynamics observed in data. Second, we stress the differential effect of energy efficiency dynamics on carbon taxes and clean energy subsidies, the latter of which are central to recent climate legislation in the United States.

In addition to the papers cited above, our work is related to the broader literature on directed technical change and the environment (e.g., Acemoglu et al., 2012, 2023; Fried, 2018). As explained in Section 2.4.3, the Leon-Ledesma and Satchi (2019) approach we adopt is similar to a directed technical change model with a fixed quantity of R&D inputs, the specification adopted in much of the recent literature. Our approach implicitly chooses the market-clearing condition for R&D inputs such that there is a closed-form solution for the long-run elasticity of substitution, which is equal to one (see also Jones, 2005; Caselli and Coleman II, 2006). This makes it possible to calibrate the model to the long-run and transition-path data. The Leon-Ledesma and Satchi (2019) approach also differs from the directed technical change literature by adding adjustment costs. This is particularly important in our setting, because it is necessary when including both the endogenous evolution of technology and a long-run elasticity equal to one. In this way, our model is leveraging directed technical change as a way to improve the quantitative modeling of transition dynamics at the cost of capturing some important features of technological progress.

2.3 Motivating Empirical Evidence

In this section, we quickly review the well-known empirical facts that motivate our study. Our energy data come from the State Energy Data System (SEDS) produced by the U.S. Energy Information Administration (EIA). We are interested in macro elasticities




Figure 2.1: Final Energy Price and Expenditure Share

Notes: Panel (a) shows the average real price of final energy and the energy expenditure share. Panel (b) shows the deviation of the real price from its deterministic exponential trend and the deviation of the energy expenditure share from its average value.

and focus on aggregate data for the United States. Appendix Section B provides a more complete description of the data. The data include quantities for different types of primary energy (e.g., coal, oil, renewables) and final energy (e.g., electricity, gasoline). The data also include prices and expenditures for fossil-fuel sources of primary energy, as well as for final energy, but do not include prices for renewable sources of primary energy, such as wind and solar.

Our analysis uses two empirical objects, the energy expenditure share of output and the real price of final energy. The energy expenditure share is measured as energy expenditure divided by the sum of final output (GDP) and energy expenditure, matching the model we present in the next section. Our measure of the energy price is the BTU-weighted average across energy sources. Hassler et al. (2021a) propose a method of calculating aggregate prices that weights the quantities of different energy types by their average prices. In Appendix Section B, we show that adopting their methodology has virtually no effect on our price series. To convert nominal to real prices, we use the GDP deflator.

Figure 2.1 plots the two empirical objects. Panel (a) focuses on long-run trends. The real price of final energy has been increasing over this period, while the expenditure share has been declining. Together, these two observations suggest that the elasticity of substitution (EoS) between energy and non-energy inputs is greater than one. Given the relatively short time period, it is hard to rule out the hypothesis that there is no trend in the share, which would suggest a long-run EoS of one. This is consistent with the econometric literature that distinguishes between short- and long-run elasticities and finds long-run elasticities at or above one (Griffin and Gregory, 1976; Pindyck, 1979; Koetse et al., 2008). The macro climate literature has interpreted the evidence as suggesting that a unitary elasticity, i.e., Cobb-Douglas production, is a good fit for the data (e.g., Golosov et al., 2014; Barrage, 2020; Hassler et al., 2021b). Since our goal is to demonstrate the importance of transition dynamics in climate-economy models, we follow the existing literature and focus on the case of a unitary long-run elasticity throughout the paper. In other words, we assume that any trend in the expenditure share is unrelated to changes in the relative price of energy and non-energy inputs.

To highlight short- and medium-run fluctuations, panel (b) plots the deviations of real energy price and the expenditure share from their deterministic exponential trends. The two series move together closely, indicating that the short-run EoS is well below one. Again, this is consistent with the econometric literature, which finds short-run elasticities that are close to zero when looking at very short time horizons (e.g., Berndt and Wood, 1975; Hassler et al., 2021a).

2.4 Model

2.4.1 Model Specification

Time is discrete and indexed by t . Each period corresponds to one year. All production is perfectly competitive.

Our model is focused on aggregate energy efficiency, i.e., substitution between energy and non-energy inputs to production. The representative final good producer has access to a CES production technology:

$$\tilde{Y}_t = \left((A_{N,t}N_t)^{\frac{\epsilon-1}{\epsilon}} + (A_{E,t}E_t)^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}, \quad (2.1)$$

where $\tilde{Y}_t$ is gross output, N_t is non-energy inputs in production (capital and labor), E_t is energy, and $\epsilon \in (0, \infty)$. The variable $A_{N,t}$ is non-energy-augmenting technology, and $A_{E,t}$ is energy-augmenting technology, which is closely related to energy efficiency. The final good producer chooses a sequence of inputs, $\{K_t, L, E_t\}_{t=0}^{\infty}$, and a sequence of technologies, $\{A_{N,t}, A_{E,t}\}_{t=0}^{\infty}$, to maximize the expected present discounted value of future profits. The price of the final good is normalized to one. We use $p_{N,t}$ to denote the price of the non-energy composite, and $p_{E,t}$ to denote the price of energy.

The representative producer can choose the two technology terms from a menu

$$\ln X_t = (1 - \eta) \ln A_{N,t} + \eta \ln A_{E,t}, \quad (2.2)$$

where $\eta \in (0, 1)$. As shown below, X_t will be a natural measure of the overall level of technology in the economy. The growth rate of X_t , $g_{X,t}$, is exogenous. When calibrating the model, we consider a stochastic process for $g_{X,t}$ and in our policy simulations, we consider a deterministic process.

Let $\theta_t \equiv A_{E,t}/A_{N,t}$ be the technology ratio. When firms adjust θ_t , they lose a fraction of total output due to quadratic adjustment costs

$$\Phi\left(\frac{\theta_t}{\theta_{t-1}}\right) = \frac{\phi}{2} \left(\frac{\theta_t}{\theta_{t-1}} - 1\right)^2. \quad (2.3)$$

Given the convex adjustment costs, firms adjust θ_t slowly over time in response to a change in relative input prices. The parameter $\phi \geq 0$ determines how quickly this adjustment occurs. Since there are adjustment costs, the final good producer looks ahead when making decisions. It discounts the future at the rate of interest.⁴

Combining production function (2.1) and the technology menu (2.2), output can be written as

$$\tilde{Y}_t = X_t \left((\theta_t^{-\eta} N_t)^{\frac{\epsilon-1}{\epsilon}} + (\theta_t^{1-\eta} E_t)^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}. \quad (2.4)$$

The overall level of technology, X_t , is exogenous and the direction of technology, θ_t , is endogenous.

An intermediate producer combines capital and labor into the non-energy input bundle using a Cobb-Douglas technology:

$$N_t = K_t^\alpha L^{1-\alpha}. \quad (2.5)$$

We assume that the size of the labor force is constant. For prices, w_t denotes the wage rate, ρ_t denotes the rental rate on capital, and $\alpha \in (0, 1)$.

Our focus is on substitution between energy and non-energy inputs. We assume that

⁴In the calibration, we use a stochastic version of the model. In this case, the firm uses the household's stochastic discount factor.

the production function for final energy is given by

$$E_t = E_{d,t}^\gamma E_{c,t}^{1-\gamma}, \quad (2.6)$$

where E_d is the quantity of dirty energy, E_c is the quantity of clean energy, and $\gamma \in (0, 1)$. Golosov et al. (2014) set the EoS between energy sources equal to 0.95, and we set it equal to one for computational convenience. In robustness analyses, we also consider a CES production function for E_t . We consider this simplified model, rather than allowing the elasticity to vary over time, because our focus is on the role of energy efficiency dynamics.

Perfect competition implies that the price of energy is

$$p_{E,t} = \tilde{\gamma} p_{c,t}^{1-\gamma} (1 + \tau_{d,t})^\gamma p_{d,t}^\gamma, \quad (2.7)$$

where $p_{c,t}$ is the cost of extracting clean energy from the environment, $p_{d,t}$ is the cost of extracting dirty energy from the environment, $\tau_{d,t} \geq 0$ is an ad valorem tax on dirty energy, and $\tilde{\gamma} = \gamma^{-\gamma} (1 - \gamma)^{\gamma-1}$ is a constant. In our policy experiments, we assume that extraction costs grow at constant exogenous rates, g_{p_c} and g_{p_d} . The budget is balanced with lump-sum rebates.

It is important to note that we do not allow for slow substitution between clean and dirty primary sources of energy in the production of energy services. Our focus is on energy efficiency, i.e., substitution between energy and non-energy inputs. Casey et al. (2026a) provide some suggestive evidence that a Cobb-Douglas production function is a reasonable fit for the long-run data and slow transition dynamics may also be important for substitution between energy sources. As they point out, a time series of aggregate clean energy prices is not available for the United States, making it impossible to apply our method to this substitution margin.

Final output (Y_t) is gross output minus energy expenditures and adjustment costs:

$$Y_t = \left(1 - \frac{\phi}{2} \left(\frac{\theta_t}{\theta_{t-1}} - 1\right)^2\right) \times \tilde{Y}_t - p_{c,t}E_{c,t} - p_{d,t}E_{d,t}. \quad (2.8)$$

This is the measure of gross domestic product (GDP) in our model.

A representative household chooses sequences of consumption (C_t) and investment (I_t) to maximize lifetime utility

$$U := \sum_{t=0}^{\infty} \beta^t \ln(C_t), \quad (2.9)$$

taking all prices and technology levels as given. Here, $\beta \in (0, 1)$ is the discount factor.

The budget constraint for the representative household is

$$C_t + I_t = w_t L + \rho_t K_t + T_t = Y_t, \quad (2.10)$$

where T_t is rebated tax revenue. Capital evolves according to

$$K_{t+1} = I_t + (1 - \delta)K_t, \quad (2.11)$$

where $\delta \in (0, 1)$ is the depreciation rate.

We now have all of the components necessary to define a competitive equilibrium.⁵

Definition 2.1. A *competitive equilibrium* is a sequence of prices, quantities, and technologies such that (1) firms maximize profits, (2) all market clearing conditions hold, and (3) individuals maximize lifetime utility.

⁵Appendix Section B presents the full set of equilibrium equations for the model.

2.4.2 The Elasticity of Substitution (EoS)

Short- versus Long-Run EoS. The key idea behind this production structure is that, following a change in energy prices, the economy behaves differently in the short and long run. In the very short run, technology is fixed and the producer can only respond to energy prices by substituting non-energy inputs for energy inputs using the CES function (2.1). At longer time horizons, the producer can also change technology terms, which makes production more flexible and leads to more substitution in response to changes in relative prices.

To further explore this intuition in terms of the elasticity of substitution between energy and non-energy inputs, we will now briefly discuss the theoretical results from Leon-Ledesma and Satchi (2019), which is the basis for our model.

They define the *short-run elasticity of substitution* as

$$\sigma_{SR} := - \frac{d \ln(\frac{E_t}{N_t})}{d \ln(\frac{p_{E,t}}{p_{N,t}})} \bigg|_{\theta_t \text{ fixed}}. \quad (2.12)$$

When climate policy or supply shocks increase the relative price of energy, σ_{SR} describes the immediate reaction of firms. The *long-run elasticity of substitution* (σ_{LR}) is the elasticity of substitution when θ_t has fully adjusted to its profit-maximizing level:

$$\sigma_{LR} := - \frac{d \ln(\frac{E_t}{N_t})}{d \ln(\frac{p_{E,t}}{p_{N,t}})} \bigg|_{\frac{\partial \Pi_t}{\partial \theta_t} = 0}, \quad (2.13)$$

where Π_t is profits earned by the final good producer in period t . When climate policy or supply shocks increase the relative price of energy, σ_{LR} describes the long-run reaction of firms.⁶

⁶The definition of σ_{LR} from Leon-Ledesma and Satchi (2019) refers to a firm that only maximizes single-period profits. In our empirical applications, we study a firm that maximizes the present discounted value of profits in an environment where the energy price is constantly changing.

Proposition 2.2. *Consider a production setting given by the production function (2.1) and the technology menu (2.2). In this case, $\sigma_{SR} = \epsilon$ and $\sigma_{LR} = 1$.*

Proof: This is a special case of Proposition 3 in Leon-Ledesma and Satchi (2019).

We provide a simple proof in Appendix Section B.

EoS and the Energy Expenditure Share. In our empirical application, we estimate the impact of energy price shocks on the energy expenditure share of output. To see how these objects are related, let σ be the EoS for some time horizon. Taking the definition of the EoS and integrating,

$$\ln \frac{E_t}{N_t} = \text{constant} + \sigma \ln \frac{p_{N,t}}{p_{E,t}}. \quad (2.14)$$

Subtracting $\ln \frac{p_{N,t}}{p_{E,t}}$ from both sides gives the following common expression:

$$\ln \frac{p_{E,t} E_t}{p_{N,t} N_t} = \text{constant} + (\sigma - 1) \ln \frac{p_{N,t}}{p_{E,t}}. \quad (2.15)$$

Now, we restrict attention to outcomes when there are no adjustment costs. This captures the long-run elasticity (after all adjustment costs are paid) and the short-run elasticity (before any adjustment costs are paid). With perfect competition, total revenue is equal to payments to energy producers plus payments to non-energy factors of production. So, we can re-write our expression as

$$\ln \frac{s_{E,t}}{1 - s_{E,t}} = \text{constant} + (\sigma - 1) \ln \frac{p_{N,t}}{p_{E,t}}, \quad (2.16)$$

where

$$s_{E,t} := \frac{p_{E,t}E_t}{Y_t + p_{E,t}E_t}, \quad (2.17)$$

is the *energy expenditure share of output*. We use this measure of the energy expenditure share for two reasons. The first is to be consistent with the data. It would be natural to measure the energy expenditure share in terms of gross output, but we would need a measure of adjustment cost expenditure to calculate gross output in a way that is consistent with our model. Unfortunately, there is no straightforward way to do this. The second reason that we choose this measure is to facilitate comparisons with the existing literature. As discussed below, the existing literature makes stark predictions about how this measure of the energy expenditure share responds to a price shock, making it easy to compare the predictions of the existing literature to our model and the data.

In our empirical application, we study how exogenous shocks to $p_{E,t}$ affect the dynamics of $s_{E,t}$, and then match the dynamics to our model. Equation (2.16) shows how these objects are related, though we will not be able to estimate it directly.⁷

2.4.3 Comparison to Existing Literature

Exogenous Technical Change (e.g., Golosov et al., 2014). In the macro climate-economy literature, it is standard to use a Cobb-Douglas (CD) production function for final output,

$$Y_t = K_t^\alpha E_t^\nu (A_t L)^{1-\alpha-\nu}.$$

⁷Since the price of output is normalized to one, the endogenous object $p_{N,t}$ is a function of only the price of energy and the production function parameters. Since we are estimating the production function parameters, we do not have a closed form equation that we can take directly to the data. But, we are still in a scenario where we are examining how one endogenous variable, $s_{E,t}$, responds to shocks in one exogenous variable, $p_{E,t}$.

This production function assumes $\sigma_{\text{SR}} = \sigma_{\text{LR}} = 1$.⁸ Thus, it has the same long-run predictions as our model, but different short- and medium-run predictions. Indeed, this production function is a special case of our model with $\epsilon = 1$ and $\phi = 0$. The CD production function implies that the energy expenditure share is $s_{E,t}^{\text{CD}} = \nu \forall t$. Thus, we can easily test the predictions of this model against the data by examining how the expenditure share responds to changes in the real energy price. The data presented in Figure 2.1 suggest that this production function will not be a good match for the data at short time horizons.

Directed Technical Change (e.g., Hassler et al., 2021a). The technology menu (2.2) defines a tradeoff between the two types of technology. This creates a close connection between our model and the literature on directed technical change and energy use. It can be conceived of as a simple way of capturing a market clearing constraint for R&D inputs. To see this, we can re-write the constraint in terms of growth rates:

$$\ln(1 + g_X) = (1 - \eta) \ln(1 + g_{A_{N,t}}) + \eta \ln(1 + g_{A_{E,t}}),$$

where $g_{J,t} := A_{J,t}/A_{J,t-1} - 1$. Here, the growth rates are the only endogenous variables. As discussed in Hassler et al. (2021a), this is a special case of the general constraint,

$$0 = G(g_{A_{N,t}}, g_{A_{E,t}}),$$

which describes models with a fixed set of R&D inputs (e.g., Acemoglu et al., 2012; Fried, 2018). In their quantitative results, Hassler et al. (2021a) assume a unit mass on R&D

⁸Models focused on short-run outcomes sometimes use lower values (e.g., Olovsson and Vestin, 2026). As noted above, models with directed technical change often use a low elasticity when technology is held fixed.

inputs that can be used in either sector.^{9,10}

Following Leon-Ledesma and Satchi (2019), we choose the log-linear constraint (2.2), because it makes the long-run elasticity equal to one, which is consistent with the data shown in Figure 2.1a. We then deviate from the usual directed technical change model by adding adjustment costs. This is especially important with the log-linear constraint. As stressed by Hassler et al. (2021a), the log-linear technology menu causes the direction of technical change to be exogenous when the growth rates can be freely chosen in each period. The adjustment costs re-introduce endogenous direction of technical change and allow the model to match both the short- and long-run elasticity of substitution. We follow Leon-Ledesma and Satchi (2019) in specifying the adjustment costs in terms of the levels of the technologies. In growth rates, the adjustment cost expression from (2.3) becomes

$$\frac{\phi}{2} \left(\frac{1 + g_{A_{E,t}}}{1 + g_{A_{N,t}}} - 1 \right)^2.$$

Thus, our model captures some of the key intuition of energy efficiency directed technical change models with fixed R&D inputs. Unlike the existing literature, we use a specific functional form for the technology constraint that allows us to set the long-run elasticity of substitution equal to one. The adjustment costs create a difference between short- and long-run outcomes. They allow the model to simultaneously have endogenous changes in the direction of technology and a long-run elasticity of one. The resulting

⁹Equation (15) in their paper captures the constraint for the functional forms used in the estimation.

¹⁰For concreteness, consider the model in Casey (2024) in which the innovation function takes the standard form:

$$A_{J,t} = \left(1 + \xi_J R_{J,t}^{1-\lambda}\right) A_{J,t-1}; \quad \sum_J R_{J,t} = 1; \quad J = N, E;$$

where $R_{J,t}$ is the share of R&D resources allocated to technology J at time t . In this setting, the constraint becomes

$$1 = \left(\frac{g_{A_{E,t}}}{\xi_E} \right)^{\frac{1}{1-\lambda}} + \left(\frac{g_{A_{N,t}}}{\xi_N} \right)^{\frac{1}{1-\lambda}},$$

a particular instance of the more general constraint.

model is simple enough that it can be easily calibrated to match the observed transition dynamics.¹¹

Relative to most of the directed technical change literature, the downside of our approach is that it is not a useful tool to study R&D policy. We specify a forward-looking representative firm, implying that there is no externality from intertemporal knowledge spillovers, and do not explicitly model payments to R&D inputs. Thus, our model cannot be used to capture the dynamic effects of R&D policy, a focus of much of the existing literature (e.g., Acemoglu et al., 2012).

2.5 Calibration

Our calibration process proceeds in three steps. In Section 2.5.1, we set a number of parameters externally. In Section 2.5.2, we match the energy efficiency dynamics in our model to the data. This is the main quantitative contribution of our calibration. In order to focus on energy efficiency, this analysis focuses on how the economy reacts to changes in final energy prices, $p_{E,t}$, without considering whether these changes are driven by changes in clean or dirty energy. Finally, Section 2.5.3 calibrates the production function for final energy. Table 2.1 presents the results of the calibration.

2.5.1 External Parameters

We set a number of parameters exogenously. We set $\beta = 0.985$ and $\delta = 0.10$, which are standard values for annual models. We choose $\eta = 0.075$ to match the average energy

¹¹Another key difference between our framework and a standard directed technical change model is that we do not impose that the growth rates of both technologies are always positive. The possibility for a decrease in technology captures the fact that a menu of technologies may exist at any given time and reflect a tradeoff. As the menu of technologies improves over time, it is always possible to improve both technologies. Following a change in prices, however, firms might choose a very different balance between technologies, leading to an overall decrease in one type of technology even as it is technologically feasible to improve both.

Table 2.1: Model Parameters

Parameter	Value	Description	Source
δ	0.10	Depreciation	Standard
β	0.985	Discount factor	Standard
η	0.075	Energy share of income	EIA
α	0.30	Capital share of income	Jones (2016)
ϵ	0.1146	Short-run EoS	IRF Matching
ϕ	5.5431	Adjustment costs	IRF Matching
γ	0.74	Dirty energy share	EIA
g_{pd}	1.9%	Dirty price growth	Calibrated
g_{pc}	0.7%	Clean price growth	Calibrated
$p_{d,0}/p_{c,0}$	0.40	Relative energy price	Calibrated

expenditure share of output, which is approximately 7.5% in the EIA data.¹² We take the remaining parameter values to match the standard facts as reported in Jones (2016). The labor share of expenditure is $(1 - \eta)(1 - \alpha)$, which we set to 65%. This yields $\alpha = 0.30$. We normalize $L = 1$.

2.5.2 IRF Matching

In this section, we discuss the empirical IRFs and how we use them to calibrate the energy demand portion of the model. Existing climate-economy models are often calibrated to long-run averages, like the average energy share of expenditure. We instead want to calibrate our model to data patterns that capture transition dynamics as well as long-run averages.

Empirical Specification. We use the annual data described above to estimate the impulse response function (IRF) of the energy expenditure share to a permanent change

¹²In the model, η is not exactly equal to the energy share of output and the difference between the two is due to adjustment costs. In our calibrated model, the adjustment costs are a small fraction of output. We therefore choose to pursue this external calibration of η , because it maintains transparency and computational simplicity. Alternatively, we could calibrate η jointly with θ and ϵ .

in the level of energy prices. To do so, we use a structural vector autoregression (SVAR). Our SVAR consists of the following two equations:

$$\Delta \ln(p_{E,t}) = c_1 + a_1 \cdot \Delta \ln(p_{E,t-1}) + b_1 \cdot s_{E,t-1} + u_1 \quad (2.18)$$

$$s_{E,t} = c_2 + a_2 \cdot \Delta \ln(p_{E,t-1}) + b_2 \cdot s_{E,t-1} + u_2. \quad (2.19)$$

To identify shocks, we make the assumption that $u_1 = \epsilon_1$ and $u_2 = d\epsilon_1 + \epsilon_2$, where ϵ_1 and ϵ_2 are orthogonal structural shocks. This specification assumes that energy price growth is not contemporaneously affected by shocks to the energy share. In other words, we assume that energy prices are predetermined for the US. This implies that energy prices do not respond to unpredicted innovations in energy expenditures within the year. This sluggishness can be explained by contracts that fix the price of energy for delivery over some period. Here, it is worth emphasizing that we consider the end-use energy price, not the energy spot price. Our IRFs focus on the dynamics following a one standard deviation shock to ϵ_1 , which we refer to as an ‘energy price shock’.

We use the detrended energy expenditure share in our main specification, following Blanchard and Quah (1989). Thus, our analysis assumes that any trend in $s_{E,t}$ is unrelated to changes in the price of energy, an assumption that is consistent with the unitary elasticity of substitution used in the existing literature (Golosov et al., 2014). Appendix Section B shows that we get similar results, although the dynamics are slower, if we do not de-trend the data. Meanwhile, detrending $\Delta \ln(p_{E,t})$ has basically no effect on the dynamics.¹³

Our baseline VAR specification uses only one lag. Appendix Section B demonstrates that the empirical IRFs are qualitatively similar with one to four lags. We also examined a number of lag selection criteria, all of which suggest one lag.

¹³Appendix B shows the time paths of $\Delta \ln(p_{E,t})$ and detrended $s_{E,t}$.

Furthermore, an underlying assumption in our baseline VAR model is that the model contains all relevant information. To test this assumption, we provide IRFs based on an augmented SVAR model with three variables in Appendix Section B. The idea is to add a third variable that measures the overall activity in the economy, which might affect the variation in both the growth rate of the energy price and the energy share of income. We use the growth rate of GDP, the growth rate of employment, and the unemployment rate as the third variable. These exercises show that the estimated dynamics are robust to adding this third variable.

Finally, we use yearly data in our main empirical analysis, as the EIA only provides data on the energy share of income at this frequency. However, the identification strategy might be considered more plausible at a higher data frequency. To test the robustness of our results, we therefore construct a quarterly dataset from monthly energy statistics from the EIA and quarterly GDP data from the BEA. We show that our qualitative results are robust to this change in data frequency. Nonetheless, we do not use this quarterly dataset in our main analysis, as a number of assumptions are made to construct the dataset. Furthermore, our quarterly dataset covers a shorter historical period due to missing data. Hence, we prefer to rely on the official annual data for our main analysis. Details on how the quarterly dataset is constructed and the consequences of using quarterly data in the SVAR analysis are provided in Appendix Section B.

Stochastic Processes for X_t and $p_{E,t}$. Our goal here is to calibrate the energy demand portion of the model. To do so, we abstract from the source of energy price variation. In addition, to match our model to the empirical IRFs, we need to consider a stochastic version of our model. Thus, we consider the case where both TFP, X_t , and final energy prices, $p_{E,t}$, are stochastic.

Table 2.2: Stochastic Processes for Calibration

Parameter	Value	Description
μ_X	0.6%	TFP drift
ρ_X	0.08	TFP persistence
Σ_X	0.01	TFP stan. dev.
μ_{p_E}	1.6%	Energy price drift
$\rho_{p_E,1}$	0.22	Energy price persistence
$\rho_{p_E,2}$	0.67	Energy price persistence
$\rho_{p_E,3}$	-0.20	Energy price persistence
$\rho_{p_E,4}$	-0.00	Energy price persistence
$\psi_{p_E,1}$	-0.12	Energy price persistence
$\psi_{p_E,2}$	-0.82	Energy price persistence
Σ_{p_E}	0.09	Energy price stan. dev.

Notes: all persistence and standard deviation parameters are rounded to two decimal places.

The growth rate of X_t , $g_{X,t}$, follows a stochastic AR(1) process

$$g_{X,t} = (1 - \rho_X)\mu_X + \rho_X g_{X,t-1} + \omega_{X,t}, \quad (2.20)$$

where $\omega_{X,t} \sim \mathcal{N}(0, \Sigma_X^2)$ is an *i.i.d.* shock and μ_X is a long-run drift parameter. The producer chooses current-period technologies prior to the realization of $\omega_{X,t}$.

The energy price growth rate is determined by the ARIMA(4, 2) process

$$\Delta \ln(p_{E,t}) = \left(1 - \sum_{i=1}^4 \rho_{p_E,i}\right) \mu_{p_E} + \sum_{i=1}^4 \rho_{p_E,i} \Delta \ln(p_{E,t-i}) + \sum_{i=1}^2 \psi_{p_E,i} \omega_{p_E,t-i} + \omega_{p_E,t}, \quad (2.21)$$

where $\omega_{p_E,t} \sim \mathcal{N}(0, \Sigma_{p_E}^2)$ is an *i.i.d.* shock and μ_{p_E} is a long-run drift parameter. It turns out that if we calibrate this ARIMA function appropriately, we are essentially able to perfectly match the empirical dynamics of $\Delta \ln(p_{E,t})$ following an energy price shock in our theoretical model.

The parameters driving the law of motion of TFP (X_t) and energy prices ($p_{E,t}$) are calibrated from data, which are presented in Table 2.2. Equation (2.20) is estimated

with total factor productivity (TFP) data covering 1970–2019,¹⁴ which yields $\mu_X = 0.6\%$, $\rho_X = 0.08$ and $\Sigma_X = 0.01$. We calibrate the drift coefficient in (2.21) to match the average annual growth rate of the energy price, resulting in $\mu_{p_E} = 1.6\%$. The remaining parameters are calibrated to match the empirical dynamics (impulse response function) of $\Delta \ln(p_{E,t})$ following an energy price shock. These dynamics are calculated using our structural VAR model described above.¹⁵

IRF Matching Methodology. In the theoretical model, we study the IRFs of $\Delta \ln(p_{E,t})$ and $s_{E,t}$ following a one standard deviation shock to $\omega_{p_E,t}$, the exogenous component of the energy price. As our calibrated ARIMA function can match the IRF for $\Delta \ln(p_{E,t})$ from the SVAR model almost exactly, we are able to consider the same energy price path in the empirical and theoretical models. The task is therefore to choose the pair ϵ and ϕ to match the IRF for $s_{E,t}$ as closely as possible. To do this, we minimize the weighted sum of squared residuals between empirical and model generated IRFs over a 21-year period. We use the inverse variance estimates from the empirical model as weights. See Appendix B for further details.

IRF Results. We study a positive and temporary shock to the growth rate of energy prices, which corresponds to a permanent change in the level of energy prices. We are interested in this type of shock, because we will use the model to study climate policies that permanently change the price of energy.

Our empirical results are presented in Figure 2.2. In all panels, the horizontal axis

¹⁴The data were originally obtained from the Federal Reserve Economic Database (FRED). Further details are provided in Appendix Section B.

¹⁵The empirical dynamics are defined as deviations from the long-run equilibrium. Hence, the ARIMA function can match these dynamics independent of μ_{p_E} . By minimizing the squared error between the SVAR and the ARIMA IRFs, we obtain the following parameter values: $\rho_{p_E,1} = 0.22$, $\rho_{p_E,2} = 0.67$, $\rho_{p_E,3} = -0.20$, $\rho_{p_E,4} = -0.00$, $\psi_{p_E,1} = -0.12$, $\psi_{p_E,2} = -0.82$, and $\Sigma_{p_E} = 0.09$.




Figure 2.2: Empirical Results and IRF Matching

Notes: Panels (a) and (b) show the path of energy prices following a one standard deviation shock to the growth rate of energy prices. This path of energy prices is used in all subsequent analyses. Panel (c) shows the path of the energy expenditure share of output in the data (solid line), our calibrated model (dashed line), and the special case of Golosov et al. (2014) (dotted line). Panel (d) shows the implications for the energy intensity of output. Panels (e) and (f) illustrate how the impulse response functions (IRFs) identify the model parameters.

shows years since the initial shock, and the vertical axis shows deviations from the initial steady state. Panels (a) and (b) show the shock to energy prices that will be used for the analysis. Panel (a) shows the dynamics of the energy price growth rate, which is estimated from the data. The solid line is the point estimate, and the dash-dotted bands are one-standard deviation (or approximately 68%) confidence intervals. The dashed line is the path of energy prices from our model, which very closely matches the data. Panel (b) translates the changes in growth rates to changes in the level of energy prices. In both cases, we show that the ARIMA(4,2) process for energy prices can exactly recreate this shock in our model.

We study the impact of this shock to energy prices in the data, our model, and older models from the literature. To represent the older literature, we focus on the influential model of Golosov et al. (2014) (henceforth, GHKT). They consider the case where $\sigma_{SR} = \sigma_{LR} = 1$. This specification is common in the literature (e.g., Nordhaus and Boyer, 2000; Hassler et al., 2016; Hassler and Krusell, 2018; Barrage, 2020; Cruz and Rossi-Hansberg, 2024) and is a special case of our model with $\phi = 0$ and $\epsilon = 1$.¹⁶

Our main findings are shown in panel (c), which plots the response of the energy expenditure share to the energy price shock. The solid line is the IRF estimated from data, the dashed line is our model with the parameters that best match the data, and the dotted line is the GHKT special case. The empirical results line up closely with the findings and interpretations in the existing literature. In the period of the shock, an increase of around 9% in energy prices leads to an approximately 0.6 percentage point increase in the expenditure share.

The methodology imposes that the expenditure share eventually returns to the original level. In other words, it imposes that the long-run elasticity is equal to one. Relative to earlier work, a key benefit of our methodology is that we can estimate the full transition

¹⁶With $\phi = 0$, the production function is the same for all $\epsilon \leq 1$.

path following an energy price shock. The transition is relatively slow, and the energy expenditure share has only returned halfway back to its original level 6 years after the shock. Importantly, however, this transition is driven by two forces, the slow reaction to the change in energy prices and the fact that energy prices are converging downward towards their new long-run level. When a carbon tax permanently raises the price of energy, the dynamics will be slower, because only the slow reaction to the change in energy prices will be relevant.

Our calibrated model matches the data closely. This is unsurprising given that the model was constructed to be able to match both the short- and long-run EoS and to include a free parameter, ϕ , that governs the speed of the transition. Our calibration yields $\epsilon = 0.1146$ ($= \sigma_{SR}$) and $\phi = 5.5431$. In contrast, the GHKT specification is clearly at odds with the data. It imposes a short-run EoS of one, which implies that the energy expenditure share is unchanged following the energy price shock. The climate literature, including GHKT, is aware that the short-run EoS is a poor match for the data, but argues that it accurately captures medium- and long-run dynamics (see also, Hassler et al., 2016, 2021b; Hassler and Krusell, 2018; Barrage, 2020). Our empirical results cast some doubt on this assertion. The data and the GHKT approach predict very different energy expenditure shares the first decade following the shock.

Panel (d) demonstrates the implications of our findings. Panel (d) shows the path of the energy intensity—defined as $s_{E,t}/p_{E,t}$ —following the energy price shock. This object is important for climate policy, because carbon taxes increase the price of energy and induce substitution towards non-energy inputs, reducing the energy intensity of output (i.e., increasing economy-wide energy efficiency). Both the data and our model suggest that the energy intensity declines quickly by 2% following the energy price shock and stays around this level thereafter. In contrast, the GHKT approach suggests a stronger short-run response (10% drop), while the long-run effect is the same. Intuitively, our

model features a low short-run EoS. The price shock, therefore, has little effect on the demand for energy in the short run, implying a weak effect on energy intensity. In the medium and long run, two counteracting forces on the energy intensity roughly balance out in our model: (i) the energy price declines, dragging the energy quantity and thereby the energy intensity up, and (ii) the economy adjusts to higher energy prices by using increasingly more energy efficient production technologies, dragging the energy quantity and thereby the energy intensity down. The GHKT approach does not feature the second effect, hence the energy intensity increases as the energy price level declines.¹⁷

Finally, panels (e) and (f) illustrate how the IRFs identify the model parameters. Panel (e) shows the effect of changing ϵ , holding ϕ fixed. In the period of the shock, the technology terms ($A_{N,t}$ and $A_{E,t}$) are fixed, and the EoS is given by $\sigma_{SR} = \epsilon$. Thus, ϵ plays a primary role in determining the size of the initial jump. Panel (f) shows the impact of changing ϕ , holding ϵ fixed. As discussed above, ϕ does not affect the short- or long-run EoS. But, higher values of ϕ imply that it is more costly to change the technology ratio, which in turn leads to slower transition dynamics. So, ϕ is largely determined by the speed of the transition.

2.5.3 Calibrating the Production Function for E_t

We now complete the calibration of the full model by calibrating the parameters in (2.6). The data are described in greater detail in Appendix Section B.

With perfect competition, γ is equal to the ratio of expenditure on dirty sources of primary energy to total expenditure on final energy. We set $\gamma = 74\%$ based on evidence from Casey et al. (2026a). The EIA does not report the price of clean primary energy

¹⁷We calibrate the two models with the same parameter η . Since there are no adjustment costs in the GHKT special case, the energy share of expenditure differs slightly.

over this period. Production function (2.6) implies that

$$\frac{p_{d,t}}{p_{c,t}} = \frac{\gamma}{1 - \gamma} \frac{E_{c,t}}{E_{d,t}}. \quad (2.22)$$

In 2019, the base year for our policy analysis, the ratio of clean and dirty primary energy, measured in BTUs, is equal to 0.14. Together with our estimate of γ , this implies that $\frac{p_{d,0}}{p_{c,0}} = 0.40$. To set units, we normalize the initial price of clean energy to one, $p_{c,0} = 1$.

We use a similar approach to pin down the growth rates of the energy prices. The ratio of clean to dirty energy grows at an average of 1.2%/year in our data. From (2.22), this implies that the relative price of clean energy decreases by 1.2%/year. Also, (2.7) implies that, in the absence of climate policy, the growth rate of the price of final energy is $g_{p_E} = (1 + g_{p_c})^{1-\gamma}(1 + g_{p_d})^\gamma - 1$. This is equal to 1.6% in the data. Together, these two relationships give $g_{p_d} = 1.9\%$ and $g_{p_c} = 0.7\%$.

In robustness analyses, we also consider CES production functions for final energy instead of the Cobb-Douglas specification in (2.6). We consider functions of the form

$$E_t = \left(\gamma E_{d,t}^{\frac{\chi-1}{\chi}} + (1 - \gamma) E_{c,t}^{\frac{\chi-1}{\chi}} \right)^{\frac{\chi}{\chi-1}}. \quad (2.23)$$

We consider substitution elasticities of $\chi \in \{0.95, 2\}$. The lower value is chosen to exactly match Golosov et al. (2014) and the higher value is chosen to capture the results from applied micro studies such as Papageorgiou et al. (2017). We use γ —our share parameter from the Cobb-Douglas specification—as the weight parameter. To be consistent with the Uzawa (1961) steady-state growth theorem, we need to assume that $g_{p_E} = g_{p_c} = g_{p_d} = 1.6\%$. Otherwise, g_{p_E} will not be constant and the model will not have a steady state.

2.6 Policy Analyses

We now turn to investigating the role that energy efficiency dynamics play in determining the effectiveness of climate change mitigation policies. In order to study the full non-linear impacts of policy, we abstract from uncertainty and use a deterministic version of the model. Our primary goal is to compare our calibrated model to the GHKT special case. We focus on outcomes over the next eighty years. All results are shown relative to a *laissez-faire* balanced growth path, which serves as our no policy baseline.¹⁸ Our framework does not include damages from climate change.

2.6.1 Constant Carbon Tax

Main Results. We simulate the impact of a permanent and unexpected ad valorem tax on dirty energy. This simple policy instrument highlights the effects of incorporating slow transition dynamics into an otherwise standard climate-economy model. In the following section, we consider a richer policy designed to achieve a realistic climate policy objective.

For this illustrative analysis, we consider a 50% tax. Casey et al. (2026b) find that this is roughly the value of the best constant carbon tax in the GHKT special case with a climate externality calibrated to the social cost of carbon (SCC) calculations of Rennert et al. (2021).

We compare the predictions of our calibrated model to the GHKT special case ($\phi = 0$ and $\epsilon = 1$). Other than the differences in ϵ and ϕ , the two models are identical. For each model, we compare outcomes with the dirty energy tax to the *laissez-faire* equilibrium.

Figure 2.3 shows the results. Panel (a) shows the dynamics of the energy expenditure share. As expected, they are qualitatively similar to those in Figure 2.2. In the period that policy is introduced, the energy expenditure share jumps almost one-for-one with the

¹⁸Given our focus on the United States, the *laissez-faire* assumption captures the absence of large-scale national climate policy. We abstract from the many state- and regional-level climate policies.




Figure 2.3: Impacts of a 50% Carbon Tax

change in the tax-inclusive price of energy and then slowly converges back to its original level. Unlike the price shock in the empirical IRFs, the carbon tax does not converge back towards its original level. As a result, the convergence in the energy share happens only because of the slowly evolving endogenous technology choice. After ten years, the expenditure share is roughly halfway back to its original level. In the GHKT special case, by contrast, the expenditure share is unchanged.

Panel (b) shows the implications for the energy intensity of output, $s_{E,t}/p_{E,t}$, which is the inverse of economy-wide energy efficiency. The two models have identical long-run predictions, but different predictions over the first two decades.

Next we turn to discussing the environmental implications of these results. In the period that policy is introduced, there is a 4 percent reduction in final energy use in our model (panel c). By contrast, energy use immediately drops by almost 28 percent in the GHKT special case. Over time, the two models converge to the lower GHKT level.

Panel (d) translates these effects to dirty energy, rather than final energy use. The dynamics are similar, but muted because we assume the substitution between clean and dirty energy can happen instantaneously. This result has important implications for interpreting data on fossil fuel use. Even if policy has a small immediate effect on emissions, it may still lead to significant long-run reductions.

Panel (e) translates these effects to differences in cumulative dirty energy use, the variable most relevant for environmental outcomes. In the GHKT model, the tax reduces dirty energy use by roughly 35 percent, and this effect is constant over time. In our model, cumulative dirty energy use is only 26 percent below the laissez-faire level in 2040. This highlights a key implication of our findings. When examining the impact of climate policy on medium-run outcomes, it is important to account for transition dynamics. Models that are only calibrated to long-run data will overstate the impact of policy on cumulative emissions. By 2100, our model predicts a 32 percent drop, which is more similar to the

GHKT results. This similarity implies that the 80 year horizon is long enough for the medium-run dynamics to be mostly overwhelmed by the long-run outcomes in determining cumulative emissions.

For completeness, panel (f) presents the dynamics of clean energy use. While our motivation for studying energy efficiency dynamics focuses on dirty energy use, the pattern of clean energy use is quite different in the two models. In our model, the demand for final energy is essentially unchanged in the period that policy is introduced, but the price of clean energy decreases relative to dirty energy. As a result, clean energy use spikes. In the GHKT model, by contrast, there is a significant reduction in the demand for energy use. Indeed, even though there is substitution away from dirty and towards clean, overall clean energy use decreases. Casey et al. (2026b) provide a simple explanation for this result. They show that, in the GHKT framework, clean and dirty energy are gross complements. When a carbon tax decreases dirty energy use, it also decreases the return to extracting clean energy, leading to lower clean energy use. Our model shows that, while this may be true in the long run, the short- and medium-run dynamics are very different.

Finally, we turn to the macro aggregates. The two models yield similar predictions for the dynamics of output and consumption as shown in panel (h). The representative household responds by lowering the savings rate, leading to an initial spike in consumption and then a swift decline. While we defer a full welfare analysis to the next section, the results suggest that the welfare benefits of a given policy intervention are smaller in our model. The paths of consumption are similar in the two models, but the environmental benefits are much lower when accounting for the slow energy efficiency dynamics. The paths of output are also similar, as shown in panel (g).

Before turning to the robustness analyses, it is important to highlight a caveat about our results: we allow for slow substitution between energy and non-energy inputs, but not between clean and dirty sources of primary energy. As noted above, we lack the necessary

data on clean energy prices to conduct this analysis. This has two main implications for our results. First, we almost certainly overstate the total reduction in dirty energy use from the policy. In both our model and the GHKT special case, the rapid substitution towards clean energy is responsible for much of the overall reduction in dirty energy use. Second, the rapid substitution implies that the 30% increase in clean energy use we observe in the first period is likely overstated.

Robustness. Figure B.11 shows the results when we assume that final energy is produced with the CES function (2.23) instead of the Cobb-Douglas function (2.6). It shows results for $\chi = 2$, which is consistent with the work of Papageorgiou et al. (2017). As before, we compare our model to the GHKT special case, and the change is implemented in both models, implying that the specification for energy demand is still the key difference between them. In all specifications, we assume that both clean and dirty extraction costs grow at the same rate, which allows the CES specification to converge to a BGP.

The qualitative patterns are quite similar with this alternative specification, especially for results related to final energy use. The main difference with the results in Figure 2.3 is that the robustness analysis has faster substitution between primary energy sources, leading to greater overall reductions in dirty energy use in both models.

Figure B.12 presents our main results using $\chi = 0.95$. This is the actual value used in Golosov et al. (2014). The results show that our simplification of implicitly setting $\chi = 1$ in the main results has virtually no impact on the results.

Figure B.13 shows results with $\eta = 0.04$, rather than $\eta = 0.075$. This lower value matches the energy expenditure share calculated in Hassler et al. (2021a), who focus only on expenditure on fossil fuels. Our measure of expenditure focuses on final energy. Importantly, this implies that we capture the costs of transforming fossil fuels (e.g., coal) into final energy (e.g., electricity), as well as the costs paid to produce electricity from

renewable energy sources.¹⁹ The results are quite similar. Emissions reductions and GDP losses are larger when energy is a greater share of total costs.

2.6.2 Clean Energy Subsidies

In Appendix Section B, we simulate a 20% subsidy for the production of clean energy in our model. Clean energy production subsidies are the central tool in the Inflation Reduction Act, which was passed in 2022. The law has very detailed provisions for different types of energy. Following Casey et al. (2026b), we choose a 20% subsidy to capture the aggregate implications of these provisions. As they explain, Bistline et al. (2023) perform a detailed analysis of all the different components of the subsidies and find that, after firms choose optimally among options, the end result is about a 20% reduction in the cost of producing utility-scale solar and offshore wind, with a slightly smaller effect for onshore wind.

Slow transition dynamics significantly bolster the ability of clean energy subsidies to reduce cumulative carbon emissions. Unlike carbon taxes, clean energy subsidies decrease the price of energy, inducing substitution towards energy inputs and decreasing energy efficiency (Holland et al., 2009; Newell et al., 2019). As a result, these policies are more effective when energy efficiency dynamics are slow.

This is a type of rebound in energy use: the clean energy subsidy lowers the price of energy, which in turn leads to an increase in use. There is a long literature studying the role of rebound in determining how exogenous shocks and policy interventions affect energy use (e.g., Gillingham et al., 2016; Fullerton and Ta, 2020).

Our model predicts that a 20 percent subsidy for clean energy reduces cumulative carbon emissions by 3 percent over the next decade, relative to *laissez-faire*. In the

¹⁹Unfortunately, as noted above, our data capture the joint costs of creating energy from clean and dirty sources and do not provide the direct estimates of the cost of clean energy.

standard model, by contrast, the same path of subsidies increases emissions by 0.6 percent. As noted above, clean and dirty energy are gross complements in the standard GHKT framework, implying that clean energy subsidies increase dirty energy use.

It is important to note that the direction of our clean energy subsidy results could be reversed if we also considered slow substitution between clean and dirty sources of primary energy. Slow substitution between primary sources of energy decreases the effectiveness of both carbon taxes and clean energy subsidies. As noted by Casey et al. (2026a), a lack of aggregate data on clean energy prices makes it difficult to analyze the substitution between clean and dirty energy in time series data. We leave a quantitative investigation of both margins to future research.

2.6.3 Least-Cost Path Analysis

To better understand the welfare consequences of slow energy efficiency dynamics, we now consider policy designed to limit total dirty energy use:

$$\sum_{\tilde{t}=0}^{\infty} E_{d,\tilde{t}} < E^{\max}, \quad (2.24)$$

where $E^{\max}$ is an exogenous politically determined dirty energy budget. In our simplified setting, this constraint plays the role of a carbon budget. We consider the case of a per-unit tax, τ_t^u , on dirty energy use. When examining policy, we consider the case where all policy is announced in 2020.

Theoretical Set-up. Consider a social planner who maximizes lifetime utility (2.9) subject to (2.4), (2.5), (2.6), (2.8), (2.10), (2.11), and (2.24). We refer to this as the least-cost path (LCP) to meeting the policy goal. We investigate the policies that implement the LCP in a decentralized equilibrium. Given that there is a single market failure, the

efficient outcome can be implemented with a tax set equal to the marginal external cost, i.e., the shadow value of the budget constraint.

Indeed, this is just a simple Hotelling (1931) problem. As is well known, along the optimal path, the gap between the price of dirty energy and its marginal extraction cost rises at the rate of interest. If the dirty energy producer internalized the limit on dirty use, this gap would be equal to the scarcity rent included in the price they charge consumers. In our case, where the limit is not internalized, the gap is equal to τ_t^u .

Proposition 2.3. *Consider the social planner problem described above. The optimal path can be implemented with a growing per-unit dirty energy tax*

$$\tau_t^u = \beta^{-t} \Omega C_t, \quad (2.25)$$

where Ω is the shadow value of the constraint measured in utils. This implies that the growth rate of the optimal tax, $g_{\tau,t} = \tau_t^u / \tau_{t-1}^u - 1$ is equal to the rate of interest

$$1 + g_{\tau,t} = \beta^{-1} \frac{C_t}{C_{t-1}} = 1 + r_t. \quad (2.26)$$

Proof: We provide a proof in Appendix Section B.

We study a tax that grows at the rate given in (2.26) with a level chosen to meet the target we discuss in the next section.

Picking a Target. For our quantitative results, we pick a dirty energy use target that is consistent with meeting a 1.5°C target for global temperature change. Achakulwisut et al. (2023) calculates cumulative fossil energy use from 2020 to 2100 for future trajectories consistent with this goal. We focus on their scenario that keeps global temperature change

below 1.5°C with a 50% probability and limits overshoot.²⁰

There are two caveats to consider. First, Achakulwisut et al. (2023) calculate results for the entire world, whereas we calibrate our model to the US. We set our target based on the ratio of cumulative future use to flow energy use in 2020. This is consistent with a scenario where all countries reduce dirty energy proportionally. Second, in the Achakulwisut et al. (2023) scenarios, dirty energy is still used after 2100. So, we choose a Ω (equivalently, an $E^{\max}$) that matches their cumulative budget through 2100, assuming that dirty energy can still be used afterward. Dirty energy is still used heavily after 2100 in our simulations.²¹

We find that cumulative dirty energy use from 2020-2100 can be 37.5 times as large as flow dirty energy use in 2020.²² Given the negative baseline trend in dirty energy use, this implies that, relative to the no policy baseline, cumulative dirty energy use needs to decrease roughly 29% to meet the target. Holding cumulative dirty energy use constant across runs, we can compare welfare by focusing on lifetime utility (2.9).

Quantitative Results. Figure 2.4 shows the least-cost path (LCP) that meets the policy goal.²³ Again, we compare the outcomes in our model to the GHKT special case in order to investigate the importance of slow energy efficiency dynamics.

Panel (a) shows the path of $\tau_t^u/p_{d,t}$ necessary to implement the LCP. The path of $p_{d,t}$ is exogenous and the same across models. We include it simply to present the units of the tax in an interpretable form. To meet the target, the tax rate in our model must be about 16% (6 p.p.) higher than in the GHKT model through the entire period. The

²⁰See figure 1.j in their paper.

²¹Their simulations include the possibility of negative emissions technologies and point-source carbon capture and storage, which are not part of our modeling. For our dirty energy target to be consistent with a 1.5°C target, these processes need to be happening outside our model. The emissions target would be much stronger in our model if we considered a target that did not take these forces into account.

²²Given that 2020 reflects policy changes in our model, we set the constraint relative to 2019 use.

²³Figure B.14 shows that the results are similar with $\chi = 2$.

adjustment costs imply that the economy reacts more sluggishly to the dirty energy tax, which in turn implies that the tax must be higher to generate the same reductions.

Panel (b) shows the path of the energy expenditure share. The graph has the familiar shape shown in the previous figure and IRFs. It is now scaled to the time-varying shock shown in panel (a).


Panel (c) shows the path of energy intensity, which highlights a key difference between the models. As expected and shown previously, energy intensity evolves more slowly in our model. Now that cumulative energy use reductions are the same across models, however, our model must catch up. As a result, energy intensity is eventually lower in our model, reflecting the higher tax. Panel (d) shows that this same pattern holds for final energy use.

Panel (e) shows that the same pattern holds for dirty energy use. To hit the target, our model must make up for its sluggish initial reaction by achieving greater reductions later in the century. It is worth noting that we do not account for changes in the carbon-intensity of dirty energy, which has generally been decreasing over time. If this were the case, the difference in required tax rates between the two models would increase. In our model, it is hard to decrease dirty energy use in the near term, which is when it is most carbon intensive.

Panel (f) shows cumulative dirty energy use, the targeted environmental outcome. By construction, both models achieve the same cumulative dirty energy use reductions in 2100, but there are notable differences over at least the first couple of decades.


Panel (g) shows the path of GDP. Since energy is a small share of GDP, even large dirty energy taxes have a relatively small effect on GDP, which is about 19 b.p. lower in the model with slow transition dynamics. This reflects, in part, the real adjustment costs that must be paid for the representative firm to change technology.

Finally, panel (h) shows the path of consumption. The differences between the two




(a) Carbon Tax along Least-Cost Path


(b) Energy Expenditure Share




(c) Energy Intensity




(d) Final Energy




(e) Flow Dirty Energy



(f) Cumulative Dirty Energy



(g) GDP



(h) Consumption

Figure 2.4: Least-Cost Path

models are quite small, again reflecting the low energy expenditure share. To compare lifetime utility across the two models, we consider the consumption equivalent variation (CEV). For each model, we compare lifetime utility from (2.9) in the laissez-faire baseline and along the LCP. We ask how much consumption the representative household would give up in each period to stay on the baseline, rather than switching to the LCP. We then compare across the two models. The CEV in our model is 0.56% and in the GHKT special case it is 0.47%, a difference of 19% (0.09 p.p.).

Our model has important implications for policy design. If policymakers use standard models, such as the GHKT special case, they will underestimate the carbon tax needed to meet intended emissions reduction targets. However, these emissions targets can still be met without large losses to consumption. Given the slow dynamics, policymakers should expect that emissions respond slowly to changes in policy and that most reductions will not occur until well after the policy is in place.

2.7 Conclusion

Macro climate-economy models are powerful tools for studying climate policy. Indeed, Hassler et al. (2018) argue that macro climate-economy models are “the only game in town” (p. 391). Do current iterations of these models generate reasonable predictions that are useful inputs to policy decisions? Given the importance of these models, surprisingly little work has been done to compare their predictions to data.

In this paper, we compare the models to data with a focus on one important determinant of carbon emissions: economy-wide energy efficiency. To do so, we borrow an approach frequently used in the business cycle literature. We estimate impulse response functions in data and compare them to the model predictions. We then build a new model that nests the existing literature and can replicate these empirical observations.

We use the new model to study the impacts of climate policy. Our model suggests that the impact of a given carbon tax on cumulative carbon emissions will be smaller than commonly believed. Carbon taxes increase the price of energy and promote energy efficiency. The data and our model suggest that this process happens slowly over time, while older models assume that it happens instantaneously. To meet a given target for emissions reductions, carbon taxes must be higher than existing models suggest.

We have focused on a simple and transparent model to highlight our key findings regarding the dynamics of energy efficiency. We hope that our results will contribute to the broader research program of using macro climate-economy models to study climate change. Given the tractability of our model, future work can integrate our insights into full quantitative integrated assessment models (IAMs) that can be used to study optimal climate policy.

2.8 Acknowledgments

Gregory Casey gratefully acknowledges financial support from National Science Foundation Grant SES-2315408. We thank Stefan Hinkelmann, Jason Jones, Ken Kuttner, Peter Pedroni and seminar participants at SURED, AERE, AERE@EEA, AERE@MEA, Brown, and UC Santa Barbara for helpful feedback. All errors are our own.

Chapter 3

Some Macroeconomic Evidence on Substitution between Clean and Dirty Energy

3.1 Introduction

Climate change mitigation policies, like carbon taxes or clean energy subsidies, are designed to alter the relative price of fossil fuel (“dirty”) and non-carbon (“clean”) sources of energy. The effectiveness of these policies depends on the ease of substitution between different energy sources. In macroeconomic models, substitution possibilities are often captured by a single parameter: the aggregate elasticity of substitution between clean and dirty energy (σ_{CD}).

Motivated by these macroeconomic models, recent studies use techniques from applied microeconomics to estimate σ_{CD} . Given data limitations, these studies focus on substitution between clean and dirty energy in certain sectors and along certain margins. Papageorgiou et al. (2017) separately estimate the clean-dirty elasticity for the electric

power sector and non-energy industries in a cross-country panel. Jo (2025) provides related firm-level evidence for French manufacturing industries and aggregates up to industry-level elasticities. Both studies find elasticities around 2 – 3 on average.

These studies provide the best available evidence on clean-dirty substitution, and they are used to calibrate macro models (e.g., Acemoglu et al., 2023). However, these studies do not estimate the aggregate elasticity, which incorporates all of the margins affected by a change in the relative price of clean and dirty energy, including substitution between (i) clean and dirty sources of electricity, (ii) electricity and fossil fuels in all end-use sectors, and (iii) output from different end-use sectors.

In this paper, we complement this existing literature by examining σ_{CD} using aggregate time-series data from the United States. We ask: “what σ_{CD} makes standard macro climate-economy models consistent with the aggregate data?” Relative to the existing literature, the benefit of our calibration-based approach is that it directly focuses on the aggregate elasticity. The drawback is that we do not have a convincing source of exogenous variation, and our results depend on model structure.

To match the data, standard macro climate-economy models require a σ_{CD} that is close to one in the long run, but below one in the short run. The distinction between long- and short-run elasticities is rarely considered in these models. The long-run value of one is consistent with the influential work of Golosov et al. (2014), but not the more recent literature. Using a reduced-form model, we provide preliminary evidence that updating macro models to match aggregate time series data would significantly alter their predictions for the effectiveness of climate policy.

3.2 Theoretical Framework

To motivate and interpret our descriptive evidence, we use the production structure of a standard macro climate-economy model (Golosov et al., 2014). All production is perfectly competitive. Final output is given by a combination of capital (K_t), labor (L_t), and energy (E_t): $Y_t = F(K_t, L_t, E_t)$. Here, E_t represents *end-use energy* that is used directly in the production of the final good, incorporating things like electricity and gasoline. The price of the final good is normalized to one, and $p_{E,t}$ denotes the price of end-use energy.

End-use energy is a composite of clean and dirty sources of *primary energy* (e.g., coal and solar). The production function is of the constant-elasticity-of-substitution (CES) form:

$$E_t = \left(\omega^{\frac{1}{\sigma_{CD}}} E_{d,t}^{\frac{\sigma_{CD}-1}{\sigma_{CD}}} + (1-\omega)^{\frac{1}{\sigma_{CD}}} E_{c,t}^{\frac{\sigma_{CD}-1}{\sigma_{CD}}} \right)^{\frac{\sigma_{CD}}{\sigma_{CD}-1}}, \quad (3.1)$$

where $E_{c,t}$ is the total quantity of clean primary energy and $E_{d,t}$ is the total quantity of dirty primary energy. Let $p_{d,t}$ and $p_{c,t}$ be the prices of dirty and clean energy, respectively. The *elasticity of substitution* (EoS) between clean and dirty energy determines how the energy mix changes in response to changes in relative energy prices:

$$\frac{d \ln \left(\frac{E_{c,t}}{E_{d,t}} \right)}{d \ln \left(\frac{p_{d,t}}{p_{c,t}} \right)} = \sigma_{CD}. \quad (3.2)$$

As $\sigma_{CD} \rightarrow 1$, the production function becomes Cobb-Douglas.

Perfect competition and profit maximization imply that

$$\ln \frac{s_{d,t}}{1 - s_{d,t}} = \text{constant} + (1 - \sigma_{CD}) \ln \frac{p_{d,t}}{p_{c,t}}, \quad (3.3)$$

where $s_{d,t} := \frac{p_{d,t}E_{d,t}}{p_{E,t}E_t}$ is the *fossil fuel share of energy expenditure* (payments to fossil fuel sources of primary energy as a fraction of total expenditure on final energy).

3.3 Descriptive Evidence

Figure 3.1 plots the components of (3.3). To measure our two key objects, we use data from the State Energy Data System (SEDS) compiled by the U.S. Energy Information Administration (2025b). It covers 1970-2023. The EIA provides data on *end-use energy* expenditure ($p_{E,t}E_t$), expenditure on fossil fuels ($p_{d,t}E_{d,t}$), and fossil fuel use ($E_{d,t}$). Expenditure is measured in dollars, and quantities are measured in BTUs. We obtain GDP figures from the U.S. Bureau of Economic Analysis (2025a) and convert prices to real dollars using the GDP deflator. With these data, we can directly calculate $s_{d,t}$ and $p_{d,t}$. Unfortunately, the EIA does not provide a direct measure of $p_{c,t}$. It is widely established, however, that the price of clean energy has fallen over this period. We view the trend in $p_{d,t}$ as a lower bound on the trend in $\frac{p_{d,t}}{p_{c,t}}$.

Panel (a) of Figure 3.1 demonstrates that there is little trend in the fossil fuel share of energy expenditure, despite a rising relative price of dirty energy. This suggests a long-run σ_{CD} that is not far above one. Using detrended data, panel (b) shows that $s_{d,t}$ and $p_{d,t}$ move together at short time horizons, suggesting a σ_{CD} that is well below one. These findings are similar to those identified for other margins of emissions reductions. Hassler et al. (2021a) and Casey et al. (forthcoming) provide similar evidence for substitution between energy and non-energy inputs to production. Hinkelmann (2024a) provides similar evidence for the aggregate elasticity between electricity and fossil fuels.




Figure 3.1: Descriptive Evidence

3.4 Impulse Response Functions

We now use impulse response functions (IRFs) to provide further information about how σ_{CD} varies by time horizon. We estimate a structural vector autoregression (VAR) consisting of the following two equations:


$$\Delta \ln(p_{d,t}) = \text{cons}_1 + a \cdot \Delta \ln(p_{d,t-1}) + u_1 \quad (3.4)$$

$$\ln \frac{s_{d,t}}{1 - s_{d,t}} = \text{cons}_2 + b \cdot \Delta \ln(p_{d,t-1}) + c \cdot \ln \frac{s_{d,t-1}}{1 - s_{d,t-1}} + u_2. \quad (3.5)$$

To identify shocks, we make the assumption that $u_1 = \epsilon_1$ and $u_2 = d\epsilon_1 + \epsilon_2$, where ϵ_1 and ϵ_2 are orthogonal structural shocks. We de-trend all variables and interpret the results as deviations from a steady state. This specification assumes a long-run σ_{CD} of one and is meant to illustrate the transition path.

Our specification implies that fossil fuel prices evolve exogenously. In other words, they are driven by supply shocks. Our IRFs focus on the dynamics following a one standard deviation shock to ϵ_1 , which we refer to as a ‘fossil fuel price shock’. We use equation (3.3) to back out the σ_{CD} that fits the data at each time horizon.

Panel (a) of Figure 3.2 shows the paths of $\ln p_{d,t}$ and $\ln \frac{s_{d,t}}{1 - s_{d,t}}$ following the shock. In

Figure 3.2: Impulse Responses and Implied σ_{CD}

the period of the shock, both the price and expenditure share variable spike. After the price levels off, the expenditure share variable slowly converges back to its long-run level. Panel (b) of Figure 3.2 calculates the corresponding σ_{CD} as $1 - \widehat{\ln \frac{s_{d,t}}{1-s_{d,t}}} / \widehat{\ln \frac{p_{d,t}}{p_{c,t}}}$, where a hat ($\widehat{}$) indicates deviation from the steady state path and $p_{c,t}$ is assumed to be constant. If the underlying data generating process is well-captured by a single CES like equation (3.1), the implied σ_{CD} should be the same at every time horizon. Our results suggest otherwise. The implied σ_{CD} is actually negative in the period of the shock. In this simple approach, we have not transformed variables to impose constraints on the parameter values. We interpret this result as a noisy estimate of a near-zero σ_{CD} . After 10 years, the implied σ_{CD} is roughly halfway back to its long-run value of one.

3.5 Implications for Climate Policy


We now turn to investigating the consequences of our results for climate policy. How different are the implications of climate policy with $\sigma_{CD} = 1$ and $\sigma_{CD} = 2$? What are the implications of taking into account the difference between the short- and long-run elasticities? We focus on flow emissions in the short- and medium-run, as well as

cumulative emissions.

To perform this analysis, we need a model that allows σ_{CD} to differ by time horizon. We consider a simple reduced-form model where we exogenously specify the path of σ_{CD} following a shock. This is meant as a stand-in for a richer model where the elasticity changes over time due to endogenous technology adoption as in Leon-Ledesma and Satchi (2019) and Hassler et al. (2021a). In particular, we consider a model where $Y_t = L^{1-\alpha} E_t^\alpha$, and E_t is given by (3.1) except that σ_{CD} exogenously changes over time to match the path shown in panel (b) of Figure 3.2. We set $\sigma_{CD} \approx 0$ in the first period. We assume that $p_{d,t}$, $p_{c,t}$, and L are constant and $\omega = 0.5$. We set $\alpha = 0.075$ to match Casey (2024). Each period is one year.

Figure 3.3 shows the impact of a 50% *ad valorem* carbon tax. All outcomes are shown relative to the no-policy steady state. The results from three models are shown: one with a constant $\sigma_{CD} = 1$, one with a constant $\sigma_{CD} = 2$, and one with the time-varying σ_{CD} .

Panel (a) shows the path of flow dirty energy use. As expected, the long-run flow emissions are the same for the $\sigma_{CD} = 1$ and time-varying models. Fifty years after the policy, these models predict emissions reductions that are approximately 13 p.p. lower than in the $\sigma_{CD} = 2$ model. The relatively slow transition dynamics in the time-varying model imply that it generates emissions reductions that are approximately 5 p.p. lower than in the $\sigma_{CD} = 1$ model ten years after the policy. Panel (b) shows cumulative dirty energy use, which is more relevant for environmental outcomes and highlights greater differences between the time-varying and $\sigma_{CD} = 1$ models. Over shorter time horizons, the dynamics of σ_{CD} impact policy effectiveness notably. After ten years, cumulative emissions reductions are a little less than 10 p.p. lower in the model with slow dynamics. The difference across models shrinks over time and is less than 5 p.p. after 50 years.

Figure 3.3: Impact of a 50% *Ad Valorem* Carbon Tax

3.6 Conclusion

The fossil fuel share of energy expenditure has been relatively constant in the United States, even as the relative price of clean and dirty energy has changed considerably. In the short run, dirty energy prices and the dirty energy expenditure share tend to move together. To recreate these patterns, standard macro climate-economy models would need to use a long-run elasticity of substitution between clean and dirty energy that is close to one and a short-run elasticity that is well below one. Based on micro evidence, existing models generally use higher long-run elasticities and do not consider a separate short-run elasticity. The results of this analysis show that updating macro models to match aggregate time series data would significantly alter their predictions for the effectiveness of climate policy.

Some caveats are in order. First, our approach has been to match *only* the aggregate data. We hope that our work will help spur efforts to build models that can match both micro and macro data. Second, we adopt a calibration-based approach that does not provide standard errors. Third, our results are necessarily backward-looking. Technology improvements, like better batteries or carbon capture and storage (CCS) technology, could significantly increase σ_{CD} over time (e.g., Gentile, 2024). Fourth, we have focused only

on substitution between aggregate clean and dirty energy composites. Similar techniques could be used to study aggregate substitution between more granular types of energy or between different categories of energy, such as intermittent versus dispatchable sources (Mehrotra, 2024).

3.7 Acknowledgments

Gregory Casey gratefully acknowledges funding from National Science Foundation Grant SES-2315408. We thank Chris Papageorgiou and numerous seminar participants for helpful feedback. All errors are our own.

Appendix A

Appendix for Chapter 1

Empirical Evidence: Additional Results

Table A.1: Summary statistics of regression sample

	Mean	SD	p25	p50	p75
CO ₂ Emissions:					
Emissions Intensity (tCO ₂ /k\$)	0.146	0.413	0.025	0.041	0.081
Emissions (million tCO ₂)	1.321	6.614	0.021	0.103	0.484
Marginal Products:					
log(MRPK)	-0.861	1.150	-1.556	-0.872	-0.181
Balance Sheet Variables:					
Sales (million \$)	9856	29353	592	2310	7040
Firm Age	35	24	17	30	50
Total Assets (million \$)	11775	32767	799	2772	8541
PPE/Assets	0.250	0.224	0.083	0.172	0.352
Debt/Assets	0.279	0.247	0.109	0.248	0.391
Intangibles/Assets	0.225	0.215	0.030	0.168	0.370
Cash/Assets	0.188	0.212	0.040	0.107	0.252

Notes: The statistics are pooled based on 15,830 firm-year observations

	Dep. Var.: log(<i>MRPK</i>)					
	Scope 1 Emissions Intensity (1)	Intensity (2)	Value-Added MRPK (3)	MRPK (4)	Level Emissions Intensity (5)	Intensity (6)
log(<i>Emissions Intensity</i>)	-0.101*** (0.017)	-0.121*** (0.017)	-0.295*** (0.021)	-0.299*** (0.021)		
<i>Emissions Intensity</i>					-0.253*** (0.059)	-0.231*** (0.051)
Firm Controls	N	Y	N	Y	N	Y
Industry × Year FE	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.304	0.334	0.336	0.389	0.301	0.327
Observations	15,830	15,830	14,964	14,964	15,830	15,830

Standard errors two-way clustered at the firm-year level are reported in parentheses.

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table A.2: Robustness: Firm-level relationship between emissions intensity and MRPK

	Dep. Var.: $\log(MRPL)$							
	Baseline		Scope 1 Emissions Intensity		Value-Added MRPL		Level Emissions Intensity	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\log(Emissions\ Intensity)$	0.058 (0.029)	0.056* (0.025)	0.040 (0.021)	0.060** (0.019)	-0.064* (0.030)	-0.059 (0.028)		
<i>Emissions Intensity</i>							0.156 (0.093)	0.128 (0.084)
Firm Controls	N	Y	N	Y	N	Y	N	Y
Industry $\times$ Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.301	0.348	0.301	0.349	0.314	0.334	0.302	0.347
Observations	15,830	15,830	15,830	15,830	14,964	14,964	15,830	15,830

Standard errors two-way clustered at the firm-year level are reported in parentheses.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table A.3: Robustness: Firm-level relationship between emissions intensity and MRPL

	Dep. Var.: <i>Cash Holdings (cash-to-asset ratio)</i>	
	(1)	(2)
$\log(Emissions\ Intensity)$	-0.040*** (0.004)	-0.037*** (0.004)
Firm Controls	N	Y
Industry $\times$ Year FE	Y	Y
Adjusted R^2	0.328	0.422
Observations	15,830	15,830

Standard errors two-way clustered at the firm-year level are reported in parentheses.
*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table A.4: Firm-level results between emissions intensity and cash holdings

	Dep. Var.: $\log(MRPK)$					
	W/o Firm Controls			W/ Firm Controls		
	(1)	(2)	(3)	(4)	(5)	(6)
$\log(Emissions\ Intensity)$	-0.196*** (0.020)	-0.099*** (0.021)	-0.184*** (0.021)	-0.198*** (0.021)	-0.105*** (0.022)	-0.146*** (0.014)
<i>Asset Tangibility</i>		-1.31*** (0.135)			-1.259*** (0.133)	
$\log(Productivity)$			0.307*** (0.015)			0.760*** (0.017)
Firm Controls	N	N	N	Y	Y	Y
Industry $\times$ Year FE	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.315	0.337	0.465	0.341	0.362	0.810
Observations	15,830	15,830	15,830	15,830	15,830	15,830

Standard errors two-way clustered at the firm-year level are reported in parentheses.

*** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$

Table A.5: Explanatory power of asset tangibility and productivity on the MRPK-emissions intensity relationship

Proof of Propositions

Proof of Lemma 1.2

Rewrite firm i 's production function in (1.4) as

$$y_i = z_i \left(\frac{k_i}{l_i} \right)^\alpha l_i \quad (\text{A.1})$$

$$= z_i \left(\frac{k_i}{l_i} \right)^{\alpha-1} k_i \quad (\text{A.2})$$

And plugging in firms' optimal capital-labor ratio, , yields

$$k_i^* = \frac{y_i}{z_i} \left(\frac{\alpha}{1-\alpha} \frac{w}{(r+\delta)(1+\tau_{ki})} \right)^{1-\alpha} \quad (\text{A.3})$$

$$l_i^* = \frac{y_i}{z_i} \left(\frac{\alpha}{1-\alpha} \frac{w}{(r+\delta)(1+\tau_{ki})} \right)^{-\alpha} \quad (\text{A.4})$$

With the firm demand function $y_i = p_i^{-\sigma} Y$, equations (A.3) and (A.4) can be rewritten as

$$k_i^* = \frac{p_i^{-\sigma}}{z_i} \left(\frac{\alpha}{1-\alpha} \frac{w}{(r+\delta)(1+\tau_{ki})} \right)^{1-\alpha} Y \quad (\text{A.5})$$

$$l_i^* = \frac{p_i^{-\sigma}}{z_i} \left(\frac{\alpha}{1-\alpha} \frac{w}{(r+\delta)(1+\tau_{ki})} \right)^{-\alpha} Y \quad (\text{A.6})$$

Combining (A.5) and (A.6) with the resource constraints (1.11) and (1.12) gives the economy's aggregate production function

$$Y = \bar{K}^\alpha \bar{L}^{1-\alpha} \frac{1}{\left(\sum_{i=1}^N \frac{p_i^{-\sigma}}{z_i} \left(\frac{1}{1+\tau_{ki}} \right)^{1-\alpha} \right)^\alpha \left(\sum_{i=1}^N \frac{p_i^{-\sigma}}{z_i} \left(\frac{1}{1+\tau_{ki}} \right)^{-\alpha} \right)^{1-\alpha}} \quad (\text{A.7})$$

Denote P as the price of the final good. The Dixit-Stiglitz structure of the economy implies $P = (\sum_{i=1}^N p_i^{1-\sigma})^{\frac{1}{1-\sigma}}$. Substituting into the expression of p_i in (1.8) and normalizing $P = 1$, the firm i 's output price can be rewritten as:

$$p_i = \frac{\frac{(1+\tau_{ki})^\alpha}{z_i(1-\tau_e\omega_i)}}{\left(\sum_{i=1}^N \left(\frac{(1+\tau_{ki})^\alpha}{z_i(1-\tau_e\omega_i)} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}} \quad (\text{A.8})$$

Finally, plug (A.8) into (A.7), and the effective aggregate TFP term can be expressed as

$$\text{TFP} = \frac{\left(\sum_{i=1}^N \left(\frac{(1-\tau_e \omega_i) z_i}{(1+\tau_{ki})^\alpha} \right)^{\sigma-1} \right)^{\frac{\sigma}{\sigma-1}}}{\left(\sum_{i=1}^N \frac{(1-\tau_e \omega_i)^\sigma z_i^{\sigma-1}}{(1+\tau_{ki})^{1+\alpha(\sigma-1)}} \right)^\alpha \left(\sum_{i=1}^N \frac{(1-\tau_e \omega_i)^\sigma z_i^{\sigma-1}}{(1+\tau_{ki})^{\alpha(\sigma-1)}} \right)^{1-\alpha}} \quad (\text{A.9})$$

□

Proof of Proposition 1.3

Taking natural logarithms on both sides of equation (1.14) gives

$$\begin{aligned} \log \text{TFP} &= \frac{\sigma}{\sigma-1} \log \left(\sum_{i=1}^N \left(\frac{(1-\tau_e \omega_i) z_i}{(1+\tau_{ki})^\alpha} \right)^{\sigma-1} \right) - \alpha \log \left(\sum_{i=1}^N \frac{(1-\tau_e \omega_i)^\sigma z_i^{\sigma-1}}{(1+\tau_{ki})^{1+\alpha(\sigma-1)}} \right) \\ &\quad - (1-\alpha) \log \left(\sum_{i=1}^N \frac{(1-\tau_e \omega_i)^\sigma z_i^{\sigma-1}}{(1+\tau_{ki})^{\alpha(\sigma-1)}} \right) \\ &= \frac{1}{\sigma-1} \log N + \frac{\sigma}{\sigma-1} \log \mathbb{E} \left(\frac{(1-\tau_e \omega_i) z_i}{(1+\tau_{ki})^\alpha} \right)^{\sigma-1} - \alpha \log \mathbb{E} \left(\frac{(1-\tau_e \omega_i)^\sigma z_i^{\sigma-1}}{(1+\tau_{ki})^{1+\alpha(\sigma-1)}} \right) \\ &\quad - (1-\alpha) \log \mathbb{E} \left(\frac{(1-\tau_e \omega_i)^\sigma z_i^{\sigma-1}}{(1+\tau_{ki})^{\alpha(\sigma-1)}} \right) \end{aligned} \quad (\text{A.10})$$

Since $\{\log z_i, \log(1+\tau_{ki}), \log(1-\tau_e \omega_i)\}$ are jointly normally distributed, the terms in the expectation operators are lognormally distributed. Applying the mean of lognormal distribution simplifies the above expression as

$$\begin{aligned} \log \text{TFP} &= \frac{1}{\sigma-1} \log N + \mathbb{E}(\log z_i) + \frac{\sigma-1}{2} \text{Var}(\log z_i) - \frac{\alpha(1-\alpha) + \alpha^2 \sigma}{2} \text{Var}(\log(1+\tau_{ki})) \\ &\quad - \frac{\sigma}{2} \text{Var}(\log(1-\tau_e \omega_i)) + \sigma \alpha \text{Cov}(\log(1-\tau_e \omega_i), \log(1+\tau_{ki})) \end{aligned} \quad (\text{A.11})$$

Taking all else as given, $\log \text{TFP}$ is maximized where

$$-\frac{\sigma}{2} \text{Var}(\log(1-\tau_e \omega_i)) + \sigma \alpha \text{Cov}(\log(1-\tau_e \omega_i), \log(1+\tau_{ki})) \quad (\text{A.12})$$

is maximized. When the capital and output wedges τ_{ki} and $\tau_e \omega_i$ are assumed to be close to 0, the above term can be rewritten as:

$$-\frac{\sigma}{2} \text{Var}(\omega_i) \tau_e^2 - \sigma \alpha \text{Cov}(\tau_{ki}, \omega_i) \tau_e \quad (\text{A.13})$$

Since $\text{Cov}(\tau_{ki}, \omega_i) < 0$, the above quadratic function in τ_e is maximized at $\bar{\tau}_e = \frac{-\alpha \text{Cov}(\tau_{ki}, \omega_i)}{\text{Var}(\omega_i)} > 0$.

□

Proof of Proposition 1.4

The aggregate emissions and economy-wide emissions intensity of output can be expressed as:

$$E = \sum_{i=1}^N \omega_i p_i y_i = Y \sum_{i=1}^N \omega_i p_i^{1-\sigma} \quad (\text{A.14})$$

$$\frac{E}{Y} = \sum_{i=1}^N \omega_i p_i^{1-\sigma} \quad (\text{A.15})$$

Plugging in the expression for p_i from equation (A.8) and simplifying terms using a procedure similar to that in *Proposition 1.3*, one can write the natural logarithm of aggregate emissions intensity as:

$$\log \left(\frac{E}{Y} \right) = \mathbb{E}(\log \omega_i) + \frac{1}{2} \text{Var}(\log \omega_i) + \alpha(1 - \sigma) \text{Cov}(\log \omega_i, \log(1 + \tau_{k_i})) \quad (\text{A.16})$$

After removing all capital wedges, the term $\text{Cov}(\log \omega_i, \log(1 + \tau_{k_i}))$ changes from negative to zero. This implies that with $\sigma > 1$, $\log \left(\frac{E}{Y} \right)$ would decrease. In other words, the economy-wide emissions intensity of output would fall following the wedge-free reform.

□

Quantitative Model: Additional Details

Firm's Profit Maximization Problem

Within each period, the monopolistically competitive firm takes its demand function $y = p^{-\sigma} D$ into account and maximizes its profits as follows:

$$\pi(z, a, \mu) = \max_{p, k, l} (1 - \tau_e \mu^\epsilon) p y - w l - (r + \delta) k \quad (\text{A.17})$$

$$= \max_{k, l} (1 - \tau_e \mu^\epsilon) D^{\frac{1}{\sigma}} y^{\frac{\sigma-1}{\sigma}} - w l - (r + \delta) k \quad (\text{A.18})$$

subject to the borrowing constraint $k \leq \kappa a$ and the production function $y = z k^\alpha l^{1-\alpha}$.

When the borrowing constraint is slack for firms, they would choose the unconstrained optimal capital which I denote by k^u . In contrast, when the borrowing constraint is binding, they can only achieve the constrained capital level k^c and that $k^c \leq k^u$. Thus the capital decision for a given state vector (z, a, μ) is given by:

$$k(z, a, \mu) = \begin{cases} k^u = \left(\frac{\sigma}{\sigma-1}\right)^{-\sigma} \left(\frac{\alpha}{R}\right)^{\alpha(\sigma-1)+1} \left(\frac{1-\alpha}{w}\right)^{(\sigma-1)(1-\alpha)} D z^{\sigma-1} (1 - \tau_e \mu^\epsilon)^\sigma, & k^u \leq \kappa a, \\ k^c = \kappa a, & k^u > \kappa a. \end{cases} \quad (\text{A.19})$$

The optimal labor demand for unconstrained (l^u) and constrained firms (l^c) are:

$$l(z, a, \mu) = \begin{cases} l^u = \frac{1-\alpha}{\alpha} \frac{R}{w} k^u, & k^u \leq \kappa a, \\ l^c = \left(\frac{\sigma}{\sigma-1}\right)^{\frac{-\sigma}{1+\alpha(\sigma-1)}} D^{\frac{1}{1+\alpha(\sigma-1)}} z^{\frac{\sigma-1}{1+\alpha(\sigma-1)}} \\ \quad \times \left[\frac{w}{(1-\alpha)(1-\tau_e \mu^\epsilon)} \right]^{\frac{-\sigma}{1+\alpha(\sigma-1)}} (k^c)^{\frac{\alpha(\sigma-1)}{1+\alpha(\sigma-1)}} & k^u > \kappa a. \end{cases} \quad (\text{A.20})$$

The monopolistic-competitive firm also sets a constant markup (η) of its price over marginal cost, which also depends on the bindingness of the borrowing constraint. The choice of price for unconstrained (p^u) and constrained firms (p^c) are as follows:

$$p(z, a, \mu) = \begin{cases} p^u = \frac{\sigma}{\sigma-1} \left(\frac{r+\delta}{\alpha}\right)^\alpha \left(\frac{w}{1-\alpha}\right)^{1-\alpha} \frac{1}{z(1-\tau_e \mu^\epsilon)}, & k^u \leq \kappa a, \\ p^c = \left(\frac{\sigma}{\sigma-1}\right)^{\frac{1-\alpha}{1+\alpha(\sigma-1)}} D^{\frac{\alpha}{1+\alpha(\sigma-1)}} z^{\frac{-1}{1+\alpha(\sigma-1)}} \\ \quad \times \left[\frac{w}{(1-\alpha)(1-\tau_e \mu^\epsilon)} \right]^{\frac{1-\alpha}{1+\alpha(\sigma-1)}} (k^c)^{\frac{-\alpha}{1+\alpha(\sigma-1)}} & k^u > \kappa a. \end{cases} \quad (\text{A.21})$$

Consequently, the firm's optimal profit function can be derived as:

$$\pi(z, a, \mu) = \begin{cases} \pi^u = \frac{1}{\sigma} \left(\frac{\sigma-1}{\sigma} \right)^{\sigma-1} (1 - \tau_e \mu^\epsilon)^\sigma D z^{\sigma-1} \left(\frac{\alpha}{r+\delta} \right)^{\alpha(\sigma-1)} \left(\frac{1-\alpha}{w} \right)^{(1-\alpha)(\sigma-1)}, & k^u \leq \kappa a, \\ \pi^c = (1 - \tau_e \mu^\epsilon) D^{\frac{1}{\sigma}} z^{\frac{\sigma-1}{\sigma}} (k^c)^{\frac{\alpha(\sigma-1)}{\sigma}} (l^c)^{\frac{(1-\alpha)(\sigma-1)}{\sigma}} - w l^c - (r + \delta) \kappa a, & k^u > \kappa a. \end{cases} \quad (\text{A.22})$$

Solution Algorithm for Stationary General Equilibrium

To numerically solve for a stationary equilibrium of the model economy which is in general equilibrium, I follow the algorithm below:

1. Guess an interest rate r , a wage w , an aggregate demand shifter D , and an aggregate transfer $\bar{T}$.
2. Solve each firm's problem given $(r, w, D, \bar{T})$.
3. Simulate a large number of firms ($N = 10000$ in my simulations) in the cross section for a sufficiently long period ($T = 1000$ in my simulations).
4. Compute, in the last period, the implied aggregate labor demand (L^d), aggregate net worth (A), aggregate capital (K), aggregate output (Y), and aggregate carbon tax revenues ($\tau_e E$), all expressed as the cross-sectional average quantities.
5. Check the market-clearing conditions:
 - Goods market clears: $Y = D$
 - Labor market clears: $L^d = 1$
 - Asset market clears: $A = K$ (zero net supply of aggregate debt)
 - Government's budget is balanced: $\bar{T} = \tau_e E$
6. If any of the markets does not clear, update $(r, w, D, \bar{T})$ and repeat until all markets simultaneously clear.
7. Once the solutions converge, ensuring all markets clear within a specified tolerance, a stationary general equilibrium is achieved.

Estimation of Firm-level TFP ($\tilde{z}_{it}$) in the Data

To address the downward bias in the estimated persistence of firm-level TFP that arises in short panels with firm fixed effects and a lagged dependent variable (Nickell, 1981), I construct an extended firm sample covering a longer time horizon. Specifically, I select firms that appear at least once in the Trucost database and extract their corresponding financial data from Compustat dating back to 1980, regardless of whether the corresponding emissions data are available.

Firm-level TFP is estimated on this sample following the methodology of Wooldridge (2009), which employs a system GMM approach to address potential endogeneity in production function estimation. I use Compustat data on labor, total capital, sales, and cost of goods sold. Labor is modeled as a flexible input and capital as a state variable, while the cost of goods sold proxies for unobserved productivity shocks. I implement a value-added production function with a second-degree polynomial in inputs to capture potential nonlinearities and include firm and year fixed effects to control for unobserved heterogeneity. The residual from this estimation constitutes firm-level TFP, capturing productivity variations conditional on input usage while mitigating endogeneity concerns.

To ensure sufficient observations for each firm and reduce the bias in the estimated persistence parameter, I then retain only firms with continuous TFP coverage over the 1980-2021 period and estimate the firm-level AR(1) productivity regression in (1.61).

Quantitative Results: Additional Details

Heterogeneous Firm Outcomes across Emissions Intensity




Figure A.1: Cross-sectional firm outcomes across emissions intensity deciles

Notes: The mean of log(MRPK), leverage (debt-to-assets ratio) and cash (cash-to-assets ratio) is calculated within each emissions intensity decile. In the model, emissions intensity deciles also correspond in rank to the asset tangibility deciles. The model outcomes are evaluated at the laissez-faire for both models.

Mapping Model Units to Real-World Units

To relate the abstract carbon tax units (τ_e) in the model to real-world values, I use a back-of-the-envelope calculation. I choose a unitless target to match between the model

and the data: the proportion of potential carbon tax revenues in U.S. GDP in the reference year 2020. The carbon tax rate can then be converted from model units to dollars per ton of CO₂ via:

$$\tau_e(\$/\text{tCO}_2) = \frac{\tau_e(\text{model})E(\text{model})}{Y(\text{model})} \times \frac{Y(\$)}{E(\text{t CO}_2)} \quad (\text{A.23})$$

where

$$Y(\$) = \frac{Y(\text{model})}{Y_{\tau_e=0}(\text{model})} \times (\text{U.S GDP in 2020}) \quad (\text{A.24})$$

$$E(\text{t CO}_2) = \frac{E(\text{model})}{E_{\tau_e=0}(\text{model})} \times (\text{CO}_2 \text{ emissions in 2020}) \quad (\text{A.25})$$

This procedure implicitly assumes that no universal carbon tax was in effect in 2020. In the data, I use U.S. GDP in 2020 (in 2020 dollars) and total greenhouse gas emissions in 2020 excluding the residential sector.

Calculation of CEV

To better compare welfare across states, I measure welfare loss (or gain) in terms of permanent consumption equivalents. Let variables with a bar (–) denote outcomes under the laissez-faire setting ($\tau_e = 0$) for comparison and variables without the bar denote the outcomes under a specific policy scenario. Based on the lifetime welfare measure in expression (1.60), to calculate the permanent consumption equivalent variation, I suppose there is a value of ψ such that

$$\bar{W}_0 = \sum_{t=0}^{\infty} \beta^t \left(\log(C_t) - m \frac{E_t}{1-\beta} \right), \quad (\text{A.26})$$

$$W_0 = \sum_{t=0}^{\infty} \beta^t \left(\log(\psi C_t) - m \frac{E_t}{1-\beta} \right). \quad (\text{A.27})$$

and therefore,

$$W_0 = \sum_{t=0}^{\infty} \beta^t \log(\psi) + \sum_{t=0}^{\infty} \beta^t \left(\log(C_t) - m \frac{E_t}{1-\beta} \right) \quad (\text{A.28})$$

$$= \sum_{t=0}^{\infty} \beta^t \log(\psi) + \bar{W}_0. \quad (\text{A.29})$$

The welfare relative to the laissez-faire scenario, measured by ψ , thus equals

$$\psi = \exp \left[(1-\beta) (W_0 - \bar{W}_0) \right]. \quad (\text{A.30})$$

Additional Welfare Results




Figure A.2: Impacts of carbon tax on firms' and workers' welfare

Notes: All welfare outcomes are shown as permanent consumption equivalents in percentage relative to laissez faire in each model.

With tax revenues rebated to firms, firms' welfare increases with carbon taxes in both HeFF and NoFF models, with larger gains seen in the HeFF model. In contrast, hand-to-mouth workers experience welfare losses purely due to lower equilibrium wages resulting from higher taxes, with a slightly greater decline in the NoFF model.

Appendix B

Appendix for Chapter 2

Model Analysis

Characterization of Competitive Equilibrium

Final Good Producer Problem

The firm's problem is

$$\max_{\{N_t, E_t, \theta_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \left[\prod_{\tilde{t}=0}^t \Lambda_{\tilde{t}} \right] \times \left[\left(1 - \frac{\phi}{2} \left(\frac{\theta_t}{\theta_{t-1}} - 1 \right)^2 \right) \underbrace{X_t \left((\theta_t^{-\eta} N_t)^{\frac{\epsilon-1}{\epsilon}} + (\theta_t^{1-\eta} E_t)^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}}}_{\tilde{Y}_t} - p_{N,t} N_t - p_{E,t} E_t \right], \quad (\text{B.1})$$

where θ_0 and θ_{-1} are given. The firm discounts future profits using the household's stochastic discount factor, $\Lambda_{t+1} \equiv \beta \frac{u'(C_{t+1})}{u'(C_t)} = \beta \frac{C_t}{C_{t+1}}$ with $\Lambda_0 = 1$. The expectation is redundant for the main analysis of the model. We include it to aid the exposition of the stochastic version used in the calibration. Similarly, in the deterministic model $\beta \frac{C_t}{C_{t+1}} = \frac{1}{1+r_{t+1}}$, and the firm discounts at the rate of interest. In the stochastic version, the firm does not know the rate of interest with certainty. Since it is owned by the household, it uses the household's expected discount factor.

The first-order conditions are

$$N_t : \quad \frac{p_{N,t} N_t}{\tilde{Y}_t} = \left(1 - \frac{\phi}{2} \left(\frac{\theta_t}{\theta_{t-1}} - 1 \right)^2 \right) \left(\frac{\tilde{Y}_t}{N_t} \right)^{\frac{1-\epsilon}{\epsilon}} (X_t \theta_t^{-\eta})^{\frac{\epsilon-1}{\epsilon}} \quad (\text{B.2})$$

$$E_t : \quad \frac{p_{E,t} E_t}{\tilde{Y}_t} = \left(1 - \frac{\phi}{2} \left(\frac{\theta_t}{\theta_{t-1}} - 1 \right)^2 \right) \left(\frac{\tilde{Y}_t}{E_t} \right)^{\frac{1-\epsilon}{\epsilon}} (X_t \theta_t^{1-\eta})^{\frac{\epsilon-1}{\epsilon}} \quad (\text{B.3})$$

$$\begin{aligned} \theta_{t+1} : \quad 0 = \mathbb{E}_t & \left\{ -\phi \left(\frac{\theta_{t+1}}{\theta_t} - 1 \right) \frac{\theta_{t+1}}{\theta_t} \right. \\ & + \Lambda_{t+2} \phi \left(\frac{\theta_{t+2}}{\theta_{t+1}} - 1 \right) \frac{\theta_{t+2}}{\theta_{t+1}} \frac{\tilde{Y}_{t+2}}{\tilde{Y}_{t+1}} \\ & + \left(1 - \frac{\phi}{2} \left(\frac{\theta_{t+1}}{\theta_t} - 1 \right)^2 \right) \left(\frac{\tilde{Y}_{t+1}}{X_{t+1}} \right)^{\frac{1-\epsilon}{\epsilon}} \\ & \left. \times \left[-\eta (\theta_{t+1}^{-\eta} N_{t+1})^{\frac{\epsilon-1}{\epsilon}} + (1-\eta) (\theta_{t+1}^{1-\eta} E_{t+1})^{\frac{\epsilon-1}{\epsilon}} \right] \right\}. \end{aligned} \quad (\text{B.4})$$

Combining all of the first-order conditions, we can re-write the last (B.4) as

$$\begin{aligned} \theta_{t+1} : \quad 0 = \mathbb{E}_t & \left\{ -\phi \left(\frac{\theta_{t+1}}{\theta_t} - 1 \right) \frac{\theta_{t+1}}{\theta_t} \right. \\ & + \Lambda_{t+2} \phi \left(\frac{\theta_{t+2}}{\theta_{t+1}} - 1 \right) \frac{\theta_{t+2}}{\theta_{t+1}} \frac{\tilde{Y}_{t+2}}{\tilde{Y}_{t+1}} \\ & \left. + \left(1 - \frac{\phi}{2} \left(\frac{\theta_{t+1}}{\theta_t} - 1 \right)^2 \right) [-\eta s_{N,t+1} + (1 - \eta) s_{E,t+1}] \right\}, \end{aligned} \quad (\text{B.5})$$

where $s_{J,t} = \frac{p_{J,t} J_t}{(1 - \Phi(\frac{\theta_t}{\theta_{t-1}})) \tilde{Y}_t}$ is share of output paid to factor $J = N, E$ and $s_{N,t} + s_{E,t} = 1$.

Plugging this identity into the previous equation and rearranging yields

$$\begin{aligned} \theta_{t+1} : \quad 0 = \mathbb{E}_t & \left\{ -\phi \left(\frac{\theta_{t+1}}{\theta_t} - 1 \right) \frac{\theta_{t+1}}{\theta_t} \right. \\ & + \Lambda_{t+2} \phi \left(\frac{\theta_{t+2}}{\theta_{t+1}} - 1 \right) \frac{\theta_{t+2}}{\theta_{t+1}} \frac{\tilde{Y}_{t+2}}{\tilde{Y}_{t+1}} \\ & \left. - \frac{p_{N,t+1} N_{t+1}}{\tilde{Y}_{t+1}} + (1 - \eta) \left(1 - \frac{\phi}{2} \left(\frac{\theta_{t+1}}{\theta_t} - 1 \right)^2 \right) \right\}. \end{aligned} \quad (\text{B.6})$$

Intermediate Producer Problem

The intermediate producer solves

$$\max_{K_t, L} \quad p_{N,t} K_t^\alpha L^{1-\alpha} - \rho_t K_t - w_t L, \quad (\text{B.7})$$

which gives

$$\rho_t = \alpha p_{N,t} K_t^{\alpha-1} L^{1-\alpha} \quad (\text{B.8})$$

$$w_t = (1 - \alpha) p_{N,t} K_t^\alpha L^{-\alpha}. \quad (\text{B.9})$$

Combining these results gives the usual price index:

$$p_{N,t} = \alpha^{-\alpha} (1 - \alpha)^{\alpha-1} w_t^{1-\alpha} \rho_t^\alpha. \quad (\text{B.10})$$

Final Energy Producer Problem

The final energy producer solves

$$\max_{E_{d,t}, E_{c,t}} \quad p_{E,t} E_{d,t}^\gamma E_{c,t}^{1-\gamma} - (1 + \tau_{d,t}) p_{d,t} E_{d,t} - p_{c,t} E_{c,t}, \quad (\text{B.11})$$

which yields

$$p_{c,t}E_{c,t} = (1 - \gamma)p_{E,t}E_t \quad (\text{B.12})$$

$$(1 + \tau_{d,t})p_{d,t}E_{d,t} = \gamma p_{E,t}E_t. \quad (\text{B.13})$$

Combining these results gives the final energy price index:

$$p_{E,t} = \tilde{\gamma} p_{c,t}^{1-\gamma} (1 + \tau_{d,t})^\gamma p_{d,t}^\gamma. \quad (\text{B.14})$$

where $\tilde{\gamma} = \gamma^{-\gamma}(1 - \gamma)^{\gamma-1}$.

Household Problem

The household maximizes (2.9) subject to the budget constraint given by (2.10) and (2.11), which yields the Euler equation,

$$1 = \mathbb{E}_t \left(\beta \frac{C_t}{C_{t+1}} (1 + r_{t+1}) \right), \quad (\text{B.15})$$

where

$$r_t := \rho_t - \delta, \quad (\text{B.16})$$

and the transversality condition,

$$\lim_{t \rightarrow \infty} \mathbb{E}_0 (\beta^T C_T^{-1} K_{T+1}) = 0. \quad (\text{B.17})$$

Again, the expectation is redundant for the main analysis of the model. We include it to aid the exposition of the stochastic version used in the calibration.

Proof of Proposition 2.2

To derive the results, we consider the first-order conditions from our fully-specified model, but consider the case where there are no adjustment costs ($\phi = 0$), which implies that firms maximize single-period profits.

To find the σ_{SR} , we combine (B.2) and (B.3) to get

$$E_t/N_t = (p_{E,t}/p_{N,t})^{-\epsilon} \theta_t^{\epsilon-1} \quad (\text{B.18})$$

and then compute

$$\sigma_{SR} := - \frac{d \ln(\frac{E_t}{N_t})}{d \ln(\frac{p_{E,t}}{p_{N,t}})} \bigg|_{\theta_t \text{ fixed}} = \epsilon. \quad (\text{B.19})$$

To find σ_{LR} , we evaluate the combined first-order condition (B.5) with $\phi = 0$:

$$E_t/N_t = \frac{\eta}{1-\eta} \left(\frac{p_{E,t}}{p_{N,t}} \right)^{-1}. \quad (\text{B.20})$$

and compute

$$\sigma_{LR} := -\frac{d \ln(\frac{E_t}{N_t})}{d \ln(\frac{p_{E,t}}{p_{N,t}})} = 1. \quad (\text{B.21})$$

Intensive Form

In this section, we present the intensive form of the model. We use $\hat{Z}_t \equiv Z_t/\bar{Z}_t$ to denote the detrended value for any variables Z_t and $\kappa_{Z,t} \equiv Z_t/Z_{t-1}$ to denote the ratio of current to last period level of any variable Z_t . In deriving the intensive form, we regularly make use of the approximation that $\exp(x) \approx 1 + x$ when x is small, which aids the translation of the model to the stochastic version used in calibration.

The trend for $C_t, I_t, \tilde{Y}_t, w_t$ is defined by $\tilde{Y}_t = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} (p_{c,t}^{1-\gamma} p_{d,t}^\gamma)^{\frac{-\eta}{(1-\alpha)(1-\eta)}}$. The trend for E_t is defined by $\bar{E}_t = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} (p_{c,t}^{1-\gamma} p_{d,t}^\gamma)^{\frac{\alpha(1-\eta)-1}{(1-\alpha)(1-\eta)}}$. The trend for N_t is defined by $\bar{N}_t = X_t^{\frac{\alpha}{(1-\alpha)(1-\eta)}} (p_{c,t}^{1-\gamma} p_{d,t}^\gamma)^{\frac{-\alpha\eta}{(1-\alpha)(1-\eta)}}$. The trends for K_t and θ_t are defined by $\bar{K}_t = X_{t-1}^{\frac{1}{(1-\alpha)(1-\eta)}} (p_{c,t-1}^{1-\gamma} p_{d,t-1}^\gamma)^{\frac{-\eta}{(1-\alpha)(1-\eta)}}$ and $\bar{\theta}_t = X_{t-1}^{\frac{-1}{1-\eta}} (p_{c,t-1}^{1-\gamma} p_{d,t-1}^\gamma)^{\frac{1}{1-\eta}}$ respectively. The trend for $p_{N,t}$ is defined by $\bar{p}_{N,t} = X_t^{\frac{1}{1-\eta}} (p_{c,t}^{1-\gamma} p_{d,t}^\gamma)^{\frac{-\eta}{1-\eta}}$. The trend for $p_{E,t}$ is defined by $\bar{p}_{E,t} = p_{c,t}^{1-\gamma} p_{d,t}^\gamma$. Lastly, the trends for $E_{c,t}$ and $E_{d,t}$ are defined by $\bar{E}_{c,t} = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} p_{c,t}^{\frac{-\eta(1-\gamma)-(1-\alpha)(1-\eta)}{(1-\alpha)(1-\eta)}} p_{d,t}^{\frac{-\eta\gamma}{(1-\alpha)(1-\eta)}}$ and $\bar{E}_{d,t} = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} p_{c,t}^{\frac{-\eta(1-\gamma)}{(1-\alpha)(1-\eta)}} p_{d,t}^{\frac{-\eta\gamma-(1-\alpha)(1-\eta)}{(1-\alpha)(1-\eta)}}$ respectively. Equations (B.22) to (B.36) describe the full competitive equilibrium in intensive form:

$$\exp\left(\frac{g_{X,t}}{(1-\alpha)(1-\eta)}\right) \exp\left(\frac{-\eta(1-\gamma)g_{p_{c,t}}}{(1-\alpha)(1-\eta)}\right) \exp\left(\frac{-\eta\gamma g_{p_{d,t}}}{(1-\alpha)(1-\eta)}\right) = \kappa_{\tilde{Y},t} \quad (\text{B.22})$$

$$\exp\left(\frac{-g_{X,t-1}}{1-\eta}\right) \exp\left(\frac{(1-\gamma)g_{p_{c,t-1}}}{1-\eta}\right) \exp\left(\frac{\gamma g_{p_{d,t-1}}}{1-\eta}\right) = \kappa_{\theta,t} \quad (\text{B.23})$$

$$\Lambda_{t+1} \equiv \beta \frac{\hat{C}_t}{\hat{C}_{t+1} \kappa_{\tilde{Y},t+1}} \quad (\text{B.24})$$

$$\mathbb{E}_t [\Lambda_{t+1} (1 + \rho_{t+1} - \delta)] = 1 \quad (\text{B.25})$$

$$\left(\left(\hat{\theta}_t^{-\eta} \hat{N}_t \kappa_{\theta,t+1}^\eta \right)^{\frac{\epsilon-1}{\epsilon}} + \left(\hat{\theta}_t^{1-\eta} \hat{E}_t \kappa_{\theta,t+1}^{\eta-1} \right)^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} = \hat{Y}_t \quad (\text{B.26})$$

$$\begin{aligned} \mathbb{E}_t \left\{ (1-\eta) \left[1 - \Phi \left(\frac{\hat{\theta}_{t+1}}{\hat{\theta}_t} \kappa_{\theta,t+1} \right) \right] - \frac{\hat{p}_{N,t+1} \hat{N}_{t+1}}{\hat{Y}_{t+1}} \right. \\ \left. - \Phi' \left(\frac{\hat{\theta}_{t+1}}{\hat{\theta}_t} \kappa_{\theta,t+1} \right) \left(\frac{\hat{\theta}_{t+1}}{\hat{\theta}_t} \kappa_{\theta,t+1} \right) \right. \\ \left. + \Lambda_{t+2} \Phi' \left(\frac{\hat{\theta}_{t+2}}{\hat{\theta}_{t+1}} \kappa_{\theta,t+2} \right) \left(\frac{\hat{\theta}_{t+2}}{\hat{\theta}_{t+1}} \kappa_{\theta,t+2} \right) \frac{\hat{Y}_{t+2}}{\hat{Y}_{t+1}} \kappa_{\tilde{Y},t+2} \right\} = 0 \end{aligned} \quad (\text{B.27})$$

$$\left[1 - \Phi \left(\frac{\hat{\theta}_t}{\hat{\theta}_{t-1}} \kappa_{\theta,t} \right) \right] \left(\left(\frac{\hat{\theta}_t}{\kappa_{\theta,t+1}} \right)^{1-\eta} \right)^{\frac{\epsilon-1}{\epsilon}} \left(\frac{\hat{Y}_t}{\hat{E}_t} \right)^{\frac{1}{\epsilon}} = \hat{p}_{E,t} \quad (\text{B.28})$$

$$\left[1 - \Phi \left(\frac{\hat{\theta}_t}{\hat{\theta}_{t-1}} \kappa_{\theta,t} \right) \right] \left(\left(\frac{\hat{\theta}_t}{\kappa_{\theta,t+1}} \right)^{-\eta} \right)^{\frac{\epsilon-1}{\epsilon}} \left(\frac{\hat{Y}_t}{\hat{N}_t} \right)^{\frac{1}{\epsilon}} = \hat{p}_{N,t} \quad (\text{B.29})$$

$$\hat{N}_t \kappa_{\tilde{Y},t}^\alpha = \hat{K}_t^\alpha \quad (\text{B.30})$$

$$\alpha \frac{\hat{N}_t}{\hat{K}_t} \kappa_{\tilde{Y},t} = \frac{\rho_t}{\hat{p}_{N,t}} \quad (\text{B.31})$$

$$(1-\alpha) \hat{N}_t = \frac{\hat{w}_t}{\hat{p}_{N,t}} \quad (\text{B.32})$$

$$\hat{p}_{E,t} = \tilde{\gamma} (1 + \tau_{d,t})^\gamma \quad (\text{B.33})$$

$$(1-\gamma) \hat{p}_{E,t} \hat{E}_t = \hat{E}_{c,t} \quad (\text{B.34})$$

$$\gamma \hat{p}_{E,t} \hat{E}_t = (1 + \tau_{d,t}) \hat{E}_{d,t} \quad (\text{B.35})$$

$$\hat{C}_t + \hat{K}_{t+1} - (1-\delta) \hat{K}_t \frac{1}{\kappa_{\tilde{Y},t}} + \hat{E}_{c,t} + \hat{E}_{d,t} = \left[1 - \Phi \left(\frac{\hat{\theta}_t}{\hat{\theta}_{t-1}} \kappa_{\theta,t} \right) \right] \hat{Y}_t. \quad (\text{B.36})$$

Balanced Growth Path

We use Z_{ss} to denote the steady-state value of any variable Z , which is frequently a detrended variable. The steady state of the full model is as follows:

$$\kappa_{\tilde{Y},ss} = \exp \left(\frac{g_X}{(1-\alpha)(1-\eta)} \right) \exp \left(\frac{-\eta(1-\gamma)g_{p_c}}{(1-\alpha)(1-\eta)} \right) \exp \left(\frac{-\eta\gamma g_{p_d}}{(1-\alpha)(1-\eta)} \right) \quad (\text{B.37})$$

$$\kappa_{\theta,ss} = \exp \left(\frac{-g_X}{1-\eta} \right) \exp \left(\frac{(1-\gamma)g_{p_c}}{1-\eta} \right) \exp \left(\frac{\gamma g_{p_d}}{1-\eta} \right) \quad (\text{B.38})$$

$$\Phi_{ss}(\kappa_{\theta,ss}) = \frac{\phi}{2} (\kappa_{\theta,ss} - 1)^2 \quad (\text{B.39})$$

$$\Phi'_{ss}(\kappa_{\theta,ss}) = \phi (\kappa_{\theta,ss} - 1) \quad (\text{B.40})$$

$$\Lambda_{ss} = \beta / \kappa_{\tilde{Y},ss} \quad (\text{B.41})$$

$$\rho_{ss} = \frac{1}{\Lambda_{ss}} - 1 + \delta \quad (\text{B.42})$$

$$s_{N,ss} = 1 - \eta - \frac{\Phi'_{ss}(\kappa_{\theta,ss})\kappa_{\theta,ss}(1 - \Lambda_{ss}\kappa_{\tilde{Y},ss})}{1 - \Phi_{ss}(\kappa_{\theta,ss})} \quad (\text{B.43})$$

$$s_{E,ss} = 1 - s_{N,ss} \quad (\text{B.44})$$

$$\hat{p}_{E,ss} = \tilde{\gamma}(1 + \tau_{d,ss})^\gamma \quad (\text{B.45})$$

$$\hat{\theta}_{ss} = \hat{p}_{E,ss}^{\frac{1}{1-\eta}} \kappa_{\theta,ss} \left((1 - \Phi_{ss}(\kappa_{\theta,ss}))^{\frac{\epsilon-1}{\epsilon}} s_{E,ss}^{\frac{-1}{\epsilon}} \right)^{\frac{-\epsilon}{(\epsilon-1)(1-\eta)}} \quad (\text{B.46})$$

$$\hat{p}_{N,ss} = (1 - \Phi_{ss}(\kappa_{\theta,ss})) s_{N,ss}^{\frac{-1}{\epsilon-1}} \left(\frac{\hat{\theta}_{ss}}{\kappa_{\theta,ss}} \right)^{-\eta} \quad (\text{B.47})$$

$$\hat{K}_{ss} = \left(\frac{\alpha \hat{p}_{N,ss}}{\rho_{ss}} \right)^{\frac{1}{1-\alpha}} \kappa_{\tilde{Y},ss} \quad (\text{B.48})$$

$$\hat{I}_{ss} = \left(1 - \frac{1-\delta}{\kappa_{\tilde{Y},ss}} \right) \hat{K}_{ss} \quad (\text{B.49})$$

$$\hat{N}_{ss} = \left(\frac{\hat{K}_{ss}}{\kappa_{\tilde{Y},ss}} \right)^\alpha \quad (\text{B.50})$$

$$\hat{w}_{ss} = (1 - \alpha) \hat{N}_{ss} \hat{p}_{N,ss} \quad (\text{B.51})$$

$$\hat{Y}_{ss} = \frac{\hat{p}_{N,ss} \hat{N}_{ss}}{(1 - \Phi_{ss}(\kappa_{\theta,ss})) s_{N,ss}} \quad (\text{B.52})$$

$$\hat{E}_{ss} = \frac{\hat{Y}_{ss} s_{E,ss} (1 - \Phi_{ss}(\kappa_{\theta,ss}))}{\hat{p}_{E,ss}} \quad (\text{B.53})$$

$$\hat{E}_{c,ss} = (1 - \gamma) \hat{p}_{E,ss} \hat{E}_{ss} \quad (\text{B.54})$$

$$\hat{E}_{d,ss} = \frac{\gamma \hat{p}_{E,ss} \hat{E}_{ss}}{1 + \tau_{d,ss}} \quad (\text{B.55})$$

$$\hat{C}_{ss} = (1 - \Phi_{ss}(\kappa_{\theta,ss})) \hat{Y}_{ss} - \hat{I}_{ss} - \hat{E}_{c,ss} - \hat{E}_{d,ss}. \quad (\text{B.56})$$

Intensive Form and Balanced Growth Path in Stochastic Model

In this section, we present changes in both the intensive form of model equations and steady-state values in the stochastic model, relative to the deterministic model. We model two stochastic processes for growth rates of X_t and $p_{E,t}$, and abstract away from the final energy producer's problem since our focus is on the energy efficiency dynamics in the stochastic setting. The introduction of randomness implies that the consumer and final good producer can only satisfy their first-order conditions in expectation.

In the stochastic model, the trend for $C_t, I_t, \tilde{Y}_t, w_t$ is defined by $\bar{\tilde{Y}}_t = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} p_{E,t}^{\frac{-\eta}{(1-\alpha)(1-\eta)}}$. The trend for E_t is defined by $\bar{E}_t = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} p_{E,t}^{\frac{\alpha(1-\eta)-1}{(1-\alpha)(1-\eta)}}$. The trend for N_t is defined by $\bar{N}_t = X_t^{\frac{\alpha}{(1-\alpha)(1-\eta)}} p_{E,t}^{\frac{-\alpha\eta}{(1-\alpha)(1-\eta)}}$. The trends for K_t and θ_t are defined by $\bar{K}_t = X_{t-1}^{\frac{1}{(1-\alpha)(1-\eta)}} p_{E,t-1}^{\frac{-\eta}{(1-\alpha)(1-\eta)}}$ and $\bar{\theta}_t = X_{t-1}^{\frac{-1}{1-\eta}} p_{E,t-1}^{\frac{1}{1-\eta}}$ respectively. Lastly, the trend for $p_{N,t}$ is defined by $\bar{p}_{N,t} = X_t^{\frac{1}{1-\eta}} p_{E,t}^{\frac{-\eta}{1-\eta}}$.

The following dynamic equations in intensive form need to be updated relative to their deterministic versions in Appendix Section B:

$$\exp\left(\frac{g_{X,t}}{(1-\alpha)(1-\eta)}\right) \exp\left(\frac{-\eta g_{pE,t}}{(1-\alpha)(1-\eta)}\right) = \kappa_{\tilde{Y},t} \quad (\text{B.57})$$

$$\exp\left(\frac{-g_{X,t-1}}{1-\eta}\right) \exp\left(\frac{g_{pE,t-1}}{1-\eta}\right) = \kappa_{\theta,t} \quad (\text{B.58})$$

$$\left[1 - \Phi\left(\frac{\hat{\theta}_t}{\hat{\theta}_{t-1}} \kappa_{\theta,t}\right)\right] \left(\left(\frac{\hat{\theta}_t}{\kappa_{\theta,t+1}}\right)^{1-\eta}\right)^{\frac{\epsilon-1}{\epsilon}} \left(\frac{\hat{\tilde{Y}}_t}{\hat{E}_t}\right)^{\frac{1}{\epsilon}} = 1 \quad (\text{B.59})$$

$$\hat{C}_t + \hat{K}_{t+1} - (1-\delta)\hat{K}_t \frac{1}{\kappa_{\tilde{Y},t}} + \hat{E}_t = \left[1 - \Phi\left(\frac{\hat{\theta}_t}{\hat{\theta}_{t-1}} \kappa_{\theta,t}\right)\right] \hat{\tilde{Y}}_t \quad (\text{B.60})$$

where

$$g_{X,t} = (1 - \rho_X)\mu_X + \rho_X g_{X,t-1} + \omega_{X,t} \quad (\text{B.61})$$

$$g_{pE,t} = \left(1 - \sum_{i=1}^4 \rho_{pE,i}\right) \mu_{pE} + \sum_{i=1}^4 \rho_{pE,i} g_{pE,t-i} + \sum_{i=1}^2 \psi_{pE,i} \omega_{pE,t-i} + \omega_{pE,t}. \quad (\text{B.62})$$

and the following steady-state values along the balanced growth path need to be updated relative to their values from Appendix Section B:

$$\kappa_{\tilde{Y},ss} = \exp\left(\frac{\mu_X}{(1-\alpha)(1-\eta)}\right) \exp\left(\frac{-\eta \mu_{pE}}{(1-\alpha)(1-\eta)}\right) \quad (\text{B.63})$$

$$\kappa_{\theta,ss} = \exp\left(\frac{-\mu_X}{1-\eta}\right) \exp\left(\frac{\mu_{pE}}{1-\eta}\right) \quad (\text{B.64})$$

$$\hat{\theta}_{ss} = \kappa_{\theta,ss} \left((1 - \Phi_{ss}(\kappa_{\theta,ss}))\right)^{\frac{\epsilon-1}{\epsilon}} s_{E,ss}^{\frac{-1}{\epsilon}} \quad (\text{B.65})$$

$$\hat{E}_{ss} = \hat{\tilde{Y}}_{ss} s_{E,ss} (1 - \Phi_{ss}(\kappa_{\theta,ss})) \quad (\text{B.66})$$

$$\hat{C}_{ss} = (1 - \Phi_{ss}(\kappa_{\theta,ss})) \hat{\tilde{Y}}_{ss} - \hat{I}_{ss} - \hat{E}_{ss}. \quad (\text{B.67})$$

In addition, all dynamic and steady state equations related to the final energy producer's

problem are omitted in the stochastic model.

CES Production for E_t

In the baseline deterministic model, we consider a Cobb-Douglas production for final energy E_t . In this section, we present the model variation with a CES production function for E_t , as specified in (2.23). With a CES production function, the final energy producer's problem yields the following first order conditions and final energy price index,

$$p_{E,t} \gamma \left(\frac{E_t}{E_{d,t}} \right)^{1/\chi} = p_{d,t} (1 + \tau_{d,t}) \quad (\text{B.68})$$

$$p_{E,t} (1 - \gamma) \left(\frac{E_t}{E_{c,t}} \right)^{1/\chi} = p_{c,t} \quad (\text{B.69})$$

$$p_{E,t} = [\gamma^\chi (p_{d,t} (1 + \tau_{d,t}))^{1-\chi} + (1 - \gamma)^\chi p_{c,t}^{1-\chi}]^{\frac{1}{1-\chi}} \quad (\text{B.70})$$

To be consistent with the balanced growth facts, we assume that $g_{p_{d,t}} = g_{p_{c,t}}$. The trend for $C_t, I_t, \tilde{Y}_t, w_t$ is defined by $\bar{Y}_t = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} (p_{c,t})^{\frac{-\eta}{(1-\alpha)(1-\eta)}}$. The trends for E_t and N_t are defined by $\bar{E}_t = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} (p_{c,t})^{\frac{\alpha(1-\eta)-1}{(1-\alpha)(1-\eta)}}$ and $\bar{N}_t = X_t^{\frac{\alpha}{(1-\alpha)(1-\eta)}} (p_{c,t})^{\frac{-\alpha\eta}{(1-\alpha)(1-\eta)}}$ respectively. The trends for K_t and θ_t are defined by $\bar{K}_t = X_{t-1}^{\frac{1}{(1-\alpha)(1-\eta)}} (p_{c,t-1})^{\frac{-\eta}{(1-\alpha)(1-\eta)}}$ and $\bar{\theta}_t = X_{t-1}^{\frac{-1}{1-\eta}} (p_{c,t-1})^{\frac{1}{1-\eta}}$ respectively. The trend for $p_{N,t}$ is defined by $\bar{p}_{N,t} = X_t^{\frac{1}{1-\eta}} (p_{c,t})^{\frac{-\eta}{1-\eta}}$. The trend for $p_{E,t}$ is defined by $\bar{p}_{E,t} = p_{c,t}$. Lastly, the trends for $E_{c,t}$ and $E_{d,t}$ are defined by $\bar{E}_{c,t} = \bar{E}_{d,t} = X_t^{\frac{1}{(1-\alpha)(1-\eta)}} p_{c,t}^{\frac{\alpha(1-\eta)-1}{(1-\alpha)(1-\eta)}}$.

With CES production for E_t , the following dynamic equations in intensive form need

to be updated relative to their versions in Appendix Section B:

$$\exp\left(\frac{g_{X,t}}{(1-\alpha)(1-\eta)}\right) \exp\left(\frac{-\eta g_{p_{c,t}}}{(1-\alpha)(1-\eta)}\right) = \kappa_{\tilde{Y},t} \quad (\text{B.71})$$

$$\exp\left(\frac{-g_{X,t-1}}{1-\eta}\right) \exp\left(\frac{g_{p_{c,t-1}}}{1-\eta}\right) = \kappa_{\theta,t} \quad (\text{B.72})$$

$$\hat{p}_{E,t} = \left[\gamma^\chi \left(\frac{p_{d,t}}{p_{c,t}} (1 + \tau_{d,t}) \right)^{1-\chi} + (1-\gamma)^\chi \right]^{\frac{1}{1-\chi}} \quad (\text{B.73})$$

$$\hat{p}_{E,t} \gamma \left(\frac{\hat{E}_t}{\hat{E}_{d,t}} \right)^{1/\chi} = \frac{p_{d,t}}{p_{c,t}} (1 + \tau_{d,t}) \quad (\text{B.74})$$

$$\hat{p}_{E,t} (1-\gamma) \left(\frac{\hat{E}_t}{\hat{E}_{c,t}} \right)^{1/\chi} = 1 \quad (\text{B.75})$$

$$\hat{C}_t + \hat{K}_{t+1} - (1-\delta)\hat{K}_t \frac{1}{\kappa_{\tilde{Y},t}} + \hat{E}_{c,t} + \frac{p_{d,t}}{p_{c,t}} \hat{E}_{d,t} = \left[1 - \Phi \left(\frac{\hat{\theta}_t}{\hat{\theta}_{t-1}} \kappa_{\theta,t} \right) \right] \hat{\tilde{Y}}_t. \quad (\text{B.76})$$

The following steady-state values need to be updated relative to their values in Appendix Section B:

$$\hat{p}_{E,ss} = \left[\gamma^\chi \left(\frac{p_{d,ss}}{p_{c,ss}} (1 + \tau_{d,ss}) \right)^{1-\chi} + (1-\gamma)^\chi \right]^{\frac{1}{1-\chi}} \quad (\text{B.77})$$

$$\hat{E}_{d,ss} = \hat{E}_{ss} \left(\frac{\hat{p}_{E,ss} \gamma}{\frac{p_{d,ss}}{p_{c,ss}} (1 + \tau_{d,ss})} \right)^\chi \quad (\text{B.78})$$

$$\hat{E}_{c,ss} = \hat{E}_{ss} (\hat{p}_{E,ss} (1-\gamma))^\chi \quad (\text{B.79})$$

$$\hat{C}_{ss} = (1 - \Phi_{ss}(\kappa_{\theta,ss})) \hat{\tilde{Y}}_{ss} - \hat{I}_{ss} - \hat{E}_{c,ss} - \frac{p_{d,ss}}{p_{c,ss}} \hat{E}_{d,ss}. \quad (\text{B.80})$$

Throughout the analysis with CES production for E_t , we assume that $\frac{p_{d,t}}{p_{c,t}} = \frac{p_{d,ss}}{p_{c,ss}} = 0.4$, i.e., the dirty-to-clean relative price is held fixed at its initial value.

Proof of Proposition 2.3

The first-order conditions for the social planner's problem are:

$$C_t : \quad \frac{1}{C_t} = \lambda_t, \quad (\text{B.81})$$

$$E_{c,t} : \quad \left(1 - \Phi\left(\frac{\theta_t}{\theta_{t-1}}\right)\right) \frac{\partial \tilde{Y}_t}{\partial E_t} (1 - \gamma) \frac{E_t}{E_{c,t}} = p_{c,t}, \quad (\text{B.82})$$

$$E_{d,t} : \quad \left(1 - \Phi\left(\frac{\theta_t}{\theta_{t-1}}\right)\right) \frac{\partial \tilde{Y}_t}{\partial E_t} \gamma \frac{E_t}{E_{d,t}} = p_{d,t} + \frac{\Omega}{\beta^t \lambda_t}, \quad (\text{B.83})$$

$$K_{t+1} : \quad 1 = \beta \frac{C_t}{C_{t+1}} (1 + \rho_{t+1} - \delta), \quad (\text{B.84})$$

$$\Omega : \quad \Omega \geq 0, \quad \sum_{t=0}^{\infty} E_{d,t} \leq E^{\max}, \quad \Omega \left(E^{\max} - \sum_{t=0}^{\infty} E_{d,t} \right) = 0. \quad (\text{B.85})$$

It is easy to confirm that these are just the decentralized equilibrium condition with one exception: the first-order condition for dirty energy has the extra term $\frac{\Omega}{\beta^t \lambda_t}$. This extra term is the external cost of dirty energy use. This term exists because firms do not take into account how their dirty energy use affects the policy goal, and the externality is just the shadow value on this constraint.

Note that $C_t^{-1} = \lambda_t$ is the marginal utility of consumption, as shown in the first-order condition for C_t . To implement the optimal policy, therefore, we simply require a tax, $\tau_t^u = \beta^{-t} C_t \Omega$, that grows at rate $g_{\tau,t} = \beta^{-1} \frac{C_t}{C_{t-1}}$. From (B.15), this growth rate is equal to $1 + r_t$.

Numerical Solution Method

We implement the model solution in *Dynare* (Adjemian et al., 2024). The stochastic model is solved with a first-order approximation of the model around its steady state. The deterministic model is solved with a perfect foresight solver given an initial and a terminal steady state.

With a per-unit tax, the tax-inclusive price of dirty energy only grows at a constant rate asymptotically. As a result, the BGP only exists asymptotically. When simulating the LCP, therefore, we assume that the tax becomes a constant ad valorem tax after $T_1 = 300$ years. We simulate out to a period of $T_2 = 600$ years. When calculating the CEV, we sum welfare over the first $T_3 = 300$ years. We confirm that, because of discounting, changing T_1 , T_2 or T_3 has essentially no impact on the dynamics during the period we study.

Data Sources

Unless otherwise noted, all of our data come from the State Energy Data System (SEDS) produced by the EIA (<https://www.eia.gov/state/seds/seds-data-complete.php?sid=US>). The data cover the period 1970-2021. All variable codes refer to these data. The codebook is available at <https://www.eia.gov/state/seds/seds-technical-notes-complete.php?sid=US> (see ‘Codes and descriptions’).

Data for Figure 2.1 and IRFs

The SEDS data directly report total expenditure on final energy (‘total end-use energy expenditures’, TETXV) and the BTU-weighted nominal price of final energy (‘total end-use energy average price’, TETXD). Both are nominal values. The EIA also reports nominal GDP, real GDP, and the GDP deflator in its Monthly Energy Review (MER) data series (see Table C1 at <https://www.eia.gov/totalenergy/data/monthly/>). These data originally come from the Bureau of Economic Analysis (BEA). The energy expenditure share is total expenditure on final energy divided by the sum of nominal GDP and energy expenditure, and the real price of final energy is the nominal price divided by the GDP deflator. Real values are measured in 2012 dollars.

Prior to estimating the IRFs in the model, we separately estimate the dynamics of TFP. To do so, we use data on total factor productivity at constant national prices for the United States from the Federal Reserve Bank of St. Louis (<https://fred.stlouisfed.org/series/RTFPNAUSA632NRUG>). The TFP data series covers the period 1970-2019.

Transformed Data Used in VARs




Figure B.1: Data Inputs into VAR

Panel (a) of Figure B.1 shows the raw time series of energy expenditure share as well

as a fitted linear time trend. Panel (b) shows the two data series used in our baseline analysis: the growth rate of energy prices and the detrended energy expenditure share.

Data for Section 2.5.3

To calculate total expenditure on fossil fuels, we sum total expenditure on coal (CLTCV), natural gas (NGTCV), and petroleum products (PATCV). Note that this includes coal and natural gas used as final energy. When coal and natural gas are not altered prior to their use in production, they are both primary and final energy. We divide total expenditure on fossil fuels by total expenditure on final energy to estimate γ .

To measure the total quantity of dirty use, we add total consumption of coal (CLTCB), natural gas (NGTCB), and petroleum products (PATCB). As in our main calculation for final energy, we measure quantities in BTUs for all energy sources. To measure the average price of dirty energy, we divide expenditure on fossil fuels by total usage. We convert nominal prices to real prices using the GDP deflator. To measure the total quantity of clean energy, we take the total quantity of renewable energy directly from the data (RETCB).

Comparison to Hassler et al. (2021a)




Figure B.2: Comparison to Hassler et al. (2021a) Measurement

Notes: Panel (a) shows the average real price of final energy from the SEDS data (\$2012/million BTUs) and the average real price of final energy using the HKO methodology (\$2012/million BTUs of electricity). Panel (b) shows quantity of final energy from the SEDS data (quadrillion BTUs) and using the HKO methodology (quadrillion BTUs of electricity).

Hassler et al. (2021a) suggest an alternative method of estimating aggregate energy prices that weights different types of energy by their average price. They apply their

method to fossil fuel primary energy. We recalculate our measure of the average real price of final energy and show that it is very similar to the series we use in our empirical analysis.

To do so, we consider final energy from four sources: electricity, coal, natural gas, and petroleum products. Coal and natural gas are considered final energy when they are not used to produce electricity (e.g., when they are used in industrial production or for residential heating). This analysis excludes final energy from biomass, because price data are unavailable. SEDS reports the nominal price of each energy type when it is used as a source of final energy (ESTXD, CLTXD, NGTXD, and PATXD). As above, we convert these to real values using the GDP deflator. On average, coal is 11% as expensive as electricity per BTU, while natural gas is 25% as expensive and petroleum products are 51% as expensive. The Hassler et al. (2021a) method assumes that these relative prices reflect differences in the relative efficiency of different energy sources. Then, we can calculate the total quantity of energy measured in BTUs of electricity as

$$E^{\text{HKO}} = E_e + 0.11E_c + 0.25E_g + 0.51E_p, \quad (\text{B.86})$$

where E_e is BTUs of electricity used as a final energy source (ESTXB in the SEDS data), E_c is BTUs of coal (CLTXB), E_g is BTUs of natural gas (NGTXB), and E_p is BTUs of petroleum products (PATXB).

Note that the calculation of total expenditure on final energy is not affected by how we weight energy sources. The weighted-average real price of final energy can then be calculated as real expenditure on final energy divided by E^{HKO} :

$$p_E^{\text{HKO}} = \frac{p_E E}{E^{\text{HKO}}}, \quad (\text{B.87})$$

where $p_E E$ is real energy expenditure. Panel (a) of B.2 plots the price series used in our analysis and the HKO version. While they have different units (BTUs versus BTUs of electricity), they move together very closely, indicating that our results are not driven by our method of calculating the average real price of final energy. For completeness, panel (b) plots total BTUs of final energy against E^{HKO} and shows that they also follow a similar path.

Matching Procedure

This appendix provides further details on the matching procedure. Consider the vector $\boldsymbol{\varphi} = (\epsilon, \phi)'$. Our aim is to choose the vector $\boldsymbol{\varphi}$ that ensures the best match between the IRFs of the empirical and theoretical models over p periods. Specifically, we want to match the IRFs for the energy share of income, $s_{E,t}$, as the ARIMA function ensures a match between the IRFs for energy price growth, $\Delta \ln(p_{E,t})$. Our estimated vector is denoted $\hat{\boldsymbol{\varphi}}$. We estimate $\hat{\boldsymbol{\varphi}}$ using the following weighted sum of squared residuals estimator:

$$\hat{\boldsymbol{\varphi}} = \arg \min_{\boldsymbol{\varphi}} \left(IRF^{\text{model}}(\boldsymbol{\varphi}) - IRF^{\text{empirical}} \right)' W \left(IRF^{\text{model}}(\boldsymbol{\varphi}) - IRF^{\text{empirical}} \right),$$

where $IRF^{\text{empirical}}$ and IRF^{model} are $(p \times 1)$ vectors with empirical and model generated IRFs for $s_{E,t}$, and W is a $(p \times p)$ diagonal weighting matrix containing inverse variance estimates from the empirical estimation. The weighting matrix ensures that deviations from the empirical IRF are penalized more heavily when the corresponding estimates are more precise.

Robustness

No Detrending

Figure B.3 shows IRFs from our SVAR model where only $s_{E,t}$ is detrended (baseline), no variables are detrended, and both variables are detrended. A notable difference is the slower transition dynamics for $s_{E,t}$ when $s_{E,t}$ is not detrended. Meanwhile, detrending $\Delta \ln(p_{E,t})$ has little effect on the dynamics.




Figure B.3: IRFs for a one-standard deviation shock to $\Delta \ln(p_{E,t})$ in a SVAR model with a short-run restriction and one lag with and without detrending: detrending of $s_{E,t}$ (baseline), no detrending, and detrending of both variables

Lag Selection

Table B.1 shows lag-order selection statistics based on VAR models with one to six lags. All selection statistics suggest one lag.

To investigate the importance of the lag-order, we compute IRFs from our SVAR model using a short-run restriction and one to four lags. As shown in Figure B.4, the overall patterns are similar. The models with one and two lags provide quite similar IRFs, while the IRFs become notably noisier for three and four lags.

Augmented SVAR Models

This appendix provides estimation results from SVAR models with three variables. We call these augmented SVAR models. The augmented SVAR models are given by

$$x_t = c + A_1 x_{t-1} + \dots + A_n x_{t-n} + u_t, \quad (\text{B.88})$$

Table B.1: Lag-order selection statistics for VAR model

Lags	LR	FPE	AIC	HQIC	SBIC
0	-	$1.4 \cdot 10^{-6}$	-7.77	-7.74	-7.69
1	190.81*	$2.5 \cdot 10^{-8}$ *	-11.84*	-11.75*	-11.60*
2	3.863	$2.7 \cdot 10^{-8}$	-11.74	-11.60	-11.34
3	4.20	$3.0 \cdot 10^{-8}$	-11.66	-11.45	-11.10
4	8.01	$3.0 \cdot 10^{-8}$	-11.66	-11.39	-10.94
5	4.57	$3.3 \cdot 10^{-8}$	-11.58	-11.26	-10.70
6	1.54	$3.8 \cdot 10^{-8}$	-11.44	-11.05	-10.40

Notes: The following acronyms are used: likelihood ratio (LR), final prediction error (FPE), Akaike's information criterion (AIC), Schwarz's Bayesian information criterion (SBIC), and the Hannan and Quinn information criterion (HQIC). All tests are conducted on the same effective sample.




Figure B.4: IRFs for a one-standard deviation shock to $\Delta \ln(p_{E,t})$ for different numbers of lags in a SVAR model with a short-run restriction

where $x_t = (\Delta \ln(p_{E,t}), y_t, s_{E,t})'$, y_t is the additional variable (e.g. GDP growth or unemployment level), c is a (3×1) vector of constants, the A variables are (3×3) coefficient matrices, and u_t is a (3×1) vector of error terms.

For identification, we assume that energy price growth is not contemporaneously affected by shocks to the energy share of income and the additional variable, and the additional variable is not contemporaneously affected by shocks to the energy share of income. Again, the idea is that end-use energy prices react with a lag to other innovations in the economy due to contracts and other rigidities.

We have estimated the augmented SVAR using the (i) detrended unemployment rate, (ii) employment growth, (iii) detrended employment growth, (iv) non-farm employment growth, (v) detrended non-farm employment growth, (vi) nominal GDP growth, (vii) detrended nominal GDP growth, (viii) real GDP growth, and (ix) detrended real GDP growth. Lag-order selection statistics generally support one lag, and all estimations are therefore conducted using one lag.

All estimated IRFs from the augmented models show the same overall patterns as our baseline specification with two variables. To illustrate this, Figure B.5 shows IRFs for $\Delta \ln(p_{E,t})$ and $s_{E,t}$ following a one-standard deviation shock to $\Delta \ln(p_{E,t})$ for four augmented models and our baseline specification with two variables.




Figure B.5: IRFs for a one-standard deviation shock to $\Delta \ln(p_{E,t})$ for the baseline SVAR model and four augmented models

Quarterly Data Analysis

The EIA only provides energy expenditure data at a yearly frequency. To conduct an econometric analysis at a higher frequency, we therefore need to construct our own energy expenditure estimate using available higher frequency data. This appendix first describes how a quarterly dataset is constructed using monthly data from the EIA and quarterly

data from the BEA. Thereafter, we conduct a SVAR analysis using the quarterly dataset to investigate the robustness of the results from the main analysis.

Computing Monthly Energy Expenditures

We use monthly energy statistics provided by the U.S. Energy Information Administration (2025a). These data contain information on the quantity of energy consumed and energy prices. We then compute energy expenditures at the monthly frequency using these figures. We verify our method by aggregating our monthly estimate into a yearly estimate, which we then compare to the official yearly estimate from the EIA.

Aggregate energy consumption is measured as end-use energy consumed across all end-use sectors. End-use energy consumed is primary energy plus electricity. The end-use sectors used by the EIA are the residential sector, the commercial sector, the industrial sector, and the transportation sector. However, as we lack price data on renewable energy, and since renewable energy constitutes a small part of primary energy consumption by the end-use sectors, we do not include renewable energy in our primary energy computation. Renewable energy in the electricity sector is included through electricity purchases by end-use sectors.

We start by computing end-use energy for coal, natural gas, petroleum products, and electricity for each end-use sector. We sum end-use energy to obtain an aggregate measure. We then aggregate to an annual level and compare the figures to end-use energy consumption from the SEDS database. The level deviates by around 5 percent on average, but – as shown in Figure B.6 – the variations in the two time series are very similar.




Figure B.6: Comparing yearly end-use energy consumption from the SEDS database with computations based on monthly data

We then obtain sector-specific monthly price data for the four energy types. For natural gas, we use "Natural Gas Price, Delivered to Consumers". These prices are reported for the residential sector, the commercial sector, and the industrial sector. There

are missing data for the commercial and the industrial sector, and there are no values for the transport sector. When data are unavailable, we use prices for the residential sector. For petroleum products, we use "Refiner Price of Finished Motor Gasoline to End Users" for all sectors. There are no price data on coal for the end-use sectors. We use "Cost of Coal Receipts at Electric Generating Plants". This is motivated by the fact that most coal is consumed by the electricity sector (around 90 percent in 2023 according to the U.S. Energy Information Administration (2024)). Hence, this coal price must be the dominating market price in the US coal market. Finally, we use end-use sector-specific electricity prices ("Average Price of Electricity to Ultimate Customers") when these are available. In case of missing data, we use residential electricity prices. These are, however, not available for the period January 1984 to December 1989. We, therefore, extrapolate sector-specific electricity prices linearly for this period.

Fuel-specific expenditures for each end-use sector are computed by multiplying prices and quantities. Aggregate energy expenditures are obtained by summing fuel-specific expenditures for all sectors. The monthly dataset starts in January 1982. The shortened time span compared to the yearly dataset is due to unavailable prices for petroleum products.

We then aggregate to yearly data to check the quality of our computation against official EIA data. Figure B.7 shows that our yearly estimate tracks the official numbers from the SEDS database closely.




Figure B.7: Comparing yearly energy expenditures from the SEDS database with yearly estimates based on monthly data

Constructing a Quarterly Dataset

The next step is to compute energy expenditures at the quarterly level, aligning with the highest available frequency of GDP data. Quarterly GDP data are downloaded from the website of the U.S. Bureau of Economic Analysis (2025b), and the energy share of

income at the quarterly frequency is obtained by dividing quarterly energy expenditures by quarterly GDP figures (using the gross energy share of income results in similar conclusions). Figure B.8 compares the yearly energy share of income from the SEDS database with the computed quarterly energy share. The two time series closely track each other. We compute the energy price by dividing energy expenditures by energy use.




Figure B.8: Comparing the energy share of income from the quarterly dataset and yearly data from the SEDS database

Notes: Yearly SEDS data placed in the first quarter of every year.

Finally, we obtain real prices using the GDP deflator.

SVAR Analysis

Using quarterly data may introduce short-run dynamics into the VAR model like business cycle fluctuations and seasonal variation in energy expenditures. We, therefore, need to add seasonal dummies to our VAR model and rerun specification tests.

We start by considering the VAR model

$$x_t = c + \sum_{j=1}^p A_j x_{t-j} + dS_t + u_t,$$

where x_t is a $(n \times 1)$ vector with n endogenous variables, c is a $(n \times 1)$ vector of constants, p is the lag length, A_j are $(n \times n)$ matrices, S_t is a (3×1) vector containing seasonal dummies, d is a $(n \times 3)$ matrix of season coefficients, and u_t is a $(n \times 1)$ vector containing the error terms. A seasonal dummy for quarter i equals one when it is quarter i and zero otherwise. Only three seasonal dummies are needed, as the fourth quarter is identified as the quarter where the three seasonal dummies equal zero. Lag length statistics generally support 5 lags, hence we run our quarterly SVAR model with 5 lags.

Figure B.9 shows IRFs based on quarterly data for our main specification with two variables and four augmented models. We see the same overall patterns as in the yearly analysis, although the energy price dynamics are notably faster. The dynamics for the energy share are also faster, but in line with our yearly IRFs, the main specification results in an IRF significantly different from zero for around 10 years at the 68% level (close to 9 years using quarterly data and 10 years using yearly data).




Figure B.9: IRFs for a one-standard deviation shock to $\Delta \ln(p_{E,t})$ for the baseline SVAR model and four augmented models based on quarterly data

Additional Results

Clean Energy Subsidies

In this section, we investigate the results of a clean energy subsidy, rather than a dirty energy tax. We focus on the case of a 20% subsidy, which is roughly in line with the subsidies in the Inflation Reduction Act (Bistline et al., 2023). Unlike carbon taxes, clean energy subsidies decrease the price of energy and discourage energy efficiency (e.g., Newell et al., 2019). Thus, these policies are *more* effective when accounting for the slow dynamics of energy efficiency.

Figure B.10 shows the results. As expected, panel (a) shows that subsidies cause the energy expenditure share to fall in the short run, because the elasticity of substitution is below one. As in the case of carbon taxes, the transition back to the original expenditure share happens very slowly. Panels (b) and (c) show the slow dynamics of energy intensity and final energy use.

Panels (d) and (e) show the dynamics of flow and cumulative dirty energy use, respectively. Casey et al. (2026b) show that clean energy subsidies actually *increase* emissions in the GHKT special case. While we find identical long-run results, our results suggest that clean energy subsidies are much more effective when taking into account the slow dynamics of energy use. The subsidies decrease flow dirty energy use by 5 percent in the short run and 0.1 percent in 2040. While the results still suggest limited effectiveness of these policies in reducing dirty energy use, the results also suggest that subsidies are unlikely to exacerbate environmental outcomes in the short run. More generally, the results again demonstrate the importance of accounting for slow energy transition dynamics when simulating the impacts of climate policy.

Panels (g) and (h) show the paths of GDP and consumption. The subsidy increases energy production, boosting GDP. But, it does so at the cost of decreasing consumption, at least partially offsetting the benefit of reduced emissions. This negative impact on consumption is expected, because the laissez-faire equilibrium maximizes the utility from consumption.

Robustness

In this section, we show robustness analyses in Figure B.11, Figure B.12, Figure B.13, and Figure B.14. When production is CES, we assume $g_{p_d} = g_{p_c} = 1.6\%/year$ in all specifications. This is necessary to ensure balanced growth.








Figure B.10: Impacts of a 20% Clean Energy Subsidy

Figure B.11: Impacts of a 50% Carbon Tax ($\chi = 2$ vs. Cobb-Douglas $\chi = 1$)

Figure B.12: Impacts of a 50% Carbon Tax ($\chi = 0.95$ vs. Cobb-Douglas $\chi = 1$)

Figure B.13: Impacts of a 50% Carbon Tax ($\eta = 0.04$ vs. $\eta = 0.075$)

Figure B.14: Least-Cost Path ($\chi = 2$ vs. Cobb-Douglas $\chi = 1$)

Bibliography

- Acemoglu, D., Aghion, P., Barrage, L., and Hémous, D. (2023). Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution. NBER Working Paper 31657, National Bureau of Economic Research, Cambridge, MA.
- Acemoglu, D., Aghion, P., Bursztyn, L., and Hémous, D. (2012). The environment and directed technical change. *American Economic Review*, 102(1):131–166.
- Achakulwisut, P., Erickson, P., Guivarch, C., Schaeffer, R., Brutschin, E., and Pye, S. (2023). Global fossil fuel reduction pathways under different climate mitigation strategies and ambitions. *Nature Communications*, 14(1):5425.
- Adjemian, S., Juillard, M., Karamé, F., Mutschler, W., Pfeifer, J., Ratto, M., Rion, N., and Villemot, S. (2024). Dynare: reference manual, version 6. Dynare Working Papers 80, CEPREMAP.
- Airaud, F. S., Pappa, E., and Seoane, H. D. (2025). The green metamorphosis of a small open economy. Unpublished manuscript.
- Almeida, H. and Campello, M. (2007). Financial constraints, asset tangibility, and corporate investment. *The Review of Financial Studies*, 20(5):1429–1460.
- André, F. J. and Smulders, S. (2014). Fueling growth when oil peaks: Directed technological change and the limits to efficiency. *European Economic Review*, 69:18–39.
- Atkeson, A. and Kehoe, P. J. (1999). Models of energy use: Putty-putty versus putty-clay. *American Economic Review*, 89(4):1028–1043.
- Barrage, L. (2020). Optimal dynamic carbon taxes in a climate–economy model with distortionary fiscal policy. *Review of Economic Studies*, 87(1):1–39.
- Barrage, L. and Nordhaus, W. D. (2024). Policies, projections, and the social cost of carbon: Results from the DICE-2023 model. *Proceedings of the National Academy of Sciences*, 121(13):e2312030121.
- Berndt, E. R. and Wood, D. O. (1975). Technology, prices, and the derived demand for energy. *Review of Economics and Statistics*, 57(3):259–268.

BIBLIOGRAPHY

- Bistline, J. E. T., Mehrotra, N. R., and Wolfram, C. (2023). Economic implications of the climate provisions of the Inflation Reduction Act. *Brookings Papers on Economic Activity*, 2023(1):77–182.
- Blanchard, O. J. and Quah, D. (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review*, 79(4):655–673.
- Bovenberg, A. L. and Goulder, L. H. (1996). Optimal environmental taxation in the presence of other taxes: General equilibrium analyses. *American Economic Review*, 86(4):985–1000.
- Brown, J. R., Fazzari, S. M., and Petersen, B. C. (2009). Financing innovation and growth: Cash flow, external equity, and the 1990s R&D boom. *The Journal of Finance*, 64(1):151–185.
- Buera, F. J. and Moll, B. (2015). Aggregate implications of a credit crunch: The importance of heterogeneity. *American Economic Journal: Macroeconomics*, 7(3):1–42.
- Busch, T., Johnson, M., and Pioch, T. (2022). Corporate carbon performance data: Quo vadis? *Journal of Industrial Ecology*, 26(1):350–363.
- Caggese, A. and Perez-Orive, A. (2022). How stimulative are low real interest rates for intangible capital? *European Economic Review*, 142:103987.
- Campiglio, E., Dietz, S., and Venmans, F. (2026). Optimal climate policy as if the transition matters. CESifo Working Paper 10139, CESifo, Munich.
- Capelle, D., Kirti, D., Pierri, N., and Villegas-Bauer, G. (2023). Mitigating climate change at the firm level: Mind the laggards. IMF Working Paper.
- Caselli, F. and Coleman II, W. J. (2006). The world technology frontier. *American Economic Review*, 96(3):499–522.
- Casey, G. (2024). Energy efficiency and directed technical change: Implications for climate change mitigation. *Review of Economic Studies*, 91(1):192–228.
- Casey, G., Gao, Y., and Kruse-Andersen, P. K. (2026a). Some macroeconomic evidence on substitution between clean and dirty energy. *AEA Papers and Proceedings*, 116:522–526.
- Casey, G., Gao, Y., and Kruse-Andersen, P. K. (forthcoming). Energy efficiency dynamics and climate policy. *Journal of the Association of Environmental and Resource Economists*.
- Casey, G., Jeon, W., and Traeger, C. (2026b). The macroeconomics of clean energy subsidies. *Journal of Environmental Economics and Management*, 138:103311.

BIBLIOGRAPHY

- Climate Impact Lab (2022). Data-driven spatial climate impact model user manual, version 092022-EPA. Available at https://www.epa.gov/system/files/documents/2022-11/epa_scghg_report_draft_0.pdf.
- Cruz, J.-L. and Rossi-Hansberg, E. (2024). The economic geography of global warming. *Review of Economic Studies*, 91(2):899–939.
- Cunha, I. and Pollet, J. (2020). Why do firms hold cash? Evidence from demographic demand shifts. *The Review of Financial Studies*, 33(9):4102–4138.
- Dell’Ariccia, G., Kadyrzhanova, D., Minoiu, C., and Ratnovski, L. (2021). Bank lending in the knowledge economy. *The Review of Financial Studies*, 34(10):5036–5076.
- EPA (2022). EPA external review draft of report on the social cost of greenhouse gases: Estimates incorporating recent scientific advances. *Supplementary Material for the Regulatory Impact Analysis for the Supplemental Proposed Rulemaking, “Standards of Performance for New, Reconstructed, and Modified Sources and Emissions Guidelines for Existing Sources: Oil and Natural Gas Sector Climate Review”*. Available at https://www.epa.gov/system/files/documents/2022-11/epa_scghg_report_draft_0.pdf.
- Falato, A., Kadyrzhanova, D., Sim, J., and Steri, R. (2022). Rising intangible capital, shrinking debt capacity, and the U.S. corporate savings glut. *The Journal of Finance*, 77(5):2799–2852.
- Farajpour, M. and Hinkelmann, S. (2021). Climate policies and input substitutability over time. Unpublished manuscript.
- Foster, L., Haltiwanger, J., and Syverson, C. (2008). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1):394–425.
- Fried, S. (2018). Climate policy and innovation: A quantitative macroeconomic analysis. *American Economic Journal: Macroeconomics*, 10(1):90–118.
- Fried, S., Novan, K., and Peterman, W. B. (2024). Understanding the inequality and welfare impacts of carbon tax policies. *Journal of the Association of Environmental and Resource Economists*, 11(S1):S231–S260.
- Fullerton, D. and Kim, S.-R. (2008). Environmental investment and policy with distortionary taxes, and endogenous growth. *Journal of Environmental Economics and Management*, 56(2):141–154.
- Fullerton, D. and Ta, C. L. (2020). Costs of energy efficiency mandates can reverse the sign of rebound. *Journal of Public Economics*, 188:104225.

BIBLIOGRAPHY

- Gentile, C. (2024). Relying on intermittency: Clean energy, storage, and innovation in a macro climate model. *Working Paper*.
- Gillingham, K., Rapson, D., and Wagner, G. (2016). The rebound effect and energy efficiency policy. *Review of Environmental Economics and Policy*, 10(1):68–88.
- Golosov, M., Hassler, J., Krusell, P., and Tsyvinski, A. (2014). Optimal taxes on fossil fuel in general equilibrium. *Econometrica*, 82(1):41–88.
- Gopinath, G., Kalemli-Özcan, Ş., Karabarbounis, L., and Villegas-Sanchez, C. (2017). Capital allocation and productivity in south Europe. *The Quarterly Journal of Economics*, 132(4):1915–1967.
- Goulder, L. H. (1995). Effects of carbon taxes in an economy with prior tax distortions: An intertemporal general equilibrium analysis. *Journal of Environmental Economics and Management*, 29(3):271–297.
- Griffin, J. M. and Gregory, P. R. (1976). An intercountry translog model of energy substitution responses. *American Economic Review*, 66(5):845–857.
- Hall, B. H. (2002). The financing of research and development. *Oxford Review of Economic Policy*, 18(1):35–51.
- Hassler, J. and Krusell, P. (2018). Environmental macroeconomics: The case of climate change. In *Handbook of environmental economics*, volume 4, pages 333–394. Elsevier.
- Hassler, J., Krusell, P., and Olovsson, C. (2018). The consequences of uncertainty: Climate sensitivity and economic sensitivity to the climate. *Annual Review of Economics*, 10:189–205.
- Hassler, J., Krusell, P., and Olovsson, C. (2021a). Directed technical change as a response to natural resource scarcity. *Journal of Political Economy*, 129(11):3039–3072.
- Hassler, J., Krusell, P., and Olovsson, C. (2021b). Presidential address 2020: Suboptimal climate policy. *Journal of the European Economic Association*, 19(6):2895–2928.
- Hassler, J., Krusell, P., and Smith Jr, A. A. (2016). Environmental macroeconomics. In *Handbook of macroeconomics*, volume 2, pages 1893–2008. Elsevier.
- Hinkelmann, S. (2024a). (Be-)coming clean: A model of the U.S. energy transition. Unpublished manuscript.
- Hinkelmann, S. (2024b). Electrification of U.S. aggregate production: Evidence and model. Unpublished manuscript.
- Holland, S. P., Hughes, J. E., and Knittel, C. R. (2009). Greenhouse gas reductions under low carbon fuel standards? *American Economic Journal: Economic Policy*, 1(1):106–46.

BIBLIOGRAPHY

- Hotelling, H. (1931). The economics of exhaustible resources. *Journal of Political Economy*, 39(2):137–175.
- Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing TFP in China and India. *The Quarterly Journal of Economics*, 124(4):1403–1448.
- Iovino, L., Martin, T., and Sauvagnat, J. (2023). Corporate taxation and carbon emissions. CEPR Discussion Paper No. 17915.
- Jo, A. (2025). Substitution between clean and dirty energy with biased technical change. *International Economic Review*, 66(2):883–902.
- Jones, C. I. (2005). The shape of production functions and the direction of technical change. *Quarterly Journal of Economics*, 120(2):517–549.
- Jones, C. I. (2016). The facts of economic growth. In *Handbook of Macroeconomics*, volume 2, pages 3–69. Elsevier.
- Kabir, P. and Tan, E. (2025). Green financing, capital (mis)allocation, and unintended consequences. Working Paper.
- Kim, S. (2023). Optimal carbon taxes and misallocation across heterogeneous firms. Working Paper.
- Klenow, P. J., Pasten, E., and Ruane, C. (2024). Carbon taxes and misallocation in Chile. Working Paper.
- Koetse, M. J., De Groot, H. L., and Florax, R. J. (2008). Capital-energy substitution and shifts in factor demand: A meta-analysis. *Energy Economics*, 30(5):2236–2251.
- Lanteri, A. and Rampini, A. A. (2025). Financing the adoption of clean technology. NBER Working Paper 33545, National Bureau of Economic Research, Cambridge, MA.
- Le Quéré, C., Korsbakken, J. I., Wilson, C., Tosun, J., Andrew, R., Andres, R. J., Canadell, J. G., Jordan, A., Peters, G. P., and van Vuuren, D. P. (2019). Drivers of declining CO₂ emissions in 18 developed economies. *Nature Climate Change*, 9(3):213–217.
- Leon-Ledesma, M. A. and Satchi, M. (2019). Appropriate technology and balanced growth. *Review of Economic Studies*, 86(2):807–835.
- Liu, T. and Xu, Z. (2024). The (mis)allocation channel of climate change: Evidence from global firm-level microdata. Working Paper.
- Mehrotra, N. (2024). The macroeconomics of net zero. *Working Paper*.
- Midrigan, V. and Xu, D. Y. (2014). Finance and misallocation: Evidence from plant-level data. *American Economic Review*, 104(2):422–458.

BIBLIOGRAPHY

- Moll, B. (2014). Productivity losses from financial frictions: Can self-financing undo capital misallocation? *American Economic Review*, 104(10):3186–3221.
- Newell, R. G., Pizer, W. A., and Raimi, D. (2019). US federal government subsidies for clean energy: Design choices and implications. *Energy Economics*, 80:831–841.
- Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426.
- Nordhaus, W. D. (1993). Optimal greenhouse-gas reductions and tax policy in the DICE model. *American Economic Review*, 83(2):313–317.
- Nordhaus, W. D. and Boyer, J. (2000). *Warming the world: Economic models of global warming*. MIT Press.
- Olovsson, C. and Vestin, D. (2026). Greenflation? *European Economic Review*, 184:105237.
- Papageorgiou, C., Saam, M., and Schulte, P. (2017). Substitution between clean and dirty energy inputs: A macroeconomic perspective. *Review of Economics and Statistics*, 99(2):281–290.
- Papoutsis, M., Piazzesi, M., and Schneider, M. (2022). How unconventional is green monetary policy? Working Paper.
- Parry, I. W. H., Williams III, R. C., and Goulder, L. H. (1999). When can carbon abatement policies increase welfare? The fundamental role of distorted factor markets. *Journal of Environmental Economics and Management*, 37(1):52–84.
- Peters, G. P., Andrew, R. M., Canadell, J. G., Fuss, S., Jackson, R. B., Korsbakken, J. I., Le Quéré, C., and Nakicenovic, N. (2017). Key indicators to track current progress and future ambition of the Paris Agreement. *Nature Climate Change*, 7(2):118–122.
- Pindyck, R. S. (1979). Interfuel substitution and the industrial demand for energy: An international comparison. *Review of Economics and Statistics*, 61(2):169–179.
- Pindyck, R. S. and Rotemberg, J. J. (1983). Dynamic factor demands and the effects of energy price shocks. *American Economic Review*, 73(5):1066–1079.
- Qi, J., Tang, X., and Xi, X. (2021). The size distribution of firms and industrial water pollution: A quantitative analysis of China. *American Economic Journal: Macroeconomics*, 13(1):151–183.
- Raupach, M. R., Marland, G., Ciais, P., Le Quéré, C., Canadell, J. G., Klepper, G., and Field, C. B. (2007). Global and regional drivers of accelerating CO₂ emissions. *Proceedings of the National Academy of Sciences*, 104(24):10288–10293.

BIBLIOGRAPHY

- Rennert, K., Errickson, F., Prest, B. C., Rennels, L., Newell, R. G., Pizer, W., Kingdon, C., Wingenroth, J., Cooke, R., Parthum, B., Smith, D., Cromar, K., Diaz, D., Moore, F. C., Müller, U. K., Plevin, R. J., Raftery, A. E., Ševčíková, H., Sheets, H., Stock, J. H., Tan, T., Watson, M., Wong, T. E., and Anthoff, D. (2022). Comprehensive evidence implies a higher social cost of CO₂. *Nature*, 610(7933):687–692.
- Rennert, K., Prest, B. C., Pizer, W. A., Newell, R. G., Anthoff, D., Kingdon, C., Rennels, L., Cooke, R., Raftery, A. E., Ševčíková, H., and Errickson, F. (2021). The social cost of carbon: Advances in long-term probabilistic projections of population, GDP, emissions, and discount rates. *Brookings Papers on Economic Activity*, (2):223–275.
- Sogalla, R., Wanner, J., and Watabe, Y. (2024). New trade models, same old emissions? KIEL Working Paper No. 2267.
- Traeger, C. (2023). Ace—analytic climate economy. *American Economic Journal: Economic Policy*, 15(3):372–406.
- U.S. Bureau of Economic Analysis (2025a). Table 1.1.5. Gross Domestic Product. BEA Interactive Data (iTable), National Income and Product Accounts (NIPA). Available at <https://www.bea.gov/itable/national-gdp-and-personal-income> (accessed November 28, 2025).
- U.S. Bureau of Economic Analysis (2025b). Tables 1.1.5. Gross Domestic Product. BEA Interactive Data (iTable), National Income and Product Accounts (NIPA). Available at <https://www.bea.gov/itable/national-gdp-and-personal-income> (accessed February 27, 2025).
- U.S. Energy Information Administration (2024). Annual Coal Report 2023.
- U.S. Energy Information Administration (2025a). Monthly Energy Review (January 2025). Available at <https://www.eia.gov/totalenergy/data/monthly/> (accessed February, 2025).
- U.S. Energy Information Administration (2025b). State energy data system (SEDS): 1960–2023 (complete). Released June 27, 2025. Available at <https://www.eia.gov/state/seds/seds-data-complete.php> (accessed November 7, 2025).
- Uzawa, H. (1961). Neutral inventions and the stability of growth equilibrium. *The Review of Economics and Statistics*, 28(2):117–124.
- Van der Werf, E. (2008). Production functions for climate policy modeling: An empirical analysis. *Energy Economics*, 30(6):2964–2979.
- Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters*, 104:112–114.