2024

Essays on the guidance of the direction of innovation

BOSTON UNIVERSITY

GRADUATE SCHOOL OF ARTS AND SCIENCES

Dissertation

ESSAYS ON THE GUIDANCE OF THE DIRECTION OF INNOVATION

by

ERIC ROBERT PAUL DONALD

B.S., Texas Christian University, 2017
M.A., Boston University, 2018

Submitted in partial fulfillment of the

requirements for the degree of

Doctor of Philosophy

2024

Approved by

First Reader

Pascual Restrepo, Ph.D.
Associate Professor of Economics

Second Reader

David Lagakos, Ph.D.
Associate Professor of Economics

Third Reader

Yuhei Miyauchi, Ph.D.
Assistant Professor of Economics

ACKNOWLEDGEMENTS

Getting a Ph.D. has been one of the great challenges of my life, and I have many people to thank for making it a rewarding and enjoyable experience.

I have been extremely lucky to have a great set of advisors: Pascual Restrepo, David Lagakos, and Yuhei Miyauchi. My experience as a graduate student and abilities as an economist were both greatly enhanced by their guidance and support. One often hears horror stories from graduate students about their bad advisors, and whenever I heard such a story, I could only nod along and pretend to understand the experience.

Pascual played a central role in my entire Boston University experience. We met in his macro course when I was a Master's student, and he has been an invaluable mentor to me ever since. Throughout my doctoral studies, our regular meetings were a constant source of encouragement and intellectual stimulation that often became an informal discussion of some topic of shared interest. I don't think we had a single meeting that stayed within its allotted time. I will be very proud if, in the future, people are unsurprised to hear that I was one of "Pascual's students".

David came to Boston University when I was a third-year student, and our first interaction was during the presentation of my second year paper. It was immediately clear that David had a special talent for framing a research project. His constant insistence that I address a specific question and avoid meandering exploration is advice that I will keep in mind for the rest of my career, and my papers will be much better as a result.

Yuhei has been a great mentor. In particular, he played an outsized role in my development by bringing me on as a co-author. He could have kept my role as simply an RA, but by entrusting me to contribute to his research agenda, he allowed me to develop many important skills. Learning how to collaborate with co-authors is not easily learned in the classroom, so I am very grateful for the opportunity to co-author work with such a talented economist.

I also benefited a great deal from my interactions with other Boston University faculty members, including Masao Fukui, Stephen Terry, Adam Guren, Bob King, Martin Fiszbein, and Bob Margo. Each one of them played a valuable role in my development as an economist.

I also want to thank the faculty members at Texas Christian University who encouraged me when I told them I wanted to pursue a Ph.D. in economics, including W. Charles Sawyer, Zack Hawley, and José Carrión. Unfortunately, W. Charles Sawyer passed away during my time at Boston University, but I believe he would have been proud to see one of his students become a successful research economist.

Finally, I cannot overstate my gratitude for the support of my beautiful wife, Bianca Donald, throughout this process. She kept me fed during the first year while I had to work on the weekend to complete problem sets. She encouraged me as I started in research and worried I would not find a good topic. And she was the best COVID quarantine partner I could have asked for. My success is due in large part to her, and I will always be grateful to have her as my rock.

Eric Robert Paul Donald

ESSAYS ON THE GUIDANCE OF THE DIRECTION OF INNOVATION

ERIC ROBERT PAUL DONALD

Boston University, Graduate School of Arts and Sciences, 2024

Major Professor: Pascual Restrepo, Ph.D.
Associate Professor of Economics

ABSTRACT

My dissertation studies how policy should guide, and respond to, technological change. The first chapter addresses the "guide" side of the question, while chapters two and three address the "respond to" side of the question. In the first chapter, I argue that cross-technology knowledge spillovers are critical for understanding policy's role in the transition to clean technology. I develop an endogenous growth model with clean and dirty technologies and a network of cross-technology spillovers. I derive formulas for the size and speed of technological transition, following a policy reform, which show that greater spillovers across technologies induce a faster transition but at the expense of a smaller long-run impact of policy. Such spillovers also prevent the lock-in of dirty technology. The economy's spillover structure can be summarized by a sufficient statistic matrix, which I estimate using patent citation data. Applying my model to US transportation and electricity generation, I find that cross-technology spillovers are mid-sized: they prevent lock-in but imply a slow transition with a high long-run impact of policy. I conclude by examining how cross-technology spillovers affect optimal clean innovation subsidies, deriving an innovation subsidy formula that holds for arbitrary carbon prices. Quantitatively, I find that optimal clean innovation

subsidies are small, reflecting the low centrality of clean technologies in the spillover network. Hence, a "big push" of temporary clean innovation subsidies is not warranted. In the second chapter, I start with the observation that the standard reaction to the problem of automation, by both lay people and the economics literature, follows a Pigouvian intuition: robots harm workers, so they should be taxed. I argue that this Pigouvian intuition is misguided, or at least oversimplified. As shown by the recent literature modeling automation within the task framework, capital only exerts a negative pecuniary externality on labor at the extensive margin of automation. At the intensive margin, more capital producing a task that has already been automated raises wages for everyone via capital deepening. To formalize this point, I present a model with heterogeneous agents where the Planner can tax income from capital and labor as well as target the extensive margin of automation by stipulating how much more expensive labor must be than capital before automation can occur. I show, via an envelope argument, that capital taxation should ignore automation when the extensive margin tool is set optimally. In a quantitative application to the US economy, I find that labor should be 3.4% more expensive than capital before automation can occur. In the third chapter, Masao Fukui, Yuhei Miyauchi, and I study optimal transfer policy in dynamic spatial equilibrium models with frictional migration and incomplete financial markets. A key policy trade-off is to provide consumption insurance while minimizing the distortion of migration flows. We derive a recursive formula for optimal spatial transfers that strikes this balance. We calibrate our model to U.S. states and find that the U.S. economy would benefit from increased transfers to low-income-growth states. Welfare gains from optimal transfers are substantial but smaller than in a framework abstracting from slow migration adjustment.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS

ACS	....	American Community Survey
BEA	....	Bureau of Economic Analysis
Btu	....	British Thermal Units
CES	....	Constant Elasticity of Substitution
CFS	....	Commodity Flow Survey
CO_2	....	Carbon Dioxide
CPC	....	Cooperative Patent System
CRRA	....	Constant Relative Risk Aversion
CRTS	....	Constant Returns to Scale
EIA	....	Energy Information Agency
EPA	....	Environmental Protection Agency
EV	....	Electric Vehicle
FE	....	Fixed Effects
FOC	....	First Order Condition
GDP	....	Gross Domestic Product
GHG	....	Greenhouse Gas
GtC	....	Gigatons of Carbon
IEA	....	International Energy Agency
iid	....	Independent and Identically Distributed
IPC	....	International Patent System
IPCC	....	Intergovernmental Panel on Climate Change
LHS	....	Left Hand Side
MRS	....	Marginal Rate of Substitution
NOAA	....	National Oceanic and Atmospheric Administration
OLG	....	Overlapping Generations
ppm	....	Parts Per Million
RHS	....	Right Hand Side
RICE	....	Regional Integrated Climate-Economy
RPP	....	Regional Price Parities
SCC	....	Social Cost of Carbon
SCF	....	Survey of Consumer Finance
TWh	....	Terawatt Hour
UBI	....	Universal Basic Income

CHAPTER ONE - SPILLOVERS AND THE DIRECTION OF INNOVATION: AN APPLICATION TO THE CLEAN ENERGY TRANSITION

1.1 Introduction

Decarbonizing the economy will require significant investments in clean innovation, so policymakers have set out to redirect the path of innovation toward clean technology. Indeed, the European Commission states that one of the main objectives of its emissions trading scheme is to "promote investment in innovative, low-carbon technologies". The United States, where carbon pricing is absent, is on track to spend roughly $360 billion on subsidies for clean technology via the Inflation Reduction Act, much of which is explicitly geared toward encouraging clean innovation.

This paper argues that *cross-technology knowledge spillovers* are crucial for understanding the role of policy in transitioning the economy to cleaner production technologies. Such spillovers are typically excluded from existing models, but they underlied the invention of many clean technologies that are now at the vanguard of the energy transition. For instance, the first Tesla prototype – the Mule 1 – was a combustion engine car that the engineers at Tesla reconfigured by ripping out the engine and stuffing the engine compartment full of batteries. Similarly, when researchers at Bell Labs invented the modern solar cell, they made use of conductive properties of silicon already known from previous research on semiconductors. The inventors of these clean technologies did not have to start from scratch. Instead, they could build on existing knowledge from other technologies.[1]

[1] Figures A·1 and A·2 provide illustrations of these two examples of cross-technology knowledge spillovers.

To study how cross-technology knowledge spillovers shape climate policy, I develop a general endogenous growth model with clean and dirty technologies whose key feature is a network of cross-technology knowledge spillovers. For policy instruments, I consider a price on carbon and technology-specific innovation subsidies. I show that cross-technology spillovers matter for both describing the impact of a policy reform and prescribing optimal clean innovation policy. Greater cross-technology spillovers allow for a more rapid transition to clean technology because they allow less advanced technologies to achieve catchup growth. However, by generating catchup growth, cross-technology spillovers also reduce the long-run impact of policy as they prevent any one technology from gaining too large a lead. Thus, depending on the level of cross-technology spillovers, a policy reform that favors clean technologies may generate rapid change in the short run or substantial redirection of innovation in the long run, but not both.

My main result is to provide formulas for the size and speed of technology's transition, following a policy reform, that depend on two matrices: one summarizing the network of knowledge spillovers and another summarizing substitution patterns in production. These formulas show that cross-technology spillovers provide a countervailing force to the substitutability of clean and dirty goods: the mechanism studied by Acemoglu et al. (2012). Indeed, higher substitutability of clean and dirty goods has the opposite effect of cross-technology spillovers – slowing down the transition and increasing the long-run impact of policy – by increasing the market size of more advanced technologies. My formulas summarize each of these forces in terms of sufficient statistic matrices, providing a clear mapping from the model to the data for an arbitrary number of technologies and a general class of production and spillover structures.

These formulas allow me to aggregate the two forces – spillovers and substitutability – into a single measure of *increasing returns to innovation*, which quantifies the degree to which more advanced technologies offer greater incentives for innovation. Formally, this measure is in terms of the spectral radius of a matrix that describes technology's transition path. As this measure increases, the transition slows down while the long-run impact of policy grows, and if this measure exceeds one, technology becomes path dependent, locking in an inefficient dirty equilibrium. Thus, by reducing increasing returns to innovation, cross-technology spillovers also help to prevent technological lock-in, whereas the reverse is true for the substitutability between clean and dirty goods.

Next, I examine how cross-technology spillovers shape optimal policy, in particular optimal clean innovation subsidies. I first show that in the first-best case where carbon prices are unrestricted, policy should separately correct the pollution externality and knowledge spillover externality, in keeping with the Pigou Principle. Carbon prices should reflect the social cost of carbon, rather than explicitly influencing the direction of innovation. Instead, with the pollution externality corrected, innovators need to internalize the knowledge spillovers they create. To this end, I derive a recursive formula for optimal innovation subsidies that reflects the social value of knowledge spillovers sent through the spillover network. These innovation subsidies consider both the value of technology in producing goods, as summarized by Domar weights, and the value of technology in producing innovations, as summarized by centrality in the spillover network, with the weight between the two determined by the Planner's level of patience. Just as carbon prices do not explicitly target innovation, innovation subsidies do not explicitly target the pollution associated with each technology.

Given the political difficulty of carbon pricing, I also consider a second-best extension where carbon prices are set to some external, potentially suboptimal, level. I show that the same recursive formula for innovation subsidies holds as before but with a simple adjustment that adds the social cost of carbon, net of the external carbon price, multiplied by the impact of innovation on equilibrium emissions. As a result, clean technologies that reduce equilibrium emissions are further subsidized, whereas dirty technologies that increase equilibrium emissions are punished. In this way, the Planner modifies innovation subsidies to prevent "immiserizing growth", where innovation reduces welfare in inefficient economies by exacerbating distortions (Bhagwati, 1958).

I then consider a quantitative application of my model to the transportation and electricity generation sectors in the United States. These are two sectors where production can be clearly delineated into "clean" and "dirty", and they together account for a little more than half of US emissions. Using the citation network of granted US patents as a proxy for the network of knowledge spillovers, I find that cross-technology spillovers are large enough to prevent lock-in of dirty technology for both sectors, but that the substitutability of clean and dirty goods still creates slow transitions with large long-run impacts for policy reforms. That is, my measure of increasing returns to innovation is near one from below. I validate my model by examining the progress made in clean transportation and electricity generation during the 2010s, showing that my model can match this episode as an untargeted moment. However, without cross-technology spillovers, my model makes the counterfactual prediction that dirty technologies would have become *more* entrenched in both sectors, not less.

I apply my model to the data with several quantitative exercises. First, I examine the impact of introducing a $51 carbon price (the Biden Administration's current estimate of the social cost of carbon) and a uniform clean innovation subsidy equivalent to a 30% tax credit (consistent with the subsidies in the Inflation Reduction Act). I find a large long-run effect of this policy reform, with the steady-state prevalence of clean technology increasing by 117.3% and 120.8% for transportation and electricity generation, respectively. However, the transition to this new steady-state is slow, with half-lives of convergence of 123 years and 128 years, respectively. In the absence of cross-technology spillovers, the economy would start out locked in a dirty equilibrium, and the policy reform would only be enough to induce a slow switch to clean innovation in electricity generation.

Second, I simulate the optimal path of innovation subsidies for a variety of potential restrictions on the carbon price. Conventional wisdom argues for a big push of large, temporary clean innovation subsidies to jolt the economy away from dirty technology (Acemoglu et al., 2012, 2016). Indeed, in the absence of cross-technology spillovers, this is what my model recommends: clean innovation subsidies for transport and electricity generation that start out at 369.4% and 238.8% of the baseline innovation wedge, respectively, and reduce down to baseline over the first century of policy. This prescription is motivated by the problem of technological lock-in, but according to my calibration, this is not the empirically relevant case for the US economy. Instead, I find that low, relatively stable subsidies for clean technology are optimal, and this holds quantitatively even when the carbon price is set below the social cost of carbon.

In the first-best, when carbon pollution is priced at its social marginal cost, clean transportation and electricity generation receive average innovation subsidies of only 39.8% and 59.9% of the baseline innovation wedge, respectively, over the first century of policy. This is because clean technologies are beneficial for their ability to produce goods without pollution, not necessarily for their ability to produce knowledge spillovers, and the patent data indicates that clean technologies do not have high centrality in the spillover network. The optimality of low clean innovation subsidies is also robust to the choice of social discount rate. In fact, a more patient Planner implements higher carbon prices and *lower* clean innovation subsidies because their greater concern for the future increases the weight they place on a technology's centrality in the spillover network.

Even in the second-best, where there is an additional justification for clean innovation subsidies, a big push is not warranted with an empirically disciplined spillover network. I consider two cases for restrictions on the carbon price: one where the carbon price is ten percent of the social cost of carbon and another where it is zero. In the case with a small carbon price, clean innovation subsidies shift upward by about 20-30% of the baseline innovation wedge but otherwise retain the same pattern as in the first-best. When a price on carbon is lost altogether, I find that clean innovation subsidies should start with a similar upward shift as in the small carbon price case and then rise rapidly after the first century of policy; an ongoing nudge rather than a big push.

Finally, my model allows me to quantify the welfare losses stemming from restrictions on the carbon price. A small, growing carbon price eventually pushes dirty technology out of production, allowing some dirty innovation to continue for the sake of spillover production. As a result, the welfare loss relative to the first-best is small. However, if pollution is unpriced, then decarbonization requires shutting down dirty innovation and relying exclusively on clean innovation. This leads to slow emission reductions, due to a rebound effect of clean innovation, and slow economic growth, due to the loss of spillovers from dirty technology, which together lead to far greater welfare losses. Hence, the welfare cost of restrictions on the carbon price is highly nonlinear.

1.1.1 Related Literature

This paper contributes to several strands of literature. First is the extensive literature on endogenous innovation and the climate. Typical macro models of climate innovation exclude the possibility of cross-technology knowledge spillovers, assuming that spillovers take place exclusively within technologies (Gans, 2012; Acemoglu et al., 2012, 2016, 2023; Hassler et al., 2021). As argued in the seminal work of Acemoglu et al. (2012), the substitutability of clean and dirty goods alone generates destabilizing increasing returns to innovation that must be overcome with a big push of clean innovation subsidies. Fried (2018) allows for cross-technology spillovers in a quantitative framework but does not analytically characterize their relevance for describing the impact of a policy reform or prescribing optimal clean innovation policy.

My paper also relates to the topic of technological path dependence. Several papers in this literature argue for the role of cross-technology knowledge spillovers in stabilizing technology's steady-state, including Acemoglu (2002, 2023) and Fried (2018). My formulas generalize these results by providing a test for technological path dependence in terms of a spectral radius, applicable for an arbitrary number of technologies and a broad class of production and spillover structures.

More generally, there is an extensive literature, going back to Hicks (1932), that considers the role of input prices in determining the direction of innovation. Examples of empirical research looking at the influence of dirty input prices on clean innovation include Newell et al. (1999), Popp (2002), Aghion et al. (2016, 2023), and Känzig (2023).[2] Examples of theoretical contributions from the skill-biased technological change literature include Acemoglu (1998, 2002) and Acemoglu and Restrepo (2018, 2022a). This literature has long argued that the degree of substitutability across goods is a fundamental determinant of the impact of input prices on the direction of innovation. My paper argues that cross-technology knowledge spillovers deserve a similar status, while also providing guidance for the measurement and aggregation of both forces.

My paper also relates to the vast literature on integrated assessment models that,

[2]For empirical work on the direction of innovation in a variety of other settings, see Acemoglu and Linn (2004), Acemoglu and Finkelstein (2008), Hanlon (2015), Budish et al. (2015), and Moscona and Sastry (2022).

starting with the seminal contribution of Nordhaus (1992), seeks to quantify the social harm of carbon pollution. Examples from this literature include Stern (2007), Golosov et al. (2014), Nordhaus (2017), Krusell and Smith (2022), and Barrage and Nordhaus (2024).[3] My paper provides a rich model of the response of clean innovation to climate policy and shows that innovation subsidies should include the social cost of carbon whenever pollution prices are incomplete. In doing so, I provide policymakers in the innovation space a way to make use of the information provided by integrated assessment models.

By focusing on innovation policy in the context of a spillover network, my work is closely related to the recent work of Liu and Ma (2021).[4] Specific to my context is the capacity for some technologies to enable externalities in production, and I show that innovation subsidies should treat improperly priced externalities in the same manner as a Domar weight, reflecting the contemporaneous impact of innovation on welfare. Finally, the methodological tools I use to characterize the transition path of technology following a policy reform are similar to those used by Kleinman et al. (2023). Their paper considers the transition dynamics of the distribution of capital and labor across space following a shock to regional fundamentals. I take a complementary approach to describe transition dynamics in the context of directed technological change.

The remainder of the paper is organized as follows. In Section 1.2, I describe the endogenous growth model I use throughout the paper. In Section 1.3, I consider a once-and-for-all policy reform, characterize the change in technology's steady-state,

[3]For integrated assessment models with endogenous innovation, see Nordhaus (2002) and Popp (2004).

[4]In general, knowledge spillovers have long been a central consideration of innovation policy (Arrow, 1962; Romer, 1990; Bloom et al., 2019; Akcigit et al., 2020; Bryan and Williams, 2021). Closely related is a body of empirical work focused on estimating the spillover benefits of innovation (Jones and Williams, 1998; Bloom et al., 2013; Jones and Summers, 2021; Myers and Lanahan, 2022).

and derive technology's transition path to its new steady-state. In Section 1.4, I derive optimal innovation subsidies and carbon prices in both the first- and second-best case. In Section 1.5, I describe the calibration of the spillover network and other structural parameters of my model. In Section 1.6, I perform quantitative exercises that simulate the impact of a policy reform as well as optimal climate policy in the first- and second-best. Section 1.7 concludes.

1.2 Model

This section lays out the endogenous growth model used throughout the paper. Time is discrete and indexed by t. There are J-many technologies in the economy, each of which is classified as either clean and dirty. The innovation process for each of these technologies is endogenously determined by the incentives to produce a step-ladder innovation which entitles its owner to a single period of monopoly rents.

The economic environment has two externalities, which are the focus of policy. First, inputs specific to dirty technologies, e.g. fossil fuels, generate carbon emissions, and those emissions damage future productive capacity. Second, innovation creates knowledge spillovers via a spillover network, so innovators do not consider the benefit they bestow on the future production of knowledge. To correct these externalities, I will consider a Planner with access to (i) a carbon price and (ii) an array of technology-specific innovation subsidies.

I will start by describing the production and innovation technology of this economy, followed by a description of how the direction of innovation is determined in competitive equilibrium. Throughout the paper, I will use bold notation to denote matrices. I will use i or j to index technologies.

1.2.1 Production

Each technology in this economy produces a distinct good, and these technology-specific goods aggregate into final output. Final output is produced according to

$$\mathcal{Y}_t = \Omega_t F\big(\{Y_{jt}\}\big), \tag{1.1}$$

where Ω_t is the climate damage function, and $F(.)$ is a constant returns aggregator of technology-specific goods. Climate damages Ω_t are a function of the past sequence of emissions $\{\mathcal{E}_{\hat{t}}\}_{\hat{t}\leq t}$. An important assumption embedded in this specification is that climate damages are Hicks-neutral, so they do not influence relative marginal products.

The good specific to technology j is produced according to

$$\ln(Y_{jt}) = \alpha \ln(\Lambda_{jt}) + (1-\alpha)\int_0^1 \ln(y_{j\iota t})d\iota, \tag{1.2}$$

where Λ_{jt} is an input and $y_{j\iota t}$ is an intermediate, both for technology j. That is, technology-specific goods are produced via a Cobb-Douglas with share parameter α that combines an input and a unit interval of intermediates.

Using one unit of input Λ_{jt} produces ω_j units of carbon emissions, so I will call technologies with $\omega_j = 0$ "clean" and those with $\omega_j > 0$ "dirty". Thus, total carbon emissions follow

$$\mathcal{E}_t = \sum_j \omega_j \Lambda_{jt}. \tag{1.3}$$

Dirty inputs can be thought of as fossil fuels whose use generates carbon emissions. Carbon intensities ω_j may differ across dirty technologies due to the different types of fossil fuels they use. For instance, the main fuels used to produce electricity – natural gas and coal – have different carbon intensities than that used in transportation:

gasoline.[5] Denote the overall emissions intensity of the economy by $\bar{\omega}_t \equiv \mathcal{E}_t/\mathcal{Y}_t$.

Inputs Λ_{jt} are produced using r_j units of the final good, so the economy features "round-about" production. For concreteness, one can imagine that dirty input costs correspond to the resources required to extract fossil fuels, but more generally, r_j should be thought of as representing the intrinsic ability of technology j to convert resources into a useful output. This is relevant in the case of the clean energy transition because fossil fuel-based machinery is relatively inefficient at converting primary energy into useful energy (BP, 2022).

Intermediate production is linear in labor, so

$$y_{j\iota t} = a_{j\iota t}\ell_{j\iota t}, \tag{1.4}$$

where $a_{j\iota t}$ is the labor productivity of intermediate ι for technology j. I define the technology stock of j as the geometric average of its labor productivities:

$$A_{jt} \equiv \exp\Big(\int_0^1 \ln(a_{j\iota t})d\iota\Big). \tag{1.5}$$

The economy has a fixed endowment of labor that is supplied inelastically, so the allocation of labor must satisfy

$$\sum_j \int_0^1 \ell_{j\iota t} d\iota \leq L, \tag{1.6}$$

where L is the aggregate supply of labor.[6]

[5] According to the EIA, the carbon intensity of natural gas, coal, and gasoline is 52.9, 96.1, and 70.7 kilograms of CO_2 per million Btu, respectively.

[6] I abstract from capital accumulation, which is a standard simplifying assumption in the directed technological change literature (Acemoglu, 2002; Acemoglu et al., 2012, 2016; Fried, 2018).

Finally, the resource constraint requires that

$$\mathcal{Y}_t = c_t + \sum_j r_j \Lambda_{jt}, \tag{1.7}$$

where c_t is household consumption.

1.2.2 Innovation

In this section, I describe the innovation production function. As is typical in the endogenous growth literature, new ideas are produced by combining scientists and old ideas. Because there are multiple technologies, I specify a spillover function that defines the network of cross-technology knowledge spillovers.

Innovation follows a step-ladder process à la Grossman and Helpman (1991) and Aghion and Howitt (1992), so if intermediate ι for technology j receives an innovation, we have that $a_{j\iota t}$ increases to $\gamma a_{j\iota t}$, where $\gamma > 1$ is the step size of innovation. Define z_{jt} as the mass of intermediates for technology j that receives an innovation. Plugging this into Equation (1.5), we have that technology evolves according to

$$A_{jt} = \gamma^{z_{jt}} A_{jt-1}. \tag{1.8}$$

This formulation implies that technology-specific growth rates follow

$$g_{jt} = \ln(\gamma) z_{jt}. \tag{1.9}$$

The production of innovation follows

$$z_{jt} = \chi_j s_{jt}^{\eta} \phi_j\big(\{A_{it-1}\}\big), \tag{1.10}$$

where s_{jt} is the number of scientists devoted to technology j, and $\phi_j(.)$ is a spillover function that governs how research productivity is affected by the state of technology in the economy. I assume that spillover functions are homogeneous of degree zero, so spillovers remain constant if all technologies scale together. The parameter $\eta \in (0,1)$ determines the degree of diminishing returns in research, and χ_j is a general research productivity term specific to each technology.

To describe the network of knowledge spillovers in this economy, consider the matrix of spillover elasticities.

Definition 1.1 (Spillover Network $\boldsymbol{\varphi}$)**.** *The spillover network is a $J \times J$ matrix with elements*

$$\varphi_{ijt} \equiv \frac{\partial \ln(\phi_{it})}{\partial \ln(A_{jt-1})}. \tag{1.11}$$

The spillover network $\boldsymbol{\varphi}_t$ describes the response of spillovers to changes in the knowledge stock of each technology. I will assume that $\boldsymbol{\varphi}_t$ is weakly positive off of its diagonal: $\varphi_{ijt}|_{i\neq j} \geq 0$. This guarantees that technologies receive higher spillovers when they are relatively less advanced.[7] By assuming that spillover functions are homogeneous of degree zero, i.e. $\sum_j \varphi_{ijt} = 0$, I ensure balanced growth in the presence of a fixed supply of scientists.[8] It will occasionally be helpful to refer to the *gross spillover network* $\tilde{\boldsymbol{\varphi}}_t$, which is defined in relation to the (net) spillover network according to $\boldsymbol{\varphi}_t = \tilde{\boldsymbol{\varphi}}_t - \mathbf{I}$. Thus, the rows of the gross spillover network describe an array of spillovers that sum to one, rather than zero.

The spillover network is a central factor in determining both the marginal cost and marginal benefit of innovation. The rows of the spillover network determine the marginal cost of innovation as they specify how the research productivity of each technology depends the overall state of technology in the economy. This is partic-

[7]To see this, note that because ϕ_{it} is homogeneous of degree zero, $\varphi_{ijt}|_{i\neq j} \geq 0$ implies $\frac{\partial \phi_i(\{A_{qt-1}/A_{it-1}\})}{\partial (A_{jt-1}/A_{it-1})}|_{i\neq j} \geq 0$. That is, technology i's spillovers increase as the other technologies become relatively more advanced. This implies $\varphi_{iit} \leq 0$.

[8]Allowing for $\sum_j \varphi_{ijt} < 0$ would be consistent with the fishing out feature of semi-endogenous growth models (Jones, 1995, 2022). In that case, balanced growth would require an exponential increase in the supply of scientists. The possibility that $\varphi_{iit} < 0$ introduces some notion of fishing out because it implies that additional innovations are more difficult for more advanced technologies (Bloom et al., 2020). However, if all technologies advance concurrently, then reductions in research productivity can be avoided, and balanced growth can still be achieved with a fixed supply of scientists.

ularly important for transition dynamics as less advanced technologies can receive a catchup-inducing boost to their research productivity. However, the rows of the spillover network are *not* relevant for policy because researchers fully internalize the productivity benefits they receive from spillovers.

On the other hand, the columns of the spillover matrix determine the marginal benefit of innovation as they specify how the efforts of future researchers will be enhanced by research that takes place today. Therefore, the columns of the spillover matrix are the focus of policy because the positive spillovers they mediate are not internalized by today's researchers. However, the columns of the spillover network are *not* relevant for laissez-faire transition dynamics because the external nature of spillovers implies they do not influence equilibrium behavior.

As an example, an economy with no cross-technology knowledge spillovers would have a spillover network of all zeros: $\boldsymbol{\varphi}_t = \mathbf{0}$. In that case, spillovers ϕ_{it} would reduce to a positive constant for each technology, so the spillover network would cease to play any role. Note that there would still be spillovers *within* technologies as Equation (1.8) states that, for a given technology, innovation builds multiplicatively on the existing technology stock. The absence of cross-technology knowledge spillovers is the implicit assumption in much of the climate innovation literature (e.g. Acemoglu et al. (2012)), but as argued throughout this paper, cross-technology knowledge spillovers play a critical role in shaping both transition dynamics and optimal policy.

Finally, there is a fixed endowment of scientists supplied inelastically, so the allocation of scientists must satisfy

$$\sum_j s_{jt} \leq \mathcal{S}, \tag{1.12}$$

where $\mathcal{S}$ is the aggregate supply of scientists. The fact that the aggregate supply of scientists is fixed implies that the Planner's concern is with the *composition* of scientific effort, not the total level. The starting point of most innovation policies is that innovators do not internalize the full social marginal benefit of their efforts (Arrow, 1962; Bryan and Williams, 2021). In a setting where scientists have an outside option, e.g. production work or leisure, equilibrium scientific effort is likely to be below the optimum. This has been the traditional justification of innovation subsidies (Romer, 1990; Bloom et al., 2019). Instead, when the supply of scientists is fixed and there are multiple technologies, the fact that differences in spillover creation across technologies may not track private profit incentives implies that the equilibrium composition of scientific effort may not be optimal.

1.2.3 Preferences & Policy Instruments

The Planner seeks to maximize the utility of a representative household with two main policy instruments: a carbon price, denoted by τ_t, and technology-specific innovation subsidies, denoted by ξ_{jt}. With these instruments (and two others discussed below), the Planner can implement the first-best as a competitive equilibrium by correcting the economy's externalities: carbon pollution and knowledge spillovers.

The social objective is household utility

$$\sum_{t\geq 0} \frac{1}{(1+\rho)^t} u(c_t), \tag{1.13}$$

where $u(.)$ is a concave function of consumption, and ρ is the rate of pure time preference. I will refer to the rate of pure time preference ρ as the (social) discount rate.[9]

As discussed below, producers of intermediates have market power, so I will allow

[9] A distinction can be drawn between the household's (private) discount rate and the Planner's (social) discount rate, but this difference only matters in contexts where the household is making an intertemporal decision, e.g. savings, that the Planner can influence with policy (e.g. Farhi and Werning (2007)). Because the choice to innovate depends on a one-period monopoly and I have abstracted from capital accumulation, none of the equilibrium outcomes of this economy depend on the household's private discount rate. For this reason, I can ignore the distinction between private and social discount rates and set ρ in accordance with the social discount rate.

the Planner to subsidize intermediates to correct the monopolist markup. This allows the Planner to correct every market failure and achieve the first-best, but given that market power is not the focus of this paper, this policy will play little role in the analysis. For the sake of simplicity, I will assume that all subsidies on intermediates are constant, which would be the case in both laissez-faire and the first-best. Denote by Υ the subsidy on intermediates. Finally, the Planner may levy a lump-sum tax D_t to fund(rebate) any deficit(surplus) from the use of corrective instruments.

1.2.4 Equilibrium

I make the typical assumption in endogenous growth models that innovation for a given intermediate grants exclusive ownership of the most advanced form of production for that intermediate. This exclusive ownership creates a rent which drives the incentive to innovate (Romer, 1990; Aghion and Howitt, 1992).

The final good is the numeraire in each period. The producers of intermediates are the only firms in the economy with market power, so prices of all other goods are set competitively. The final good producer solves

$$\max_{\{Y_{jt}\}} \mathcal{Y}_t - \sum_j p_{jt} Y_{jt}, \tag{1.14}$$

where p_{jt} is the price of the technology j good. This yields the condition

$$\Omega_t \frac{\partial F_t}{\partial Y_{jt}} = p_{jt}. \tag{1.15}$$

The producers of technology-specific goods solve

$$\max_{\Lambda_{jt}, \{y_{j\iota t}\}} p_{jt} Y_{jt} - (r_j + \omega_j \tau_t)\Lambda_{jt} - \int_0^1 p_{j\iota t} y_{j\iota t} d\iota \tag{1.16}$$

where $p_{j\iota t}$ is the price of intermediate ι for technology j. This yields the conditions

$$\alpha p_{jt}\frac{Y_{jt}}{\Lambda_{jt}} = r_j + \omega_j \tau_t \tag{1.17}$$

$$(1-\alpha)p_{jt}\frac{Y_{jt}}{y_{j\iota t}} = p_{j\iota t}. \tag{1.18}$$

One can see from Equation (1.17) that carbon pricing reduces demand for dirty inputs by making those dirty inputs more expensive. This, in turn, reduces the demand for all goods and intermediates associated with dirty production as well.

Given their market power, an intermediate producer internalizes their demand curve from Equation (1.18), so they solve

$$\max_{p_{j\iota t}, y_{j\iota t}} \Upsilon p_{j\iota t} y_{j\iota t} - \frac{w_{\ell t}}{a_{j\iota t}} y_{j\iota t} \;\; s.t. \tag{1.19}$$
$$(1-\alpha)p_{jt}\frac{Y_{jt}}{y_{j\iota t}} = p_{j\iota t},$$

where $w_{\ell t}$ is the wage paid to production workers. The optimal markup is infinite with a Cobb-Douglas demand system, so to pin down the markup, I will assume intermediate producers limit price one step down the productivity ladder to the competitive fringe. This yields

$$p_{j\iota t} = \frac{\gamma}{\Upsilon}\frac{w_{\ell t}}{a_{j\iota t}}. \tag{1.20}$$

That is, intermediate prices equal the subsidy-inclusive break-even price of the competitive fringe. This allows intermediate producers to achieve profit

$$\pi_{j\iota t} = \Upsilon\frac{\gamma-1}{\gamma}(1-\alpha)p_{jt}Y_{jt}. \tag{1.21}$$

We can see from Equation (1.21) that relative profits depend on relative prices of technology-specific goods, so I show in Appendix A.1.1 that relative prices follow

$$\frac{p_{jt}}{p_{Jt}} = \left(\frac{r_j + \omega_j\tau_t}{r_J + \omega_J\tau_t}\right)^{\alpha}\left(\frac{A_{jt}}{A_{Jt}}\right)^{\alpha-1}. \tag{1.22}$$

That is, relative prices are increasing in relative input costs – inclusive of the carbon price – and decreasing in relative technology.

Turning to the innovation side of the economy, I assume there is a competitive research firm that produces innovations for every technology. The purchase of an innovation in technology j allows its owner to be the productivity leader of a random intermediate for one period, so by no-arbitrage, the equilibrium price for such an innovation is equal to expected profit: $\Pi_{jt} = \int_0^1 \pi_{j\iota t} d\iota$.

The research firm solves

$$\max_{\{s_{jt}\}} \sum_j \xi_{jt} z_{jt} \Pi_{jt} - w_{st} \sum_j s_{jt} \tag{1.23}$$

where w_{st} is the wage paid to scientists. That is, the research firm produces mass z_{jt} of innovations for technology j and sells them at price Π_{jt}. Considering the no-arbitrage condition, we have that the price of innovation for technology j is proportional to market size:

$$\Pi_{jt} = \Upsilon \frac{\gamma - 1}{\gamma} (1 - \alpha) p_{jt} Y_{jt} \propto S_{jt} \mathcal{Y}_t, \tag{1.24}$$

where S_{jt} is the income share of technology j. That is, the rent one can achieve by innovating in technology j scales in accordance with total spending on the technology j good. This is a standard result in the directed technological change literature (Acemoglu, 2002; Acemoglu et al., 2012).

The research optimality condition yields

$$\left(\frac{s_{jt}}{s_{Jt}}\right)^{1-\eta} = \left(\frac{\chi_j}{\chi_J}\right) \left(\frac{\xi_{jt}}{\xi_{Jt}}\right) \left(\frac{\phi_{jt}}{\phi_{Jt}}\right) \left(\frac{\Pi_{jt}}{\Pi_{Jt}}\right), \tag{1.25}$$

which, together with the fixed supply of scientists (1.12), pins down the innovation equilibrium. Thus, the allocation of scientists, and therefore the direction of innovation, is determined by three forces: subsidies, spillovers, and market size. First, innovation subsidies have a natural influence; as policy favors one technology with higher subsidies, more scientists are devoted to that technology. Second, spillovers influence the allocation of scientists by affecting the productivity of research. As discussed in Section 1.2.2, technologies that are relatively less advanced receive a boost in research productivity.

Third, technologies with larger markets, i.e. greater income shares, have larger innovation rents, and therefore, stronger incentives to innovate. This is the channel through which the substitutability of technology-specific goods influences the direction of innovation. If technologies are substitutes(complements), then the less advanced technology will obtain a smaller(larger) innovation rent as it accounts for a smaller(larger) share of final output. Market size is also the channel through which carbon prices influence the direction of innovation. In the case of substitutes(complements), a price on carbon reduces(increases) relative market size for dirty forms of production as it increases dirty input costs.

Finally, the representative household owns all of the factors and firms in the economy and consumes their income in each period. Both the household and government must satisfy their budget constraints:

$$c_t = w_{\ell t} L + w_{st} \mathcal{S} + \Pi_t - D_t \tag{1.26}$$

$$D_t + \tau_t \mathcal{E}_t = \sum_j (\xi_{jt} - 1) z_{jt} \Pi_{jt} + \sum_j \int_0^1 (\Upsilon - 1) p_{j\iota t} y_{j\iota t} d\iota, \tag{1.27}$$

where Π_t is a dividend equal to all of the profits in the economy. Because the household owns all of the factors and firms in the economy, and net government expenditure is financed lump-sum, equilibrium consumption is equal to final output net of input costs.

Given these conditions, the competitive equilibrium is defined as follows:

Definition 1.2 (Equilibrium)**.** *Given an initial condition for technology $\{A_{j,-1}\}$ and an array of input costs $\{r_j\}$, an equilibrium consists of a sequence of carbon prices*

$\{\tau_t\}$, innovation subsidies $\{\xi_{jt}\}$, intermediate subsidies Υ, lump-sum taxes $\{D_t\}$, final output $\{\mathcal{Y}_t\}$, technology-specific goods $\{Y_{jt}\}$, technology-specific good prices $\{p_{jt}\}$, inputs $\{\Lambda_{jt}\}$, intermediates $\{y_{\jmath\iota t}\}$, intermediate prices $\{p_{\jmath\iota t}\}$, labor $\{\ell_{\jmath\iota t}\}$, production wages $\{w_{\ell t}\}$, innovation $\{z_{jt}\}$, innovation rents $\{\Pi_{jt}\}$, scientists $\{s_{jt}\}$, scientist wages $\{w_{st}\}$, technology $\{A_{jt}\}$, emissions $\{\mathcal{E}_t\}$, and household consumption $\{c_t\}$ such that:

(i) *Prices $\{\{p_{jt}, \{p_{\jmath\iota t}\}\}, w_{\ell t}\}$ and quantities $\{\mathcal{Y}_t, \{Y_{jt}, \Lambda_{jt}, \{y_{\jmath\iota t}, \ell_{\jmath\iota t}\}\}\}$ on the production side of the economy solve the profit maximization problems* (1.14), (1.16), *and* (1.19)*;*

(ii) *The labor market clears* (1.6)*;*

(iii) *Prices $\{\{\Pi_{jt}\}, w_{st}\}$ and quantities $\{z_{jt}, s_{jt}\}$ on the innovation side of the economy satisfy the no-arbitrage condition* (1.24) *and solve the profit maximization problem* (1.23)*;*

(iv) *The scientist market clears* (1.12)*;*

(v) *Technology $\{A_{jt}\}$ evolves according to* (1.8)*;*

(vi) *Emissions $\{\mathcal{E}_t\}$ follow* (1.3)*;*

(vii) *The resource constraint* (1.7) *and budget constraints* (1.26) *and* (1.27) *hold.*

Given this paper's focus on innovation, one can think about the equilibrium evolution of technology as a dynamic process with technology $\{A_{jt}\}$ as the state and scientists $\{s_{jt}\}$ as the control. In each period, the inherited state $\{A_{jt-1}\}$ pins down the control $\{s_{jt}\}$ via a spillover effect and a market size effect. The state then updates to $\{A_{jt}\}$, and the process starts over in the next period. In Appendix A.1.1, I provide a sufficient condition for the existence and uniqueness of the economy's equilibrium path, given some initial condition.

1.2.5 Example Production & Spillover Structure

The production and spillover structure I have described thus far involves minimal assumptions, so I will now sketch an example economy where I take a stance on the functional forms of production and spillovers. I will assume this structure when I calibrate and simulate my model in Sections 1.5 and 1.6, but the theoretical results of Sections 1.3 and 1.4 apply to the more general setup described above.

Production is divided into Θ-many sectors, each of which has a clean and dirty form of production. Final output follows

$$\mathcal{Y}_t = \Omega_t \Big(\sum_{\theta} \nu_\theta^{\frac{1}{\lambda}} E_{\theta t}^{\frac{\lambda-1}{\lambda}} \Big)^{\frac{\lambda}{\lambda-1}}, \tag{1.28}$$

where $E_{\theta t}$ is the output of sector θ. That is, final output is a complements CES over sectors with elasticity of substitution $\lambda < 1$ and sector share parameters $\{\nu_\theta\}$.

Output at the sector level follows

$$E_{\theta t} = \Big(Y_{\theta c t}^{\frac{\sigma-1}{\sigma}} + Y_{\theta d t}^{\frac{\sigma-1}{\sigma}} \Big)^{\frac{\sigma}{\sigma-1}}, \tag{1.29}$$

where $Y_{\theta c}$ and $Y_{\theta d}$ are the clean and dirty forms of production for sector θ, respectively. That is, sectoral output is a substitutes CES over clean and dirty production with elasticity of substitution $\sigma > 1$. This production structure nests the setup of Acemoglu et al. (2012) when there is a single sector.

I will assume the final sector of the economy Θ only has a clean form of production. My application considers the transportation and electricity generation sectors, but because these sectors make up only a small portion of the economy, the final sector, which I will refer to as the "general" sector, closes the model while abstracting from emissions that originate outside of transportation or electricity generation.[10]

On the innovation side of the economy, I will assume that research productivity follows $\chi_j = \chi \nu_{\theta(j)}^{-\eta}$, where $\theta(j)$ is the sector associated with technology j. Thus, the

[10] Motor vehicle output and electricity generation accounted for 2.6% and 1.8% of US GDP in 2021, respectively.

production of innovation follows

$$z_{jt} = \chi \left(\frac{s_{jt}}{\nu_{\theta(j)}} \right)^{\eta} \phi_{jt}. \tag{1.30}$$

Following Acemoglu and Restrepo (2022b), I show in Appendix A.1.2 that one can interpret the sectoral CES shares ν_θ as the share of tasks in the economy performed by sector θ. Thus, the relevant notion of effective scientific effort in this case is the number of scientists per economic task associated with a technology.

The spillover functions follow

$$\phi_{it} = \frac{\prod_j A_{jt-1}^{\tilde{\varphi}_{ij}}}{A_{it-1}}, \tag{1.31}$$

which implies that the spillover network is constant. The numerator represents an idiosyncratic Cobb-Douglas aggregator of all knowledge stocks in the economy, and the denominator represents the recipient's own knowledge stock. The Cobb-Douglas exponents define the gross spillover network $\tilde{\boldsymbol{\varphi}}$, so as discussed above, the (net) spillover network follows mechanically from $\boldsymbol{\varphi} = \tilde{\boldsymbol{\varphi}} - \mathbf{I}$. For example, the case of no cross-technology knowledge spillovers corresponds to a gross spillover network equal to the identity matrix: $\tilde{\boldsymbol{\varphi}} = \mathbf{I}$. This specification for spillovers was studied in the recent work by Liu and Ma (2021).

Finally, household utility follows the common CRRA specification

$$u(c_t) = \frac{c_t^{1-\vartheta} - 1}{1 - \vartheta}, \tag{1.32}$$

where $1/\vartheta$ is the intertemporal elasticity of substitution.[11]

[11]This specification implies a consumption discount rate of $\rho + \vartheta g_{ct}$ via the Ramsey equation, where g_{ct} is the growth rate of consumption.

1.3 Policy's Impact on the Direction of Innovation

This section characterizes the dynamic response of the direction of innovation to a policy reform. Specifically, I will imagine there is a once-and-for-all policy reform $\{d\tau, \{d\xi_j\}\}$.[12] I first characterize how technology's steady-state changes in response to the policy reform using what I call the *amplification matrix*. Next, I characterize the transition path of technology towards – or away from – the new steady-state using what I call the *transition matrix*.[13]

Both my amplification matrix and transition matrix are composed of the same two sufficient statistic matrices: one describing the network of knowledge spillovers and another describing substitution patterns in production. These two matrices show how the two forces, spillovers and substitutability, determine both the degree to which the policy reform will change technology's steady-state and the speed at which technology will converge to its new steady-state. These matrices are sufficient statistics in the sense that, conditional on estimating each matrix, one does not need to know the underlying structure of production and spillovers to first-order characterize the impact of policy on the direction of innovation.

Taking an eigendecomposition of the transition matrix, I quantify the degree of increasing returns to innovation in terms of the spectral radius of the transition matrix. I show that increasing returns determine both the long-run impact of policy and the speed of technology's transition, but that the two are inversely related. Finally, the spectral radius of the transition matrix provides me with a condition for technological path dependence, where the path of technology depends discontinuously on its initial condition. I conclude the section with a simple, two-technology example that

[12] As discussed in Section 1.4, time-invariant policy is not necessarily optimal. However, I focus on time-invariant policy in this section to analytically characterize how the direction of innovation would respond to a given policy regime, given unlimited time.

[13] My analysis does not require technology to be in a previous steady-state at the time of the policy reform.

illustrates the arguments made throughout the section.

1.3.1 Sufficient Statistic Matrices

I start by defining the two sufficient statistic matrices that will be relevant throughout this section. These matrices summarize the two main forces of my model: (i) substitution patterns in production and (ii) the network of cross-technology knowledge spillovers.

First, consider the substitution matrix:

Definition 1.3 (Substitution Matrix $\mathbf{\Sigma}$)**.** *The substitution matrix is a $J-1 \times J-1$ matrix with elements*

$$\Sigma_{ijt} \equiv \frac{\partial \ln \left(Y_{it}/Y_{Jt}\right)}{\partial \ln \left(p_{Jt}/p_{jt}\right)}. \tag{1.33}$$

The substitution matrix $\mathbf{\Sigma}_t$ describes how relative demand responds to changes in relative prices. As the relative price of good j changes, this induces a change in relative equilibrium demand for good i, as described by Σ_{ijt}. For both relative prices and quantities, good J is the base good.

The elements of the substitution matrix $\mathbf{\Sigma}_t$ are similar to elasticities of substitution. Indeed, the diagonal elements of $\mathbf{\Sigma}_t$ are elasticities of substitution between goods i and J, but for off-diagonal elements, I am allowing the change in relative price and change in relative quantity to refer to different goods. We then have the usual interpretation that diagonal elements above one describe substitutes, while diagonal elements below one describe complements. As an illustration, if the output aggregator were a standard CES with elasticity of substitution σ, we would have $\mathbf{\Sigma}_t = \sigma\mathbf{I}$. Therefore, the demand responses described by the substitution matrix $\mathbf{\Sigma}_t$ provide a general characterization of substitution patterns in production.

In the case of the nested-CES production structure described in Section 1.2.5, I show in Appendix A.1.3 that $\mathbf{\Sigma}_t$ is a block-diagonal matrix that depends on the elasticities of substitution of both nests, as well as income shares within sectors. When

I calibrate and simulate my model in Sections 1.5 and 1.6, I will use this closed-form solution for the substitution matrix to estimate $\boldsymbol{\Sigma}_t$ using existing estimates of the relevant elasticities of substitution.

Second, consider the spillover matrix:

Definition 1.4 (Spillover Matrix $\boldsymbol{\Phi}$)**.** *The spillover matrix is a $J-1 \times J-1$ matrix with elements*

$$\Phi_{ijt} \equiv \frac{\partial \ln\left(\phi_{Jt}/\phi_{it}\right)}{\partial \ln\left(A_{jt-1}/A_{Jt-1}\right)}. \tag{1.34}$$

The spillover matrix $\boldsymbol{\Phi_t}$ describes how relative spillovers change in response to changes in relative technology stocks. As before, the base technology is J. The fact that spillover functions are homogeneous of degree zero implies that one can map from the spillover network $\boldsymbol{\varphi}_t$ to the spillover matrix $\boldsymbol{\Phi}_t$ using

$$\Phi_{ijt} = \varphi_{Jjt} - \varphi_{ijt}. \tag{1.35}$$

Therefore, we can interpret the spillover matrix $\boldsymbol{\Phi}_t$ as a matrix of relative spillover elasticities.[14] That is, a positive value for Φ_{ijt} implies that the elasticity of ϕ_{Jt} with respect to technology j is greater than that of ϕ_{it}, and a negative value implies the reverse. Moreover, the mapping of Equation (1.35) implies that any estimate of the spillover network provides an estimate of the spillover matrix as well.

The spillover matrix $\boldsymbol{\Phi}_t$ provides a general description of an economy's spillovers for a wide class of spillover functions. In particular, the diagonal of $\boldsymbol{\Phi}_t$ provides a measure of the degree to which technologies receive spillovers from the *other* technologies in the economy. To see this, note that

$$\Phi_{iit} = \varphi_{Jit} - \varphi_{iit} = \sum_{j \neq i} \left(\varphi_{ijt} - \varphi_{Jjt}\right), \tag{1.36}$$

so the diagonal elements of $\boldsymbol{\Phi}_t$ will be large when technologies receive high spillovers

[14]To see this, note that if ϕ_{it} is homogeneous of degree zero, then its derivatives are homogeneous of degree negative one. Thus, we have $\frac{\partial \phi_{it}(\{A_{jt-1}/A_{Jt-1}\})}{\partial (A_{jt-1}/A_{Jt-1})} \frac{A_{jt-1}/A_{Jt-1}}{\phi_{it}(\{A_{jt-1}/A_{Jt-1}\})} = \frac{\partial \phi_{it}(\{A_{jt-1}\})}{\partial A_{jt-1}} \frac{A_{jt-1}}{\phi_{it}(\{A_{jt-1}\})}$.

from their peers, relative to the spillovers received by the base technology. When technologies receive minimal spillovers via the spillover network, i.e. $\sum_{j\neq i}\varphi_{ijt}$ approaches 0, less advanced technologies will experience only a minor enhancement in their research productivity. Therefore, the spillover matrix $\mathbf{\Phi}_t$ is informative of the extent to which cross-technology knowledge spillovers will stimulate catch-up growth.

For illustration, consider the two spillover networks of Figure 1·1.

Figure 1·1: Example Spillover Networks

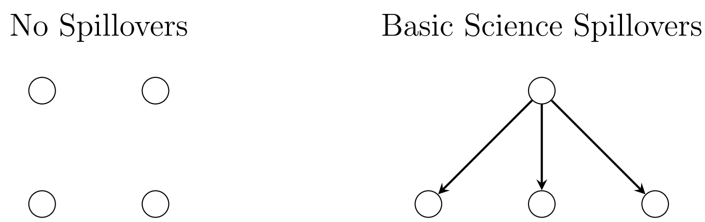

In the first network, the "No Spillovers" network, there are no cross-technology knowledge spillovers, so $\boldsymbol{\varphi}_t = \mathbf{0}$. In this case, the spillover matrix is also zero $\mathbf{\Phi}_t = \mathbf{0}$ to reflect the irrelevance of cross-technology spillovers in this economy. In the second network, the "Basic Science Spillovers" network, there is one technology that creates all of the spillovers in the economy. Assuming these spillovers are uniform, the spillover matrix follows $\mathbf{\Phi}_t = \varsigma\mathbf{I}$, where ς is the size of the spillover sent to each downstream technology in the economy. Thus, in this example, the diagonal of the spillover matrix reflects the strength of cross-technology spillovers in this economy.[15]

1.3.2 Steady-State Impact of Policy

I will now characterize how the balanced growth steady-state of this economy changes in response to a policy reform.

Definition 1.5 (Balanced Growth Steady-State)**.** *A balanced growth steady-state is an equilibrium in which every technology experiences uniform growth at rate g.*

[15] In these two stylized cases, the value of $\mathbf{\Phi}_t$ is independent of the choice of base technology, but this is not true in general.

I will describe the steady-state of technology in terms of relative technology, rather than technology levels, because relative technology is constant along a balanced growth path. Denote technology j's relative technology stock by

$$\bar{A}_{jt} \equiv \frac{A_{jt}}{A_{Jt}}, \tag{1.37}$$

where the base technology is J. Therefore, we can characterize the steady-state of this economy as a fixed point in a dynamic process for relative technology. In Appendix A.1.4, I provide conditions for the existence and uniqueness of the economy's steady-state.

The following proposition characterizes the steady-state impact of a policy reform in terms of an amplification matrix.

Proposition 1.1. *A policy reform $\{d\tau, \{d\xi_j\}\}$ induces first-order changes in the steady-state level of relative technology according to*

$$d\ln(\bar{A}_{ss}) = \eta\boldsymbol{\mathcal{M}}\big[d\ln(\Xi) - \alpha\big(\boldsymbol{\Sigma} - \boldsymbol{I}\big)d\ln(\mathcal{R})\big], \tag{1.38}$$

where $d\ln(\Xi_j) \equiv d\ln\big(\xi_j/\xi_J\big)$ and $d\ln(\mathcal{R}_j) \equiv d\ln\big((r_j + \omega_j\tau)/(r_J + \omega_J\tau_t)\big)$ are the changes to relative innovation subsidies and relative input costs induced by the policy reform, respectively. The amplification matrix $\boldsymbol{\mathcal{M}}$, evaluated at the steady-state, follows

$$\boldsymbol{\mathcal{M}} = \big[\boldsymbol{\Phi} - \eta(1-\alpha)\big(\boldsymbol{\Sigma} - \boldsymbol{I}\big)\big]^{-1}. \tag{1.39}$$

The proof of Proposition 1.1 can be found in Appendix A.1.4. Proposition 1.1 defines the amplification matrix $\boldsymbol{\mathcal{M}}$, which describes the role of spillovers and substitutability in shaping the steady-state impact of a policy reform. I will assume that the amplification matrix $\boldsymbol{\mathcal{M}}$ does not flip the sign of policy impacts, which I show to be the empirically relevant case in Section 1.6.1. Furthermore, I provide an economic interpretation for the case when the amplification matrix *does* flip the sign of policy impacts in Section 1.3.4.

Increases in relative subsidies Ξ raise the prevalence of favored technologies in steady-state, while higher carbon prices reduce the prevalence of dirty technologies in steady-state, insofar as goods are substitutes. The amplification matrix $\boldsymbol{\mathcal{M}}$ describes the full general equilibrium impact of these policy reforms. First, cross-technology knowledge spillovers dampen the impact of policy by increasing the research productivity of technologies disfavored by policy. The catchup growth this generates prevents any technology from gaining too large of a lead. Next, when goods are substitutes, technologies that are favored by policy increase their market size, further increasing the equilibrium incentive to innovate. The opposite is true when goods are complements because technologies favored by policy have their market size reduced. This channel shows up through the substitution matrix net of the identity matrix $(\mathbf{\Sigma} - \mathbf{I})$ because one is the boundary point between complements and substitutes. Intuitively, the spillover matrix decreases the "size" of the amplification matrix as it enters with a positive in the "denominator". In contrast, the substitution matrix increases the "size" of the amplification matrix as it enters with a negative in the "denominator".

1.3.3 Technology's Transition Path

Having examined how policy affects technology in the economy's long-run steady-state, I can now describe the transition path technology takes to its new steady-state. To determine the forces that are of first-order importance in governing this transition, I will linearize the equilibrium conditions around the steady-state to derive a transition matrix.

I will describe the transition path of technology in terms of the log deviation of relative technology from steady-state

$$\bar{\mathcal{A}}_t \equiv \ln(\bar{A}_t) - \ln(\bar{A}_{ss}). \tag{1.40}$$

That is, relative technology converges to its steady-state as $\bar{\mathcal{A}}_t$ converges to zero.

The following proposition provides a first-order characterization of the transition path of technology in terms of a transition matrix.

Proposition 1.2. *Given $\bar{\mathcal{A}}_0$, the transition path of technology follows, to a first-order, the linear process*

$$\bar{\mathcal{A}}_t \approx \boldsymbol{\mathcal{J}} \bar{\mathcal{A}}_{t-1}, \tag{1.41}$$

where the transition matrix $\boldsymbol{\mathcal{J}}$, evaluated at the steady-state, follows

$$\boldsymbol{\mathcal{J}} = \big[(1-\eta)\boldsymbol{I} - g\eta(1-\alpha)\big(\boldsymbol{\Sigma} - \boldsymbol{I}\big)\big]^{-1}\big[(1-\eta)\boldsymbol{I} - g\boldsymbol{\Phi}\big]. \tag{1.42}$$

The proof of Proposition 1.2 can be found in Appendix A.1.6. The main object of interest in Proposition 1.2 is the transition matrix $\boldsymbol{\mathcal{J}}$. This matrix shows that the transition path of technology is governed by the same two sufficient statistic matrices. Spillovers influence the transition matrix by enhancing the research productivity of less advanced technologies. As stated above, the diagonal of the spillover matrix $\boldsymbol{\Phi}$ provides a measure of the degree to which less advanced technologies can use spillovers to achieve catchup growth.

Substitutability influences the transition matrix through a market size effect that reduces the profits of technologies that are less advanced than their substitutes. Conversely, when technologies are less advanced than their complements, they achieve higher profits. These effects are summarized by the substitution matrix net of the identity matrix $(\boldsymbol{\Sigma} - \mathbf{I})$. As before, the identity matrix is subtracted from the substitution matrix because one is the boundary point between complements and substitutes. Furthermore, substitutability shows up through an inverse matrix because it follows the logic of a Leontief inverse. A small change in the inherited technological state $\{\bar{A}_{jt-1}\}$ induces a change in the realized technological state $\{\bar{A}_{jt}\}$. The realized technological state is what governs market size, so the change in market sizes induces yet another change in the realized technological state, and so on. This sequence of ripple effects generates a geometric series that culminates in an inverse matrix which describes the cumulative impact of the market size effect.[16]

1.3.4 Degree of Increasing Returns to Innovation

In this section, I eigendecompose the transition matrix and show that the degree of increasing returns to innovation can be quantified in terms of the transition matrix's spectral radius. This governs whether technology converges to its new steady-state quickly, slowly, or in an extreme case, not at all. Furthermore, I show that the amplification matrix can be written in terms of the eigenvalues of the transition matrix, with the implication that the steady-state impact of policy and the speed of technology's transition are inversely related.

The speed of convergence quantifies the degree of increasing returns to innovation. When increasing returns are high, convergence will be slow as technologies that are more advanced in the initial condition than in the steady-state will only slowly lose their advantage. Conversely, when increasing returns are low, convergence will be fast as any advantage enjoyed in the initial condition, relative to the steady-state, will quickly disappear. In what follows, I derive the speed of convergence from the transition matrix using spectral analysis.

First, use Equation (1.41) of Proposition 1.2 to iterate forward from technology's initial log deviation from steady-state

$$\bar{\mathcal{A}}_t \approx \boldsymbol{\mathcal{J}}^t \bar{\mathcal{A}}_0. \tag{1.44}$$

Intuitively, we can see that the speed of transition depends on how quickly the transition matrix $\boldsymbol{\mathcal{J}}$ shrinks technology's log deviation from steady-state $\bar{\mathcal{A}}_t$. If $\boldsymbol{\mathcal{J}}$ were a

[16]Formally, we have that

$$\left[\mathbf{I} - \frac{g\eta(1-\alpha)}{1-\eta}(\boldsymbol{\Sigma} - \mathbf{I})\right]^{-1} = \sum_{n \geq 0} \left(\frac{g\eta(1-\alpha)}{1-\eta}(\boldsymbol{\Sigma} - \mathbf{I})\right)^n = \mathbf{I} + \left(\frac{g\eta(1-\alpha)}{1-\eta}(\boldsymbol{\Sigma} - \mathbf{I})\right) + \left(\frac{g\eta(1-\alpha)}{1-\eta}(\boldsymbol{\Sigma} - \mathbf{I})\right)^2 + \dots, \tag{1.43}$$

which is proportional to the inverse matrix contained in $\boldsymbol{\mathcal{J}}$.

matrix of zeros, technology would converge immediately, whereas if $\boldsymbol{\mathcal{J}}$ were the identity matrix, technology would remain in its initial state forever. Thus, the speed at which less advanced technologies can catch up to their peers, and therefore the degree of increasing returns to innovation, depends on the "size" of the transition matrix. As we will see, the proper notion of a matrix's "size" is in terms of the magnitude of its eigenvalues, i.e. its spectral radius. I will focus on the case where the transition matrix $\boldsymbol{\mathcal{J}}$ has $J-1$ distinct real eigenvalues $\{\kappa_j\}$. I verify numerically that this is the empirically relevant case under the calibration described in Section 1.5. Consider the eigendecomposition of the transition matrix

$$\boldsymbol{\mathcal{J}} = \boldsymbol{Q}\boldsymbol{D}(\kappa)\boldsymbol{Q}^{-1}, \tag{1.45}$$

where $\boldsymbol{D}(\kappa)$ is a diagonal matrix whose diagonal elements are the eigenvalues of the transition matrix, and $\boldsymbol{Q}$ is an invertible matrix whose columns are the eigenvectors of the transition matrix. This eigendecomposition follows from the fact that any real square matrix has a Jordan decomposition (Galor, 2007).[17] Plugging the eigendecomposition into Equation (1.44), we have

$$\bar{A}_t \approx \boldsymbol{Q}\boldsymbol{D}(\kappa)^t\boldsymbol{Q}^{-1}\bar{A}_0. \tag{1.46}$$

From this, we can see that the speed of convergence is governed by the eigenvalues $\{\kappa_j\}$ of the transition matrix. As time progresses, technology's transition path is driven by the geometric decay of the eigenvalues. Eigenvalues slightly above zero will rapidly shrink to zero, signifying fast convergence, while eigenvalues slightly below one will gradually shrink to zero, signifying slow convergence.

[17] The basic arguments of this section hold for the more general case of repeated or complex eigenvalues because the main properties of Equation (1.45) hold for the Jordan decomposition of any real square matrix: $\boldsymbol{\mathcal{J}} = \boldsymbol{Q}\boldsymbol{D}\boldsymbol{Q}^{-1}$, where $\boldsymbol{D}$ is the Jordan normal form of $\boldsymbol{\mathcal{J}}$, and $\boldsymbol{Q}$ is some invertible matrix. Therefore, the convergence properties of technology always depend on the convergence properties of the Jordan normal form of $\boldsymbol{\mathcal{J}}$. See Lemma 2.8 of Galor (2007).

The eigenvalues of the transition matrix determine the speed of convergence, but which eigenvalues play the most important role quantitatively? To answer this question, consider the linear combination of eigenvectors which recovers the initial state of technology

$$\beta = \boldsymbol{Q}^{-1}\bar{\mathcal{A}}_0. \tag{1.47}$$

That is, β is a linear projection of the initial state $\bar{\mathcal{A}}_0$ onto the eigenvectors of the transition matrix. In other words, one can regress $\bar{\mathcal{A}}_0$ on $\boldsymbol{Q}$ to obtain β.[18] I will refer to the eigenvectors of the transition matrix as *eigenstates*. These eigenstates represent specific states of technology that correspond to the eigenvectors of the transition matrix.[19] Since $\boldsymbol{Q}$ forms a basis, any state of technology can be expressed as a linear combination of these eigenstates. Moreover, each eigenstate has a convergence rate determined by its corresponding eigenvalue. For instance, if the initial state of technology is proportional to the jth eigenstate, then β will load exclusively on the jth dimension, and technology will geometrically converge to the steady state with a decay rate of κ_j. More generally, the loading of the initial state on the eigenstates determines the speed of convergence as a weighted average of the eigenvalues. The following proposition provides a formal statement of how the spectral properties of the transition matrix determine technology's transition.

Proposition 1.3. *Suppose the transition matrix $\boldsymbol{\mathcal{J}}$ has $J-1$ distinct real eigenvalues $\{\kappa_j\}$. Then the transition path of technology follows*

$$\bar{\mathcal{A}}_t \approx \sum_j \kappa_j^t \beta_j Q_j, \tag{1.48}$$

where Q_j is the jth eigenstate of the transition matrix, and β is the linear combination

[18] Using the standard regression formula, we have $\beta = (\boldsymbol{Q}'\boldsymbol{Q})^{-1}\boldsymbol{Q}'\bar{\mathcal{A}}_0$, but since $\boldsymbol{Q}$ forms a basis, this formula simplifies to $\beta = \boldsymbol{Q}^{-1}\bar{\mathcal{A}}_0$. Thus, such a regression would have zero mean squared error.

[19] The "eigenstates" of my paper are analogous to the "eigenshocks" of Kleinman et al. (2023). In both cases, these terms describe particular configurations of state variables that (i) correspond to eigenvectors of a transition matrix and (ii) determine the eigenvalues that are most relevant for convergence speeds.

of the eigenstates that recovers the initial state of technology $\beta = \boldsymbol{Q}^{-1}\bar{\mathcal{A}}_0$.

Proposition 1.3 is derived through matrix manipulation of the eigendecomposition (1.45) of the transition matrix. It demonstrates that technology's transition depends on three factors: eigenvalues, eigenstates, and the initial state's loading on the eigenstates. The initial state of technology is composed of a linear combination of the eigenstates, and as time progresses, each eigenstate component converges toward the steady-state based on the geometric decay of its respective eigenvalue. Consequently, the overall speed of convergence is determined by each eigenvalue according to the extent to which the initial state loads on its respective eigenstate.

In summary, the speed of convergence depends on a weighted average of the transition matrix's eigenvalues, with the weights determined by the initial state's loading on the eigenstates. The largest eigenvalue will tend to have an outsized influence on the speed of convergence, so the spectral radius provides a summary measure of increasing returns to innovation. Thus, the formula for the transition matrix explains the determinants of increasing returns to innovation. Combining Equations (1.42) and (1.45), we have

$$\big[(1-\eta)\mathbf{I} - g\eta(1-\alpha)\big(\mathbf{\Sigma} - \mathbf{I}\big)\big]^{-1}\big[(1-\eta)\mathbf{I} - g\mathbf{\Phi}\big] = \boldsymbol{Q}\boldsymbol{D}(\kappa)\boldsymbol{Q}^{-1}. \tag{1.49}$$

Therefore, cross-technology knowledge spillovers and substitution patterns in production determine the degree of increasing returns to innovation through their influence on the eigenvalues of the transition matrix. Cross-technology spillovers accelerate convergence by stimulating catchup growth. Conversely, substitutability slows convergence by reducing profits for less advanced technologies. Intuitively, the spillover matrix decreases the "size" of the transition matrix as it enters with a negative in the "numerator", while the substitution matrix increases the "size" of the transition matrix as it enters with a negative in the "denominator".

Measuring increasing returns with the spectral radius of the transition matrix allows me to derive a condition for technological path dependence, a central topic in the directed technical change literature (Acemoglu et al., 2012; Aghion et al., 2016, 2019). When the direction of innovation is path dependent, the path of technology depends discontinuously on its initial state. We can see this case play out by considering the eigenvalues of the transition matrix. Technological convergence slows down as the eigenvalues of the transition matrix approach unity, and once one of the eigenvalues goes outside of the unit circle, the steady-state becomes unstable. The following corollary states this point formally.

Corollary 1.1. *The direction of innovation is path dependent if the spectral radius of the transition matrix* $\max|\kappa_j|$ *exceeds one.*

Corollary 1.1 is an application of a well-known result in dynamical systems establishing the local instability of steady-states (Galor, 2007).[20] It provides a sufficient condition for technological path dependence, which hinges on the spectral radius of the transition matrix. By examining Equation (1.48), we can see that if the initial state of technology loads on an eigenstate with an eigenvalue exceeding one, then technology will geometrically diverge from the steady-state in the direction of this unstable eigenstate. Consequently, the dynamic path for technology will depend discontinuously on the initial state, as a flip in the sign of the loading on the unstable eigenstate would result in divergence along the same ray but in the opposite direction.

My second corollary establishes the relationship between the amplification matrix and the eigenvalues of the transition matrix.

Corollary 1.2. *The amplification matrix follows*

$$\boldsymbol{\mathcal{M}} = g\boldsymbol{Q}\boldsymbol{D}(1-\kappa)^{-1}\boldsymbol{Q}^{-1}\big[(1-\eta)\boldsymbol{I} - g\eta(1-\alpha)\big(\boldsymbol{\Sigma} - \boldsymbol{I}\big)\big]^{-1}. \tag{1.50}$$

[20] In particular, see Theorem 4.8 of Galor (2007). The more general statement pertains to the modulus, rather than the absolute value, of the eigenvalues of the transition matrix, allowing for the possibility of complex eigenvalues.

The proof of Corollary 1.2 can be found in Appendix A.1.7. This result establishes that increasing returns to innovation increase the long-run impact of policy reforms through the term $\boldsymbol{D}(1-\kappa)^{-1}$. This implies that the long-run impact of policy and the speed of technology's transition are inversely related. When increasing returns are high ($\kappa_j \to 1$), the steady-state impact of policy will be large, but the transition will be slow. Conversely, when increasing returns are low ($\kappa_j \to 0$), the steady-state impact of policy will be small, but the transition will be fast. By favoring more advanced technologies, increasing returns allow technologies that are favored by policy to gain larger leads over their peers in the long run. However, the same mechanism leads to sluggish transitions by preventing less advanced technologies from catching up. For instance, high cross-technology knowledge spillovers reduce increasing returns and allow rapid transitions by facilitating catchup growth, but those same spillovers prevent technologies from gaining a significant advantage in the long run.

In describing the policy impacts of Proposition 1.1, I assumed the amplification matrix does not flip the sign of policy impacts. For example, clean innovation subsidies or carbon pricing reduce the prevalence of dirty technology in the steady-state. By combining Corollaries 1.1 and 1.2, we can interpret this assumption in terms of the degree of increasing returns to innovation. As we can see from Equation (1.50), eigenvalues that are much larger than one will tend to make the amplification matrix flip the sign of policy impacts. This is because, in the case of path dependence, the steady-state defines the boundary of the different basins of attraction. Thus, policy reforms that favor clean technology now shrink the dirty basins of attraction. For instance, an increase in the carbon price now *increases* the prevalence of dirty technology in the steady-state, implying that fewer initial conditions will land in dirty basins of attraction.

Finally, to measure convergence speeds in units of time, I will consider half-lives of convergence. That is, the number of periods required to halve the log distance from steady-state. Each eigenstate has its own half-life, which follows

$$t_{ES_j}^{(1/2)} = \left\lceil \frac{\ln(1/2)}{\ln(\kappa_j)} \right\rceil, \tag{1.51}$$

where $\lceil . \rceil$ is the ceiling function. Thus, if the initial state of technology is proportional to the jth eigenstate, the half-life of convergence for each technology will equal $t_{ES_j}^{(1/2)}$. More generally, as stated in Proposition 1.3, each technology will have its own half-life governed by the initial state's loading on the eigenstates. Half-lives are scale independent, so the overall distance of technology's initial state from its steady-state will not affect this metric of convergence.[21] However, the location of the initial state relative to the steady-state matters in that it influences the loading on the eigenstates β.

1.3.5 Two Technology Example

To illustrate the arguments of this section, I will now consider a simple example where there are only two technologies. I will make the parametric assumptions of Section 1.2.5 and imagine that there is a single sector with clean and dirty technology, denoted by c and d. In this case, there is only one relative technology $\bar{A}_t = A_{ct}/A_{dt}$. This setup nests the canonical model of Acemoglu et al. (2012). Throughout this section, I will assume a mild regularity condition to ease exposition.[22]

When there are only two technologies, the various matrices reduce to scalars, and

[21] To be exact, the half-life of convergence for each technology will be the same for initial states $\bar{A}_0$ and $k\bar{A}_0$, where $k > 0$.

[22] Specifically, I assume $1 - \eta > g\Phi$ and $1 - \eta > g\eta(1-\alpha)(\Sigma - 1)$, both of which are easily satisfied under empirically plausible parameter choices (see Section 1.5). The main reason for the ease with which these conditions are satisfied is that empirically plausible growth rates are on the order of $g = 0.02$, so this shrinks the right-hand side of both inequalities.

the two sufficient statistic matrices become

$$\Sigma = \sigma \tag{1.52}$$

$$\Phi = \varphi_{cd} + \varphi_{dc}. \tag{1.53}$$

That is, Σ is the elasticity of substitution between clean and dirty production, and Φ is the sum of cross-technology knowledge spillovers in the economy. Consistent with Proposition 1.1, the steady-state for relative technology follows

$$\ln(\bar{A}_{ss}) = \eta \frac{\ln(\Xi) - \alpha(\Sigma - 1)\ln(\mathcal{R})}{\Phi - \eta(1-\alpha)(\Sigma - 1)}, \tag{1.54}$$

where, as before, $\Xi \equiv \xi_c/\xi_d$ is the relative innovation subsidy and $\mathcal{R} \equiv r_c/(r_d + \omega_d \tau)$ is the relative input price (inclusive of the carbon price). From this, we can see that the amplification matrix follows

$$\mathcal{M} = \frac{1}{\Phi - \eta(1-\alpha)(\Sigma - 1)}. \tag{1.55}$$

As discussed in Section 1.3.2, the amplification matrix is decreasing in the level of cross-technology spillovers Φ and increasing in the elasticity of substitution Σ. Applying Proposition 1.2, we have that the transition matrix follows

$$\mathcal{J} = \frac{(1-\eta) - g\Phi}{(1-\eta) - g\eta(1-\alpha)(\Sigma - 1)}. \tag{1.56}$$

The transition matrix is also decreasing in the level of cross-technology spillovers Φ and increasing in the elasticity of substitution Σ, as argued in Section 1.3.4. In accordance with Proposition 1.3, technology's transition path is determined by the geometric decay, or expansion, of $\mathcal{J}^t$, so the speed of convergence increases as the transition matrix shrinks. In particular, the half-life of convergence follows

$$t^{(1/2)} = \left\lceil \frac{\ln(1/2)}{\ln(\mathcal{J})} \right\rceil. \tag{1.57}$$

To consider the possibility of path dependence, we can apply Corollary 1.1 and ask if $\mathcal{J}$ is outside the unit circle. This will be the case whenever

$$\Phi < \eta(1-\alpha)(\Sigma - 1). \quad (1.58)$$

That is, the direction of innovation is path dependent if the substitutability of clean and dirty technology is high relative to the size of cross-technology spillovers. This effect is also modulated by the curvature of research effort η and the importance of labor in production $(1-\alpha)$. As research effort in each period experiences steeper diminishing returns $(\eta \downarrow)$, this reduces the incentive to allocate greater research resources to the more advanced technology. As labor becomes less important in production $((1-\alpha)\downarrow)$, this reduces the competitive advantage of the more advanced technology as the factor that benefits from the productivity differential is less important.

Finally, the two forces of spillovers and substitutability have an intuitive influence: if the two technologies are highly substitutable relative to catchup-inducing spillovers, then the less advanced technology will not be profitable enough to attract the innovation effort necessary to overcome their disadvantage. Instead, the more advanced technology will enjoy greater innovation and further solidify its lead. As this process continues, the more advanced technology will strengthen its lead until, in the limit, all production and innovation will be devoted to the more advanced technology. This is the path dependence scenario considered by Acemoglu et al. (2012), so Inequality (1.58) provides a test for their path dependence hypothesis.

Finally, by applying Corollary 1.2, we can see that the long-run impact of policy and the speed of technology's transition are inversely related as

$$\mathcal{M} = \frac{g(1-\mathcal{J})^{-1}}{(1-\eta) - g\eta(1-\alpha)(\Sigma - 1)}. \quad (1.59)$$

This shows that increasing returns to innovation, as measured by $\mathcal{J}$, slow the transition while at the same time increasing the long-run impact of policy. Equation (1.59) also shows that the steady-state is unstable if and only if the amplification matrix $\mathcal{M}$ is negative. As discussed in Section 1.3.4, in the case of path dependence, the steady-state defines the boundary of the two basins of attraction. If the economy starts with dirty technology relatively more advanced than the steady-state value (i.e. $\bar{A}_0 < \bar{A}_{ss}$), then innovation will continue to favor dirty technology. Thus, policies that favor clean technology now lower steady-state relative technology $\bar{A}_{ss}$ and shrink the dirty basin of attraction.

1.4 Optimal Policy

We have seen the response of the direction of innovation to an arbitrary policy regime, so I will now characterize optimal climate innovation policy. I start with a first-best problem where the Planner's instruments are unrestricted. In that case, the two externalities, carbon pollution and knowledge spillovers, are corrected separately, in keeping with the Pigou Principle. In particular, the fact that clean technology can produce goods without pollution does not warrant an additional innovation subsidy when carbon pollution is priced properly. Instead, innovation subsidies reward technologies for their ability to produce knowledge spillovers, and I derive a recursive formula for optimal innovation subsidies that corrects the knowledge spillover externality created via the spillover network.

I then consider a realistic second-best problem where the Planner cannot control the price of carbon pollution. In that case, the Planner must choose innovation subsidies when there is some externally imposed carbon price. I show that the adjustment to optimal innovation subsidies is simple; the same recursive formula holds as in the first-best but with a simple adjustment for the wedge between the external and efficient carbon price. Now clean technologies *do* receive additional innovation subsidies for their ability to produce goods without pollution as the Planner must correct both externalities with a single instrument.

1.4.1 First-Best

I start with the first-best planning problem. The Planner has a complete set of instruments, so I will consider a first-best primal problem and back out the corrective instruments that implement the Planner's allocation as an equilibrium.

Definition 1.6 (First-Best Planning Problem)**.** *The Planner solves*

$$\max_{\{c_t,\{\Lambda_{jt},\{\ell_{j\iota t}\},A_{jt},s_{jt}\}\}} \sum_{t\geq 0} \frac{1}{(1+\rho)^t} u(c_t) \quad s.t. \tag{1.60}$$

$$\mathcal{Y}_t = c_t + \sum_j r_j \Lambda_{jt}$$

$$L = \sum_j \int_0^1 \ell_{j\iota t} d\iota$$

$$A_{jt} = \gamma^{\chi_j s_{jt}^{\eta} \phi_{jt}} A_{jt-1}$$

$$\mathcal{S} = \sum_j s_{jt}.$$

That is, the Planner picks the technologically feasible allocation that will maximize household utility. Note that climate damage from past carbon emissions is embedded in $\mathcal{Y}_t$. It will be helpful to define the Planner's intertemporal marginal rate of substitution $R_{t+1} \equiv (1+\rho)u'_t/u'_{t+1}$. This is the Planner's discount factor over consumption goods. The following proposition describes first-best carbon prices and innovation subsidies.

Proposition 1.4. *To implement the first-best allocation as an equilibrium, the Planner separately corrects the externalities from carbon pollution and knowledge spillovers. They set the carbon price according to*

$$\tau_t = -\sum_{\hat{t}\geq t} \prod_{\hat{s}=1}^{\hat{t}-t} \frac{1}{R_{t+\hat{s}}} \mathcal{Y}_{\hat{t}} \frac{\partial \ln(\Omega_{\hat{t}})}{\partial \mathcal{E}_t}, \tag{1.61}$$

which is the social cost of carbon. Next, they set innovation subsidies according to the recursion

$$\xi_{jt}\Pi_{jt} = (1-\alpha)S_{jt}\mathcal{Y}_t + \frac{1}{R_{t+1}}\big[\xi_{jt+1}\Pi_{jt+1} + \sum_i \xi_{it+1}\Pi_{it+1}g_{it+1}\varphi_{ijt+1}\big]. \tag{1.62}$$

The proof of Proposition 1.4 can be found in Appendix A.1.8. Several comments are in order. First, the Planner corrects the two externalities in the economy with distinct instruments. This principle of separating externalities has been argued for elsewhere in the climate innovation literature (Acemoglu et al., 2012, 2016; Golosov et al., 2014). Clean innovators are blind to the climate benefit of their innovations in laissez-faire, but once the carbon price has been set equal to the social cost of carbon (1.61), innovation rents will be properly adjusted to reward clean technology. Innovation subsidies do not need to provide additional support to clean innovation as such.

Instead, innovation subsidies correct the knowledge spillover externalities mediated through the spillover network, and in doing so, they ignore the pollution externalities associated with each technology. Innovation subsidies set according to the recursive formula of Equation (1.62) guarantee that the private reward for innovation $\xi_{jt}\Pi_{jt}$ equals the social value of innovation. This social value includes two terms. First, there is a contemporaneous benefit to production from having more advanced technology. This effect is captured by the $(1-\alpha)S_{jt}\mathcal{Y}_t$ term, which reflects the Domar weight of technology à la Hulten's Theorem (Domar, 1961; Hulten, 1978). Next, the terms in the brackets represent the spillover benefits realized in the following period. Tomorrow's innovators will build on top of the knowledge stock they inherit, represented by the $\xi_{jt+1}\Pi_{jt+1}$ term, and their research productivity will be affected by knowledge stocks via the spillover network $\boldsymbol{\varphi}_{t+1}$, represented by the $\sum_i \xi_{it+1}\Pi_{it+1}g_{it+1}\varphi_{ijt+1}$ term. Therefore, the innovation subsidies of Proposition 1.4 provide a general formula for innovation policy in the presence of a spillover network as it relies on minimal parametric assumptions about the structure of production and spillovers. Finally, the Planner sets the intermediate subsidy Υ to close the monopoly markup γ and the lump-sum tax D_t to balance the government's budget (1.27) in each period.

As I argue throughout this paper, knowledge spillovers are critical for understanding both the impact and efficiency of policy in shaping the direction of innovation. However, we can see from Proposition 1.4 that the role of spillovers differs in these two cases. For describing the impact of policy, the spillovers technologies *receive* are what matter. These spillovers allow less advanced technologies to achieve catchup growth, and as argued in Section 1.3, this shapes both the steady-state impact of policy and the speed with which technology converges to its steady-state. Put differently, the rows of the spillover matrix $\boldsymbol{\varphi}_t$ are what matter for describing the impact of policy, and this can be seen from the fact that the diagonal of the spillover matrix $\boldsymbol{\Phi}_t$ (1.36) is composed of row sums. For prescribing optimal policy, the spillovers technologies *send* are what matter. This is because the recipients of spillovers already internalize the boost to their research productivity, so policy needs to reflect the spillover externalities that profit signals ignore. It is the columns of the spillover matrix $\boldsymbol{\varphi}_t$ that matter for efficiency, and this can be seen from the fact that the spillover benefit term of Equation (1.62) is composed of column sums.

To provide further intuition for the determinants of optimal innovation subsidies, I will consider steady-state innovation policy. My results mirror those of Liu and Ma (2021), but I will consider a more general structure for production and spillovers, albeit with a narrow focus on steady-state policy. I will describe the steady-state in terms of innovation subsidies multiplied by income shares $\tilde{\xi}_{jt} \equiv \xi_{jt} \cdot S_{jt}$ as this value converges to a constant even in cases where income shares of some technologies go to zero. For instance, ever-increasing carbon prices may drive the income share of dirty technologies asymptotically to zero.[23] It will be helpful to define the Planner's growth-adjusted intertemporal marginal rate of substitution $\tilde{R}_t \equiv R_t/(1+g_{yt})$, where g_{yt} is the growth rate of output. The following corollary characterizes steady-state innovation policy.

Corollary 1.3. *In steady-state, innovation subsidies multiplied by income shares satisfy*

$$\tilde{\xi}' = \frac{1}{\gamma - 1} S' \Big[(1 - \tilde{R}^{-1}) \boldsymbol{I} - g \tilde{R}^{-1} \boldsymbol{\varphi} \Big]^{-1}. \tag{1.63}$$

The proof of Corollary 1.3 can be found in Appendix A.1.9. This result states that steady-state innovation policy depends on both the spillover network $\boldsymbol{\varphi}$ and the Planner's level of patience, as measured by their growth adjusted discount factor $\tilde{R}$. To better understand the intuition, note that the inverse matrix of Equation (1.63) follows the logic of a Leontief inverse. Today's innovations expand output according to income shares S and create spillovers for the following period. Iterating forward with Equation (1.62), these spillovers beget further innovations tomorrow, which generate their own spillovers in the following period, and so on. The sequence of ripple effects generates a geometric series that sums to the inverse matrix of Equation (1.63).[24]

An informative case is when there are no cross-technology spillovers $\boldsymbol{\varphi} = \mathbf{0}$. Then, all innovation subsidies are set equal to $1/(\gamma - 1)(1 - \tilde{R}^{-1})$, which I call the *baseline innovation wedge*. This wedge reflects the difference between the private marginal benefit of innovation, $(\gamma - 1)$ times the contemporaneous income growth from innovation, and the social marginal benefit of innovation, a permanent increase in income from better technology, when there is no spillover network. In our setting,

[23] This is what happens in the steady-state under the parametric assumptions of Sections 1.2.5 and 1.5.2.

[24] Formally, we can write

$$\big[(1 - \tilde{R}^{-1})\mathbf{I} - g\tilde{R}^{-1}\boldsymbol{\varphi}\big]^{-1} = \sum_{n \geq 0} \Big(\frac{\mathbf{I} + g\boldsymbol{\varphi}}{\tilde{R}}\Big)^n = \mathbf{I} + \Big(\frac{\mathbf{I} + g\boldsymbol{\varphi}}{\tilde{R}}\Big) + \Big(\frac{\mathbf{I} + g\boldsymbol{\varphi}}{\tilde{R}}\Big)^2 + \ldots, \tag{1.64}$$

so the Planner considers a full geometric series of spillover effects in setting steady-state innovation policy.

it is natural to think about innovation policy net of the baseline innovation wedge $\hat{\xi} \equiv (\gamma - 1)(1 - \tilde{R}^{-1})\tilde{\xi}$ because this term controls the *composition* of research effort. Without a spillover network, the laissez-faire composition of research effort, i.e. $\hat{\xi} = S$, would be efficient in steady-state because the lack of cross-technology spillovers implies innovation rents are proportional to the social value of innovation.

For the more general case with a spillover network, manipulating Equation (1.63), $\hat{\xi}$ solves

$$(1 - \tilde{R}^{-1})\big(\hat{\xi}' - S'\big) - g\tilde{R}^{-1}\hat{\xi}'\boldsymbol{\varphi} = 0. \tag{1.65}$$

This shows the role of the Planner's level of patience. A Planner that is completely myopic, i.e. $\tilde{R}^{-1} \to 0$, will allow the composition of research effort to be set entirely by profit signals: $\hat{\xi} = S$. This maximizes the contemporaneous benefit of innovation by focusing exclusively on the immediate impact on output. A Planner that is perfectly patient, i.e. $\tilde{R}^{-1} \to 1$, will set innovation policy with an exclusive focus on the spillover benefit of innovation. They set innovation policy according to

$$\hat{\xi}'\boldsymbol{\varphi} = \hat{\xi}'\big(\tilde{\boldsymbol{\varphi}} - \mathbf{I}\big) = 0, \tag{1.66}$$

where $\tilde{\boldsymbol{\varphi}}$ is the gross spillover network discussed in Section 1.2.2. That is, $\hat{\xi}$ solves $\hat{\xi}'\tilde{\boldsymbol{\varphi}} = \hat{\xi}'$, so $\hat{\xi}$ is a measure of eigenvector centrality for the gross spillover network. In this case, the Planner prioritizes technologies that are central in the spillover network because these technologies are best able to produce spillovers over time.

In summary, technology stocks are used in the production of both physical goods and ideas, and Equation (1.65) states that steady-state innovation policy strikes a balance between these two benefits according to the Planner's degree of patience. For plausible values of $\tilde{R}^{-1}$, the spillover network will be the primary determinant of innovation policy.

1.4.2 Second-Best: Innovation Subsidies with Incomplete Carbon Pricing

Having considered optimal policy with an unrestricted set of instruments, I will now consider a second-best case where the Planner cannot set the price on carbon. I view this as a realistic case given that pricing carbon has proved to be politically difficult.

Let $\{\hat{\tau}_t\}$ denote a sequence of externally imposed carbon prices. For example, these external carbon prices may simply be zero if carbon pricing is politically impossible. The loss of policy instruments implies that the Planner must now satisfy an incentive compatibility constraint. To keep focus on innovation, I will assume that intermediate subsidies continue to close the monopolist markup $\Upsilon = \gamma$.

Definition 1.7 (Second-Best Planning Problem)**.** *The Planner solves*

$$\max_{\{c_t,\{A_{jt},s_{jt}\}\}} \sum_{t\geq 0} \frac{1}{(1+\rho)^t} u(c_t) \;\; s.t. \tag{1.67}$$

$$\mathcal{Y}_t = c_t + \sum_j r_j \Lambda_{jt}$$

$$A_{jt} = \gamma^{\chi_j s_{jt}^{\eta} \phi_{jt}} A_{jt-1}$$

$$\mathcal{S} = \sum_j s_{jt}$$

$$\{\Lambda_{jt}, \{\ell_{j\iota t}\}\} = \operatorname{argmax} \left[\mathcal{Y}_t - \sum_j r_j \Lambda_{jt} - \hat{\tau}_t \mathcal{E}_t\right] \;\; s.t. \;\; L = \sum_j \int_0^1 \ell_{j\iota t} d\iota.$$

The final constraint is an incentive compatibility constraint which reflects the private optimizing behavior of producers. The equilibrium of this economy can be represented as a simple revenue maximization problem, and the Planner must now choose an innovation allocation knowing that the choice of technology will affect production choices in equilibrium. The Planner still has a complete set of policy instruments on the innovation side of the economy, so they can back out the innovation subsidies that implement their desired allocation of scientists as an equilibrium. The following proposition characterizes those innovation subsidies.

Proposition 1.5. *When the Planner can no longer control the carbon price, they set innovation subsidies according to the recursion*

$$\xi_{jt}\Pi_{jt} = (1-\alpha)S_{jt}\mathcal{Y}_t - (\tau_t - \hat{\tau}_t)\mathcal{E}_t \frac{\partial \ln(\mathcal{E}_t)}{\partial \ln(A_{jt})} + \frac{1}{R_{t+1}}\left[\xi_{jt+1}\Pi_{jt+1} + \sum_i \xi_{it+1}\Pi_{it+1}g_{it+1}\varphi_{ijt+1}\right], \tag{1.68}$$

where τ_t follows the social cost of carbon formula of Equation (1.61)*.*

The proof of Proposition 1.5 can be found in Appendix A.1.10. The first thing to note is that second-best innovation subsidies are remarkably similar to those of the first-best; the only difference is a new term that reflects the wedge between the external and efficient carbon price. The social cost of carbon is no longer properly internalized in production decisions, so to make up for this distortion, the Planner adjusts the direction of innovation in favor of clean technology.[25] This adjustment multiplies the wedge between the external and efficient carbon price, total emissions, and the elasticity of equilibrium emissions with respect to technology. For clean technology, this emissions elasticity will tend to be negative, implying an increase in clean innovation subsidies. Conversely, improvements in dirty technology will tend to increase emissions in equilibrium, so dirty innovation subsidies will be reduced.[26] As one might expect, this formula recovers the first-best when the external carbon price $\hat{\tau}_t$ happens to be set at the efficient level τ_t.

The similarity between first- and second-best innovation subsidies implies that the latter inherit the properties discussed in Section 1.4.1. All else equal, technologies are rewarded for their importance in production and their ability to create spillovers, as measured by Domar weights and eigenvector centrality respectively, with the weight between the two determined by the Planner's level of patience. Now that carbon

[25] I will assume that $\tau_t \geq \hat{\tau}_t$ in my discussion of Proposition 1.5 as this is the empirically relevant case, but the result does not formally require this to be true.

[26] This opens the possibility that innovation subsidies are negative. In that case, the allocation of scientists will be a corner solution with zero scientists allocated to technologies with negative innovation subsidies. Formally, $\xi_{jt} \leq 0 \Rightarrow s_{jt} = 0$. Indeed, if all technologies lead to sufficiently large increases in equilibrium emissions, the Planner may want to stop growth altogether.

pollution is improperly priced, the flow value of innovation includes the degree to which technological change exacerbates the economy's distortions. Thus, one can interpret the innovation subsidies of Equation (1.68) as adjusting for the possibility of "immiserizing growth", where innovation can reduce welfare in inefficient economies by further deepening distortions (Bhagwati, 1958, 1968). By this view, Proposition 1.5 provides a general way of thinking about innovation policy in the presence of distortions. The flow value of innovation that guides policy must include both Domar weights and the effect on distortions to prevent technological change from reducing welfare.

1.5 Calibration

In this section, I describe my calibration of the structural parameters of my model. My application is the transportation and electricity generation sectors in the United States. Throughout this section, I make the parametric assumptions for production and spillovers described in Section 1.2.5. I conclude the section with a model validation exercise that tests my model's ability to match the advances made by clean technology in both sectors from 2010 to 2021. I show that my model is able to match the evolution of technology in this period when it includes the spillover network but fails to do so when the spillover network is shut down.

A time period represents one year. Longer time periods, such as one representing five years, imply counterfactually rapid growth of clean transportation from 2010 to 2021. Finally, any data series referenced in this section that is published in nominal terms is adjusted using a US consumer price index that takes 2012 as its base year. Thus, all dollar values referenced throughout the paper are in 2012 dollars.

1.5.1 Spillover Network

In this section, I describe how I calibrate one of the central objects in my model: the spillover network. I follow Liu and Ma (2021) in using the citation network of patents as a proxy for the spillover network.

For my sample of patents, I take granted US patents from PatentsView.[27] First, I need to assign patents to my five technology classes: clean transportation, dirty transportation, clean electricity generation, dirty electricity generation, and general technology. To do so, I make use of the International Patent Classification (IPC) and Cooperative Patent Classification (CPC) systems.[28] IPC codes pertaining to both clean and dirty transportation come from Aghion et al. (2016), who identify IPC codes for the transportation sector in their empirical study of directed innovation in the auto industry.[29] I also include categories from CPC subclass Y02T that pertain to transportation. For clean transportation, this includes electromobility (Y02T64-72), efficient charging and discharging systems (Y02T10/92), charging of electric vehicles (Y02T90/10-16), hydrogen technology in transport (Y02T90/40), and hybrids (Y02T10/62). For dirty transportation, this includes internal combustion engine efficiency (Y02T10/12) and engine management (Y02T10/40).

For clean electricity generation, I take directly from the CPC subclass Y02E related to GHG reducing innovations in electricity generation. This includes renewables (Y02E10), nuclear (Y02E30), and energy storage (Y02E60/10-16). I also include in-

[27]I am abstracting from international patents in the construction of my spillover network, but I do not think this introduces bias for the following two reasons. First, US patents primarily cite other US patents. Indeed, 70% of US patent citations reference other US patents (Liu and Ma, 2021). Second, citation shares are scale independent, so for the exclusion of international patents to introduce bias, citations to international patents would have to systematically differ in their distribution across technology classes.

[28]The CPC system is a simple extension of the IPC system. While there is some discordance between the two systems, I manually verify that all of the IPC codes that I use map directly to the same CPC codes with the same meaning. For instance, B60L denotes electric vehicles in both the IPC and CPC system.

[29]For the full list of IPC codes pertaining to both clean and dirty transportation, see Table 1 of Aghion et al. (2016).

tegration of photovoltaics in buildings (Y02B10/10). IPC codes pertaining to dirty electricity generation come from Lanzi et al. (2011), who identify a list of IPC codes associated with fossil fuel technologies for electricity generation.[30] Finally, I define general technology patents as those patents that are not in any of the above categories. Table 1.1 gives a representative, but not exhaustive, list of my patent classification codes.

Table 1.1: Sampling of Patent Classification Codes

Transportation		Electricity Generation	
Description	IPC/CPC Codes	Description	IPC/CPC Codes
Clean		Clean	
Electric Vehicles	B60L	Renewables	Y02E10
Hybrids	Y02T10/62	Nuclear	Y02E30
Hydrogen Fuel Cells	H01M8	Energy Storage	Y02E60/10-16
Dirty		Dirty	
Internal Combustion Engines	F02B	Steam Engine Plants	F01K
Controlling Combustion Engines	F02D	Gas-Turbine Plants	F02C
Supplying Combustion Engines	F02M	Steam Generation	F22
Cylinders for Combustion Engines	F02F	Combustion Apparatus	F23

Notes: IPC codes for both clean and dirty transport patents come from Aghion et al. (2016), CPC codes for clean electric patents come directly from CPC subclass Y02E, and IPC codes for dirty electric patents come from Lanzi et al. (2011). General patents are those not classified as pertaining to either transport or electricity.

With patents placed into technology classes, I estimate the gross spillover network $\tilde{\boldsymbol{\varphi}}$ as the proportion of technology i's patent citations that reference patents in technology j. That is, I set gross spillovers according to

$$\tilde{\varphi}_{ij} = \frac{cites_{ij}}{\sum_k cites_{ik}}, \tag{1.69}$$

where $cites_{ij}$ is the number of citations from patents in technology class i that reference patents in technology class j. This approach leverages the widely accepted view that patent citations are a metric of knowledge spillovers (Jaffe et al., 1993; Hall et al., 2005). The (net) spillover network then follows mechanically from $\boldsymbol{\varphi} = \tilde{\boldsymbol{\varphi}} - \mathbf{I}$, and because Equation (1.69) implies that the rows of $\tilde{\boldsymbol{\varphi}}$ mechanically sum to one, this

[30]For the full list of IPC codes pertaining to dirty electricity generation, see Table A1 of Lanzi et al. (2011).

guarantees that the spillover function described in Equation (1.31) is homogeneous of degree zero.

Figure 1·2 contains a heat map that illustrates the gross spillover network $\tilde{\boldsymbol{\varphi}}$. It excludes the row pertaining to the general technology as almost all of the general technology's patent citations reference another general technology patent.[31] Figure 1·2 allows us to see several intuitive patterns in the spillover network. First, all of the climate-relevant technology classes extensively cite both themselves and the general technology. Second, the transportation sector has relatively high within-sector spillovers when compared with the electricity generation sector. As suggested by the story of Tesla's prototype, both clean and dirty vehicles can learn from each other because they are fundamentally similar machines. This does not appear to be the case in the electricity generation sector where, for example, solar panels and coal-fired power plants generate electricity via entirely different means. Instead, clean electricity has an abnormally high citation share on the general technology, consistent with the solar panel inventors from Bell Labs building on knowledge from the ICT sector. Finally, clean and dirty technologies each form their own distinct spillover clusters. That is, dirty technologies of both sectors cite each other, and the same is true of clean technologies but to a lesser extent. Each of these features of the spillover network are relevant for convergence dynamics, as I will discuss in Section 1.6.1.

1.5.2 Climate Module

To simulate optimal climate policy, I must add structure to my model specifying (i) how a sequence of emissions influences the Earth's climate system and (ii) how warming of the Earth's climate system impacts the economy's productive capacity. To this end, I follow Golosov et al. (2014) in my description of the carbon cycle and

[31] Specifically, 98.6% of general technology patent citations reference another general technology patent.

Figure 1·2: Heat Map of Gross Spillover Network $\tilde{\varphi}$

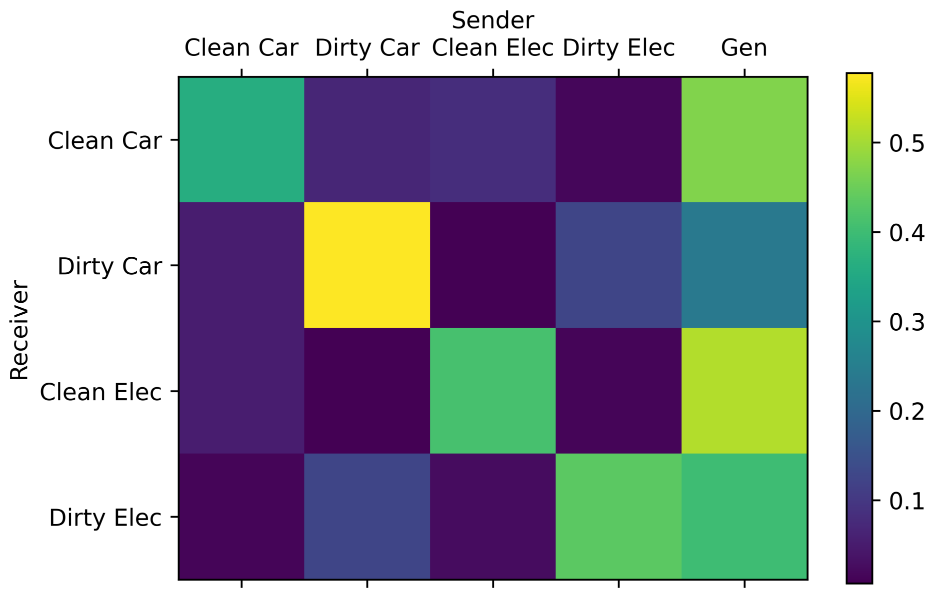

Notes: The y-axis denotes technologies receiving spillovers, while the x-axis denotes technologies sending spillovers. Granted US patents come from PatentsView.

economic damage function.

Atmospheric carbon concentrations evolve according to

$$\mathcal{C}_t = \sum_{\hat{t}=1800}^{t} \left(\psi_p + (1-\psi_p)\psi_0\psi^{t-\hat{t}}\right)\mathcal{E}_{\hat{t}} + \bar{\mathcal{C}}, \tag{1.70}$$

where $\bar{\mathcal{C}}$ is the pre-industrial level of atmospheric carbon concentration. I select the year 1800 as the starting point of industrialization.[32] This specification of the carbon cycle provides a mapping from past carbon emissions to the current level of atmospheric carbon concentrations. In particular, it states that fraction $(\psi_p + (1-\psi_p)\psi_0\psi^{t-\hat{t}})$ of carbon emitted at time $\hat{t} \leq t$ will remain in the atmosphere at time t. This remaining carbon has both a permanent and transitory component. The permanent component, fraction ψ_p, will remain in the atmosphere forever. The

[32] One could debate the exact start date of industrialization, but given that carbon emissions were trivially small before the year 1800, the choice is insignificant for present purposes.

remainder, fraction $(1-\psi_p)$, is transitory. For the transitory component, a fraction $(1-\psi_0)$ exits the atmosphere within a period and is absorbed by the biosphere or surface oceans. Next, the remainder of the transitory component, fraction ψ_0, decays geometrically according to ψ. As argued in Archer (2005) and Golosov et al. (2014), this relatively simple specification of the Earth's climate system provides a good approximation of the complex relationship between carbon emissions and atmospheric carbon concentrations. In particular, it captures the slow process by which the deep oceans absorb carbon from the atmosphere.[33]

The pre-industrial level of atmospheric carbon concentration $\bar{\mathcal{C}}$ was 596.4 gigatons of carbons.[34] Throughout the paper, I list carbon quantities in terms of gigatons of carbon (GtC). For example, US greenhouse gas emissions in 2021 were equivalent to 1.7 GtC. I will, however, list carbon prices in terms of dollars per ton of CO_2 as this is the convention. To map back and forth, one can note that a ton of CO_2 contains 12/44 tons of carbon.

To calibrate my climate module, I start by setting $\psi_p = 0.2$, in line with the 2007 IPCC report estimate that 20% of carbon emissions will remain in the atmosphere after thousands of years. For the remaining two parameters $\{\psi_0, \psi\}$, I perform a similar exercise as Acemoglu et al. (2016) and pick these values to match the relationship between carbon emissions and atmospheric carbon concentrations throughout the 1960-2020 period, selecting $\psi_0 = 0.377$ and $\psi = 0.994$ as the values that minimize

[33] This approach is an alternative to that of Nordhaus (2017), which models the Earth's climate system as consisting of three carbon reservoirs: the atmosphere, the upper oceans and biosphere, and the deep oceans. In that model, carbon concentrations of the three reservoirs evolve according to a system of linear difference equations. As discussed in Archer (2005), Archer et al. (2009), and Golosov et al. (2014), the linear, three reservoir specification implies a counterfactually rapid absorption of atmospheric carbon into the deep oceans. Instead, the approach taken here is better able to match the depreciation of atmospheric carbon concentrations, while also reducing the number of state variables needed to describe the climate. The implication of this difference in modeling choice is longer-lived, and therefore larger, impacts of anthropogenic carbon emissions on the climate.

[34] Atmospheric carbon concentrations are often stated in terms of parts per million (ppm) of CO_2. One ppm of CO_2 is equivalent to 2.13 GtC (O'Hara, 1990). This conversion was used to map the widely-accepted pre-industrial 280 ppm of CO_2 to GtC: $596.4 = 280 \times 2.13$.

the distance between the model's predictions and the data. A major advantage of Equation (1.70) is that it allows for the state of the climate to be written in terms of a two-dimensional recursion. The first dimension $\mathcal{C}_{1t}$ represents the permanent component of atmospheric carbon, and it follows

$$\mathcal{C}_{1t} = \psi_p \mathcal{E}_t + \mathcal{C}_{1t-1}. \tag{1.71}$$

The second dimension $\mathcal{C}_{2t}$ represents the transitory component of atmospheric carbon, and it follows

$$\mathcal{C}_{2t} = (1 - \psi_p)\psi_0 \mathcal{E}_t + \psi \mathcal{C}_{2t-1}. \tag{1.72}$$

The sum of these two components is therefore atmospheric carbon concentration: $\mathcal{C}_t = \mathcal{C}_{1t} + \mathcal{C}_{2t}$. Writing the state of the climate recursively allows for a simpler representation of the climate system as one does not need to keep track of the entire history of emissions. This requires the selection of an initial condition $\{\mathcal{C}_{1t_0}, \mathcal{C}_{2t_0}\}$. For any t_0, such as 1960 in the case of my calibration exercise, I set $\mathcal{C}_{1t_0} = \sum_{\hat{t}=1800}^{t_0} \psi_p \mathcal{E}_{\hat{t}} + \bar{\mathcal{C}}$ and $\mathcal{C}_{2t_0} = \hat{\mathcal{C}}_{t_0} - \mathcal{C}_{1t_0}$, where $\hat{\mathcal{C}}_{t_0}$ is the observed level of atmospheric carbon concentrations. Data on carbon emissions dating back to 1800 comes from Our World in Data, and data on atmospheric carbon concentrations comes from the NOAA's Mauna Loa Observatory.[35] The results of my calibration exercise can be seen in Figure 1·3.

Finally, atmospheric carbon concentrations create damage to production via

$$\Omega_t = \exp\big(- \varrho(\mathcal{C}_t - \bar{\mathcal{C}})\big), \tag{1.73}$$

where $\varrho > 0$ is a scale parameter determining the sensitivity of output to the climate.[36] This specification implies that the semi-elasticity of final output with respect

[35] Carbon emissions from land use change are only available going back to 1850, so I extrapolate the linear trend in those emissions from 1850 to 1950 back to 1800.

[36] One can also map from atmospheric carbon concentrations to temperature increases with the equation $\Delta T_t = \Gamma \ln(\mathcal{C}_t/\bar{\mathcal{C}})/\ln(2)$, where Γ represents the climate sensitivity parameter. That is, a doubling of atmospheric carbon concentrations over pre-industrial levels leads to warming of Γ

Figure 1·3: Climate Model Matches Historic Relationship of Emissions and Carbon Concentrations

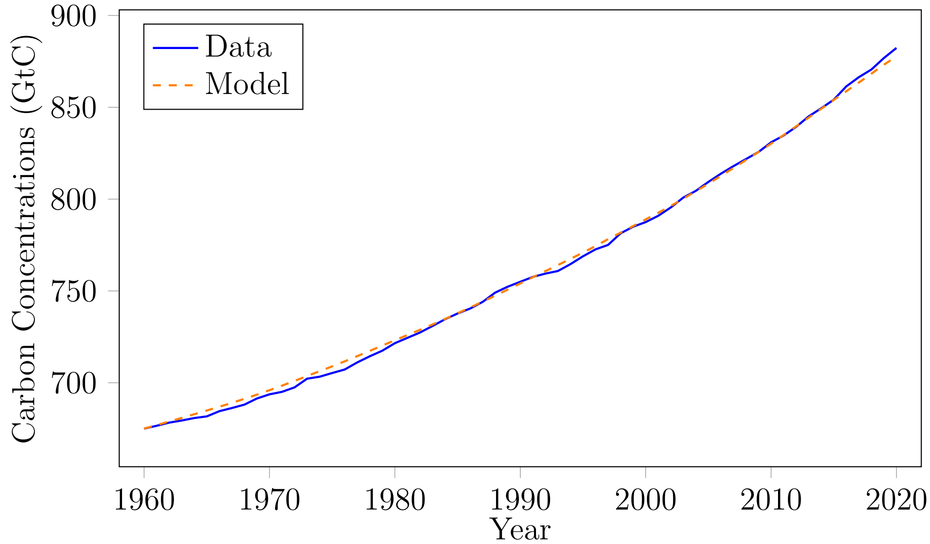

Notes: Match of model predicted atmospheric carbon concentrations with data when $\psi_p = 0.2$, $\psi_0 = 0.377$, and $\psi = 0.994$. Carbon emissions come from Our World in Data, and atmospheric carbon concentrations come from the NOAA's Mauna Loa Observatory.

to atmospheric carbon concentrations is constant: $\frac{\partial \ln(\mathcal{Y}_t)}{\partial \mathcal{C}_t} = \varrho$. Assuming a constant semi-elasticity of damages with respect to atmospheric carbon concentrations does not allow for severe non-linearities in the climate system, e.g. tipping points, but this functional form for damages is consistent with the approach taken in much of the climate literature (Nordhaus, 1992, 2017; Stern, 2007; Krusell and Smith, 2022; Barrage and Nordhaus, 2024). In my choice of ϱ, I take the same value as Golosov et al. (2014) and set $\varrho = 5.3 \times 10^{-5}$. However, given the considerable uncertainties surrounding the damage that will result from climate change, I will also consider an alternative damages parameter that is four times as large, roughly in line with the "catastrophic damages" scenario considered by Golosov et al. (2014).[37]

degrees Celsius. The standard value for Γ is 3, but there remains considerable uncertainty over the value of Γ. The IPCC deems 1.5-4.5 to be the "likely" range of Γ, but this leaves a very real chance of much higher values (Wagner and Weitzman, 2016).

[37] I abstract from the consideration of risk in optimal climate policy. There are substantial uncertainties regarding the climate system's response to emissions, as well as the sensitivity of human

1.5.3 Structural Parameters & Initial Conditions

In this section, I describe my calibration strategy for the remaining structural parameters of my model $\{\sigma, \lambda, \{\nu_\theta, \omega_{\theta d}\}, L, \mathcal{S}, \gamma, \chi, \alpha, \{r_j\}, \eta, \vartheta, \rho\}$ as well as the initial condition for technology $\{A_{jt_0}\}$. In Appendix A.2.1, I provide details explaining how the equilibrium conditions of my model, and therefore the calibration moments discussed below, can be written in terms of just technology, policy instruments, and structural parameters. Table 1.2 provides a summary of my parameter choices.

Table 1.2: Parameter Choice Summary

Parameter	Description	Value	Source
σ	Clean/Dirty ES	1.86	Papageorgiou et al. (2017)
λ	Cross-Sector ES	0.1	Hassler et al. (2021)
ν_{car} ν_{elec}	CES Shares	0.029 0.025	See Text
γ	Innovation Step Size	1.07	Acemoglu et al. (2023)
α	Input Share	0.4	Standard Barrage (2020)
$r_{\theta d}/r_{\theta c}$	Relative Dirty Input Price	2.25	BP (2022)
η	Research Elasticity	0.5	See Text
ϑ	Inverse Intertemporal ES	1	Standard
ρ	Rate of Pure Time Preference	0.001 0.015	Stern (2007) Nordhaus (2017)

For the within-sector elasticity of substitution between clean and dirty forms of production, I set $\sigma = 1.86$ in line with the average estimate from Papageorgiou et al. (2017).[38] The evidence for this estimate comes from the electricity generation sector,

civilization to changes in the Earth's climate, so a large literature explicitly considers the role of risk and insurance in determining the social cost of carbon (Weitzman, 2009; Lemoine and Traeger, 2014; Gillingham et al., 2018; Cai and Lontzek, 2019). A thorough consideration of climate uncertainty is beyond the scope of this paper, but I will capture some elements of the insurance motive for climate policy by considering the possibility of catastrophic damages as an extension.

[38]See Table 3 of Papageorgiou et al. (2017).

but I set it as the value for the transportation sector as well in part because it is of a similar magnitude to the estimate of Lanzi and Sue Wing (2011) who look at the energy sector more broadly. For the cross-sector elasticity of substitution, I set $\lambda = 0.1$ in line with the evidence from Hassler et al. (2021) who find that energy is a near perfect complement with capital/labor at annual frequencies. Across the literature, it is common to model energy as a complement to other factors (Van der Werf, 2008; Böhringer and Rutherford, 2009). For instance, Fried (2018) selects a near-zero elasticity of substitution between energy and other factors, making production almost Leontief.

For the sectoral CES shares, I normalize the sum of the shares to one and pick $\{\nu_{car}, \nu_{elec}\}$ so that the income shares of transportation and electricity generation in the laissez-faire steady-state match the average values of those shares in the data from 2000 to 2020. Data on revenue for the electricity generation sector comes from the US Energy Information Agency (EIA).[39] Data on motor vehicle output and GDP come from the US Bureau of Economic Analysis (BEA). This approach yields $\nu_{car} = 0.029$ and $\nu_{elec} = 0.025$ to match average income shares of 2.9% and 2.3%, respectively.

I normalize the supply of scientists and workers to $\mathcal{S} = 1$ and $L = 100$, respectively. These are both without loss normalizations, but they match the fraction of the labor force that works in R&D in the US (Jones and Vollrath, 2013). Following Acemoglu et al. (2023), I set the step size of innovation to $\gamma = 1.07$ to match the profit share of the petroleum and coal products, durable manufacturing, and wholesale trade sectors.[40] This is a similar value to those found in other models of step-ladder innovation (Acemoglu et al., 2016; Akcigit and Kerr, 2018). I set research productivity χ so that growth in the laissez-faire steady-state is 2% per year.

For the income share of inputs, I set $\alpha = 0.4$ to match estimates of the labor

[39] Specifically, Forms EIA-826, EIA-861, and EIA-861M.

[40] Acemoglu et al. (2023) match the weighted average profit share for these sectors from 2004 to 2014 using the Quarterly Financial Report from the US Census Bureau.

share. The standard economy-wide value for the labor share is 0.67, but Barrage (2020) estimates a labor share in electricity and resource production of 0.403. Given my interest in the latter, I set my labor share $1-\alpha$ as a weighted average of the two, with about one quarter weight on the lower estimate from electricity and resource production. For input costs, I normalize all of the clean input costs to one: $r_{\theta c}=1$. I then set dirty input costs to match the relatively low efficiency of fossil fuel-based machinery at converting primary energy into useful energy. To compare the energy conversion efficiency of fossil fuels with non-fossil sources of energy, the BP Statistical Review of World Energy report assumes a thermal equivalent efficiency factor that rises linearly from 36% in 2000 to 45% in 2050 (BP, 2022). That is, in 2050 45 TWh of electricity produced from solar panels will require the burning of 100 TWh worth of coal, down from 125 TWh in 2000.[41] Therefore, I set the dirty input costs to $r_{\theta d}=2.25\approx 1/0.45$.

For my choice of the elasticity of innovation with respect to scientists η, I rely on a body of empirical evidence from the innovation literature. One body of studies considers the elasticity of patents with respect to R&D expenditures and generally finds a value of about 0.5 (Griliches, 1990; Hall and Ziedonis, 2001; Blundell et al., 2002). Another considers the elasticity of R&D expenditure with respect to the tax price of research and generally finds an elasticity of about unity (Hall, 1993; Hall and Van Reenen, 2000; Bloom et al., 2002; Wilson, 2009). As I discuss in Appendix A.2.1, both sets of findings correspond to a value of $\eta=0.5$ in my model. Similar conclusions were reached by Akcigit and Kerr (2018), Acemoglu et al. (2018), and Bloom et al. (2021).

[41] This efficiency factor is consistent with the US Department of Energy estimate that electric vehicles convert more than 77% of the electricity they pull from the grid into kinetic energy, while gasoline-powered vehicles can only convert 12-30% of the chemical energy in gasoline into kinetic energy. Given that gasoline-powered vehicles are a more mature technology, this superior conversion rate of electric vehicles likely reflects an intrinsic difference between the two technologies.

For the preference parameters, I set the inverse intertemporal elasticity of substitution ϑ equal to the standard value of unity, which implies a logarithmic period utility function. For the rate of pure time preference ρ, I consider both the Stern rate of 0.1% per year and the Nordhaus rate of 1.5% per year (Stern, 2007; Nordhaus, 2017). The former is the standard value for those who believe ethical considerations should drive the choice of discount rate, while the latter is the standard value for those who think the discount rate should match the rate of return on capital. There has been substantial debate in the climate literature over the appropriate choice of social discount rate (Stern, 2007; Nordhaus, 2007; Dasgupta, 2008; Barrage, 2018).[42] This debate is not the focus of this paper, but given the long-term impact of both climate and innovation policy, it is important to understand how the choice of social discount rate shapes the optimal path of climate innovation policy.

For the initial condition of technology at time t_0, I first set within-sector relative technology $\{A_{\theta c t_0}/A_{\theta d t_0}\}$ to match the clean quantity share of each sector. I target clean quantity shares, as opposed to income shares, because of data availability, but as I show in Appendix A.2.1, there is a straightforward mapping between quantity shares and income shares. Data for quantity shares in the transportation sector comes from the US Department of Energy's Transportation Energy Data Book which lists the share of new light vehicles in the US that were hybrids, plug-in hybrids, and EVs from 1999 to 2021 (Davis and Boundy, 2022).[43] Data on the share of electricity generated from non-fossil sources comes from the EIA.[44] Next, I set cross-sector relative technology $\{A_{\theta d t_0}/A_{\Theta c t_0}\}$ to match the income share of each sector.

[42]It is worth noting that moral philosophers universally reject the idea that utility should be discounted simply because it occurs at a later date (Parfit, 1984; Ord, 2020). The only justification for discounting future utility that is widely accepted amongst philosophers is the risk of extinction, so it is acceptable to discount future utility in accordance with the probability that it may not happen. Indeed, this is the justification given by Stern (2007) for his choice of discount rate.

[43]In particular, see Table 6.02 of the Transportation Energy Data Book.

[44]Specifically, Form EIA-923.

In both cases, I include the clean input subsidies discussed below. Note that I am using data on sectoral income shares for my calibration of both the CES shares and the initial condition for technology. In the case of the CES shares, I am matching the average empirical income shares from 2000 to 2020 when technology is in steady-state, which guarantees that the income shares of these two small sectors will remain plausible as technology evolves in my simulations. In the case of the initial condition of technology, I am matching the empirical income shares at a given point in time to back out appropriate values for technology. Finally, I pin down the absolute level of technology by normalizing output to 100.

For my carbon intensity parameters $\{\omega_{\theta d}\}$, I match the emissions of transportation and electricity generation in 2021. Data on US GHG emissions by sector comes from the US Environmental Protection Agency (EPA). I find that the carbon intensity of electricity generation is about 65.2% larger than that of transportation. My optimal policy simulations track emissions from the US transportation and electricity generation sectors, but given that climate change is a global phenomenon, I need to select a sequence of emissions for the remainder of the world economy. To this end, I set outside emissions equal to the optimal emissions path of the 2010 RICE model, subtracting a little more than half of US emissions (Nordhaus, 2010).[45]

1.5.4 Model Validation

In this section, I test my model's ability to replicate the dynamics of clean technology from 2010 to 2021. During this period, clean technology in both transportation and electricity generation made substantial gains, leading to increases in the clean quantity share in both sectors.[46] As shown in Figure 1·4, the percentage of new light

[45]To be exact, I subtract 56.9% of US emissions as this is the average contribution of transportation and electricity generation to US emissions from 2000 to 2020.

[46]Solar photovoltaics and onshore wind are now cheaper than fossil fuels in terms of the private lifetime cost of new generation (Roser, 2020). These two forms of renewable electricity generation saw price declines of 89% and 70% from 2009 to 2019, respectively.

vehicles that were hybrid or electric rose from 2.4% to 9.8%, while the percentage of electricity generated from non-fossil sources rose from 30.1% to 39%. I show that my model can match these clean advances in both sectors, in spite of the fact that no model parameters was selected to target this change.[47] However, when I shut down the spillover network, my model predicts path dependence, implying that the 2010s would have seen dirty technology gain an increasing lead in both sectors.[48] Thus, the nascent clean energy transition of the 2010s is difficult to reconcile with endogenous growth models that exclude cross-technology knowledge spillovers.

To argue that technological path dependence is inconsistent with the 2010s advance in clean technology, it must be the case that US climate policy was too weak to push these sectors out of dirty basins of attraction. To account for US climate policy during this time, I allow for subsidies to both innovation and clean inputs. Further details on the selection of these subsidies are in Appendix A.2.2. For the clean input subsidies, the clean form of production in sector θ receives an input subsidy of $\bar{\xi}_{\theta c}$, so the effective input price becomes $(1-\bar{\xi}_{\theta c})r_{\theta c}$. For the transportation sector, I set $\bar{\xi}_{car,c}$ to match federal spending on the plug-in electric vehicle tax credit as a proportion of total spending in the transportation sector.[49] This yields a subsidy of $\bar{\xi}_{car,c} = 0.012$. For the electricity generation sector, I set $\bar{\xi}_{elec,c} = 0.3$ in line with the 30% investment tax credit given to renewable electricity generation during this period.

In selecting innovation subsidies, I focus on relative innovation subsidies because the fixed supply of scientists implies that only relative innovation subsidy matter, as one can see in the research equilibrium condition (1.25). That is, subsidies that apply to innovation across the board, such as the R&D tax credit, will not affect the

[47]The match of the model to the data over this period did inform my selection of the number of years in a time period, as I discuss at the start of Section 1.5.

[48]By shutting down the spillover network, I mean setting the spillover network to a matrix of zeros: $\boldsymbol{\varphi} = \mathbf{0}$. This is equivalent to setting the gross spillover matrix to the identity: $\tilde{\boldsymbol{\varphi}} = \mathbf{I}$.

[49]Data on federal spending on the plug-in electric vehicle tax credit during this period comes from US Congressional Research Service Report IF11017.

composition of scientists across technologies. To this end, I take data on public R&D spending by technology from the International Energy Agency (IEA). I then select innovation subsidies to match this spending as a proportion of total R&D spending, which I take from the BEA. This procedure involves simulating the equilibrium path of technology, so I conduct a separate calibration of innovation subsidies for the version of the model where I shut down the spillover network.

Using the procedure described in Section 1.5.3, I take 2010 as the start year to set my initial condition for technology. I then simulate the equilibrium path of technology with and without the spillover network and plot the implied dynamics of clean quantity shares in Figure 1·4.

Figure 1·4: Model with Spillovers Matches 2010s Advance of Clean Technology

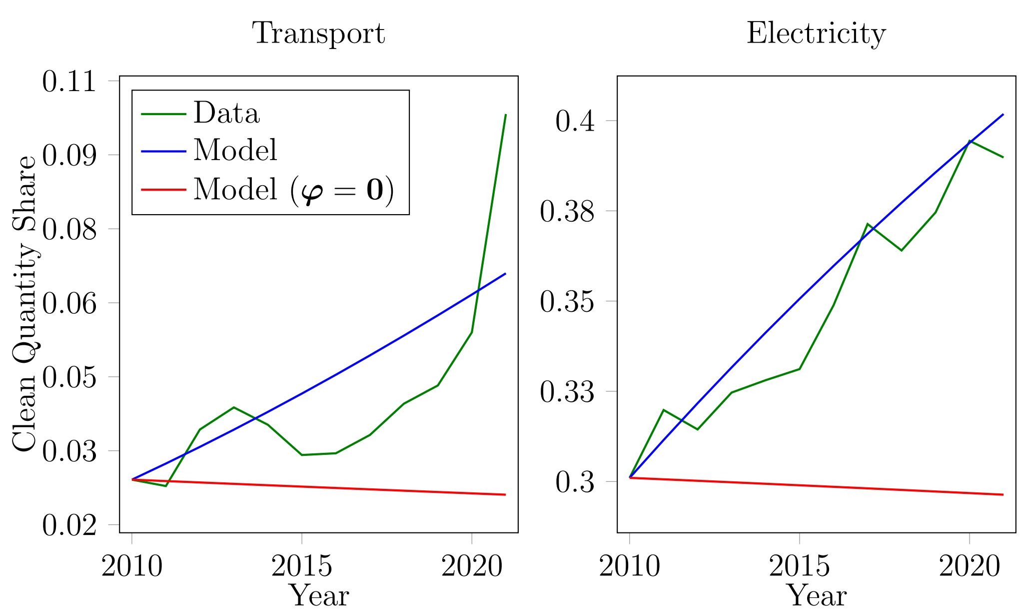

Notes: For transportation, the clean quantity share is the proportion of new light vehicles that are hybrid or electric, whereas for electricity generation, it is the proportion of electricity generated from non-fossil sources. Data for these two series come from Davis and Boundy (2022) and the EIA.

Comparing the simulations with the data, one can see that the model with spillovers broadly matches, as an untargeted moment, the empirical evolution of clean quantity shares in both sectors. This is especially true in the electricity generation sector. The simulated clean quantity shares start at the same value as the data in 2010 by construction, and by 2021, transportation and electricity generation reach clean quantity shares of 6.6% and 40.2% in the simulation, respectively. By comparison, these sectors had achieved clean quantity shares of 9.8% and 39% by 2021 in the data. When I shut down the spillover network, the model makes the counterfactual prediction that clean technology would have fallen even further behind during this period. Without cross-technology spillovers, the clean quantity shares of transportation and electricity generation would have fallen to 2.1% and 29.6% by 2021. Therefore, a model without cross-technology spillovers cannot match the qualitative direction of the data.

Using Propositions 1.2 and 1.3, we can better understand the role of cross-technology spillovers in shaping the path of clean technology during this period. When the model includes the spillover network, the spectral radius of the transition matrix is equal to 0.993, indicating that technology is converging toward its steady-state. The model predicts that starting in 2010, transportation and electricity generation will converge halfway to their steady-states in 79 and 100 years, respectively. However, when I shutdown the spillover network, the spectral radius rises to 1.01, inducing path dependence in accordance with Corollary 1.1. In that case, the half-life of convergence of each sector becomes infinite as the economy diverges from its interior steady-state and moves deeper into the dirty basins of attraction. Thus, the model's ability to explain the nascent clean energy transition of the 2010s hinges on the inclusion of knowledge spillovers.

1.6 Quantitative Exercises

In this section, I conduct three quantitative exercises. First, I examine the impact of introducing a carbon price and clean innovation subsidy. I do so for several levels of cross-technology spillovers to demonstrate their importance in determining the impact of the policy reform, both in the long-run and transition. To make use of the theoretical results in Section 1.3 at the level of the sector, I define sector-specific relative clean technology $\bar{B}_{\theta t} \equiv A_{\theta ct}/A_{\theta dt}$, which is log-linear in relative technology $\{\bar{A}_{jt}\}$.[50] To quantify convergence speeds, I will also describe half-lives of convergence at the level of the sector $t_\theta^{(1/2)}$.

Second, I simulate first-best climate innovation policy. This exercise highlights the distinction between the spillovers a technology *receives* and the spillovers it *sends*. Clean technologies are not particularly central in the spillover network, so they do not receive favorable innovation subsidies. Instead, clean technologies are beneficial for their ability to produce physical goods without pollution, but this attribute is already favored by the carbon price.

Finally, I simulate second-best innovation subsidies, where the price on carbon pollution is incomplete. I show that a small, growing carbon price can recover most of the welfare gains of the first-best. However, if pollution is unpriced, then the Planner must decarbonize with clean innovation alone. This leads to slow emission reductions, due to a rebound effect of clean innovation, and slow economic growth, due to the loss of spillovers from dirty technology. Together, these result in massive welfare losses relative to the first-best.

1.6.1 Impact of Policy Reform

In this section, I examine the impact of introducing a policy reform. I consider a \$51 carbon price, the Biden Administration's current estimate of the social cost of

[50] To make use of Propositions 1.2 and 1.3 at the level of the sector, I can define

$$\bar{\mathcal{B}}_{\theta t} \equiv \ln(\bar{B}_{\theta t}) - \ln(\bar{B}_{\theta,ss}) = \bar{\mathcal{A}}_{\theta ct} - \bar{\mathcal{A}}_{\theta dt}. \tag{1.74}$$

Thus, there is a linear mapping from $\bar{\mathcal{A}}_t$ to $\bar{\mathcal{B}}_t$ which allows me to examine transition dynamics at the level of the sector.

carbon, as well as a uniform clean innovation subsidy equivalent to a 30% tax credit.[51] I pick the clean innovation subsidy to be quantitatively similar to the investment tax credits in the Inflation Reduction Act, but it should be noted that tax credits for production have different effects from tax credits for innovation. Consistent with my theoretical analysis, the policy reform's impact, in both the long-run and transition, depends critically on the level of cross-technology spillovers through its mediation of increasing returns to innovation.

I will assume the policy reform begins after research decisions are sunk in 2021, and for the sake of simplicity, I will make comparisons with laissez-faire. Table 1.3 summarizes the impact of the policy reform for three cases that differ in their level of cross-technology spillovers. First, I shut down the spillover network. Second, I take my calibrated spillover network. Third, I double the level of cross-technology spillovers in the economy. These three cases follow from a single parameterization which scales the spillover network by ζ. Variation in ζ controls the level of cross-technology spillovers by multiplying the off-diagonal components of the spillover network by factor ζ.[52] For example, the double spillovers case sets $\zeta = 2$, doubling the off-diagonal row sums of the spillover network. The third panel of Table 1.3 summarizes the difference between the three cases in terms of their degree of increasing returns to innovation, as measured by the spectral radius of the transition matrix.

The first panel of Table 1.3 considers the long-run impacts of the policy reform. The first two rows make use of Proposition 1.1 and show the first-order change in steady-state relative clean technology by sector. With calibrated spillovers, we have an increase of 117.32% for transportation and 120.8% for electricity generation.[53]

[51] This implies an innovation subsidy for clean transportation and electricity generation of $\xi_c = 1/(1-0.3) \approx 1.43$.

[52] One can also think in terms of the spillover matrix $\mathbf{\Phi}$, which measures the level of cross-technology spillovers in the economy. Multiplying the spillover network by ζ does the same for the spillover network, giving us $\zeta\mathbf{\Phi}$.

[53] Formally, these changes in relative technology are listed in log points.

Together, the carbon price and induced increase in clean technology increase within-sector clean income shares and reduce the economy's emissions intensity in steady-state.

Table 1.3: Impact of Policy Reform

	No Spillovers	Calibrated Spillovers	Double Spillovers
Long-Run Impacts			
Relative Clean Technology by Sector			
$\%\Delta\bar{B}_{car}$	0%	+117.32%	+33.87%
$\%\Delta\bar{B}_{elec}$	$+\infty$%	+120.82%	+38.66%
Clean Income Shares by Sector			
ΔS_c^{car}	0 pp	+12.54 pp	+7.44 pp
ΔS_c^{elec}	+100 pp	+14.88 pp	+9.54 pp
Emissions Intensity			
$\%\Delta\bar{\omega}$	-74.44%	-74.55%	-57.41%
Transitional Impacts			
Half-Lives of Convergence by Sector			
$t_{car}^{(1/2)}$	–	123 years	20 years
$t_{elec}^{(1/2)}$	–	128 years	25 years
Carbon Emissions by Year			
$\%\Delta\mathcal{E}_{2035}$	-43.7%	-50.8%	-53.49%
$\%\Delta\mathcal{E}_{2060}$	-46.55%	-57.01%	-55.23%
Degree of Increasing Returns to Innovation			
Spectral Radius			
$\max\lvert\kappa_j\rvert$	1.01	0.994	0.977

Notes: Impact of introducing a carbon price at the Biden Administration's estimate of the SCC ($51) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Changes in relative technology are listed in log points. For path dependent economies, long-run impacts refer to corner, rather than interior, steady-states.

The other two cases illustrate the role of cross-technology spillovers in determining the long-run impact of policy reforms. When the spillover network is shut down, the spectral radius of the transition matrix increases from 0.994 to 1.01, inducing path dependence. In that case, policy can only have a long-run impact on technology if it pushes the economy into one of its clean basins of attraction, but when it does so, the long-run impact is transformative. Table 1.3 shows the policy reform is strong enough to do so for electricity generation, but not for transportation.[54] This difference stems

[54] As discussed in Corollary 1.2, in the case of path dependence, policy reforms push the economy

from the fact that clean technology starts out with a greater disadvantage, relative to its dirty rival, in transportation. Conversely, when cross-technology spillovers are doubled, the spectral radius of the transition matrix decreases from 0.994 to 0.977, substantially reducing increasing returns to innovation. Now, the catchup growth generated by cross-technology spillovers reduces the long-run impact of policy because clean technologies cannot gain a large lead without enhancing research productivity in dirty technologies. Thus, the long-run change in relative technology for both sectors is reduced substantially, along with the change in clean income shares and emissions intensity.

The second panel of Table 1.3 considers the transitional impacts of the policy reform. Using Proposition 1.2, I can characterize the speed of transition to the new steady-state.[55] For my calibrated spillover network, increasing returns to innovation are strong enough to create a slow transition but are not so strong that they generate path dependence. Following the policy reform, transportation and electricity generation converge halfway to their steady-states in 123 and 128 years, respectively. The policy reform's impact on both input prices and the path of technology reduces emissions by about half in 2035 and 2060, relative to the laissez-faire level in those years.[56]

As before, the other two cases illustrate how cross-technology spillovers shape the speed with which policy can influence the direction of innovation. When the spillover network is shut down, the economy doesn't converge to an interior steady-state, mak-

into clean basins of attraction by expanding those clean basins of attraction to include more initial conditions. Indeed, the policy reform of Table 1.3 *reduces* relative clean technology in the interior steady-state by 101.27 and 116.1 log points for transportation and electricity generation, respectively, thereby increasing the set of initial conditions that lead to clean growth. This expansion includes the initial condition for electricity generation but is too small to include that of transportation.

[55] Given that these results involve linearizing my model, one may be concerned that my results are not informative about such a large policy reform. To address this concern, I compare the linearized transition path to the full simulation in Figure A·3 and show that the two do not differ significantly.

[56] The emission reductions presented here abstract from differences in climate damages brought about by the different emission paths.

ing the half-lives of convergence undefined. Moreover, the long-run emission reductions from pushing electricity generation into its clean basin of attraction are achieved only slowly because clean electricity cannot enjoy catchup growth. This is reflected in the weak emission reductions in the transition. Conversely, when cross-technology spillovers are doubled, the transition speeds up substantially, with comparatively quick half-lives of convergence of 20 years and 25 years. In the near term, the more rapid technological transition allows for greater emission reductions, but this lead is eventually lost as high cross-technology spillovers allow dirty technologies to stay relatively advanced.

Figure 1·5: Technology Path

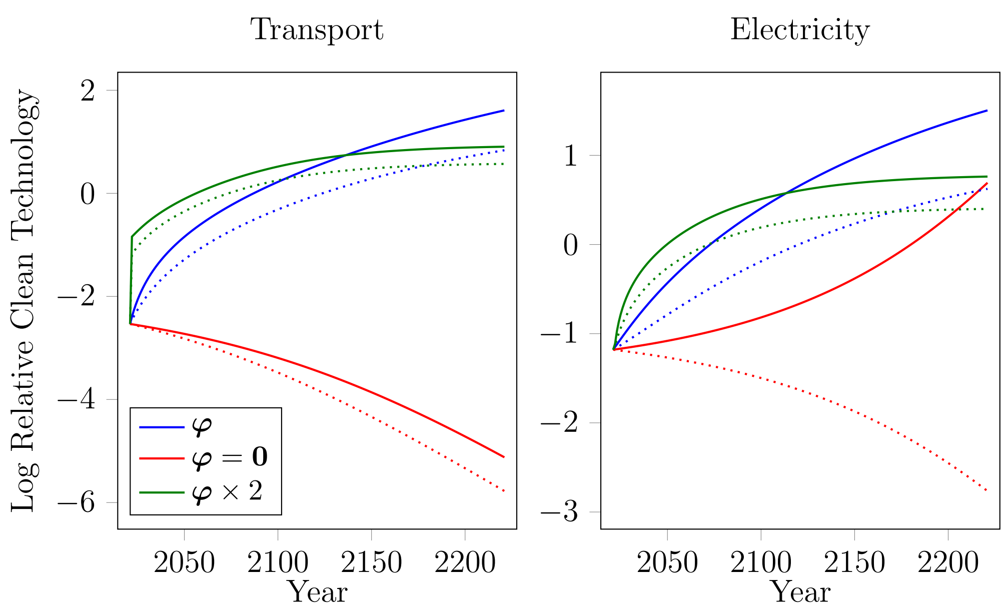

Notes: Impact of introducing a carbon price at the Biden Administration's estimate of the SCC ($51) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Dotted lines indicate laissez-faire paths.

Figures 1·5 and 1·6 provide further detail on the impact of the policy reform. The dotted lines indicate the paths that would have been taken in the absence of reform. With calibrated spillovers, the policy reform allows clean technologies to slowly gain a lead over their dirty counterparts as the economy transitions to a new, cleaner steady-state. This induced improvement in clean technology, as well as the carbon price, leads to reductions in total emissions over time, relative to laissez-faire.[57] Figure A·4 shows the impact of the policy reform on clean income shares over time, which follows a similar pattern as relative clean technology.

When the spillover network is shut down, the policy reform enables a slow shift toward clean technology for electricity generation while only delaying the advance of transportation into its dirty basin of attraction. The high increasing returns to innovation, to the point of path dependence, bifurcate the impact of policy. For large enough policy reforms, like in the case of electricity generation, the shift in the direction of innovation will be slow, but in the long-run, the impact will be massive. This contrasts sharply with the case where cross-technology spillovers are doubled. High spillovers allow for a rapid response to the policy reform, but the overall transition eventually peters out due to the small long-run change.

These differences showcase how both the long-run and transitional impacts of a policy reform are shaped by cross-technology spillovers. High cross-technology spillovers allow policy to shape the direction of innovation quickly by allowing less advanced technologies to achieve catchup growth. However, the same mechanism implies smaller long-run effects of policy as any advantage granted to some technologies will trickle down to others via spillovers. These effects are summarized by the degree of increasing returns to innovation, as measured by the spectral radius of the transition matrix. Low increasing returns allow for rapid transitions in response to policy reforms, but as argued in Corollary 1.2, the forces that prevent increasing returns also reduce the long-run impact of policy by preventing any technology from gaining a significant advantage.

[57] Absolute emissions still rise steadily in the long-run because a fixed carbon price is eventually overcome by economic growth.

This section has focused on cross-technology spillovers as a source of variation in increasing returns to innovation, but similar results hold for substitution patterns in production. In Figure A·5, I show how the steady-state impact of the policy reform varies in the level of cross-technology spillovers ζ and the elasticity of substitution between clean and dirty goods σ. As I have argued, the steady-state impact reduces as ζ increases. Similarly, the steady-state impact increases with the elasticity of substitution σ because higher substitutability allows clean goods to gain greater market share when favored by policy. In both cases, higher increasing returns to innovation increase the steady-state impact of policy while at the same time increasing the time needed for policy to take effect by slowing the transition. Figure A·6 illustrates how the degree of increasing returns to innovation varies in ζ and σ. As I have argued throughout this paper, cross-technology spillovers reduce increasing returns, with technology becoming path dependent if cross-technology spillovers are reduced to 66.1% of their calibrated value. Similarly, substitutability raises increasing returns, with technology becoming path dependent if the elasticity of substitution σ rises to 2.29.

The EPA has recently proposed increasing the federal SCC to \$190. I examine the impact of introducing a carbon price of \$190 in Table A.2 as well as Figures A·7, A·8, and A·9. The results are qualitatively similar but quantitatively larger. In fact, a \$190 carbon price is still not large enough to shift transportation into its clean basin of attraction when the spillover network is shut down. Figures A·10 and A·11 show the carbon prices and clean innovation subsidies necessary to shift each technology into its clean basin of attraction when the spillover network is shut down. With a clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$), the electricity generation sector would shift to clean innovation without a carbon price, while the transportation sector would require a carbon price of \$248. With a carbon price

Figure 1·6: Pollution Path

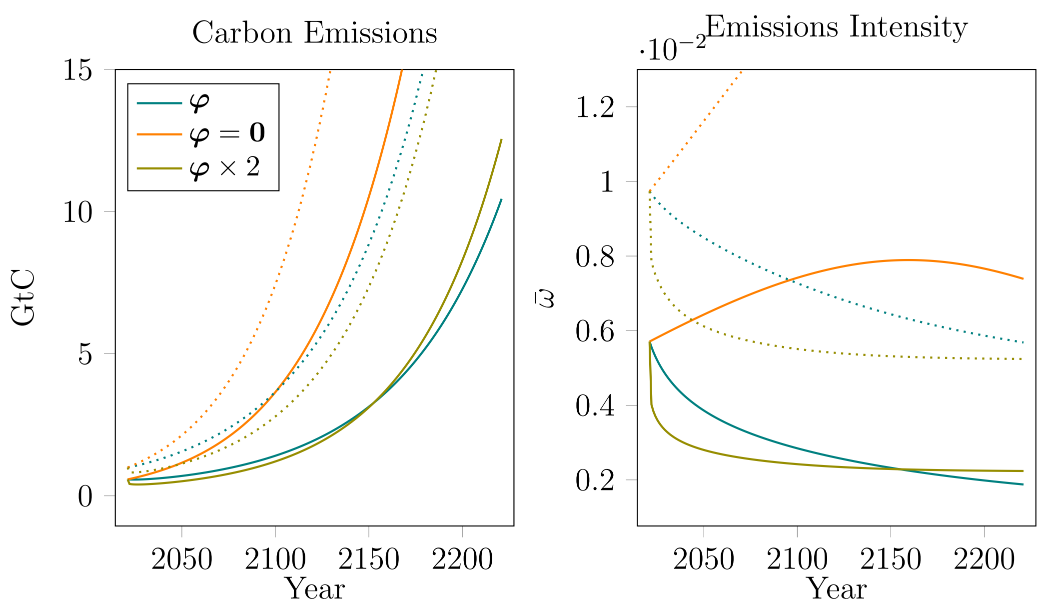

Notes: Impact of introducing a carbon price at the Biden Administration's estimate of the SCC ($51) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Dotted lines indicate laissez-faire paths.

at the Biden Administration's estimate of the SCC ($51), the electricity generation sector would shift to clean innovation with a clean innovation subsidy of $\xi_c = 1.03$, while the transportation sector would require an innovation subsidy of $\xi_c = 1.62$.[58]

Figure 1·7 shows the half-lives of convergence for each eigenstate and sector of the economy.[59] Proposition 1.3 allows me to unpack the substitution patterns and connections in the spillover network that drive each sector's convergence speed. Each eigenstate represents a region of state space with a distinct convergence speed, and Figure 1·7 shows that the slowest eigenstate drives the speed of convergence for both sectors. That is, technology's initial condition loads primarily on the eigenstate representing the region of state space where convergence is especially slow. This eigenstate represents initial conditions where dirty technology has an economy-wide advantage

[58]These innovation subsidies are equivalent to tax credits of 2.9% and 38.3%, respectively.

[59]I perform the same exercise for the case where cross-technology spillovers are doubled in Figure A·12.

Figure 1·7: Half-Lives of Convergence

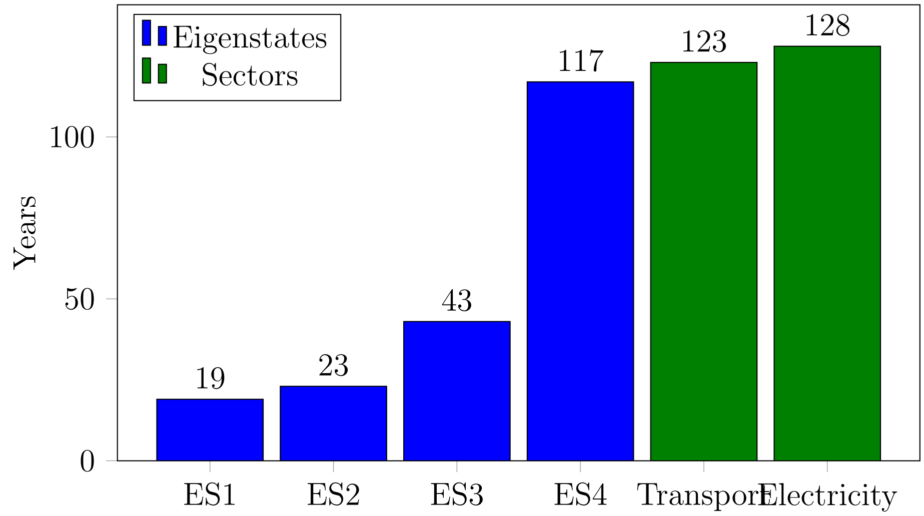

Notes: Transition speeds following the introduction of a carbon price at the Biden Administration's estimate of the SCC ($51) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$).

over clean technology; exactly the initial condition that policymakers must contend with.

What about this initial condition leads to slow convergence? First, the substitutability of clean and dirty goods makes switching the predominant technology in any sector slow as the small market size of the laggard will reduce innovation rents. However, this is not the whole story because the third eigenstate, which is substantially faster, represents initial conditions where dirty technology is advanced in one sector and clean technology is advanced in another. Moving from dirty to clean technology in both sectors concurrently is especially slow because, as argued in Section 1.5.1, clean and dirty technologies both form their own spillover clusters. Thus, catchup growth for clean technologies is reduced because clean technologies receive spillovers from their peers who are themselves trying to catch up. In spite of the forces that generate a slow transition, both sectors avoid path dependence. Overall, cross-technology knowledge spillovers are large enough to stimulate sufficient catchup growth for clean technologies. In particular, both clean technologies receive high spillovers from outside of the transportation and electricity generation sectors. Without those spillovers, the spectral radius would increase to 1.007, inducing path dependence.[60]

Another noteworthy feature of Figure 1·7 is that transportation converges a bit more quickly than electricity generation, with electricity generation taking about 4.1% longer to halve its log distance from steady-state. This difference is due to the higher within-sector spillovers in transportation. Another feature of the spillover network visible in Figure 1·2 is that clean transport receives spillovers from dirty transport, while there are almost no spillovers from dirty electricity to clean electricity. This is consistent with the spillover anecdotes from my introduction, where electric vehicles receive spillovers from dirty vehicles, but solar panels receive spillovers from the ICT sector. To quantify the importance of within-sector spillovers, Figure A·13 shows half-lives of convergence when within-sector spillovers are shut down. Without these spillovers, the spectral radius rises to 0.998, and the half-lives for transportation and electricity generation rise to 449 years and 390 years, respectively. The loss of within-sector spillovers substantially slows the overall speed of convergence, but more importantly, the sectors switch their convergence speed rankings, signifying that transportation's more rapid transition stems from within-sector spillovers.

Finally, one may wonder whether the degree of increasing returns to innovation depends on the size of the policy reform. Figure A·14 shows that the eigenvalues of the transition matrix, and therefore the spectral radius, are not sensitive to the size of the policy reform. Thus, variation in transition speed across different policy reforms comes from differences in the loading of the initial condition on the eigenstates β, rather than the size of the eigenvalues $\{\kappa_j\}$ themselves.

[60]To shut down spillovers from the general sector, I transfer all of the general spillover elasticities to the diagonal of the gross spillover network $\tilde{\boldsymbol{\varphi}}$. I take a similar approach to shutdown within-sector spillovers, transferring all of the within-sector spillover elasticities to the diagonal.

1.6.2 First-Best Policy

In this section, I simulate the policy path that generates an optimal clean transition, starting in 2022.[61] As I have emphasized, optimal policy considers the spillovers a technology sends, rather than receives, because they create a benefit that innovators fail to internalize. Clean technologies create relatively few spillovers, so for this reason, clean innovation subsidies play a small role in the clean transition. Instead, carbon emissions are brought down primarily through pricing pollution.

Figure 1·8: First-Best Policy Path

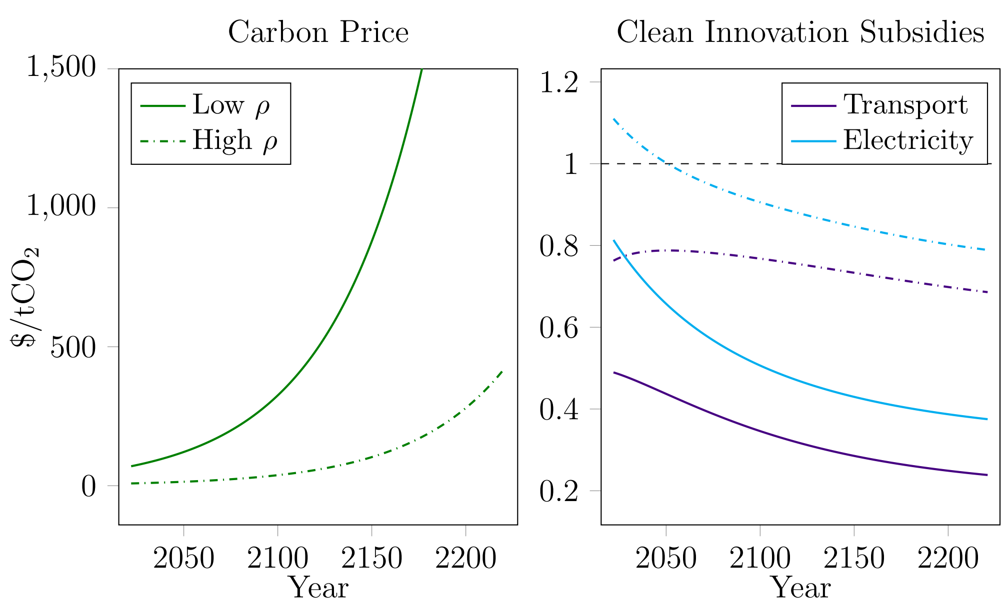

Notes: Innovation subsidies are listed as a fraction of the baseline innovation wedge.

Figure 1·8 shows the first-best policy path for two values of the discount rate. The left panel displays the optimal carbon price. As expected, a lower discount rate increases the carbon price by increasing the importance of future climate damages. The difference is substantial, with the less patient Planner taking more than a century to reach the more patient Planner's initial carbon price.

[61] I provide further details on my solution method in Appendix A.2.3.

The right panel displays clean innovation subsidies as a fraction of the baseline innovation wedge. Strikingly, clean technologies do not receive favorable innovation subsidies. The only case where a clean technology receives an innovation subsidy above the baseline innovation wedge is clean electricity for the first few decades of policy when the discount rate is high. As shown in Proposition 1.4, the primary function of innovation subsidies is to correct the positive knowledge spillover externality. Clean technologies are valued for their ability to produce goods without pollution, not necessarily for their ability to produce spillovers. On the latter metric, they don't score particularly well. Clean transport's eigenvector centrality in the spillover network is barely more than a tenth of its steady-state income share (0.64 vs. 5.19%), while clean electricity's eigenvector centrality is slightly less than a quarter of its steady-state income share (0.91 vs. 3.7%).[62] This is why, in both cases, the innovation subsidies for clean electricity are uniformly higher than those for clean transport; the former is the superior clean technology for producing spillovers. Alternatively, dirty transport and electricity have eigenvector centralities of 1.14 and 0.83, respectively, relative to steady-state income shares of zero.

Moreover, a lower discount rate *reduces* clean innovation subsidies because, as shown in Corollary 1.3, a more patient Planner places more weight on spillover creation, rewarding technologies more for their ability to produce ideas via spillovers, rather than their ability to produce physical goods. The more patient Planner is more concerned about future climate damages, but this shows up in their higher carbon price. With the social cost of carbon properly priced, there is no reason for them to give additional support to clean technologies as such. For both values of the discount rate, clean innovation subsidies decline toward their steady-state value over time. This slope stems from the fact that clean income shares increase to their steady-state

[62] I have scaled eigenvector centrality to sum to 100, making it comparable to a percent. The steady-state income shares I have listed come from the low discounting allocation.

value over time. Thus, the private reward for clean innovation in early periods is low relative to the forward-looking social value, warranting a front-loaded path for clean innovation subsidies.

Figure 1·9: First-Best Technology Path

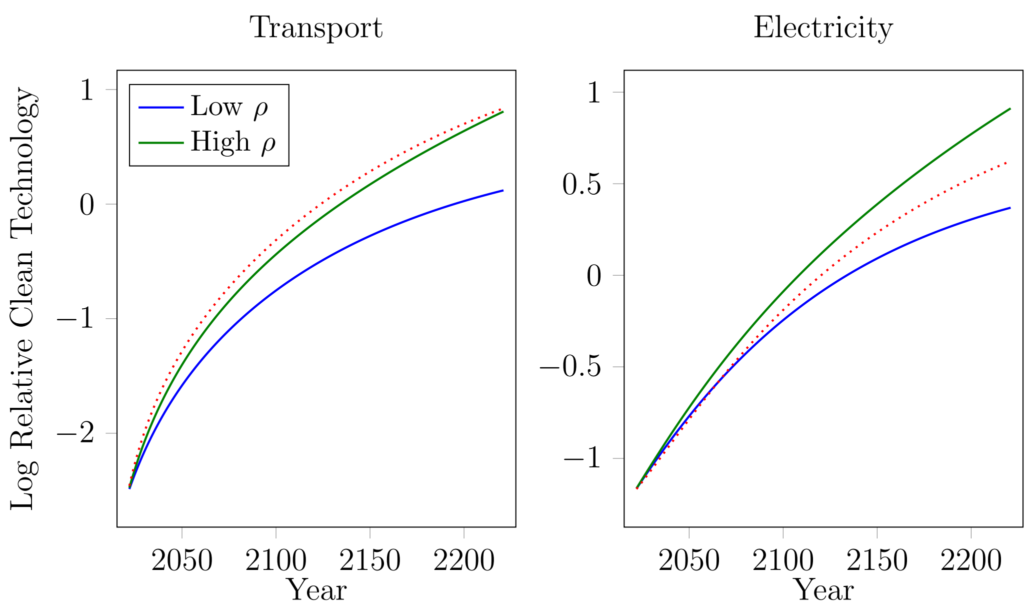

Notes: Dotted lines indicate laissez-faire paths.

Figures 1·9 and 1·10 provide further details on the first-best allocation. The dotted lines indicate laissez-faire paths. For the path of relative clean technology, optimal policy does not unambiguously favor clean technology. For transportation, the laissez-faire path actually has the highest relative clean technology, and for both transportation and electricity generation, the patient Planner ends up with the lowest relative clean technology. This again reflects the relatively low spillovers generated by clean technologies. Because clean technologies are actually disfavored by innovation subsidies, they do not gain as much of a lead over their dirty counterparts. This is especially true for transportation, where the clean technology is particularly bad at generating spillovers, and for the more patient Planner, who puts greater value on the generation of spillovers.

Nonetheless, as shown in Figure A·15, the steadily rising price on carbon allows clean technology to slowly take over production in both sectors. For both levels of discounting, within-sector clean income shares rise to steady-state values of one, but this happens more quickly for the more patient Planner, who imposes a higher price on carbon. Clean technology starts at a greater disadvantage in the transportation sector, but catchup growth via the spillover network allows clean transport to achieve a similar within-sector income share as clean electricity by 2200.

Finally, the carbon price substantially reduces carbon emissions, relative to laissez-faire. The more patient Planner immediately reduces carbon emissions by about half and steadily reduces them thereafter, whereas the less patient Planner allows carbon emissions to rise for a little more than a century. In both cases, the economy continues to grow, leading to steady reductions in emissions intensity. Together, these figures show that the Planner does not reduce emissions by abandoning dirty technologies. Instead, they push dirty technology out of production by raising the price on carbon, while at the same time continuing research in dirty technologies for the sake of spillovers.

In Figure A·16, I plot the temperature increases associated with first-best policy. Due to my focus on the transportation and electricity generation sectors in the US, most of the emissions that determine the global increase in temperature come from outside of my model. In spite of this, US climate policy still has a substantial influence on the level of warming beyond 2100. Warming peaks around 3 degrees Celsius or 3.2 degrees Celsius, depending on the discount rate, but in laissez-faire, warming continues unabated as the economy grows.

To further examine the role of spillovers in optimal policy, I shut down the spillover network in Figure 1·11. In that case, carbon prices are largely unchanged, but clean

Figure 1·10: First-Best Pollution Path

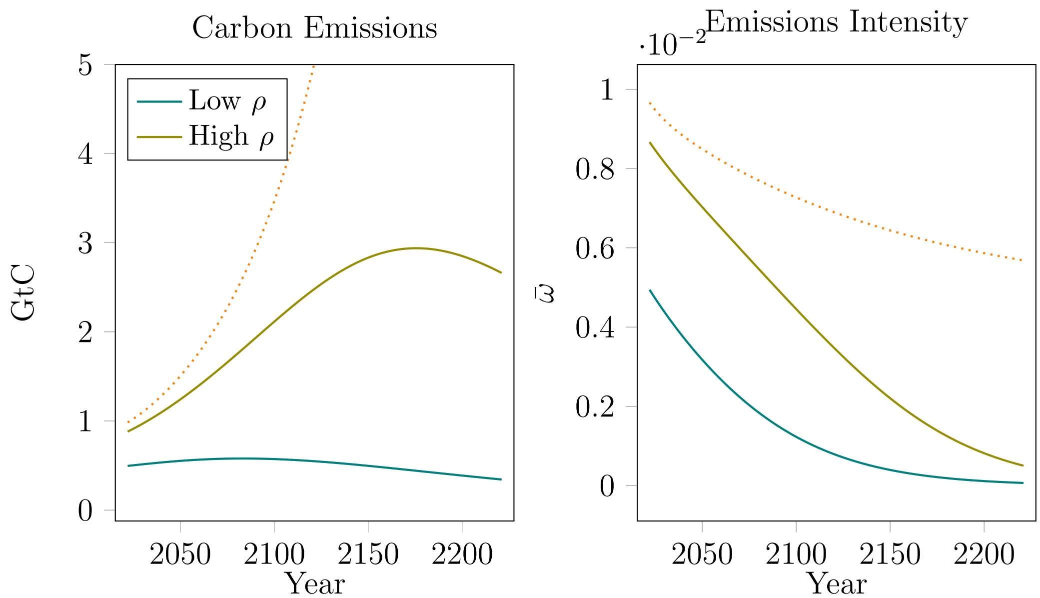

Notes: Dotted lines indicate laissez-faire paths.

technology receives a "big push" with massive, yet temporary, innovation subsidies. Now innovation subsidies for clean transport and electricity start at more than triple and double the baseline innovation wedge, respectively, and slowly decrease down to the baseline level over the first century of policy. This is in sharp contrast to innovation policy in the case of calibrated spillovers because, without a spillover network, clean technologies are beneficial in terms of both their ability to produce goods without pollution and their ability to produce spillovers. Before every technology could create spillovers on clean technologies, but now that only clean technologies can create clean spillovers, the usefulness of clean technology in production grants an exclusive value to their spillovers as well.

Finally, the damages we can expect from a warmer climate are highly uncertain, so I consider the possibility of substantially higher damages in Figure A·17. For that case, I quadruple the damage parameter ϱ, which substantially increases the social cost of carbon. Interestingly, despite the increased cost of climate change, clean

Figure 1·11: First-Best Policy Path (No Spillovers)

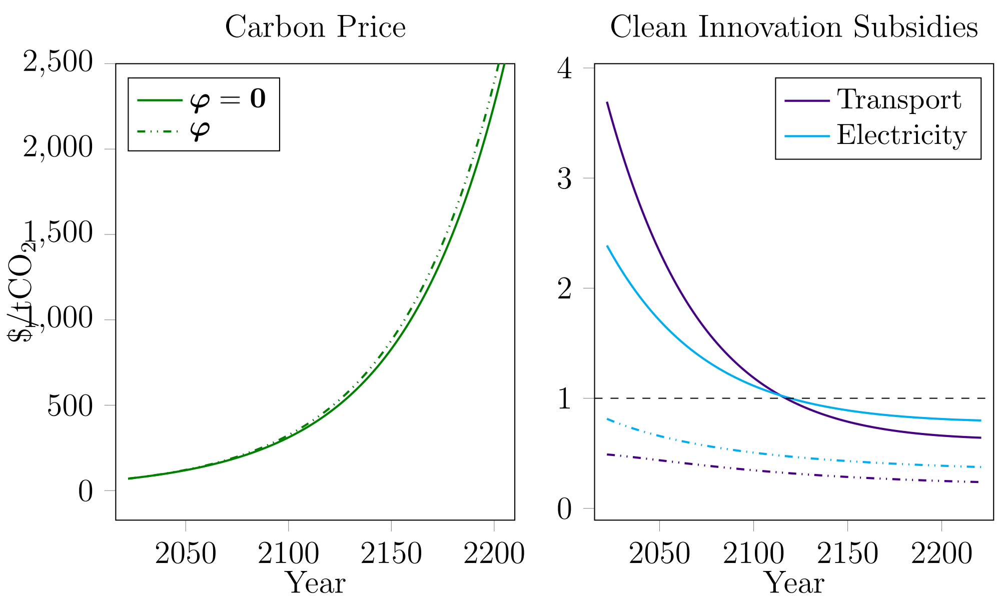

Notes: Optimal policy when the spillover network is shut down. Dashed-dotted lines represent the benchmark calibration. Both cases use the low discount rate. Innovation subsidies are listed as a fraction of the baseline innovation wedge.

innovation subsidies are reduced even further. The higher price on carbon leads to higher clean income shares, raising the private return to clean innovation. What does not change is the relatively low centrality of clean technologies in the spillover network. Thus, the social value of clean innovation increases by less than the private return, implying a reduction in innovation subsidies.

1.6.3 Second-Best Innovation Subsidies

In this section, I simulate second-best innovation subsidies, where carbon prices are incomplete. I consider two cases: one where the external carbon price is zero and another where the external carbon price is ten percent of the social cost of carbon. I show that these cases differ substantially as even a small, but growing, price on carbon allows the Planner to achieve welfare similar to that of the first-best, but without a price on carbon, society suffers massive welfare losses relative to the first-best. As before, the policy simulation starts in 2022.[63] In the main text, I take my preferred specification of low discounting, but I consider high discounting as a robustness check.

Figure 1·12: Second-Best Policy Path

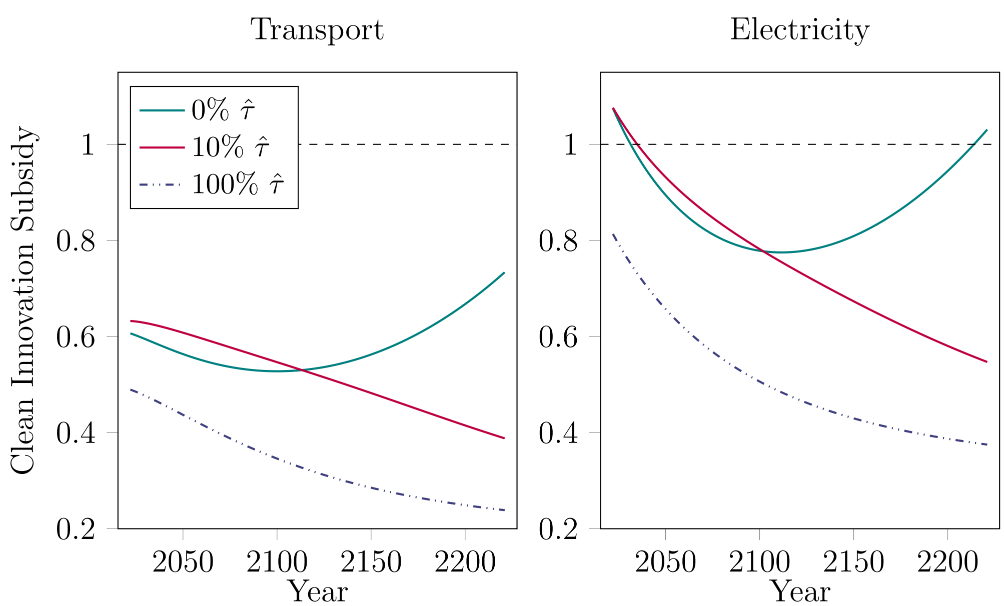

Notes: Policy paths use the low discount rate. The external carbon price $\hat{\tau}$ is a proportion of the social cost of carbon, so 100% is the first-best. Innovation subsidies are listed as a fraction of the baseline innovation wedge.

Figure 1·12 displays second-best innovation subsidies, as a fraction of the baseline innovation wedge, with the external carbon price $\hat{\tau}_t$ set as a proportion of the social cost of carbon. As shown in Proposition 1.5, innovation subsidies adjust to accommodate the distortion of incomplete carbon pricing. The dashed-dotted lines reference the first-best policy path. First, we can see that the case with $\hat{\tau}_t$ set to ten percent of the social cost of carbon is qualitatively similar to the first-best, with clean innovation subsidies shifted upwards by about 20-30% of the baseline innovation wedge. In that case, rising carbon prices eventually push dirty technology out of production, so the Planner can still develop dirty technologies to some degree for the sake of spillovers.

[63] Appendix A.2.4 contains further details on my solution method.

This contrasts sharply with the case where carbon pollution is entirely unpriced. In that case, clean innovation subsidies are similar to the ten percent case for the first century of policy, but beyond that, they begin to rise precipitously. Moreover, dirty innovation is immediately shut down. Without a carbon price, the only way for the Planner to push dirty technology out of production is to shut down dirty innovation altogether, forgoing spillovers in the process. One should note that the higher clean innovation subsidies of the second-best do not justify a big push for clean innovation. Even in the second-best, the inclusion of cross-technology spillovers implies a different prescription for clean innovation policy.

Figure 1·13: Second-Best Technology Path

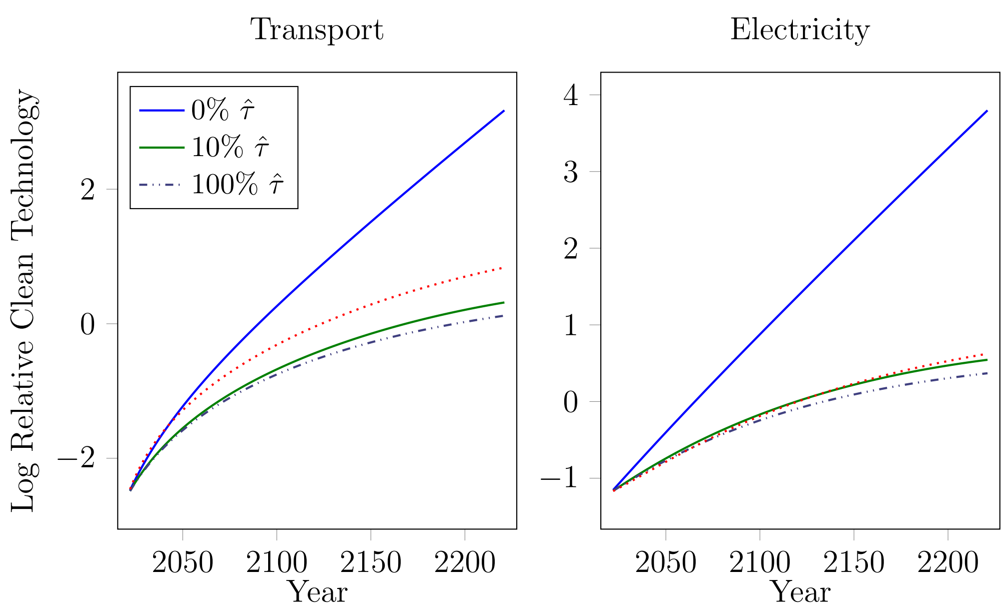

Notes: Policy paths use the low discount rate. The external carbon price $\hat{\tau}$ is a proportion of the social cost of carbon, so 100% is the first-best. Dotted lines indicate laissez-faire paths.

Figures 1·13 and 1·14 provide further details on second-best policy. The dotted lines indicate laissez-faire paths. The difference between the two second-best cases emerges clearly in the path of relative clean technology. Again, a small, growing carbon price allows the Planner to continue dirty innovation for the sake of spillovers, so relative clean technology is only marginally higher than in the first-best. However, without a carbon price, the Planner must drive relative clean technology to infinity to push dirty production out of the economy. This opens a vast difference in relative clean technology for the zero case, relative to the others.

A small, growing carbon price leads to an overshooting of emissions relative to the first-best but still allows a substantial reduction in emissions compared to laissez-faire. However, when the carbon price is lost altogether, emissions are only slightly below their level in laissez-faire for the first few centuries. This is because increases in clean technology lead to expansions in output, which increase demand for dirty inputs in turn. That is, reducing emissions via clean technological growth alone is ineffective because the substitution from clean to dirty is dampened by a rebound effect. Thus, a small carbon price goes a long way in preventing technological advance from translating into increases in emissions. For more detail, Figure A·18 shows the path for clean income shares for each case of second-best policy.

In Figure A·19, I show the path of temperature increases associated with each second-best policy path. As before, most global emissions come from outside of the model, but US climate policy has a substantial influence on the level of warming beyond 2100. As suggested by the path of emissions, the ten percent case leads to a fair amount of additional warming, with maximum warming of about 3.5 degrees Celsius, roughly half a degree higher than in the first-best. However, the zero case leads to disastrous warming of about 10 degrees Celsius and rising by 2500, underscoring the power of even a small carbon price.

These simulations illustrate the main tradeoff of second-best policy. Technology stocks are an input into the production of both physical goods and new ideas. As shown in the first-best, clean technologies should dominate production because they can produce goods without pollution, but in the production of ideas, dirty tech-

Figure 1·14: Second-Best Pollution Path

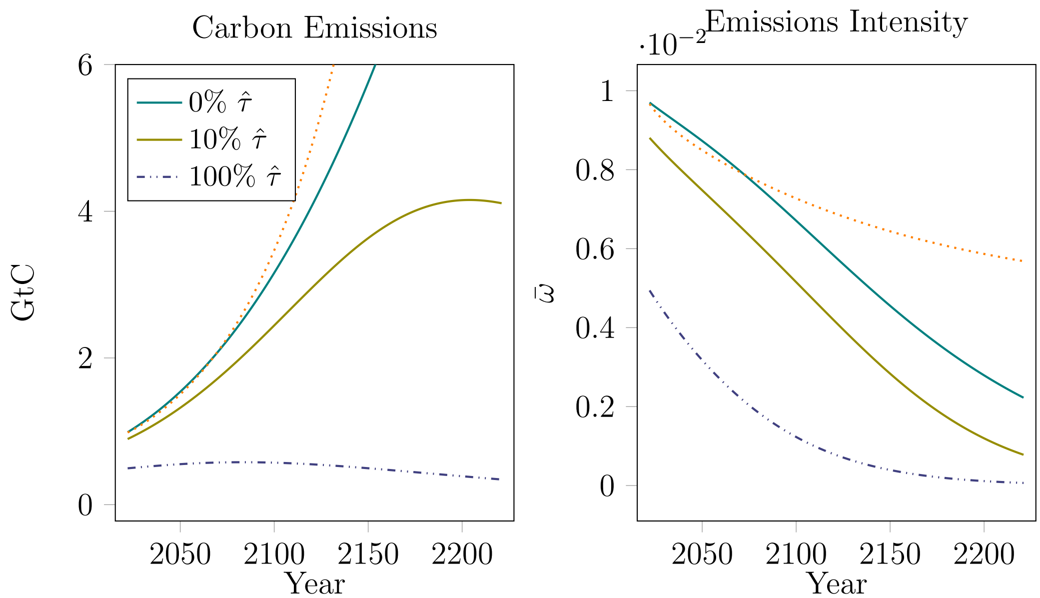

Notes: Policy paths use the low discount rate. The external carbon price $\hat{\tau}$ is a proportion of the social cost of carbon, so 100% is the first-best. Dotted lines indicate laissez-faire paths.

nologies are generally better due to their higher centrality in the spillover network. When the Planner can only control innovation subsidies, they must pick a path for technology that compromises between the two. A growing carbon price ensures that dirty technology is eventually pushed out of production, allowing dirty innovation to continue without increasing emissions. However, when pollution is unpriced, any dirty innovation will raise emissions, so the Planner must stop developing the dirty technology altogether.

In the end, decarbonizing by only improving clean technology leads to both slow emission reductions from a rebound effect and slow growth from the loss of spillovers.[64] For this reason, the welfare loss of moving from the first-best to the second-best

[64] In the steady-state, the loss of dirty innovation leads growth to stop altogether. This stems from using a Cobb-Douglas spillover function in my simulations. Put differently, when one or more technologies are lost, the model switches from endogenous growth to semi-endogenous growth. If the spillover function had an elasticity of substitution above one, growth could continue without dirty innovation, though the abrupt stop in dirty innovation would still dampen growth.

increases sharply when carbon pollution is unpriced. Pricing carbon pollution at ten percent of the social cost of carbon leads to a consumption equivalent loss of 1.52%, implying that a small, growing carbon price can achieve most of the welfare of the first-best. However, if carbon pollution is unpriced, the consumption equivalent loss is almost 100%. Clearly, the clean technological transition requires both instruments so as to decouple the production of physical goods and new ideas.

Figure A·20 shows second-best innovation policy in the case of high discounting. The conclusions are broadly similar. Clean innovation subsidies are higher than in the first-best in both the zero and ten percent case, and those of the zero case grow substantially over time. However, the difference between the zero and ten percent case is much less stark for the less patient Planner. For instance, dirty innovation is now allowed to continue for about another 150 years in the zero case. This reduced difference stems from the fact that the higher emissions and slower growth of the zero case are both costs that materialize in the future. A less patient Planner puts less weight on these costs, with a consumption equivalent loss of 0.31% for the ten percent case and 0.92% for the zero case, relative to the first-best. Therefore, the cost of incomplete carbon pricing depends critically on the Planner's degree of patience.

1.7 Conclusion

One of the main goals of climate policy is to redirect innovation from dirty to clean technology. Without such redirection, decarbonization will be prohibitively expensive. This paper has argued that cross-technology knowledge spillovers are critical for understanding the role of policy in the transition to clean technology. Such spillovers matter for both describing the impact of policy reforms and prescribing optimal clean innovation policy. They enable a more rapid transition to clean technology by allowing less advanced technologies to achieve catchup growth, but this same mechanism reduces the long-run impact of policy reforms. Formally, I have shown that two forces – cross-technology spillovers and substitution patterns in production – govern both the size and speed of technological transition following a policy reform and that each of these forces can be summarized in terms of a sufficient statistic matrix. Together, these forces govern the degree of increasing returns to innovation, which can be quantified in terms of a spectral radius.

For optimal policy, I have shown that innovation subsidies should focus on spillover creation whenever carbon prices correct the pollution externality. That is, clean technologies should not be subsidized simply because they can create goods without pollution; they should be subsidized insofar as they produce spillovers. However, I have shown that a similar innovation subsidy formula holds in the presence of arbitrary, potentially suboptimal, prices on carbon pollution, but with an adjustment for the equilibrium impact of innovation on emissions.

To take my model to the data, I performed several quantitative exercises, focusing on transportation and electricity generation in the US. First, I showed that cross-technology knowledge spillovers are strong enough to prevent the lock-in of dirty technology, but that the substitutability of clean and dirty goods still leads to a degree of increasing to returns to innovation near one. Thus, a realistic carbon price and clean innovation subsidy bring about a substantial redirection of innovation toward clean technology in the long run, but with a slow transition. Second, I showed, contrary to conventional wisdom, that clean innovation should receive little subsidy because clean technologies do not produce particularly large spillovers, as measured by their centrality in the spillover network. This holds quantitatively even when carbon prices are set below the social cost of carbon. Finally, I quantified the welfare cost of restrictions on the carbon price and showed that a small, growing carbon price can recover most of the welfare of the first-best. However, if carbon pollution is unpriced, then decarbonization requires shutting down dirty innovation and relying exclusively on clean innovation, with disastrous consequences for welfare and the climate.

Economists have long understood, going back to Hicks (1932), that substitution patterns in production are important whenever the direction of innovation is endogenous. This paper has argued that cross-technology knowledge spillovers are deserving of a similar status. I have made this argument in the context of climate change, but I believe it has more general applicability. Just as the importance of substitutability has been applied to a variety of contexts – including skill-biased technological change, automation, and more – the importance of cross-technology spillovers could be applicable in a wide range of settings as well.

CHAPTER TWO - OPTIMAL TAXATION WITH AUTOMATION: NAVIGATING CAPITAL AND LABOR'S COMPLICATED RELATIONSHIP

2.1 Introduction

The last few decades of income growth in the United States have not borne widely shared fruits. Unlike income growth in the mid-twentieth century, the benefits have primarily accrued to the top of the income distribution (Bailey and Danziger, 2013). This growth has been so uneven that many groups have seen reduced wages, a phenomenon that is difficult to explain with traditional, skill biased technological change models (Acemoglu and Autor, 2011). Furthermore, this stagnation of wage growth has happened alongside a reduction in the labor share of national income (Karabarbounis and Neiman, 2014). Automating technologies, and their displacement effect on labor, are now seen as one of the prime suspects for these changes (Acemoglu and Restrepo, 2020). To put the matter in historical context, Frey (2019) has argued that today's automating technologies represent a return of the labor replacing technologies of the Industrial Revolution, distinct from the labor enabling technologies of the first half of the twentieth century. This paper will remain agnostic as to the exact extent that automating technologies can account for the adverse trends described above. The goal, instead, is to answer the question: How should policymakers address inequality brought on from the displacement effect of automation?

If enough policy tools are available, the answer to the above question need not involve any disruption to the production process, the extreme case being the Second

Welfare Theorem. But, if policy tools are more limited, there may be reason to disrupt production to reduce the inequality consequences of displacement from automation. This paper will argue that the primary trade off faced when limiting the extent of automation in production is the vacillating relationship between labor and capital. Reflecting a longstanding view of macro-economics, capital deepening on the infra margin of automation, where automation has already taken place, benefits labor by making it more productive. However, the friendly relationship between capital and labor becomes adversarial at the extensive margin of automation, where displacement occurs. The Planner must therefore balance these two forces when distorting the level of automation in production.

To formalize the above point, this paper will consider a heterogeneous-agents task model with a Ramsey problem inspired by Acemoglu et al. (2020). In this setting, the Planner will have access to linear factor taxes, a uniform lump-sum transfer, and a policy variable that I will call an automation threshold rule. This threshold rule will stipulate the degree to which the cost of employing capital must be below the cost of employing labor before a firm is allowed to automate and use capital. With such a rule, the Planner is able to exclusively target the extensive margin of automation because the requirement that capital must be much better at a task before that task can be automated only binds when the automation decision is near the margin. Moreover, the threshold rule does not have an inframarginal effect on tasks away from the margin of automation, where capital poses no threat to labor.

The analytical utility of the model comes from the comparison of the threshold rule with the capital tax. Capital taxes can also influence the extensive margin of automation by increasing the user cost of capital, but they do so by reducing the economy-wide supply of capital, which also affects capital deepening on the infra margin. The main result of this paper will be to show that, conditional on setting the threshold rule optimally, capital taxation is no longer a useful policy tool for controlling the level of automation. Put differently, once labor has been protected along the extensive margin, there is no reason to distort the use of capital along the infra margin beyond the need to raise revenue. In practice, real-world policy makers are unlikely to have a tool that precisely targets the extensive margin of automation, and nothing else. Instead, this result should be interpreted as a formal expression of the desirability of targeting the extensive margin of automation in a limiting case where the extensive and infra margin can be considered separately. In particular, the comparison of two extreme policies, one which perfectly targets the extensive margin and the other which is indiscriminate, allows a vantage point to understand the trade offs at play when distorting the extent of automation in production.

The reasoning behind this result will also draw heavily on the notion of "so-so technologies" developed in Acemoglu and Restrepo (2018). The margin of automation is, by definition, the place where capital's comparative advantage is the smallest, so distorting the task allocation near this margin has little efficiency cost. The tasks sent back to labor will be the exact tasks where switching between capital and labor has the lowest impact on output. These marginal tasks are "so-so" in the sense that their automation elicits small productivity improvements while still reducing wages of exposed workers through a displacement effect. Labor may still be highly productive in these tasks, and thus loath to lose them; it is only the productivity difference between capital and labor that is small.

The main theoretical results will be derived in a static setting where different households can supply distinct types of labor. Each of the various labor types are subject to their own endogenously determined margin of automation. Because automation is endogenous, the Planner must take into account the reverberating effect on the level of automation from each decision. To restate the result described in the above paragraphs, with an optimal threshold rule, the optimality condition for the capital tax ignores this automation effect. This is not, however, the case for the labor tax. Furthermore, if the Planner is unable to use a threshold rule, the optimal capital tax must take into consideration its effect on automation. This is an important secondary case because of the practical difficulty of exclusively targeting the extensive margin of automation. Furthermore, the optimal lump-sum transfer is not directly related to the level of automation, implying that a UBI need not have any special relationship with automation per se.

Policy considerations that pertain to capital are naturally dynamic, so it is important to consider whether the insights of a simplified static setting extend to a more general dynamic environment. To address this question, I develop a discrete-time dynamic extension with heterogeneous agents in a deterministic, generalized OLG setting. The OLG setup is general in the sense that it allows an arbitrary set of generational groups and lifespans. In this dynamic setting, it is again the case that optimal capital taxation can ignore automation effects if and only if there is an optimal threshold rule. Furthermore, the dynamic model will show that the choice of optimal threshold rule is static: the threshold variable can be solved out in each period with only that period's variables. The dynamic extension does, however, have some major shortcomings. It does not consider a margin for directed technological change via research, nor does it consider the possibility of human capital investment that allows a worker to change type. The result that setting optimal threshold rules is a sequence of static problems may not be robust to the inclusion of these considerations.

With a theoretical framework in place, I turn to a static, heterogeneous-agents numerical exercise to quantify the optimal magnitude of the various policy parameters as well as their relative significance for welfare. One advantage of this paper's structural model is that, with proper calibration, it allows for the computation of the global optimum. The most important empirical moment for this exercise to match is the extent to which different occupations are exposed to automation, so I use estimates from Webb (2019) of occupation-level exposure to both robots and software to differentiate the effect of automation across heterogeneous types of labor. The other relevant moment reflecting the economy-wide substitutability between capital and labor is an estimate of the elasticity of substitution from Karabarbounis and Neiman (2014). Otherwise, the remaining moments for calibration will be chosen to simulate the 2010 American economy. This includes a wealth distribution which matches the substantially greater inequality in the wealth distribution relative to the income distribution.

The results of this numerical exercise imply that substantial redistribution via a lump-sum transfer is the most important margin for welfare improvement. Moving to optimal policy is welfare equivalent to roughy a 5% increase in everyone's consumption, and much of this improvement is driven by straightforward redistribution. In particular, optimal policy entails a lump-sum transfer equal to about 17% of per-capita consumption which is funded by similar taxes on both factors of about 33%. The similar magnitude of capital and labor taxes is due to both the substantial level of inequality in the wealth distribution and the similarity of observed elasticities for capital and labor. If either of these features are relaxed, the optimal capital tax drops significantly. Of less importance is the threshold rule, which is set optimally at about 3.4%; a level not sensitive to the ownership distribution or elasticity of capital. This suggests that the practical difficulty of targeting the extensive margin of automation is not of great welfare consequence and that the welfare insults from automation are not that large relative to other sources of inequity in the United States.

This paper is related to intersecting lines of research in both macro and public finance. We have recently made great advances in our understanding of automation through the insights afforded by the task framework; Acemoglu and Restrepo (2018) being canonical work in this area. The literature on the optimal policy response to automation is still fairly young (Guerreiro et al., 2022; Thuemmel, 2022; Costinot and Werning, 2023), and I hope this paper can contribute to that literature by incorporating some of the knowledge gained from the task framework. The work most similar to my own in the automation policy literature is Acemoglu et al. (2020), for the simple reason that my paper borrows heavily from their approach. The conclusions, however, are quite different. In their model, the best available tax system involves production efficiency; capital and labor are taxed (at similar rates) and the level of automation is left undistorted. This is in stark difference to my model's recommendation that a threshold rule be used to distort the level of automation away from production efficiency. The source of this difference is the number of agents. Acemoglu et al. (2020) consider a single type of labor, so they satisfy the conditions of the Diamond and Mirrlees (1971a,b) production efficiency result: all goods and factors are taxed at their own linear rate. Because there is only one type of labor, the requirement that all workers face the same tax schedule does not restrict different inputs to the same tax. In a heterogeneous-agents model, the assumption of a uniform labor tax schedule violates the production efficiency conditions. Thus, production inefficiency can be optimal when there are multiple types of workers. Nevertheless, there are variants of the Acemoglu et al. (2020) model where the authors consider distortions of the level of automation, so I will discuss how these cases relate to my approach in greater detail later in the paper.

Another important, more developed, literature with which this paper interacts is that on optimal capital taxation (Atkinson and Stiglitz, 1976; Chamley, 1986; Naito, 1999; Straub and Werning, 2020). The primary result of this paper, that automation concerns need not affect the capital tax when the extensive margin of automation has been targeted optimally, can be viewed as complementary to past capital taxation results that did not consider an automation margin. If the extensive margin of automation can be controlled separately, our past thinking about capital taxation can again become applicable even in the face of automation concerns. It is only in the absence of such targeting that one needs to update their views of the capital tax to account for automation. Thus, the primary contribution of this paper to the capital taxation literature is that the standard equity/efficiency considerations motivating the analysis of capital taxation can be separated from a desire to control the level of automation.

The distinction between the extensive and infra margin of automation presents an important conceptual caveat to the common intuition among both laypeople and some economists that taxes on capital and universal, unconditional cash transfers are the primary policy tools with which to address automation. In an influential 2017 interview, Bill Gates claimed "the robot that takes your job should pay taxes". In addition, Democratic primary candidate Andrew Yang ran on a platform of UBI to help workers displaced by automation. More importantly, the technical literature on the subject has almost exclusively revolved around determining optimal taxes on automating capital. In Costinot and Werning (2023), the authors develop a general framework for the taxation of new technologies. They consider two technologies, old and new, where the former employs all labor. Each of the goods traded by the new technology can be linearly taxed at a good-specific rate. With this framework, the authors derive optimal tax formulas for robots comprised of measurable sufficient statistics. Another paper in this literature, Guerreiro et al. (2022), considers a model with two types of workers, routine and non-routine, and derives optimal robot taxes along a transition path where the cost of producing robots steadily decreases. Both of these papers find positive robot taxes as a means to dampen the extent of automation. In addition, they both consider universal lump-sum transfers, either explicitly or implicitly with general, Mirrleesian income tax functions. Of course, sufficiently targeted taxes may allow for targeting of the extensive margin of automation, but any application of taxes should consider raising revenue on the infra margin, and slowing automation on the extensive margin, separately. This is all said keeping in mind the possibility that deliberate targeting of the extensive margin of automation may be of little welfare consequence, which can only be examined in a calibrated model.

Finally, the reasoning of this paper has general application to the trade-offs imposed on society by *creative destruction*. We know that creative destruction is the engine of long-run prosperity, but there is now ample evidence that productivity enhancing changes can create uncompensated losers (Autor et al., 2013; Acemoglu and Restrepo, 2020). Of course, the theoretical possibility of compensating the losers with the surplus is not a systemic framework for dealing with this problem in a second-best world. This paper deploys the insights of the task model to provide a helpful principle for thinking about this problem. Indeed, the trade-offs facing policymakers in their response to creative destruction are similar to those in the case of automation (which is really just a form of creative destruction). A new technology is only harmful to a disrupted incumbent at the extensive margin of displacement; once displacement has taken place, productive use of new technology raises output and provides the incumbent with more demand in their remaining tasks. As before, the most efficient place to reduce the disruption of a productive new technology is at the marginal task, where the comparative advantage of the new technology is the smallest, while the gains from novelty are best realized in tasks where the old way of doing things is at a clear disadvantage. With this framework, the goal is to deploy the new technology where it is most useful to society and hold it back when disruption provides little reward.

The remainder of the paper is organized as follows. Section 2.2 will present the static model and derive the primary theoretical conclusions of this paper. Section 2.3 will present a dynamic extension that examines the robustness of the static results in an OLG environment. Section 2.4 outlines the numerical strategy and derives quantitative optimal policy results. Section 2.5 concludes. The technical details and proofs, as well as more details describing the construction of the numerical model, will be contained in the Appendix.

2.2 Static Model

2.2.1 The Environment

Households

We will have N-many households, or labor types, each of which will select consumption and labor supply to maximize their utility function subject to budget constraints

$$u_i(C_i, L_i) \tag{2.1}$$

Therefore, given after-tax wages $(1-\tau_L)W_i$ and capital rents $(1-\tau_K)R$, each household satisfies a labor FOC and budget constraint

$$(1-\tau_L)W_i = -u_{i,l}/u_{i,c} \tag{2.2}$$

$$C_i = (1-\tau_L)W_iL_i + n_iD + \omega_i\big((1-\tau_K)RK - \phi(K)\big) \tag{2.3}$$

where D is a UBI, n_i is the number of members of household i, and households own exogenous fraction ω_i of an investment bank with cost function $\phi(K)$ that invests according to

$$(1-\tau_K)R = \phi'(K) \tag{2.4}$$

Note that we have chosen to place taxes on the supply side, but as is well known, this does not matter for economic incidence. Furthermore, we will have $\sum_i \omega_i = 1$.

Tasks

Each type will have their own continuum of tasks to reflect differential exposure to automation. Denote by $X_i \subset \mathbb{R}$ the tasks to which each type is exogenously capable of supplying labor. We will have

$$X_i = (\underline{x}_i, \overline{x}_i] \tag{2.5}$$

Each type of worker can have some of their tasks automated. Denote the set of automated tasks for type i by $(\underline{x}_i, x_i]$, i.e. automation is coming from the left for each worker. Note that I have chosen to give the threshold task to capital, but this is not of practical importance.

Therefore, the tasks given to each factor can be denoted by

$$\Omega_{L_i} = (x_i, \overline{x}_i] \tag{2.6}$$

and

$$\Omega_K = \cup_i (\underline{x}_i, x_i] \tag{2.7}$$

Production

Each factor J will have its own general productivity parameter $A_J \in \mathbb{R}$ and task-specific productivity function $\alpha_J : X_J \to \mathbb{R}$. The output of each task will be produced with perfect substitutes:

$$y(x) = A_K \alpha_K(x) k(x) + A_{L_i} \alpha_{L_i}(x) \ell_i(x) \quad x \in X_i \tag{2.8}$$

A price taking representative firm produces output according to a CES task aggregator

$$F(X,L,K)=\Big(\int\limits_{\Omega_K}\big(A_K\alpha_K(x)k(x)\big)^{\frac{\sigma-1}{\sigma}}dx+\sum_i\int\limits_{\Omega_{L_i}}\big(A_{L_i}\alpha_{L_i}(x)\ell_i(x)\big)^{\frac{\sigma-1}{\sigma}}dx\Big)^{\frac{\sigma}{\sigma-1}} \quad (2.9)$$

where $X \equiv (x_1,\ldots x_N)$ and $L \equiv (L_1,\ldots L_N)$. We will also impose factor market clearing, so $\int_{\Omega_k}\kappa(j)dj = K$ and $\int_{\Omega_{L_i}}\ell_i(j)dj = L_i\ \forall i$.

As described in Appendix B.1.1, the further assumption of cost minimization across tasks and the ideal price index as the numeraire allows us to derive a CES production function in terms of aggregate factor supplies:

$$F\big(X,L,K\big)=\Big(\Lambda_K^{\frac{1}{\sigma}}\big(A_KK\big)^{\frac{\sigma-1}{\sigma}}+\sum_i\Lambda_{L_i}^{\frac{1}{\sigma}}\big(A_{L_i}L_i\big)^{\frac{\sigma-1}{\sigma}}\Big)^{\frac{\sigma}{\sigma-1}}, \quad (2.10)$$

where $\Lambda_K \equiv \int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx$ and $\Lambda_{L_i} \equiv \int_{\Omega_{L_i}}\alpha_{L_i}(x)^{\sigma-1}dx$. From this form of the production function, we can interpret variation in the task allocation as eliciting variation in the CES shares.

Endogenous Task Allocation

We will have that the factor with the lowest marginal cost is the one assigned to a given task, so $\forall i$

$$\Omega_{L_i}=\Big\{x\in X_i\Big|\frac{W_i}{A_{L_i}\alpha_{L_i}(x)}<\frac{R(1+\theta)}{A_K\alpha_K(x)}\Big\} \quad (2.11)$$

while Ω_K encompasses all other tasks.

Here, θ is the threshold rule to be determined by the tax authority. In words, firms are not allowed to automate tasks unless the marginal cost of labor is $\theta \cdot 100$ percent greater than that of capital. To see this, consider that automated tasks will satisfy:

$$\frac{R(1+\theta)}{A_K\alpha_K(x)}\leq\frac{W_i}{A_{L_i}\alpha_{L_i}(x)}$$
$$\Rightarrow 1+\theta\leq\frac{MC_{L_i}}{MC_K}$$

Furthermore, with the assumption that $\frac{A_{L_i}\alpha_{L_i}(x)}{A_K\alpha_K(x)}$ is strictly increasing and continuous in x $\forall i$ (i.e. labor has a comparative advantage in higher ordered tasks) we will have that the task sets are partitioned by the thresholds of automation defined by X.

It is worth reiterating that θ is not a tax on capital. Figure 1 shows the various curves determining the automation decision of the firms. Firms pay according to the marginal cost curves that are shaded, so the marginal user cost of capital is not increased by $(1+\theta)$. The unshaded line just stipulates what tasks firms are allowed to automate: those where the unshaded line is still below the marginal cost of labor. This then shifts the level of automation down to x_i as more tasks are regulated back to labor. This figure also provides intuition for the functioning of the threshold rule. With no threshold rule, the shaded marginal cost curves would determine the task allocation and capital would automate any task where it was at all cheaper than labor. With a threshold rule, firms may only automate tasks where the comparative advantage of capital is especially strong. This guarantees that the tasks given back to labor are the exact ones where the productivity gains from automation are the smallest. The resultant deadweight loss triangle from the threshold rule is also visible.

Figure 2·1: Determination of the Automation Threshold

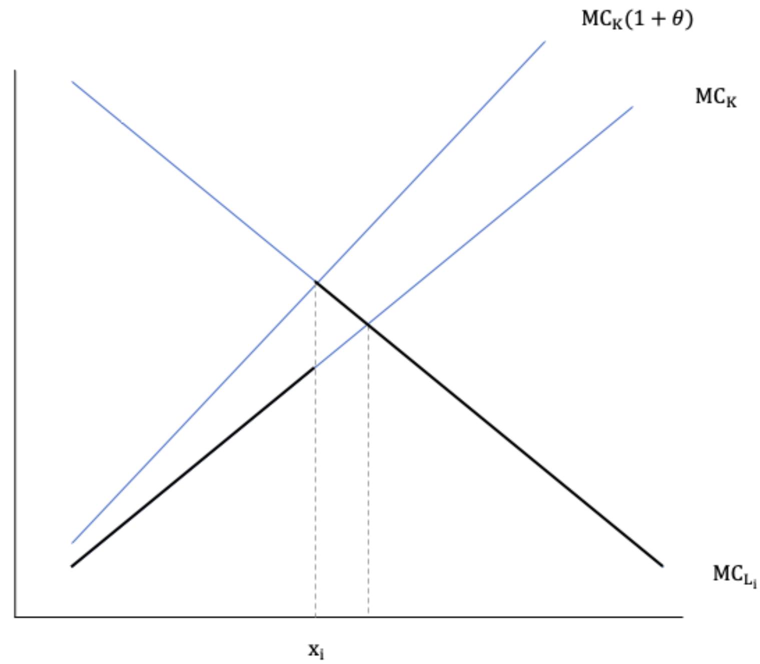

The endogenous task allocation gives us task thresholds that are a function of (L, K, θ). To derive $X(L, K, \theta)$, we can note that this choice of automation defines a level set:

$$\left\{(X, L, K, \theta) \in \mathbb{R}^{2N+2} \middle| \frac{F_K(1+\theta)}{A_K \alpha_K(x_i)} = \frac{F_{L_i}}{A_{L_i}\alpha_{L_i}(x_i)} \quad \forall i\right\} \tag{2.12}$$

The Implicit Function Theorem can then be used to derive $X(L, K, \theta)$. The derivation of the Jacobian of this function is detailed in Appendix B.1.2. The most important feature of this Jacobian is that $\frac{\partial X}{\partial K} \propto \frac{\partial X}{\partial \theta}$.

We can then define a production function that has an endogenous task allocation built into its internal machinery:

$$\tilde{F}(L, K, \theta) \equiv F(X(L, K, \theta), L, K) \tag{2.13}$$

Note that for factor J, we have

$$\tilde{F}_J = F_J + \sum_i \frac{\partial F}{\partial x_i}\frac{\partial x_i}{\partial J}$$

However, in equilibrium, factor prices will be F_J, not $\tilde{F}_J$. When a firm takes factor prices as exogenous, the task allocation will be determined entirely by factor prices, the threshold rule, and the schedule of productivities. Conditional on this task allocation, firms then set marginal products equal to factor prices. $X(L, K, \theta)$ is, therefore, a general equilibrium object whose derivatives are not internalized by firms.

Government

The government selects linear factor taxes, a threshold rule, and a UBI to maximize the social welfare function:

$$W(u_1, u_2, \ldots, u_N) \tag{2.14}$$

The government must also satisfy its own budget constraint:

$$G + \sum_i n_i D = \sum_i \tau_L W_i L_i + \tau_K R K \tag{2.15}$$

where G is some exogenous level of government spending.

The total demand for output in this economy is then

$$Y = \sum_i C_i + G + \phi(K) \tag{2.16}$$

so satisfaction of the resource constraint requires $\tilde{F} = Y$.

Equilibrium

For a given set of policy choices $\big(\tau_L, \tau_K, D, \theta\big)$, an equilibrium consists of factor prices $\big(\{W_i\}, R\big)$, an allocation of factors, output, and consumption $\big(\{L_i, C_i\}, K, Y\big)$, and an assignment of factors to tasks $\big(\{\Omega_{L_i}\}, \Omega_K\big)$ such that:

(i) Each household maximizes utility subject to their budget constraint.

(ii) The investment bank and representative firm maximize profit.

(iii) Factor markets clear while the resource constraint is satisfied.

Note that satisfaction of the government's budget constraint also follows from these criteria.

2.2.2 Ramsey Problem

Setup

To determine optimal policy, the Planner will solve the following Ramsey problem:

$$\max_{\{C_i\},\{L_i\},K,D,\theta,\tau_L,\tau_K} W(u_1, u_2, \ldots, u_N) \;\; s.t.$$

$$\big[\text{IC1}\big] : \; (1-\tau_L)F_{L_i} = -u_{i,l}/u_{i,c} \; \forall i$$

$$\big[\text{IC2}\big] : \; (1-\tau_K)F_K = \phi'(K)$$

$$\big[\text{GBC}\big] : G + \sum_i n_i D = \sum_i \tau_L F_{L_i} L_i + \tau_K F_K K$$

$$\big[\text{BC}\big] : C_i = (1-\tau_L)F_{L_i}L_i + n_i D + \omega_i\big((1-\tau_K)F_K K - \phi(K)\big) \; \forall i$$

The first two constraints are implementation constraints reflecting the requirement that the Planner respect the decentralized behavior of households and firms. Specifically, the first term reflects the optimization criteria of the labor market, while the second term reflects the optimization criteria of the capital market. The next two constraints are just budget constraints for both the government and households. Note that the satisfaction of these budget constraints implies the satisfaction of the resource constraint. Effectively, the Planner is picking the welfare maximizing policy and allocation that is consistent with market equilibrium.

A feature of this problem that is worth emphasizing is that all of the different types of labor face the same labor tax. This is important because, in addition to being realistic, it prevents the use of the Diamond and Mirrlees (1971a,b) production efficiency result, which would trivially require that $\theta = 0$. The production efficiency result states that if all goods and factors can be linearly taxed at their own rate, there would be no reason to disrupt the transformation of inputs into outputs. In the context of automation, all of the distributional consequences from the displacement of labor could be dealt with via taxes on the factors that benefit and subsidies on those that don't. This is distinct from the Second Welfare Theorem, which would use idiosyncratic lump-sum taxes to correct the adverse distributional consequences of automation while also maintaining production efficiency. Here, because we are in a second-best environment with more limited policy tools, there may be reason to violate production efficiency as a corrective for labor displacement.

Optimal Policy

From the above Ramsey problem, we can derive a Lagrangian as described in Appendix B.1.3. In the results I describe below, λ_i will refer to the multiplier on the household i's budget constraint, μ to the multiplier on the government budget constraint, and η_i to the multiplier on the requirement that household i faces the same labor tax as household 1. This last constraint enforces universality of the labor tax schedule. Also, denote by $\varepsilon_K \equiv \frac{\phi'(K)}{\phi''(K)K}$ the elasticity of capital supply.

Proposition 2.1. *When the threshold rule is optimally chosen, the optimal capital tax and UBI satisfy*

$$\frac{\tau_K}{1-\tau_K} = \frac{\sum_i \omega_i(\mu - \lambda_i)}{\mu}\frac{1}{\varepsilon_K} \tag{2.17}$$

$$\frac{\sum_i \lambda_i n_i}{\sum_i n_i} = \mu \tag{2.18}$$

The proof of this proposition is provided in Appendix B.1.4. The first thing to notice in these equations is that there is no effect from the change in automation thresholds. Put differently, these are the same optimality conditions one would derive if the production function had been a CES with fixed, rather than endogenous, shares. For the UBI, this is not mysterious given that it is a lump-sum transfer. However, capital still does have an effect on the automation thresholds as the supply of capital determines the user cost of capital, but this effect is ignored by the Planner because the change in thresholds from capital are proportional to those from the threshold rule. Therefore, the First Order Condition on the threshold rule implies that $\frac{\partial \mathcal{L}}{\partial X}' \frac{\partial X}{\partial K} \propto \frac{\partial \mathcal{L}}{\partial X}' \frac{\partial X}{\partial \theta} = 0$. Herein lies the primary result of this paper: The optimal capital tax may ignore the effect of capital on the level of automation when the threshold rule has been chosen optimally.

The interpretation of these equation is otherwise fairly straightforward. Capital taxes should be higher when there is adverse wealth inequality: if those with large shares of the wealth have low private value of funds relative to the social value. With utility functions that are concave in consumption, we would interpret this as those with low consumption having little wealth. This equity effect is also tempered by a typical inverse elasticity rule. Turning to the UBI, transfers should be made until the social value of funds is equal to the average private value of funds. One qualifier to this proposition is that automation can have an indirect effect on the capital tax and UBI through the distribution of welfare, i.e. by changing the λ vector. For instance, if automation leads to drastic inequality, this may imply a higher capital tax and UBI. But that would be a response to the downstream consequences of automation, rather than an upstream attempt at controlling the level of automation itself.

Proposition 2.2. *The optimality condition for the threshold rule satisfies*

$$\mu \sum_i \frac{\partial F}{\partial x_i} \hat{\alpha}_i^{-1} = \sum_{i \neq 1} \frac{\eta_i}{\sigma} \frac{F_{L_1}}{F_{L_i}} \Big(\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_i} \alpha_{L_i}(x)^{\sigma-1} dx} \hat{\alpha}_i^{-1} - \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_1} \alpha_{L_1}(x)^{\sigma-1} dx} \hat{\alpha}_1^{-1} \Big) \tag{2.19}$$

Here, $\hat{\alpha}_i$ is an object from the $X(L, K, \theta)$ Jacobian which is defined in equation (B.3). The proof of this proposition is contained in Appendix B.1.5. Equation (2.19) should be interpreted as the simple equality of marginal cost with marginal benefit. The LHS is the benefit from increased production due to more efficient automation, while the RHS is the equity cost of further automation. Because the LHS represents a deviation from production efficiency, it will converge to zero as θ does the same. To gain further intuition for equation (2.19), it is worth considering the FOC that would result if the Planner could pick an automation threshold by fiat, as shown in equation (2.20).

$$\mu \frac{\partial F}{\partial x_i} = \frac{\eta_i}{\sigma} \underbrace{\frac{F_{L_1}}{F_{L_i}}}_{\text{a}} \underbrace{\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_i} \alpha_{L_i}(x)^{\sigma-1} dx}}_{\text{b}} \quad \forall i \neq 1 \tag{2.20}$$

The main intuition of this equation is that we don't want to take tasks from poor people where they really excel. If the wages of household i are relatively low (a) or the task being considered for automation is productive relative to those in their task allocation (b), then the automation of this task must be justified with a large increase in output, as represented by the LHS. Furthermore, note that the elasticity of substitution is in the denominator of the RHS. To see the importance of this, consider the case of perfect substitutes across tasks in production. If tasks are linearly substitutable, workers are indifferent as to the number of tasks they are allocated; they pay no cost when they are squeezed into a smaller number of tasks by automation. Finally, note that the analog of equation (2.20) for the first household is similar but for a few adjustments that account for their status as the reference household. We can now see that equation (2.19) is just the weighted average of threshold FOCs where the weighting is the sensitivity of a threshold to automation.

I will leave the optimality condition of the labor tax to Appendix B.1.6 as it is somewhat cumbersome. The most important feature of this optimality condition is that, unlike the capital tax, the labor tax must take into consideration the effect of labor supply on automation thresholds. This is perhaps unsurprising given that the root cause for deviating from production efficiency is the inability to differentiate labor taxes by type.

Optimal Policy Without a Threshold Rule

A requirement for Proposition 2.1 was that the effect of capital on automation thresholds envelopes out due to the threshold rule FOC. However, what if the Planner is unable to target the extensive margin of automation? This may be due to the political influence of those that stand to benefit from automation, the policy maker's inability to carry out task-level rules, or simple incompetence on the part of policy makers. Whatever the reason, it is worth understanding what optimal capital taxation looks like when it is the only available tool for influencing the extensive margin of automation. A related question, considered in the numerical exercise of Section 2.4, is whether there is a substantial welfare cost in losing the ability to separately affect the extensive and infra margin of automation.

Let us consider a case where the automation thresholds are always set to facilitate production efficiency:

$$X^m(L, K) \equiv \underset{X \in \prod_i X_i}{\operatorname{argmax}} F(X, L, K) \tag{2.21}$$

The level set that determines the automation thresholds is then

$$\left\{(X, L, K) \in \mathbb{R}^{2N+1} \middle| \frac{F_K}{A_K \alpha_K(x_i)} = \frac{F_{L_i}}{A_{L_i} \alpha_{L_i}(x_i)} \ \ \forall i \right\} \tag{2.22}$$

which is equivalent to requiring $\theta = 0$. Thus, if we denote the level set of (2.12) by A

and that of (2.22) by B, we have

$$\forall (X, L, K) \in B, \;\; (X, L, K, 0) \in A$$

which implies that the implicit functions are as such:

$$X^m(L, K) = X(L, K, 0)$$

We then have that the partial derivatives of X^m are the same as those of X when θ is evaluated at zero. The maximization problem is then the same as before but with $\theta = 0$, rather than as a variable under control.

Proposition 2.3. *In the absence of a threshold rule, the optimal capital tax satisfies*

$$\begin{aligned} \frac{\tau_K}{1-\tau_K} =& \frac{\sum_i \omega_i(\mu - \lambda_i)}{\mu} \frac{1}{\varepsilon_K} \\ &+ \underbrace{\frac{c}{\mu} \frac{1}{\phi'(K)K} \sum_{i \neq 1} \frac{\eta_i}{\sigma} \frac{F_{L_1}}{F_{L_i}} \Big(\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_i} \alpha_{L_i}(x)^{\sigma-1} dx} \hat{\alpha}_i^{-1} - \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_1} \alpha_{L_1}(x)^{\sigma-1} dx} \hat{\alpha}_1^{-1} \Big)}_{a} \end{aligned} \tag{2.23}$$

Here, c is another object from the $X(L, K, \theta)$ Jacobian which is defined in equation (B.2). The proof of this proposition is in Appendix B.1.7. In the absence of a threshold rule, the capital tax must take on a role in controlling automation, so it is additively adjusted by the term (a). This additional term is an equity effect that takes the same form as the RHS of the threshold rule FOC in equation (2.19). The efficiency effect of automation (the LHS of (2.19)) is absent because of the satisfaction of production efficiency. We can now see that the common intuition, that capital must be taxed to keep labor competitive, does hold under plausible circumstances.

The increase of the capital tax will aid in the control of automation by increasing the user cost of capital, re-allocating some tasks to labor at the margin. However, this tax rise will reduce the entire supply of capital. This will lead to an inframarginal effect by reducing capital intensity in tasks that are far from the margin of automation;

tasks unrelated to displacement from automation due to capital's clear comparative advantage.

The optimality conditions for the UBI and labor tax do not substantively change when the threshold rule is removed. This is, again, unsurprising in the case of the UBI as it is just a transfer that should equate the value of funds across entities. The labor tax still considers the effect of labor supply on the automation thresholds; this effect is just different due to the new equilibrium where task thresholds are chosen in a decentralized manner. The effect of automation on production is absent due to production efficiency, but otherwise, the optimality condition for the labor tax is the same as in the case with the threshold rule. The labor tax optimality condition in the case with no threshold rule can be seen in equation (B.13) of Appendix B.1.7.

2.2.3 Implementation

The question is naturally raised as to how a tax authority could target the extensive margin of automation in practice. Acemoglu et al. (2020) come down on the side of output maximizing automation in their main analysis, but when they consider a secondary case where capital taxes are politically restricted to be low, they find it optimal to distort the level of automation in favor of labor. Their reasoning is similar to that of this paper: automation near the margin induces small productivity gains, so a small reduction in automation generates first-order increases in wages with only second-order decreases in output. To implement control over the threshold of automation, they consider task-specific capital taxes that would make the use of capital prohibitively expensive in the tasks the Planner would like to give back to labor. Task-specific capital taxes are also allowed in the setup of Costinot and Werning (2023), but their optimal tax formulas do not explicitly consider the interplay between a general capital tax and distorting the level of automation through task-specific taxes.

The policy tool considered in this paper, a threshold rule, is, in fact, a weaker instrument than task-specific capital taxation. If one has access to task-specific capital taxes, they can precisely control the threshold of automation by making any given task prohibitively expensive for the dispreferred factor. In that case, each threshold could be chosen optimally, whereas in the case of a threshold rule, the restriction to a single policy instrument requires that the threshold only be optimized on average, as shown in equation (2.19). Furthermore, the main result of this paper, that the choice of the optimal capital tax can ignore automation effects, would follow trivially if the thresholds of automation could be controlled by fiat. One does not need to worry about capital's effect on automation if automation can be controlled with complete precision using another instrument.

Models with task-specific capital taxes therefore nest the possibility of a threshold rule as they allow an arbitrary choice of automation thresholds. However, whereas a threshold rule requires the specification of a single variable, task-specific capital taxes require the specification of infinitely-many tax rates. Furthermore, a defining feature of the threshold rule is that it leaves capital undistorted in tasks that are already automated. This allows capital to be undisturbed when acting as a q-complement for labor. For task-specific capital taxes to have this feature, they would have to be specified so precisely that no one pays them in equilibrium. This is, of course, mathematically coherent, but not terribly plausible. Finally, while it is obvious that task-specific capital taxation would make controlling automation with a general capital tax unnecessary, it is perhaps surprising that, despite being a weaker instrument, the threshold rule has the same power to nullify this role for general capital taxation.

Despite the conceptual simplicity of a threshold rule, it is not at all obvious how one would implement such a rule. In practice, a more realistic goal would be taxes that place a fixed cost on an initial act of automation. For instance, a tax on the installation of automating technology would select against so-so automation and would not affect intensive usage conditional on adoption. It would be difficult to precisely determine which capital investments are an act of initial automation and which are inframarginal, but even without perfect targeting, it is helpful for policy makers to understand that the goal is to target the extensive margin of automation while avoiding the infra margin. Note that this is not an argument for a zero capital tax; a zero capital tax would be a knife-edge for equation (2.17). It is just to say that revenue raising and control of the extent of automation in production should be viewed as separate policy goals.

2.3 Dynamic Extension

2.3.1 The Environment

Generations

We will consider a set of generational groups denoted by $\mathcal{G}$ and indexed by g. Generational groups will still be partitioned into labor types, and each generational group will supply a single type of labor for the entirety of their life. Let group g live from period t_g to period $t_g + T_g$. This group will pick sequences of consumption(or savings) and labor supply to maximize the value function

$$V_g(C_{g,t_g}, L_{g,t_g}, C_{g,t_g+1}, L_{g,t_g+1}, \dots C_{g,t_g+T_g}, L_{g,t_g+T_g}) \tag{2.24}$$

subject to a sequence of intertemporal budget constraints

$$a_{g,t+1} + C_{g,t} = \big(1 + (1 - \tau_{K,t})R_t - \delta\big)a_{g,t} + (1 - \tau_{L,t})W_{i,t}L_{g,t} + D_t \tag{2.25}$$

Groups that are alive in the initial period will inherit wealth level $a_{g,0}$, while those born later will come into the world with no wealth. In their final period, a group will consume all of its wealth. Capital taxes are applied to rents, rather than net returns, but this is unimportant as there is only a single type of capital.

Denote by $\mathcal{G}_t$ the set of generational groups alive in time t and by $\mathcal{G}_i$ the set of generational groups of type i. Thus, $\mathcal{G}_{i,t}$ is the set of generational groups of type i alive at time t. We will still have N-many types of workers in each period, and I will assume $\mathcal{G}_{i,t} \neq \emptyset \; \forall t$, i.e. there is always some worker of each type. We will have $L_{i,t} = \sum_{g \in \mathcal{G}_{i,t}} L_{g,t}$. One point of distinction is that the number of generational groups of a given labor type now stands in for the variable determining the size of a labor type in the static model.

The optimization conditions for each of these generations will then be

$$(1 - \tau_{L,t})W_{i,t} = \frac{-V_{g,l_t}}{V_{g,c_t}} \;\; \forall i, t; \; \forall g \in \mathcal{G}_{i,t} \tag{2.26}$$

and

$$V_{g,c_{t-1}} = \big(1 + (1 - \tau_{K,t})R_t - \delta\big)V_{g,c_t} \;\; \forall g \in \mathcal{G}; \; \forall t \in \{t_g + 1, \ldots t_g + T_g\} \tag{2.27}$$

Finally, denote by

$$\mathcal{G}_t^B \equiv \left\{ g \middle| t_g = t \right\}$$

the generations being born at the onset of period t and

$$\mathcal{G}_t^D \equiv \left\{ g \middle| t_g + T_g = t \right\}$$

the generations dying at the conclusion of period t.

Note that this setup, with proper choice of variables, can give both garden-variety two-period OLG models with old and young for each type, as well as infinitely-lived households of each type where the generational component becomes moot. In fact, this setup is general to all deterministic overlapping generations models. It is also worth noting that equation (24) allows for a fairly general set of value functions including, but not limited to, recursive preferences.

Tasks & Production

Each type of labor will again have their own continuum of tasks where they face a time-specific level of automation. That is, all generations in $\mathcal{G}_i$ may supply labor to the interval of tasks X_i and must face a sequence of automation thresholds $\{x_{i,t}\}_t$.

Factors will have potentially time-varying general productivity parameters $A_{J,t} \in \mathbb{R}$ and task-specific productivity functions $\alpha_{J,t} : X_J \to \mathbb{R}$.

Factors are again assigned to tasks where they have the lowest marginal cost. The resulting task thresholds can be described in each time period by the following level set:

$$\left\{(X_t, L_t, K_t, \theta_t) \in \mathbb{R}^{2N+2} \middle| \frac{F_{K,t}(1+\theta_t)}{A_{K,t}\alpha_{K,t}(x_{i,t})} = \frac{F_{L_i,t}}{A_{L_i,t}\alpha_{L_i,t}(x_{i,t})} \;\; \forall i\right\} \tag{2.28}$$

That is, the task allocation is entirely pinned down by intra-temporal objects. Therefore, the implicit function for the thresholds will remain as it was in the static model.

Production will remain as in the static model with a representative, price-taking firm producing output in each period. Output will again be produced using a CES task aggregator where production in each task is perfect substitutes.

Government

The government will select a sequence of linear factor taxes, threshold rules, and UBI to maximize the social welfare function

$$W\left(\left\{V_g\right\}\right) \tag{2.29}$$

subject to a sequence of budget constraints

$$G_t + D_t|\mathcal{G}_t| + \left(1 + R_t - \delta\right)B_t = B_{t+1} + \sum_{g\in\mathcal{G}_t} \tau_{K,t} R_t a_{g,t} + \sum_i \sum_{g\in\mathcal{G}_{i,t}} \tau_{L,t} W_{i,t} L_{g,t} \tag{2.30}$$

Here the government can issue bonds which pay the market rate on capital to satisfy no arbitrage. Note that B_t signifies debt, so it has the opposite sign from household assets. Furthermore, we will have $K_t + B_t = \sum_{g \in \mathcal{G}_t} a_{g,t}$

Finally, the resource constraint for a given period of this economy requires

$$\tilde{F}_t + (1-\delta)K_t = \sum_{g \in \mathcal{G}_t} C_{g,t} + G_t + K_{t+1} \tag{2.31}$$

Equilibrium

For a given sequence of policy choices $\left\{\tau_{L,t}, \tau_{K,t}, D_t, \theta_t\right\}_t$, an equilibrium consists of a sequence of factor prices $\left\{\{W_{i,t}\}_i, R_t\right\}_t$, a sequence of allocations of factors, output, and consumption $\left\{\{L_{g,t}, C_{g,t}\}_g, K_t, Y_t\right\}_t$, and a sequence of assignments of factors to tasks $\left\{\{\Omega_{L_i,t}\}_i, \Omega_{K,t}\right\}_t$ such that:

(i) Each generational groups maximizes utility subject to their budget constraints.

(ii) The representative firm maximizes profit in each period.

(iii) Factor markets clear and the resource constraint is satisfied in each period.

2.3.2 Dynamic Ramsey Problem

The formulation of the dynamic Ramsey problem, and subsequent Lagrangian, is described in Appendix B.2.1. I will assume that a binding, unanticipated sequence of policy choices will be credibly announced at the onset of initial period.

Optimal Policy

Proposition 2.4. *Given an optimal sequence of threshold rules, the optimal sequence for household savings and the UBI must satisfy*

$$\lambda_{g,t+1}\big(1 + (1-\tau_{K,t+1})F_{K,t+1} - \delta\big) + \mu_{t+1}\tau_{K,t+1}F_{K,t+1} = \lambda_{g,t} \tag{2.32}$$

$$\frac{\sum_{g\in\mathcal{G}_t}\lambda_{g,t}}{|\mathcal{G}_t|}=\mu_t \tag{2.33}$$

for each time period and generational group.

We can now see that the main result of Proposition 2.1 holds in the dynamic environment as neither of these equations consider an automation effect. Again, we have the same optimality conditions that would arise in an environment with fixed-share CES production. The reasoning is the same: capital affects the automation threshold, but this effect is proportional to that of the threshold rule. The UBI is just a lump-sum transfer. The proof of this proposition is contained within Appendix B.2.2

The interpretation of equation (2.33) is again that the social value of funds should equal the private value of funds. For equation (2.32), we have a simple equality of marginal cost and marginal benefit. This household must lose a unit of funds in period t to gain an after-tax return in the following period, while the public purse gains revenue from the investment.

Proposition 2.5. *The optimal sequence of threshold rules satisfies*

$$\mu_t\sum_i\frac{\partial F_t}{\partial x_{i,t}}\hat{\alpha}_{i,t}^{-1}=\sum_{i\neq 1,g\in\mathcal{G}_{i,t}\setminus\{g_t'\}}\frac{\eta_{g,t}}{\sigma}\frac{F_{L_1,t}}{F_{L_i,t}}\Big(\frac{\alpha_{L_i,t}(x_{i,t})^{\sigma-1}}{\Lambda_{L_i,t}}\hat{\alpha}_{i,t}^{-1}-\frac{\alpha_{L_1,t}(x_{1,t})^{\sigma-1}}{\Lambda_{L_1,t}}\hat{\alpha}_{1,t}^{-1}\Big) \tag{2.34}$$

in each period.

The first thing to note is that equation (2.34) is just the static optimality condition from Proposition 2.2, adjusted with time subscripts. We then have that the Planner can treat the choice of sequence of threshold rules as though it were a series of static problems to be considered within each time period. This is because the thresholds of automation for a given period are only a function of variables from that period, including the threshold rule.

It remains an open question whether this result is robust to the explicit modeling of endogenous technological change or human capital investment. For example, if technology is determined in part by the research effort of firms, those research efforts would take into account the expectation of future control of the extensive margin of automation. One would not want to invent an automating technology if they know policy makers will prevent them from using it. A similar arguments applies to the modeling of human capital investments. Workers may decide not to train their way into a different occupation if policy makers dampen the effect of automating technology on relative wages. I believe the main value of this result is not to demonstrate, beyond a shadow of a doubt, that the control of the extensive margin of automation is a static problem. Instead, it clarifies what sort of economic processes might make the control of the extensive margin of automation a dynamic decision. Capital accumulation will not suffice.

The remaining policy choices for the dynamic environment are derived in Appendix B.2.3. These include debt policy and the initial capital tax. I do not derive the labor tax as this would require additional separability assumptions on the generations' value functions whereas the goal of this section is to derive the results of Propositions 2.4 and 2.5 in as general an environment as possible. The most important lesson from these additional policy choices is that the Planner could use debt to control the level of automation. Because government debt crowds out capital investment, the choice of debt affects the supply, and thus user cost, of capital, which in turn affects the level of automation. However, much like household savings, this margin of control over automation is unnecessary when aa threshold rule is optimally employed. For the initial capital tax, the surprise nature of the policy announcement opens the well-known problem that optimization requires wealth confiscation in the initial period. The heterogeneous-agents feature of the model complicates this story slightly, as shown in equation (B.17).

Optimal Policy Without a Threshold Rule

I will now consider the set of optimal policy choices when the Planner is unable to target the extensive margin of automation. As stated before, there is no guarantee that political institutions will be able to achieve such targeting. Without a threshold rule to distort automation decisions, the level of automation will be chosen in each period to maximize profits, which will beget production efficiency. This is again equivalent to solving the same problem but with the added constraint that $\theta_t = 0$ for all time periods.

Proposition 2.6. *Given a sequence of decentralized automation decisions, the optimal sequence of household savings must satisfy*

$$\begin{aligned} \lambda_{g,t}\big(1 + (1-\tau_{K,t})F_{K,t} - \delta\big) + \mu_t \tau_{K,t} F_{K,t} = \lambda_{g,t-1} \\ + \frac{c}{K_t} \sum_{i \neq 1} \sum_{g \in \mathcal{G}_{i,t} \setminus \{g'_t\}} \frac{\eta_{g,t}}{\sigma} \frac{F_{L_1,t}}{F_{L_i,t}} \Big(\frac{\alpha_{L_i,t}(x_{i,t})^{\sigma-1}}{\Lambda_{L_i,t}} \hat{\alpha}_{i,t}^{-1} - \frac{\alpha_{L_1,t}(x_{1,t})^{\sigma-1}}{\Lambda_{L_1,t}} \hat{\alpha}_{1,t}^{-1} \Big) \end{aligned} \tag{2.35}$$

for each time period and generational group.

As in Proposition 2.3, the optimality condition determining the capital tax now has an additional term that reflects the automation effect. This is the term on the second line of equation (2.35). We still have the interpretation that marginal cost equals marginal benefit, but now the marginal cost includes the equity consequences of automation brought about by increasing the supply of capital in the period following the savings decision. Put differently, the cost of a household saving is no longer just their forgone consumption, but the effect that act of saving has on the level of automation in tomorrow's economy. The proof of this proposition follows from that of Proposition 2.4 in Appendix B.2.2 but without the cancellation of automation effects.

As in the static model, the lack of a threshold rule does not affect the optimality condition for the sequence of UBI payments. It does, however, affect the optimal choice of debt, as discussed in Appendix B.2.3. Without automation, the Planner would like to use the crowding out effect of government debt to curtail automation. This is in-keeping with the unifying theme of this paper: The Planner should influence the level of automation by changing the supply of capital if and only if they do not have access to control over the extensive margin of automation.

2.4 Numerical Results

To pin down the quantitative implications of the static model of Section 2.2, I will simulate the optimization problem in an environment calibrated to the 2010 American economy. This section will outline the calibration procedure and present optimal policy in an environment with heterogeneous exposure to automation from robotic and software technology.

2.4.1 Calibration

Functional Forms

I will use a utilitarian social welfare function and additively separable utility.

$$W(u_1, u_2, \ldots, u_N) = \sum_i n_i u_i \tag{2.36}$$

$$u_i = \frac{\left(C_i/n_i\right)^{1-\gamma} - 1}{1-\gamma} - \frac{\left(L_i/n_i\right)^{1+1/\varepsilon_L}}{1+1/\varepsilon_L} \tag{2.37}$$

That is, we have CRRA utility of per-person consumption and iso-elastic disutility of per-person labor effort. One can imagine that each household is an occupational guild, so n_i is the number of people in an occupation. Occupations will be broken into the ten occupation partition used in Acemoglu and Autor (2011).

I will set the capital cost function to a scaled, iso-elastic function:

$$\phi(K) = \frac{\xi}{1 + 1/\varepsilon_K} K^{1+1/\varepsilon_K} \tag{2.38}$$

I will set $X_i = (0, 100]$ for each occupation. Following Acemoglu et al. (2020), I will use the following functional form for factor's task-specific productivity schedules: $\alpha_{L_i} = x^{\zeta_i \nu_i}$ and $\alpha_K(x) = x^{\zeta_i(\nu_i - 1)}$ when evaluated over X_i. Here, we can interpret ζ_i as determining labor's comparative advantage and ν_i as determining labor's absolute advantage in higher order tasks. Therefore, we have relative task-specific productivity schedules

$$\frac{\alpha_{L_i}(x)}{\alpha_K(x)} = x^{\zeta_i} \tag{2.39}$$

Data on occupation-specific exposure to automating technologies will come from Webb (2019). This data gives scores quantifying the degree of similarity between the tasks in an occupation's job description and those in the text of patents for the relevant technology. These scores provide percentiles of exposure to software and robotics. I will assume that ζ takes the linear form

$$\zeta_i = \Gamma_0 + \Gamma_1 \cdot exp_i^r + \Gamma_2 \cdot exp_i^s \tag{2.40}$$

where exp_i^r and exp_i^s denote the percentiles of exposure to robots and software.

Finally, I will assume that wealth shares take the form

$$\omega_i = \left(S_{L_i}\right)^{\Theta} \Big/ \sum_j \left(S_{L_j}\right)^{\Theta} \tag{2.41}$$

where S_{L_i} is the income share of group i.

This equation can be derived from $\omega_i = \left(y_{L_i}\right)^{\Theta} \Big/ \sum_j \left(y_{L_j}\right)^{\Theta}$ where y_{L_i} is labor income. The parameter Θ then determines the degree to which the distribution of wealth is more concentrated than the distribution of income. If Θ was equal to one, the distribution of wealth would exactly equal the distribution of income. But if the rich had disproportionately higher levels of wealth, we would have Θ greater than one. This setup is valuable because we need to calibrate to a mapping from income shares into wealth shares, not the empirical wealth distribution itself. Because I have collapsed the economy into occupational groups, it would be inappropriate to target the observed wealth distribution of the United States. The goal is to be able to translate the income shares of my occupational groups into wealth shares using a mapping that is empirically realistic.

Assigned Parameters

Data on the size, wages, and labor supply $\{n_i, W_i, L_i\}$ of occupations are taken from the ACS 2010 5yr on IPUMS. Further details on the cleaning procedure for this data are found in Appendix B.3.1.

Table 2.1 contains the values chosen for each of the model's parameters. Relative risk aversion is set to $\gamma = 1$, which is a conventional choice in the literature. The Frisch elasticity of labor is set to $\varepsilon_L = 0.82$; the value found in Chetty et al. (2011)'s review of labor supply elasticities. Acemoglu et al. (2020) find effective factor taxes of $\tau_L = 0.255$ and $\tau_K = 0.1$ for the United States in 2010. Jakobsen et al. (2020) use administrative wealth records in Denmark to find long-run elasticities of taxable wealth with respect to the after-tax rate of return of 0.77 for the moderately wealthy and 1.15 for the very wealthy, so I will take the rough midpoint of these estimates with an elasticity of unity: $\varepsilon_K = 1$.

A CES task aggregator allows an aggregate factor-level CES interpretation (with the same elasticity of substitution) when the task allocation is held constant. However, when the task allocation is allowed to change, this introduces an additional margin of substitution. Therefore, short-run estimates of the elasticity of substitution, in time frames that do not allow for changes in the task allocation, can be viewed as reflecting the elasticity of substitution across tasks. Conversely, long-run estimates of the elasticity of substitution can be viewed as reflecting both the inherent substitutability of tasks and shifts in the allocation of tasks. To be specific, Oberfield and Raval (2021) estimate an elasticity of substitution of $\sigma = 0.7$ for the US manufacturing sector; i.e. factors are complements in production. However, Karabarbounis and Neiman (2014) estimate that factors are substitutes with $\Sigma = 1.33$. The latter derive this estimate by looking at reductions in the labor share stemming from secular declines in the cost of capital, which allows time for changes in the task allocation to play out. I will assume changes in the task allocation are the source of this additional substitutability which bumps the observed elasticity of substitution past one.

Table 2.1: Parameter Choices

Parameter	Name	Value	Source
Assigned Parameters			
γ	Relative Risk Aversion	1	Convention
ε_L	Frisch Elasticity	0.82	Chetty et al. (2011)
ε_K	Elasticity of Capital Supply	1	Jakobsen et al. (2020)
τ_L	Labor Tax	0.255	Acemoglu et al. (2020)
τ_K	Capital Tax	0.1	Acemoglu et al. (2020)
σ	Short-Run Elasticity of Substitution	0.7	Oberfield and Raval (2021)
Parameters used for Calibration			
S_K	Capital Share of Income	0.4	Convention
Σ	Long-Run Elasticity of Substitution	1.33	Karabarbounis and Neiman (2014)
β_r	Robot Effect on Log Wages	-0.28	Webb (2019)
β_s	Software Effect on Log Wages	-0.11	Webb (2019)
$\Delta \ln R$	Log Reduction in Price of Automating Capital	-2	Hubmer (2023)
Calibrated Parameters			See Text
Θ	Wealth Share Convexity Parameter	1.88	
Γ_0	Comparative Advantage Constant	7	
Γ_1	Comparative Advantage Robot Coefficient	-0.037	
Γ_2	Comparative Advantage Software Coefficient	-0.015	

Calibrated Parameters

To calibrate Θ, I will use the 2010 Survey of Consumer Finances to compute labor income shares and wealth shares by decile. These deciles are computed using the weighting, labor income, and net worth variables contained within the SCF Summary Extract dataset. As a secondary confirmation that the SCF shares capture the salient features of US wealth inequality, the top decile wealth share is similar to that obtained by Saez and Zucman (2016). Taking the functional form assumed in equation (2.41), I pick Θ to minimize the distance between observed and predicted wealth shares. The result can be seen in Figure B·1.

The most salient feature of the model for describing exposure to automation is the choice of $\Gamma \in \mathbb{R}^3$ as this determines the level and heterogeneity of the substitutability of capital and labor. The three values will be determined jointly by targeting the long-run elasticity of substitution and two regression estimates of the effect of exposure on changes in log wages. Despite the joint determination, one can naturally think of Γ_0 as determining the economy-wide elasticity of substitution between capital and labor and $\Gamma_{1,2}$ as determining the degree to which slightly more exposure affects the sensitivity of wages to automation induced by reductions in the user cost of capital.

The identifying variation Karabarbounis and Neiman (2014) consider in their estimate of the long-run elasticity of substitution is a reduction in the cost of producing capital, so I will consider a change in equilibrium prices and quantities brought on from a shift in the price of producing capital, ξ. The details of this procedure are described in Appendix B.3.2.

Next, the effect of occupation-level automation exposure on the sensitivity of wages to automation comes from regressions contained within an old version of Webb (2019). These regressions estimate the effect of exposure to both robots and software on long differences in log wages from 1980 to 2010 at the industry-occupation-level. If we assume that the estimated effect of exposure entirely reflects sensitivity to automation induced by reductions in the user cost of capital, we can view the regression coefficients as approximations of how much more a given change in the log price of capital would affect log wages through the channel of automation if exposure were slightly increased. An additional data point necessary for this pair of calibrating equations is the log change in the price of capital in the period of study. The model in question only has one type of capital, but in mapping parameters to the data, the relevant capital price is that at the margin of automation. I will therefore set $\Delta \ln R = -2$, which is the observed reduction in the price of computer-powered equipment and software from 1982 to 2012 (Hubmer, 2023).

Finally, for each guess of Γ within the calibration process, labor and capital's general productivities and the absolute advantage parameters, $(\{A_{L_i}, \nu_i\}, A_K)$, will be chosen to match income shares, normalize the initial automation thresholds at fifty, and match a capital to output ratio of three. Capital's production cost, ξ, will be chosen to scale the model to match an initial level of output of one hundred. The government share of income is trivially determined by factor tax rates and income shares, which pins down the level of required government spending in the status quo.

2.4.2 Results

The results of this numerical experiment are contained in Table 2.2. In addition to optimal policy choices, the table displays three ways of comparing the optimal allocation with the status quo. The first is consumption equivalence, or the constant multiple Φ that must be applied to everyone's consumption in the status quo to achieve the same level of social welfare as in the optimum. The second is the percent change in output between the status quo and optimum. The third is the percent change in total labor supply between the status quo and optimum. Finally, the UBI is displayed as a proportion of per-capita consumption in the optimum.

To compare the interactions and relative welfare consequence of each policy instrument, I consider optimal policy with and without the threshold rule and UBI. One can immediately see that simple redistribution is of greater significance for welfare than targeting the extensive margin of automation. For reference, this analysis implies an optimal lump-sum transfer in 2010 of roughly \$5,500. One caveat to keep in mind is that optimal taxes and transfers address all of the sources of inequities in the American economy, including automation, whereas the threshold rule can only address displacement from automation. This suggests that the inequality brought on by automation has, thus far, not been of great welfare significance relative to other sources of inequity. Figure B·2 suggests why this may be the case: Exposure to robots and software is not a phenomenon that exclusively affects those with low wages. The occupations of this figure are arrayed in descending order of hourly wage. Indeed, exposure to robots and software seems to particular afflict the middle-skilled, a relationship commonly associated with wage polarization (Acemoglu and Autor, 2011).

Comparing columns (1) and (2) as well as (3) and (4), one can see that the ability to target the extensive margin of automation allows for higher welfare with reduced taxes on capital and greater taxes on labor. The reduction in the capital tax is not surprising in light of Proposition 2.3. With an optimal threshold rule, the capital tax does not need to play a role in holding down automation. Furthermore, with automation held in check, labor taxes can be increased without posing too much of a threat to labor's competitiveness. However, the across the board reduction in the labor tax relative to the capital tax does not appear motivated to keep labor competitive in the task allocation, as in Acemoglu et al. (2020). The relative magnitudes of the factor taxes remain roughly the same even if the Planner is able to target the extensive margin of automation, suggesting that tax rates are not chosen to influence the extent of automation.

Table 2.2: Numerical Results

	Threshold Rule & UBI (1)	UBI (2)	Threshold Rule (3)	Taxation Only (4)
Policy Choices				
τ_L	0.327	0.323	0.146	0.137
τ_K	0.329	0.336	0.287	0.298
θ	0.034	0.0	0.054	0.0
UBI	0.168	0.168	0.0	0.0
Comparative Statics				
Φ	1.0458	1.0457	1.0185	1.0183
$\%\Delta Y$	-0.1241	-0.1242	-0.0334	-0.0332
$\%\Delta L$	-0.0851	-0.0853	0.0385	0.0385

Notes: This table presents the policy choices from local maxima of the Planner's problem. Comparative statics are in reference to the 2010 American status quo, which is assumed to have no threshold rule or lump-sum transfer.

Instead, the marked increase in the taxation of capital from the status quo rate of $\tau_K = 0.1$ is due to both the extreme inequality of capital ownership and the similar supply elasticities of capital and labor. Because capital income goes to agents with already high consumption, and because capital is not particularly elastic, it is welfare improving to levy substantial taxes on capital. Table B.1 derives optimal policy with differing parameters governing the wealth distribution and elasticity of capital supply. In particular, the optimal capital tax is substantially reduced if the wealth distribution is set equal to the income distribution or if the supply elasticity of capital is doubled. However, the optimal threshold rule is not sensitive to these changes, reflecting a different motivation for controlling the extensive margin of automation than for taxing capital.

In the absence of a UBI, reductions in the labor tax become a useful instrument for redistribution as the poorest members of society earn almost all of their income from supplying labor. Furthermore, a higher threshold rule is warranted because the displacement of automation is worse in the absence of any lump-sum redistribution.

We can see that all of the changes considered lead to fairly small welfare gains. This is in-keeping with the well-understood logic of Harberger (1954): Welfare changes are small near the optimum as the FOCs are close to being met. An additional reason for the small utility gain is that the entire economy has been partitioned into ten occupational groups. Because each member of these groups eats the average level of consumption, I have tacitly assumed a great deal of risk sharing that reduces the need for redistribution. A UBI is a form of social insurance, but such insurance would be of less value if you are perfectly insured within an occupational union. There is a similar logic behind why Lucas (2003) found small welfare losses from the business cycle. His calculations involved a representative agent, which may mask the true intertemporal variance in consumption elicited by the business cycle in a world with incomplete markets. Aggregating households into occupations has a similar effect here; the cost of inequality-inducing automation, or inequality generally, is lower as the true variation in consumption is masked. Business cycles impose costs by varying consumption across time, while the static cost of inequality comes from varied consumption in the cross-section. In both cases, the aggregation of households dampens the level of variance in consumption.

While the total change in social welfare is small, the welfare effects for each occupation can be more substantial. The policy recommendations of Table 2.2 involve income losses for some and are, as such, not Pareto improvements. This can be seen in Figure 2·2, which contains the occupation-specific consumption equivalence for the full policy optimum of column (1). The computation is the same as for the aggregate consumption equivalence, except now a multiple on consumption is picked for each occupation to match their specific change in utility. The occupations are arrayed in descending order of hourly wage. As you can see, it is primarily the low wage workers that benefit from the policy change, and their benefit is a great deal larger than that seen when considering social welfare.

Figure 2·2: Consumption Equivalence by Occupation

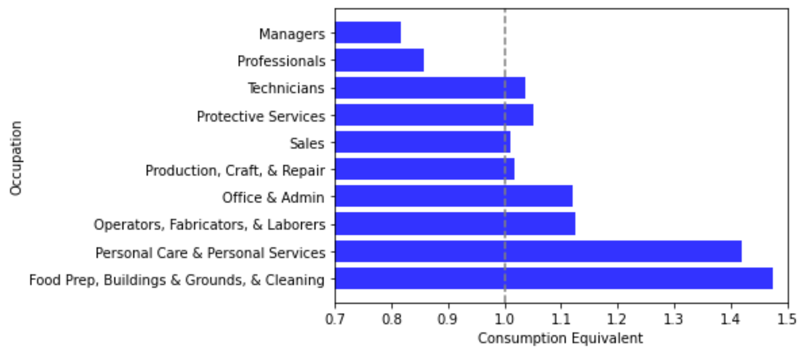

2.5 Conclusion

This paper has considered the optimal policy response to automating technology in a second-best environment. Contrary to the common intuition that capital taxes must be raised to keep labor competitive, capital taxes need only serve as a revenue-gathering instrument when control of the extensive margin of automation is optimal. Capital taxation has no further role to play in manipulating the level of automation in the productive process. This result comes from the complicated relationship between capital and labor in production. Capital deepening on the infra margin of automation is beneficial for labor, but displacement on the extensive margin of automation hurts labor. The Planner would, therefore, like to consider these margins separately. Furthermore, the efficiency value of targeting the extensive margin of automation comes from an important intuition of the task model: the marginal automated task will only be slightly more productive when performed by capital. However, labor's absolute productivity in these automated tasks can be quite high, so the reduction in displacement can have a first-order effect on wages. Reducing the extent of automation at the margin therefore has a second-order effect on output but a first-order effect on the distribution of income. If the Planner stipulates that automation can only happen when capital has a decisive advantage, the tasks regained by labor will be the exact ones where the distortion has the smallest effect on output.

The primary results of this paper were derived in Section 2.2 in a static Ramsey problem with minimal assumptions on either utility functions or the Planner's objective. Preferences for redistribution, the concavity of utility in consumption, and the separability of consumption and labor all play no role in the applicability of capital taxation to the control of automation. The ability to separately consider a general capital tax and the extensive margin of automation depends primarily on the endogenous determination of the task allocation. Because the effect of capital on the level of automation is proportion to that of the threshold rule, the FOC of the threshold causes the effect of further capital supply on the extent of automation to envelope out. This result was shown in Section 2.3 to extend to a dynamic environment where the determination of the sequence of optimal threshold rules can be solved as a sequence of static problems in each period. Distortions of the extensive margin of automation do create production inefficiency, so the desirability of such distortions depends crucially on policy tools being too limited to satisfy the Production Efficiency Theorem of Diamond and Mirrlees (1971a,b).

To apply this framework in an empirically realistic environment, Section 2.4 calibrated the static model to the American economy in 2010. In such an environment, much of the available welfare gain comes from simple, lump-sum redistribution funded by similar taxes on both capital and labor. The ability to target the extensive margin does improve welfare, but much of the inequity that the Planner would like to address does not appear to stem from automation.

The intuitions underlying this paper may have more general applicability to controlling the forces of creative destruction. Indeed, the core trade off in second-best automation policy extends to creative destruction: the victims of displacement benefit from more intensive use of the new technology along the infra margin. Policies pertaining to creative destruction should therefore focus themselves on targeting the extensive margin of displacement. For the sake of clarity, it is worth noting that this logic does not extend to the displacement of fossil fuel technology. For example, in the comparison of solar panels and coal, we are interested in reducing the use of coal on both the extensive and infra margin. In that case, a general tax is more appropriate because we are addressing a cost, in this case an externality, that has the same sign wherever we use coal.

Finally, it is worth stressing that this paper derives its theoretical results in a limit case where the extensive margin can be considered independently of a general tax, which indiscriminately affects both the infra and extensive margin. This setting allows a vantage point from which to illustrate the two opposing influences capital exerts on labor. In practice, the lesson for policy makers is that concerns about displacement from automation should remained focused squarely on the extensive margin of automation.

CHAPTER THREE - OPTIMAL DYNAMIC SPATIAL POLICY

3.1 Introduction

In recent decades, there has been a notable divergence in the trajectories of economic activity among different regions in the United States (Moretti, 2012). Regions specialized in the manufacturing sector have been particularly affected by the increasing import competition (Autor et al. 2013, Caliendo et al. 2019) and the rise of automation (Acemoglu and Restrepo 2020). Looking ahead, the evolving natural resource landscape, influenced by anticipated climate change, is expected to further reshape regional economic activity (Cruz and Rossi-Hansberg, 2023; Hanson, 2023; Arkolakis and Walsh, 2023).

The divergent trajectories of regional economic activity have sparked discussions among policymakers and researchers regarding place-based policies. These discussions revolve around the need to support declining regions or those facing adverse conditions. Critics argue that such policies inadvertently incentivize individuals to stay in, or relocate to, unproductive regions (Glaeser and Gottlieb, 2008; Kline and Moretti, 2014b; Fajgelbaum and Gaubert, 2020). According to this perspective, place-based policies are only justified if they address market failures by internalizing productivity spillovers or other local externalities. The disincentive effect is potentially more problematic if agents make forward-looking migration decisions anticipating future policies. On the other hand, proponents of place-based policies highlight the slow and frictional nature of people's migration responses. They argue that these policies can effectively provide insurance against adverse regional shocks without distorting migration decisions (Bartik, 2020; Gaubert et al., 2021b). Despite the contentious debate, there is currently limited theoretical and empirical understanding of how place-based policies should strike a balance between these two perspectives, particularly in the context of dynamically evolving regional economic activity.

This paper addresses this ongoing debate by examining optimal place-based transfer policy using dynamic quantitative spatial equilibrium models that account for frictional migration and incomplete financial markets. Our main theoretical result is a recursive formula for the optimal spatial transfers that balances the benefit of providing consumption insurance while minimizing distortions from migration responses. We use this theory to assess the optimality of the observed patterns of spatial transfers across U.S. states since 1980.

We begin by establishing a model environment that captures the dynamic evolution of spatial economic activity under forward-looking migration decisions. Locations may differ in their fundamental components, including productivity and amenities, as well as bilateral trade and migration frictions. These fundamental components can arbitrarily evolve over time, capturing the heterogeneous trajectories of regional economic activity. These fundamental factors, along with localized production and amenity externalities, shape the dynamics of wages and amenities in equilibrium. We model forward-looking migration decisions as a dynamic discrete choice problem following the approach of Artuç et al. (2010), Kennan and Walker (2011), and Caliendo et al. (2019). However, unlike these previous papers, which assume a specific functional form of idiosyncratic preference shocks, our framework allows for flexible functional forms and arbitrary correlation across potential destinations.

Using this framework, we study the problem of optimal spatial transfers across locations and over time. The key restriction of the policy is that the Planner cannot directly regulate migration flows (i.e., directly specify the location decisions based on the individual idiosyncratic preference shocks). Therefore, the optimal transfer must overcome the dynamic moral hazard problem associated with agents' migration decision. Given this restriction, the optimal transfers balance two forces. On the one hand, due to incomplete financial markets, the Planner aims to equalize the dispersion in marginal utility through transfers. On the other hand, the Planner seeks to use transfers to incentivize migration towards areas that improve aggregate efficiency.

Our main theoretical result is a recursive expression of the optimal spatial transfers that strikes this balance. Our formula equates the marginal benefit of providing consumption insurance and smoothing, which depends on the marginal utility from a transfer in each location and time, with the marginal efficiency loss, which depends on a fiscal externality arising from distorting migration decisions. Importantly, the fiscal externality depends not only on contemporaneous economic conditions and externalities, but also on future fiscal externalities in a recursive manner. This implies that transfer policies are influenced not only by current economic activity but also by its dynamic evolution. For example, in response to persistent negative shocks in a region, the Planner gradually decreases transfers to provide dynamic incentives for relocation.

To provide further intuition, we consider a special case with a complete financial market. In this case, marginal utility is equalized across agents in different locations at every period. Therefore, the only market failures the Planner aims to address are production and amenity externalities. In this scenario, a simple period-by-period Pigouvian externality correction is sufficient to achieve efficiency. Notably, since future externalities are fully addressed by future transfer policies, contemporaneous policies do not need to directly respond to the future evolution of economic activity and externalities. Therefore, despite considering a dynamic environment, our optimal spatial transfer formula aligns with those of the static spatial equilibrium environment (Fajgelbaum and Gaubert 2020).

We show that our results can be generalized in many different directions. For one thing, we consider the case with a dynamic form of agglomeration externality, i.e., local population affects not only contemporaneous but also future productivity, as well as cross-region agglomeration spillovers (e.g., Desmet et al., 2018; Allen and Donaldson, 2020; Peters, 2022; Cai et al., 2022). Our expression remains to hold by simply incorporating these forces into the recursive expression for the fiscal externality. For another, we also show that our framework can be extended to the case with heterogeneous types of households with respect to preference, productivity, and migration cost (e.g., Fajgelbaum and Gaubert, 2020; Rossi-Hansberg et al., 2019).

Equip with this theoretical framework, we calibrate our model to the evolution of economic activity across U.S. states to assess the optimality of the observed patterns of spatial transfers. In particular, we calibrate our model to exactly replicate the evolution of pre- and post-tax-and-transfer income (hereafter post-tax income in short), consumption expenditure, and population movement across the U.S. states since 1980. Using the calibrated model, we resolve the model by replacing these spatial transfers with the optimal ones prescribed by our framework.

In the data, we observe that the net public transfer, defined as the difference between post- and pre-tax income, is positive for low-income states and negative for high-income states, consistent with the previous research emphasizing the role of spatial transfers for redistribution (e.g., Gaubert et al., 2021a). We also find that these public transfers flow from high-income-*growth* states to low-income-*growth* states. These results are consistent with the observation that public transfers play an important role in mitigating regional income shocks as documented by Asdrubali et al. (1996). Interestingly, we also find that private transfers, defined by consumption expenditure minus post-tax income, tend to flow from high-income states to low-income states, and particularly from high-income-*growth* states to low-income-*growth* states. These results are consistent with the observation that there are pre-existing mechanisms aside from public transfers that mitigate regional income shocks as documented by Asdrubali et al. (1996). Our calibration takes into account the presence of these non-public regional risk-sharing mechanisms.

Our calibrated model reveals substantial and systematic differences between the optimal spatial transfers prescribed by our theoretical framework and the observed ones. In particular, in our baseline specification, the slope between the log consumption expenditure and log pre-tax income should be 0.52 under optimal spatial transfers in 2000, instead of 0.63 as observed in the data. These results indicate that optimal transfers entail more redistribution toward low-income states. Moreover, the slope between the growth rate of consumption expenditure and the growth rate of pre-tax income should be 0.07 from 1980 to 2000 under optimal spatial transfers, instead of 0.41 as observed in the data. Therefore, optimal transfers entail more redistribution toward low-income-growth states. Such large redistribution is desirable in our context because of the large and persistent location decisions across U.S. states as observed in data. Therefore, the benefits of providing additional consumption insurance outweigh the costs of distorting migration decisions.

We find that implementing the optimal transfer policy from 1980 onwards leads to a 0.05 percent increase in consumption equivalent welfare. While this is a significant gain considering the size of the U.S. economy and population, the welfare gains are smaller compared to the range of estimates provided using a static spatial framework (Fajgelbaum and Gaubert, 2020; Rossi-Hansberg et al., 2019). This difference arises because migration is frictional and it takes time for the effects and gains of the policy to materialize. Therefore, using a dynamic framework is important not only to understand the structure of the optimal policy but also to understand the welfare gains resulting from the policy.

Our paper is related to several strands of literature. First, our paper contributes to the literature on place-based policy. A majority of the academic research on place-based policies has focused on its role in ensuring aggregate efficiency by correcting localized externalities (e.g., Glaeser and Gottlieb, 2008; Kline and Moretti, 2014b; Fajgelbaum and Gaubert, 2020; Rossi-Hansberg et al., 2019). More recent work highlights the redistributive motive of place-based policy. Gaubert et al. (2021b) advocate the role of place-based policy as a tool for redistribution across different types of households. Davis and Gregory (2021) focus on its role in equalizing the dispersion of marginal utility within types and discuss a similar policy trade-off as in our paper in a static partial equilibrium environment.[1] To date, little effort has been devoted to analyzing place-based policy in a dynamic environment with heterogeneous and divergent regional economic activity with slow and frictional migration responses (Bartik, 2020). We contribute to this literature by characterizing optimal policies that strike the trade-off between consumption insurance and efficiency in a dynamic general equilibrium environment.

Second, our paper contributes to the literature on the dynamic evolution of spatial economic activity with population mobility. In particular, our framework builds on the approach of modeling forward-looking migration decisions as a dynamic discrete choice problem (Artuç et al., 2010; Kennan and Walker, 2011; Caliendo et al., 2019). Such a framework is used and extended to study various aspects of the dynamics of regional economic activities, such as the regional incidence of trade shocks (Caliendo et al., 2019), spatial development and growth (Desmet et al., 2018; Allen and Donaldson, 2020; Kleinman et al., 2023; Cai et al., 2022), regional depopulation and aging (Giannone et al., 2023b), climate change and natural disasters (Balboni, 2019, Cruz, 2021, Pang and Sun, 2023, Bilal and Rossi-Hansberg, 2023), and the role of financial frictions (Giannone et al., 2023a; Greaney, 2023; Dvorkin, 2023). Despite

[1]The insight that marginal utility is not equalized in spatial equilibrium model dates back to Mirrlees (1972).

the popularity of the framework and the breadth of the applications, there have only been limited attempts to characterize optimal policy in this framework. The closest contribution is by Cameron et al. (2007), who show that market equilibria are efficient in a setting where there are no trade frictions, households are risk neutral, and there are no externalities. We show that there is room for policy intervention in a more general setting and explicitly characterize the optimal spatial transfers.

3.2 Model

We consider a world consisting of J locations. Time is discrete and indexed by t. There is a unit measure of households in the economy. Each location features a competitive labor market. In each location, competitive firms produce differentiated goods that can be traded across locations subject to iceberg trade costs. The productivity and amenity levels of each location are influenced by localized agglomeration externalities.

Households are forward-looking, have perfect foresight, and optimally decide where to move given some initial distribution of labor across locations. Households face costs to move across locations and experience an idiosyncratic shock that affects their moving decision. This formulation of forward-looking migration decisions closely follows the approach of Artuç et al. (2010) and Kennan and Walker (2011) in partial equilibrium and Caliendo et al. (2019) in general equilibrium. However, unlike these previous studies, which assume specific functional forms for idiosyncratic preference shocks, our framework allows for flexible functional forms and arbitrary correlation across potential destinations.

In the following sections, we first outline how households make decisions based on equilibrium prices, wages, and local amenities. We then describe how these prices, wages, and amenities are determined in competitive equilibrium.

3.2.1 Household Decisions

In each period t, households consume a basket of goods produced in various locations. We denote the consumption of goods produced in location i and consumed by households in location j as q_{ijt}.[2] The bundle of products consumed by households is denoted as $c_{jt} = c_{jt}\big(\{q_{ijt}\}\big)$, where $c_{jt}\big(\cdot\big)$ represents a function with constant returns to scale.[3] Additionally, households derive utility from location-specific amenities, denoted as a_{jt}. Amenity a_{jt} is influenced by both exogenous factors (e.g., rivers, beaches) and agglomeration externalities, as elaborated in the next section. We denote the flow utility of households in period t in location j by $u\left(c_{jt}, a_{jt}\right)$.

In each period t, households in location j provide one unit of labor at a wage rate of w_{jt}. In addition to labor income, households receive a location-specific public transfer denoted as τ_{jt}. This transfer aims to capture various forms of subsidies and taxes that are specific to different locations and times. The goal of Section 3.3 is to characterize the optimal choice of $\{\tau_{jt}\}$ by the Planner. Furthermore, households also receive private transfers s_{jt}. We include s_{jt} to account for transfers occurring across regions, encompassing elements such as factor income owned in other locations, pure altruistic transfers (e.g., charities), or payments from private insurance associated with regional income shocks. Together, the household consumption is given by

$$c_{jt} = \frac{w_{jt} + \tau_{jt} + s_{jt}}{P_{jt}}, \tag{3.1}$$

[2]In Appendix C.2.2, we provide an extension that incorporates heterogeneous household types with varying income, consumption, and migration decisions even conditional on location and time.

[3]Although we do not explicitly model nontradables (e.g., services, housing) or intermediate inputs for production, these factors can be captured by reinterpreting $c_{jt}(\cdot)$ as "reduced factor demand" in Adao et al. (2017), who show that any trade models under perfect competition can be represented as a factor exchange economy.

where P_{jt} is the ideal price index of the consumption bundle given by

$$P_{jt} = \min_{\{q_{ijt}\}} \big[\sum_i p_{ijt} q_{ijt} \big] \;\; s.t. \;\; c_{jt}\big(\{q_{ijt}\}\big) \geq 1, \tag{3.2}$$

where p_{ijt} is the price of good produced in i sold in location j at period t. The consumption of each variety is also given by

$$\{q_{ijt}\} = \arg\min_{\{q_{ijt}\}} \big[\sum_i p_{ijt} q_{ijt} \big] \;\; s.t. \;\; c_{jt}\big(\{q_{ijt}\}\big) \geq c_{jt}. \tag{3.3}$$

A key feature of our environment is the presence of dispersion in the marginal utility of consumption, $\partial u_{jt}/\partial c_{jt}$, across locations and over time. In particular, if agents are risk averse, i.e., the utility function is concave, the Planner can enhance households' welfare by offering consumption smoothing through public transfers $\{\tau_{jt}\}$. In other words, a key rationale for the policy is to address the incompleteness of the financial market. We further explore this aspect when analyzing the optimal transfer policy in Section 3.3.

At the end of each period t, households make decisions regarding their next place of residence in the following period. Relocating from location j to location k incurs migration costs $\tilde{\chi}_{jkt}$ in utility. Additionally, each household experiences an idiosyncratic preference shock ε_{jkt} when considering a move from location j to location k in period t. We assume that the vector of idiosyncratic shocks $\varepsilon_{jt} = [\varepsilon_{j1t}, \ldots, \varepsilon_{jJt}]$ follows the cumulative distribution function $G_{jt}(\cdot)$. Our only requirements for $G_{jt}(\cdot)$ are that it is continuous and differentiable. It is worth noting that this specification encompasses the common parametric assumption where ε_{jkt} is assumed to follow an i.i.d. extreme-value distribution (Artuç et al., 2010, Kennan and Walker, 2011, and Caliendo et al., 2019). By allowing for flexible functional forms and arbitrary correlation of shocks across destinations, our model captures flexible substitution patterns of migration choices.[4] Given these assumptions, households' idiosyncratic lifetime utility from residing in location j at time t is given by

$$v_{jt} = u\left(c_{jt}, a_{jt}\right) + \max_{k} \Big[\beta V_{kt+1} - \tilde{\chi}_{jkt} + \varepsilon_{jkt}\Big], \tag{3.4}$$

where β represents the discount factor, and $V_{jt} \equiv \mathbb{E}(v_{jt})$ denotes the expected lifetime utility from residing in location j at time t. Note that the uncertainty of v_{jt} arises because the sequence of future preference shocks has not yet been realized at the time of making movement choices in period t. Note also that households anticipate future public transfers and equilibrium wages when making location decisions, as captured by V_{jt}. The fraction of households that move from location j to location k in period t is given by

$$\mu_{jkt} = Pr\Bigg(k = \arg\max_{k} \Big[\beta V_{kt+1} - \tilde{\chi}_{jkt} + \varepsilon_{jkt}\Big]\Bigg). \tag{3.5}$$

With these movement flows, the population in each location, denoted as L_{jt}, follows the law of motion

$$L_{jt} = \sum_{i} L_{it-1}\mu_{ijt-1}. \tag{3.6}$$

Before proceeding to the description of the equilibrium, we provide a useful reformulation of the aforementioned dynamic migration decision:

Lemma 3.1. *Given $\{V_{kt+1}\}$, the expected lifetime utility and aggregate behavior of this economy's heterogeneous households at time t coincide with the solution to the*

[4]Our formulation does not account for the case where individual preference shocks are serially correlated, as explored by Howard and Shao (2022). In such cases, the Planner must infer the previous idiosyncratic shocks conditional on location and time, adding additional complication to the optimal transfer policy problem. In an extreme scenario where preference shocks are perfectly correlated over time and observable, the model is isomorphic to our extension with multiple types of households as described in Appendix C.2.2.

following representative agent recursive formulation:

$$V_{jt} = u(c_{jt}, a_{jt}) + \max_{\{\mu_{jkt}\}_k} \Big[-\psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1} \Big] \qquad (3.7)$$
$$s.t. \quad \sum_k \mu_{jkt} = 1,$$

where $\psi_{jt}(\cdot)$ is an exogenous function that depends on the migration costs $\{\tilde{\chi}_{jkt}\}$ and the distribution of idiosyncratic preference shocks $G_{jt}(\cdot)$.

The proof and the exact expression for $\psi_{jt}(\cdot)$ is in Appendix C.1.1. This lemma demonstrates that the dynamic discrete choice problem faced by agents with heterogeneous idiosyncratic shocks can be effectively represented as a representative agent's problem, wherein the migration flows μ_{jkt} are directly chosen. This formulation builds upon prior research that formulated static discrete choice problems as representative agent problems (e.g., Anderson et al., 1992, Dubé et al., 2022), but we extend it to a dynamic context and explicitly derive the corresponding direct (as opposed to indirect) utility function under arbitrary preference shocks. While this lemma simply provides an alternative expression for the dynamic migration decision, it significantly aids in formulating the optimal transfer problem. It is worth noting that in the special case where $G_{jkt}(\cdot)$ follows an i.i.d. Type-I Extreme Value distribution with dispersion parameter θ, we have $\psi_{jt}\big(\{\mu_{jkt}\}\big) = \frac{1}{\theta}\sum_k \mu_{jkt} \ln\big(\mu_{jkt}/\chi_{jkt}\big)$, and migration flows follow the familiar expression that $\mu_{ijt} = \frac{\chi_{ijt}\exp(\beta V_{jt+1})^{\theta}}{\sum_k \chi_{ikt}\exp(\beta V_{kt+1})^{\theta}}$, where we define $\chi_{ijt} \equiv -\exp(\tilde{\chi}_{ijt})$. We refer to the functions $\psi_{ijt}(\cdot)$, which capture both the idiosyncratic shocks and migration costs, as "migration frictions".

3.2.2 Competitive Equilibrium

We now discuss how wages, amenities, price indices, and production are determined in competitive equilibrium.

In each location, there exists a representative and competitive producer that utilizes labor as its sole factor of production. Firms operate under a linear production technology with labor productivity denoted as z_{jt}. As we specify below, we allow this productivity to be influenced by both exogenous location-specific factors and agglomeration externalities. When goods are shipped from location i to location j during period t, they incur iceberg trade costs represented by $\kappa_{ijt} \geq 1$. Specifically, to enable residents of location j to consume q_{ijt} units of goods, $\kappa_{ijt} q_{ijt}$ units of goods must be dispatched from location i. We also normalize within-region trade cost such that $\kappa_{iit} = 1$. Given perfect competition in the labor and product markets, the price of good produced in i and sold in j at period t is given by

$$p_{ijt} = \kappa_{ijt} w_{it} / z_{it}. \tag{3.8}$$

The goods market clearing condition at each location and time is given by

$$\sum_j \kappa_{ijt} q_{ijt} L_{jt} = z_{it} L_{it}. \tag{3.9}$$

The productivity and amenities of locations are influenced by both location-specific factors and agglomeration externalities driven by the local population size:

$$z_{jt} = Z_{jt}(L_{jt}), \qquad a_{jt} = A_{jt}\left(L_{jt}\right). \tag{3.10}$$

Here, $Z_{jt}(\cdot)$ and $A_{jt}(\cdot)$ represent functions that are specific to the location and time and depend on the population size. It is worth noting that households do not internalize the impact of their migration decisions on the productivity and amenities of their region. Previous studies have highlighted that, in a static context, addressing these externalities is one important motivation for implementing place-based policies (Glaeser and Gottlieb, 2008; Kline and Moretti, 2014a; Fajgelbaum and Gaubert, 2020). In Section 3.3, we show how these agglomeration externalities influence the optimal spatial transfers.[5]

We assume that public transfers $\{\tau_{jt}\}$ are set to satisfy the government budget constraint:

$$\sum_j \tau_{jt} L_{jt} = 0. \tag{3.11}$$

We also assume that private transfers $\{s_{jt}\}$ have zero net supply:

$$\sum_j s_{jt} L_{jt} = 0. \tag{3.12}$$

Given these conditions, the competitive equilibrium is defined as follows.

Definition 3.1 (Competitive Equilibrium)**.** *Given fundamentals for labor productivity and amenities* $\{Z_{jt}(\cdot), A_{jt}(\cdot)\}$*, migration frictions* $\{\psi_{ijt}(\cdot)\}$*, iceberg trade costs* $\{\kappa_{ijt}\}$*, and the initial population distribution* $\{L_{j0}\}$*, the competitive equilibrium of this economy with transfers consists of movement flows* $\{\mu_{ijt}\}$*, continuation values* $\{V_{jt}\}$*, population distributions* $\{L_{jt}\}$*, consumption decisions* $\{c_{jt}, \{q_{ijt}\}\}$*, prices* $\{P_{jt}, \{p_{ijt}\}\}$*, amenities* $\{a_{jt}\}$*, public transfers* $\{\tau_{jt}\}$*, private transfers* $\{s_{jt}\}$*, interest rates* $\{R_t\}$*, government bonds* $\{B_t\}$*, and wages* $\{w_{jt}\}$ *such that (i) Households' consumption decisions* $\{c_{jt}, \{q_{ijt}\}\}$ *are given by Equations* (3.1) *and* (3.3)*, (ii) Movement flows* $\{\mu_{ijt}\}$ *and continuation values* $\{V_{jt}\}$ *solve the representative household problem* (3.7)*, (iii) Population distributions* $\{L_{jt}\}$ *follow the law of motion* (3.6)*, (iv) Goods prices* $\{P_{jt}, \{p_{ijt}\}\}$ *satisfy Equations* (3.2) *and* (3.8)*, (v) Goods markets clear by Equation* (3.9)*, (vi) Amenities* $\{a_{jt}\}$ *and productivity* $\{z_{jt}\}$ *satisfy Equation* (3.10)*, (vii) Government budget constraint holds* (3.11)*, and (viii) Private transfers* $\{s_{jt}\}$ *are in zero net supply* (3.12)*.*

3.3 Optimal Dynamic Spatial Transfers

In this section, we define and characterize the solution to the optimal spatial transfer problem.

[5]Some studies have examined more general forms of externalities where past population size, as well as population sizes in other locations, generate externalities on contemporaneous productivities and amenities (e.g., Desmet et al., 2018; Allen and Donaldson, 2020; Peters, 2022; Cai et al., 2022). In Appendix C.2.1, we demonstrate how our expression for optimal spatial transfers can be extended to such environments.

3.3.1 Planning Problem

This section sets up a planning problem for the optimal transfer policy. At $t = 0$, households are distributed across locations with measure $\{L_{j0}\}$. We consider a planning problem that sketches the Pareto frontier across households in different initial locations.[6]

Definition 3.2 (Planning Problem)**.** *Given Pareto weights $\{\varkappa_i\}$ attached to individuals in location i in period 0, the Planner solves*

$$\max_{\left\{\{q_{kjt},\mu_{jkt},p_{kjt}\},c_{jt},L_{jt+1},V_{jt},P_{jt},a_{jt},B_{t+1},R_{t+1},w_{jt},\tau_{jt},s_{jt}\right\}_{t\geq 0}} \quad \sum_i \varkappa_i L_{i0} V_{i0}$$

$$s.t. \qquad V_{jt} = u(c_{jt}, a_{jt}) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1} \tag{PK}$$

$$\{\mu_{jkt}\} \in \arg \max_{\{\mu_{jkt}\}:\sum_k \mu_{jkt}=1} u(c_{jt}, a_{jt}) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1} \tag{IC}$$

(3.1)-(3.3), (3.6)-(3.12).

The set of solutions to this problem given an arbitrary Pareto weights $\{\varkappa_i\}$ defines the Pareto frontier.

A crucial constraint for the Planner is that transfers $\{\tau_{jt}\}$ cannot be contingent on the idiosyncratic preference shocks $\{\varepsilon_{ijt}\}$. This restriction is evident in the incentive compatibility constraints for migration decisions (IC). If transfers were allowed to depend on idiosyncratic shocks ε_{ijt}, the Planner could directly manipulate migration flows by providing negative consumption unless individuals choose the desired location. This restriction gives rise to a dynamic moral hazard problem, where the Planner must utilize spatial transfers $\{\tau_{jt}\}$ to influence migration flows while satisfying (IC) and the promise-keeping constraint (PK) that specifies the evolution of households' continuation values.

We view this restriction as capturing a realistic constraint for place-based policies.

[6] Note that our solution encompasses any solution to a concave and differentiable social welfare function by appropriately setting the Pareto weights.

Firstly, in most democratic countries, regulating mobility would violate constitutional rights. Secondly, even in contexts where there are quotas for migration flows, such as the Fukou policies in China, it is unrealistic to assume that the government can observe an individual's idiosyncratic preferences for location decisions. If the government cannot do so, the policy may lead to misallocation of people with respect to these idiosyncratic shocks.[7,8]

We also assume that the Planner can perfectly commit to future policies. Commitment is relevant in our setting because once a location choice has been made, it becomes a sunk decision for households. Therefore, without commitment, the Planner would expropriate all income and equalize marginal utility in each period.[9]

3.3.2 Characterizing Optimal Spatial Transfers

This section focuses on solving for the optimal sequence of spatial transfers. To accomplish this, we follow the primal approach in the public finance literature to focus on a relaxed planning problem where the Planner picks an incentive-compatible sequence of consumption and migration flows, and show that the solution to the original problem coincides with the original one. The relaxed planning problem is defined as follows.

Definition 3.3 (Relaxed Planning Problem)**.** *Given Pareto weights* $\{\varkappa_i\}$ *attached to*

[7]See Glaeser and Luttmer (2003) for evidence that rent control policy functioned as an effective quota for the residential population in New York City, and it may have led to misallocation of residents in terms of willingness to pay to reside in the city.

[8]In Appendix C.1.3, we characterize the solution to the first-best problem and demonstrate that the transfer policies cannot achieve the first-best allocation.

[9]This mirrors the commitment problem in the capital taxation literature (Kydland and Prescott, 1977; Chari and Kehoe, 1990; Klein et al., 2008; Farhi et al., 2012), where investment decisions are sunk at the beginning of each period and the Planner has an incentive to expropriate the entire capital stock to facilitate redistribution.

individuals in location i in period 0*, the Planner solves*

$$\max_{\{\{q_{kjt},\mu_{jkt}\},L_{jt+1},V_{jt}\}_{t\geq 0}} \sum_i \varkappa_i L_{i0} V_{i0}$$

$$s.t. Z_{it}(L_{it})L_{it} = \sum_j \kappa_{ijt} q_{ijt} L_{jt}, \qquad L_{jt+1} = \sum_i L_{it}\mu_{ijt} \tag{RC}$$

$$V_{jt} = u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1} \tag{PK}$$

$$\{\mu_{jkt}\} \in \arg \max_{\{\mu_{jkt}\}:\sum_k \mu_{jkt}=1} u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1}. \tag{IC}$$

The relaxed planning problem differs from the original in that the Planner directly specifies consumption and migration flows without considering the equilibrium determination of prices. Therefore, the Planner's choice set is larger in the relaxed problem than in the original. However, we show in Appendix C.1.2 that there exists supporting prices and transfers which implement the solution to the relaxed problem as a competitive equilibrium. Thus, the solution to the relaxed planning problem is also a solution to the original planning problem.

This relaxed planning problem also highlights the fundamental trade-off faced by the Planner. In the presence of a location experiencing a negative productivity shock $Z_{jt}(\cdot)$, the concave utility function implies higher marginal utility, creating an incentive for the Planner to increase consumption c_{jt}. However, this action increases migration flows μ_{ijt} into the location experiencing negative productivity shock through the constraint (IC). The optimal transfer policy must strike a balance between these two forces.

To characterize the solution to the planning problem, it is useful to introduce the concept of the fiscal externality in location j in period t: Δ_{jt}. Formally, we define Δ_{jt} as the Lagrange multiplier associated with the population law of motion in (RC) in the dual of the Relaxed Planning Problem 3.3. In simple terms, Δ_{jt} represents the marginal benefit to the government budget of relocating one unit of population to location j in period t. In Appendix C.1.2, we demonstrate that Δ_{jt} can be recursively expressed as follows:

$$\Delta_{jt} = w_{jt} - P_{jt}c_{jt} + w_{jt}\gamma_{jt}^{z} + \frac{\partial_a u_{jt}}{\partial_c u_{jt}} P_{jt}A_{jt}\gamma_{jt}^{a} + \frac{1}{R_{t+1}}\sum_k \mu_{jkt}\Delta_{kt+1}, \tag{3.13}$$

where $\gamma_{jt}^{Z} \equiv \frac{\partial \log z_{jt}}{\partial \log L_{jt}}$ and $\gamma_{jt}^{A} \equiv \frac{\partial \log a_{jt}}{\partial \log L_{jt}}$ represent the elasticities of the agglomeration externalities for productivity and amenity, respectively; and $\partial_c u_{jt} \equiv \frac{\partial u_{jt}}{\partial c_{jt}}$ and $\partial_a u_{jt} \equiv \frac{\partial u_{jt}}{\partial a_{jt}}$ represent the marginal utility of consumption and amenity, respectively; and R_t is the marginal rate of substitution for the Planner between period t and $t-1$. In this expression, the first two terms, $w_{jt} - P_{jt}c_{jt}$, capture the net resources utilized by households in location j and time t, accounting for labor productivity minus consumption expenditure. The next two terms correspond to the agglomeration externalities for productivity and amenity. Finally, the last term captures future fiscal externalities. Specifically, locating an individual in location j generates fiscal externalities in the future, depending on the probability of migration from j to other locations.

It is worth noting that, if $\Delta_{jt} > \Delta_{kt}$, the Planner can enhance the total budget by relocating individuals from location k to location j. However, the Planner may choose to maintain this gap if the marginal utility from consumption is lower in location j. Our main proposition characterizes the consumption levels $\{c_{jt}\}$ that characterize this trade-off.

Proposition 3.1. *If a competitive equilibrium with transfers solves the Planner's problem, then the following condition is satisfied for all locations j and time $t \geq 0$:*

$$-\sum_i L_{it}\sum_k \Delta_{kt+1}\frac{1}{P_{jt+1}}\frac{\partial \mu_{ikt}}{\partial c_{jt+1}} = \sum_i L_{it}\mu_{ijt}\left(\frac{\beta R_{t+1}\partial_c u_{jt+1}/P_{jt+1}}{\partial_c u_{it}/P_{it}} - 1\right), \qquad \forall j, t \tag{3.14}$$

where Δ_{jt} *is defined recursively by Equation (3.13), and* R_{t+1} *is the marginal rate of substitution for the Planner between period* $t+1$ *and* t*, implicitly determined to satisfy the government budget constraint (3.11).*

The proof of Proposition 3.1 can be found in Appendix C.1.2. This proposition implicitly specifies the consumption levels $\{c_{jt+1}\}$ that equalize the marginal cost and benefit of increasing transfers τ_{jt+1}. On the left-hand side of Equation (3.14), we summarize the marginal cost of increasing τ_{jt+1} by distorting migration decisions. An incremental increase in τ_{jt+1} leads to a consumption increase of $1/P_{jt+1}$. This, in turn, affects the migration shares to destinations k by $\frac{1}{P_{jt+1}} \frac{\partial \mu_{ikt}}{\partial c_{jt+1}}$. It is important to note that the migration patterns change not only for destination j but also for all other destinations $k \neq j$. For each destination k, a unit change in migration share is associated with a fiscal externality Δ_{kt+1}, which encompasses the net resource cost and agglomeration externality in location k at period $t+1$ (Equation (3.13)). Summing up these effects across all origin locations i, multiplied by L_{it}, yields the total efficiency loss resulting from the marginal increase in transfers τ_{jt+1}. Note that the efficiency loss would be zero if fiscal externalities Δ_{jt} were equalized across locations since a marginal reallocation of the population across locations would have no effect on the government's budget.

On the right-hand side of Equation (3.14), we summarize the marginal benefit of increasing τ_{jt+1} from consumption insurance. Marginal transfers in τ_{jt+1} provide consumption smoothing benefits if the marginal utility of income in location j at $t+1$, $\partial_c u_{jt+1}/P_{jt+1}$, is relatively high compared to that of the origin location i at t, $\partial_c u_{it}/P_{it}$, adjusted by the marginal rate of intertemporal substitution: βR_{t+1}. It is worth noting that if perfect consumption smoothing is achieved, i.e., the usual Euler equation is satisfied so that $\beta R_{t+1} \partial_c u_{jt+1}/P_{jt+1} = \partial_c u_{it}/P_{it}$, the right-hand side of Equation (3.14) becomes zero, implying no benefit from public transfers for consumption insurance and smoothing. We will revisit this scenario in the next section when discussing a special case of a complete financial markets.

The trade-off between providing insurance and incentivizing efficient migration, as presented in Proposition 3.1, resonates with the incentive-insurance trade-off emphasized in the unemployment insurance literature (Baily, 1978 and Chetty, 2006) and the dynamic labor contract literature (Burdett and Coles, 2003 and Balke and Lamadon, 2022). In the unemployment insurance literature, the trade-off arises between the benefits of unemployment insurance and the fiscal externality resulting from changes in the probability of transitioning out of unemployment. In the dynamic labor contract literature, firms face a trade-off in tenure contracts between providing consumption smoothing and altering workers' effort and churning probability.

Proposition 3.1 further highlights the forward-looking nature of optimal spatial transfers. The recursive nature of fiscal externality, as presented in Equation (3.13), demonstrates that altering migration patterns not only generates fiscal externalities in the present but also has implications for the future. In Appendix C.3, we show through a numerical simulation that these dynamic considerations lead to frontloading transfers to locations with negative permanent shocks, i.e., provide a large transfer and gradually decrease transfers to provide additional incentive for out-migration.

It is worth noting that the expression in Proposition 3.1 does not depend on the Pareto weights $\{\varkappa_{jt}\}$. This is because the Pareto weights are encapsulated in the consumption levels in period t. Conditional on those values, Proposition 3.1 specifies the growth rates of the marginal utility of consumption at every period after $t = 0$ irrespective of the Pareto weights. The initial period consumption level at $t = 0$, in turn, is determined to simply equalize the marginal utility of consumption up to Pareto weights, i.e., $\varkappa_i \partial_c u_{i0}/P_{io} = \varkappa_j \partial_c u_{j0}/P_{j0}$ for all i, j, since the migration decision occurs after the consumption decision within a period.

Finally, we turn our attention to the construction of the public transfers τ_{jt} necessary to achieve the optimal consumption levels identified in Proposition 3.1. Once we have determined the optimal consumption values, we can establish the corresponding optimal public transfers by utilizing the households' budget constraint (3.1). Specifically, the optimal transfers are given by the expression:

$$\tau_{jt} = P_{jt}c_{jt} - w_{jt} - s_{jt}. \tag{3.15}$$

If private transfers $\{s_{jt}\}$ are exogenous, this equation provides a direct characterization of the optimal public transfers. However, if private transfers s_{jt} respond to equilibrium conditions (including $\{\tau_{jt}\}$), the determination of $\{\tau_{jt}\}$ becomes a fixed-point problem that needs to be solved.

3.3.3 Special Case: Complete Financial Markets

In this subsection, we explore a special case of our model that assumes a complete financial market, providing further insight into Proposition 3.1.

There are two settings that attain a complete financial market in our setting. The first case involves the introduction of Arrow-Debreu securities that pay out for each location, time, *and idiosyncratic preference shocks of each individual.* Note that, without the availability of securities contingent on idiosyncratic shocks, the financial market is still incomplete. We formally present this model extension in Appendix C.1.4. The second case assumes risk neutrality (constant marginal utility $\partial_c u_{jt}$ across locations and time) and equalized consumption prices across locations ($P_{it} = P_{jt}$ for all $i, j \in J$). It is important to note that risk neutrality alone does not guarantee the completeness of financial market due to differential growth rates of prices faced by individuals moving across locations. In both scenarios, we assume that private transfers s_{jt} are set to zero for simplicity, as they would introduce additional distortions in migration flows under the complete financial market.[10]

[10]If s_{jt} is not equal to zero, the expression for τ_{jt} in Corollary 3.1 is simply replaced by $\tau_{jt} + s_{jt}$,

We acknowledge that both of these settings rely on unrealistic and highly specific assumptions. Therefore, the purpose of this discussion is not to present empirically relevant cases, but rather to elucidate the intuition behind our main result in Proposition 3.1.

In the scenario of a complete financial market, there are no benefits from providing additional consumption smoothing. Consequently, the sole objective of policy becomes correcting production and amenity agglomeration externalities. More formally, in the case of complete financial markets, the right-hand side of Proposition 3.1 becomes zero. To equate the left-hand side with zero, it is sufficient to equalize fiscal externalities Δ_{jt} across locations for each period t. The Planner can achieve this via a Pigouvian correction, i.e., setting transfers to offset the agglomeration externalities that enter as a part of fiscal externality (3.13), as demonstrated in the following corollary.

Corollary 3.1. *Under complete financial markets, the optimal set of public transfer* $\{\tau_{jt}\}$ *is given by*

$$\tau_{jt} = w_{jt}\gamma_{jt}^{z} + \frac{\partial_a u_{jt}}{\partial_c u_{jt}} P_{jt} A_{jt} \gamma_{jt}^{a} - K_t, \tag{3.16}$$

where K_t *is a lump-sum set to satisfy the government budget constraint (3.11).*

The proof of Corollary 3.1 is in Appendix C.1.4. As mentioned earlier, it is intuitive that the optimal transfer only depends on agglomeration externalities due to the absence of objectives related to providing consumption insurance. An interesting result is that future fundamentals and externalities do not directly influence the optimal transfers. This is because future transfer policies fully address future externalities. Therefore, despite the consideration of a dynamic environment, our optimal spatial transfer formula recovers that of a static environment (Fajgelbaum

meaning that the net amount of public and private transfers serve as a Pigouvian correction.

and Gaubert 2020).[11,12]

3.3.4 Extensions

In this section, we lay out several extensions to our optimal policy framework.

Dynamic and Cross-Regional Externalities First, we consider an extension with a dynamic form of agglomeration externality, i.e., local population affects not only contemporaneous but also future productivity, as well as cross-region agglomeration spillovers, as pursued by Desmet et al. (2018); Allen and Donaldson (2020); Peters (2022); Cai et al. (2022). In particular, we assume that the productivity in location j in time t takes the form: $Z_{jt} = g_{jt}\big(\bar{Z}_{jt}, \{L_{kt}\}\big)$, where $g_{jt}\big(\,\cdot\,\big)$ is an arbitrary function, $\{L_{kt}\}$ is the profile of population size across locations (including those outside location j), and $\bar{Z}_{jt}$ is a fundamental component of the productivity. $\bar{Z}_{jt}$, in turn, evolves according to $\bar{Z}_{jt} = h_{jt}\big(\{\bar{Z}_{kt-1}\}, \{L_{kt-1}\}\big)$, where $h_{jt}\big(\,\cdot\,\big)$ is an arbitrary function, which can depend on the profiles of fundamental productivity in the previous period and the profiles of population sizes in the previous period. This specification allows changes in fundamental productivity to follow a dynamic, endogenous process that is itself dependent on the population distribution. For instance, one could allow the production of knowledge to have a spatial component. Furthermore, we can allow for the possibility of a lagged agglomeration effect.

We show in Appendix C.2.1 that the basic features of Proposition 3.1 carry over to this case. The optimization condition of Equation (3.14) continues to hold with straightforward adjustments accounting for the additional externalities. The only

[11] See Golosov et al. (2014) for a related result in the context of a pollution externality in a dynamic general equilibrium setting. When dynamic externalities arise, such that the current population distribution affects future externalities, Corollary 3.1 can be adjusted to incorporate these dynamic externalities, as further discussed in Section 3.3.4.

[12] It is also worth noting that, in the special case of complete financial markets, the transfers of Corollary 3.1 implement the utilitarian first-best allocation as a competitive equilibrium.

substantive change is that the fiscal effects of movement choice, previously summarized by Δ_{jt} in Equation (3.13), must now account for a recursive term that describes the dynamic effect of altering fundamental productivity. Nonetheless, the basic logic of the formula, where the Planner equalizes the distortion cost with the consumption smoothing benefit of a marginal transfer, is unchanged with the inclusion of a general, dynamic relationship between labor productivity and the population distribution across space.

Heterogeneous Household Types In Appendix C.2.2, we provide an extension that incorporates heterogeneous household types with varying income, consumption, and migration decisions even conditional on location and time, as pursued by Fajgelbaum and Gaubert (2020) and Rossi-Hansberg et al. (2019) in a static environment. The basic features of Proposition 3.1 carry over to this case, except that the fiscal externality Δ_{jt} is now specific to each household type. In particular, the type-specific Δ_{jt} depends on the cross-type agglomeration spillovers in this extension.

Origin-Destination Specific Transfers In the baseline model, we have restricted ourselves to location-specific transfers. In some countries or contexts, one may argue for a richer implementation of transfers that depends on households' past location as well. In Appendix C.2.3, we extend our model to an environment where transfers can depend on households' origin and destination location. We show that the basics of our formula remain unchanged in such cases except that now the formula holds for each origin-destination pair.

3.4 Data and Calibration

To take the model to the data, we use as an empirical setting the distribution of economic activity across U.S. states since 1980. This section discusses our main data sources and how we calibrate our model to data. We also document salient patterns of spatial transfer across U.S. states over time.

3.4.1 Data

Our quantitative analysis focuses on the 48 contiguous U.S. states plus the District of Columbia, excluding Alaska and Hawaii.

Pre-Tax Income, Taxes, and Public Transfers We follow Gaubert et al. (2021a) to obtain the measure of average pre-tax income, taxes and transfer in each state from the Bureau of Economic Analysis (BEA)'s Regional Account. Pre-tax income equals wages, employer-provided benefits, proprietors' income, dividends, interest, and rent and excludes capital gains and thereby corporate retained earnings. Taxes include all major federal, state, and local taxes except sales taxes. Transfer includes social security payment including Social Security Disability Insurance, as well as all major government transfers including Social Security and Medicaid. For brevity, in what follows, we refer to the pre-tax income plus the net tax and transfers as "post-tax income".

Consumption Expenditure We obtain personal consumption expenditure from Bureau of Economic Analysis (BEA)'s Regional Account. Unlike the aforementioned data on income, taxes, and public transfers, this data only start from 1997. In the subsequent section, we explain how we impute these values before 1997.

Consumer Price Index We first obtain Regional price parities (RPPs) from BEA to measure the differences in price levels across states in 2008. We then use the state-level inflation rates from Hazell et al. (2022) to impute the price indices since 1980.

Migration Flows We use data on bilateral five-year migration flows between U.S. states from the U.S. population census from 1980-2000 and from the American Community Survey (ACS) after 2000. We define a period in the model as equal to five years to match these observed data. We interpolate between census decades to obtain five-year migration flows for each year of our sample period. Following Kleinman et al. (2023), to take account of international migration to each state and fertility/mortality differences across states, we adjust these migration flows by a scalar for each origin and destination state, such that origin population in year t pre-multiplied by the migration matrix equals destination population in year $t+1$, as required for internal consistency.

Trade Flows In the U.S., the Commodity Flow Survey (CFS) reports the flow of manufacturing goods shipped between CFS zones in the US every five years. There are two key limitations in this data set. First, the CFS only started in 1997. Second, the CFS does not report the service trade flows. To overcome this data limitation, we follow the approach of Fajgelbaum and Gaubert (2020) to use the gravity equation of trade flows to find the unique estimates of trade flows between states that are consistent with the actual distance between states, observed pre- and post-tax income, under the assumption that trade frictions are isoelastic with respect to distance.[13] In our baseline analysis, we set the elasticity of distance to trade flow as $\psi = -1$. In Section 3.5, we show that our results are robust to alternate assumptions of this elasticity.

3.4.2 Observed Patterns of Spatial Transfers over Time

Before proceeding with the analysis for the optimal transfers, we present salient patterns of spatial transfer across U.S. states over time. In particular, we demonstrate

[13] See Appendix C.5.1 for further details on this calibration procedure.

that both public and private transfers involve positive net transfers from high-income states to low-income states, as well as from high-income-growth states to low-income-growth states.

To investigate this question, we run the following regression:

$$\log(w_{jt} + \tau_{jt}) = \beta^{\tau} \log(w_{jt}) + \xi_t^{\tau} + \eta_j^{\tau} + \epsilon_{jt}^{\tau}. \tag{3.17}$$

Here, j is the state, t is year, w_{jt} is pre-tax income, τ_{jt} is the net transfer (so that $w_{jt} + \tau_{jt}$ is the post-tax income), ξ_t^{τ} is the year fixed effect, η_j^{τ} is the state fixed effect, and ϵ_{jt}^{τ} is the residual. In our analysis below, we investigate the results both without and with state fixed effects; the former specification informs the patterns of tax-and-transfer between low- and high-income states, and the latter specification informs those between slower- and faster-income growth states. We run this regression from 1997 to 2019 at an annual frequency.

Table 3.1: Observed Transfers vs Income

	log (Post-Tax-and-Transfer Income)				log (Consumption Expenditure)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
log (Pre-Tax-and-Transfer Income)	0.74***	0.75***	0.81***	0.77***				
	(0.02)	(0.01)	(0.04)	(0.02)				
log (Post-Tax-and-Transfer Income)					0.89***	0.52***	0.41***	0.45***
					(0.08)	(0.06)	(0.07)	(0.08)
Observations	1173	1173	459	575	1173	1173	459	575
Adjusted R^2	0.992	0.999	0.999	0.999	0.951	0.995	0.996	0.995
Year FE	X	X	X	X	X	X	X	X
State FE		X	X	X		X	X	X
Sample			Before 2005	Low Income			Before 2005	Low Income

Notes: This table reports the results of the regression specification (3.17) for Columns (1)-(4) and (3.18) for Columns (5)-(8). The sample of this regression is the U.S. states from 1997 to 2019 at an annual frequency.

Columns (1)-(4) of Table 3.1 present the results. Column (1) starts with the specification without state fixed effects. The result shows that a one percent increase in the pre-tax income is associated with a 0.74 percent increase in the post-tax income. These results are consistent with Gaubert et al. (2021a), who document that there is a net transfer from high-income states to low-income states. In Column (2), we present the results with the specification with state fixed effects. The result indicates that a state experiencing a one percent increase in the growth rate of pre-tax income is associated with a 0.75 percent increase in the growth rate of post-tax income. These results are consistent with the observation that public transfer plays a role in mitigating regional income shocks as documented by Asdrubali et al. (1996). These results are robust to solely focusing on before 2005 (Column 3) and by focusing on low-income states (Column 4).

Columns (5)-(8), in turn, focus on the patterns of private transfers using the following regression specification:

$$\log(w_{jt} + \tau_{jt} + s_{jt}) = \beta^s \log(w_{jt} + \tau_{jt}) + \xi_t^s + \eta_j^s + \epsilon_{jt}^s, \tag{3.18}$$

where the dependent variable $(w_{jt} + \tau_{jt} + s_{jt})$ is the consumption expenditure (so that s_{jt} is defined as consumption expenditure minus post-tax income). Column (5) starts with the specification without state fixed effects. The result shows that a one percent increase in the post-tax income is associated with a 0.89 percent increase in the post-tax income. Therefore, while the magnitude is smaller compared to the public transfer, private transfer flows from high-income to low-income states. In Column (6), we present the results with the specification with state fixed effects. The result indicates that a state experiencing a one percent increase in the growth rate of post-tax income is associated with a 0.52 percent increase in the growth rate of consumption expenditure. These results are consistent with the observation that there are pre-existing mechanisms aside from public transfers that mitigate regional income shocks as documented by Asdrubali et al. (1996). These results are robust to solely focusing before 2005 (Column 7) and by focusing on low-income states (Column 8).

To further illustrate the patterns of public and private transfers, in Figure 3·1, we present the patterns of pre-tax income, post-tax income, and consumption expenditure, for 2000 (Panel a) and the growth from 1980-2000 (Panel b). Since we only observe the consumption expenditure after 1997, we use regression (3.18) to predict the patterns of private consumption (and private transfer as its difference from the observed post-tax income). More specifically, we use the estimated regression coefficients β^s and state fixed effects η_j^s to predict the value of consumption expenditure, where we adjust ξ_t so that the net private transfer is zero, i.e., $\sum_j s_{jt} L_{jt} = 0$. Panel (a) shows that, in 2000, states with one percent higher pre-tax income are associated with 0.76 percent higher post-tax income and 0.63 percent higher consumption expenditure. Panel (b) shows that states with one percent higher pre-tax income growth are associated with 0.78 percent higher post-tax income growth and 0.41 percent higher consumption expenditure.

Overall, the evidence suggests that both public and private transfers play a role in mitigating the differences in nominal income levels across states, as well as the differences in the growth rates of nominal income across years. In what follows, we assess to what extent these patterns align with the optimal transfers through the lens of our framework.

3.4.3 Parametric Model

To facilitate our quantitative analysis, we parameterize our model using common specifications found in the existing literature. We specify a utility function that is additively separable between consumption and amenity, where the consumption component follows a constant relative risk aversion (CRRA) form:

$$u(c_{jt}, a_{jt}) = \begin{cases} \frac{c_{jt}^{1-\rho}}{1-\rho} + \log a_{jt} & \rho \neq 1 \\ \log c_{jt} + \log a_{jt} & \rho = 1. \end{cases} \tag{3.19}$$

Figure 3·1: Observed Transfers in the U.S. Over Time

(a) 2000 **(b)** Change from 1980-2000

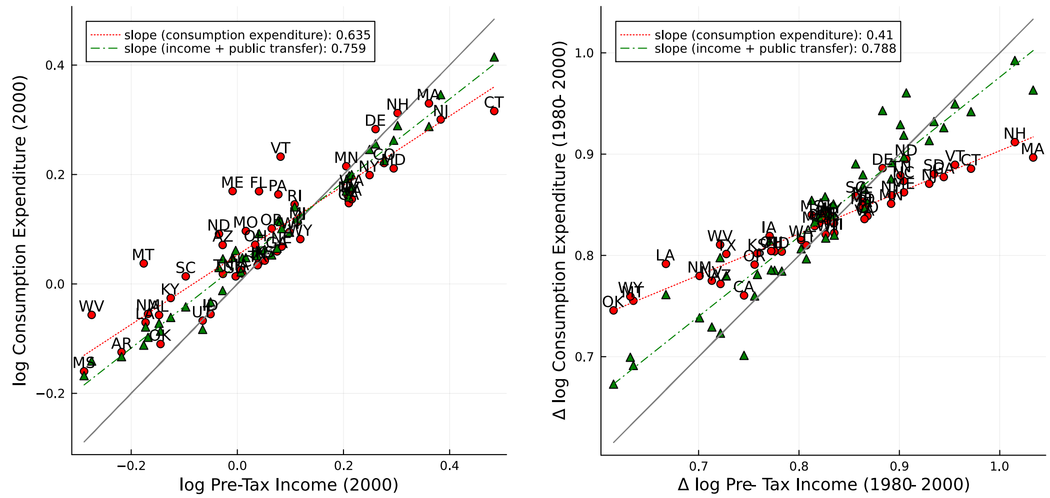

Notes: These figures plot the patterns of pre-tax income, post-tax income, and consumption expenditure, for 2000 (Panel a) and the growth from 1980-2000 (Panel b). To construct the consumption expenditure in 1980, we use the estimated regression coefficients β^s and state fixed effects η_j^s to predict the value of consumption expenditure, where we adjust ξ_t so that the net private transfer is zero, i.e., $\sum_j s_{jt} L_{jt} = 0$.

We further assume that the consumption bundle takes the form of a constant elasticity of substitution aggregator across varieties produced in different locations:

$$c_{jt}\left(\{q_{ijt}\}\right) = \left(\sum_i q_{ijt}^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}. \tag{3.20}$$

We assume that idiosyncratic preference shocks for migration decisions follow the i.i.d. type-I extreme value distribution following Artuç et al. (2010) and Caliendo et al. (2019), so the migration friction function $\psi_{it}(\cdot)$ is given by

$$\psi_{it}\left(\{\mu_{ijt}\}\right) = \frac{1}{\theta}\sum_j \mu_{ijt} \log\left(\mu_{ijt}/\chi_{ijt}\right). \tag{3.21}$$

We assume that productivity and amenity spillovers take an iso-elastic form such that

$$Z_{jt}\left(L_{jt}\right)=\bar{Z}_{jt}L_{jt}^{\gamma^{Z}}, \quad A_{jt}\left(L_{jt}\right)=\bar{A}_{jt}L_{jt}^{\gamma^{A}}, \tag{3.22}$$

where γ^Z and γ^A are parameters that govern the spillover elasticities, and $\bar{Z}_{jt}$ and $\bar{A}_{jt}$ are productivity and amenity fundamentals.[14]

3.4.4 Calibration

To take our model to the data, we use the parametric model discussed in the previous section. Doing so requires the values of structural parameters $\{\beta, \rho, \theta, \sigma, \gamma^A, \gamma^Z\}$ and location fundamentals, including productivities $\{\overline{Z}_{jt}\}$, amenities $\{\overline{A}_{jt}\}$, bilateral migration costs $\{\chi_{ijt}\}$, and bilateral trade costs $\{\kappa_{ijt}\}$. We discuss the calibration of these parameters in turn.

Structural Parameters $\{\beta, \rho, \theta, \sigma, \gamma^A, \gamma^Z\}$

Table 3.2 summarizes our baseline calibration of the structural parameters. We set the 5-year discount factor equal to the conventional value of $\beta = (0.95)^5$. We set the relative risk aversion to $\rho = 1$, which corresponds to logarithmic utility. We assume a value for the migration elasticity of $\theta = 0.5$, which is within the range of parameters estimated and calibrated in the literature (Kleinman et al., 2023; Caliendo et al., 2019; Cruz, 2021). We set the agglomeration spillovers for amenity and productivity to $\gamma^Z = \gamma^A = 0.05$, which is within the range of parameters found in the literature (e.g., Melo et al., 2009). For all parameters, we provide sensitivity analysis and discuss how our quantitative results are affected by the choices of these parameter values.

[14] In Appendix C.4, we list up the system of equations that determine the equilibrium and optimal allocation in this parametric setting.

Table 3.2: Baseline Parameter Calibration

Parameter	Description	Value
β	Discount Factor	0.95^5
ρ	Relative Risk Aversion	1
θ	Migration Elasticity	0.5
σ	Trade Elasticity	5
γ^A	Amenity Spillover Elasticity	0.05
γ^P	Productivity Spillover Elasticity	0.05

Fundamentals and Transfers

We now discuss how we calibrate productivities $\{\overline{Z}_{jt}\}$, amenities $\{\overline{A}_{jt}\}$, bilateral migration costs $\{\chi_{ijt}\}$, bilateral trade costs $\{\kappa_{ijt}\}$, public transfers $\{\tau_{jt}\}$, and private transfers $\{s_{jt}\}$ from the data on trade flows, pre- and post-tax income, prices, and migration flows. Since we only observe the data until 2015, we need to make an assumption regarding the evolution of fundamentals and transfers after 2015. Below, we show how to calibrate fundamentals to exactly replicate the observed data between 1980 to 2015.[15]

Productivity $\{\overline{Z}_{jt}\}$ To recover productivity, we first back out the variety-specific prices sold within the region using the CES demand in equation (3.20), $p_{jjt}^{1-\sigma} = P_{jt}^{1-\sigma} x_{jjt} / \sum_{\ell} x_{\ell jt}$, where P_{jt} is the observed consumer price index and x_{ijt} is the nominal trade flow from state i to j in period t. Furthermore, using equation (3.8) and the normalization of within-region trade cost at one ($\kappa_{iit} = 1$), we can back out $z_{jt} = w_{jt}/p_{jjt}$. Together with the parametric assumption on the productivity spillover (3.22), we can back out $\overline{Z}_{jt} = z_{jt} L_{jt}^{-\gamma^Z}$. We also assume that the exogenous component of productivity remains constant after 2015, i.e., $\overline{Z}_{jt} = \overline{Z}_{jT}$ for all $t > T$.

[15]Our approach of recovering the fundamentals to exactly fit the observed data patterns is isomorphic to the "dynamic-hat algebra" approach in Caliendo et al. (2019). Explicitly recovering fundamentals has the additional benefit of clarifying identifiable objects from the data, as well as clarifying the possibility of relaxing parametric assumptions, such as the non-log utility in consumption.

Bilateral Trade Costs $\{\kappa_{ijt}\}$ Using the CES demand in equation (3.20), the ratio of nominal trade flows satisfies $x_{ijt}/x_{jjt} = (\kappa_{ijt} z_{it} w_{it})^{1-\sigma} / (z_{jt} w_{jt})^{1-\sigma}$. Using this equation, together with the value of $\{z_{jt}\}$ obtained in the previous procedure and the observed pre-tax income $\{w_{jt}\}$, we recover $\{\kappa_{ijt}\}$. We also assume that bilateral trade costs remain constant after 2015, i.e., $\kappa_{ijt} = \kappa_{ijT}$ for all $t > T$.

Amenities $\{\overline{A}_{jt}\}$ and Bilateral Migration Frictions $\{\chi_{ijt}\}$ We start by introducing the normalization of these variables such that $\overline{A}_{jt} = 1$ and $\chi_{jjt} = 1$. Note that these normalizations do not affect the subsequent estimation and simulation of the optimal spatial transfers.[16] We back out off-diagonal elements of migration frictions, $\{\chi_{ijt}\}$ in two steps. We first back out those in the last data period, $t = T$, using the migration flow in period T, $\{\mu_{ijT}\}$, and the future evolution of fundamentals and transfers. Here, we also assume that bilateral migration frictions do not change after 2015, such that $\chi_{ijt} = \chi_{ijT}$ for all $t > T$. Once we obtain $\{\mu_{ijT}\}$, we sequentially back out $\{\chi_{ijt}\}$ for $t < T$ backward starting from $t = T - 1$ using migration flows $\{\mu_{ijt}\}$ in each period. Appendix C.5.2 describes the details of this procedure.

Public Transfers $\{\tau_{jt}\}$ For the time period when we observe the data, we calibrate them as the difference between pre-tax and post-tax income. To calibrate anticipated public transfers for future periods ($t > T$), we assume that the marginal rate of public transfer with respect to income is unchanged from T. In particular, we assume that $\tau_{it} = \varphi_{it} w_{it} L_{it} + \varrho_t$, where φ_{it} is the marginal rate of public transfer and ϱ_t is a lump-sum component to satisfy the government budget constraint ($\sum_i \tau_{it} L_{it} = 0$). We assume that marginal rate of public transfer is unchanged from the last observed

[16]More specifically, under the assumption of additively separable utility (3.19), we cannot separately identify amenity and migration frictions. Similarly, we cannot separately identify the level of migration frictions *from* location j at period t and the migration frictions *to* location j in $t-1$. Note that, although these variables cannot be separately identified, the normalization does not affect any equilibrium predictions.

period, such that $\varphi_{it} = \varphi_{iT}$ for all $t > T$.

Private Transfers $\{s_{jt}\}$ For the time period that we observe the data, we calibrate them as the difference between post-tax income and consumption expenditure. Outside these periods, we assume that private transfers are determined based on time-invariant state-specific components and the fluctuation of post-tax income. More specifically, we use the estimated regression coefficients β^s and state fixed effects η_j^s from regression (3.18) to predict the value of consumption expenditure, where we adjust ξ_t so that the net private transfer is zero, i.e., $\sum_j s_{jt}L_{jt} = 0$. We also use the same relationship between private transfer and the post-tax income when we solve for optimal tranfers.

Pareto Weights $\{\varkappa_j\}$ Our framework presented in Section 3.3 allows us to trace the Pareto frontier using an arbitrary set of Pareto weights attached for individuals in each location in the initial period. For quantitative illustration, we use a particular value for $\{\varkappa_j\}$ given its high dimensionality. In our baseline specification, we calibrate the Pareto weights such that the optimal consumption patterns in the initial data period perfectly coincide with those observed in the data. Namely, by normalizing the Pareto weights in location 1 by $\varkappa_1 = 1$, we set the Pareto weights for j, $\varkappa_j = \frac{\partial_c u_{10}/P_{10}}{\partial_c u_{j0}/P_{j0}}$, so that the marginal utility from transfers is equalized across locations in the initial period up to Pareto weights. We also present how our results are affected by alternatively assuming a utilitarian Planner, i.e., $\varkappa_j = 1$ for all location j.

3.5 Optimal Dynamic Spatial Transfers in the U.S.

In this section, we use our calibrated model to compute the optimal spatial transfers in the U.S. since 1980. In particular, using the calibrated model in Section 3.4.4, which exactly matches the observed spatial transfers in the data between 1980-2015, we solve the model by replacing these observed public transfers with the optimal ones prescribed under Proposition 3.1 (see Appendix C.4 for the full system of equations). In this section, we assess the differences between optimal and observed transfers, population distribution, and welfare gains from implementing the optimal transfer.

3.5.1 Optimal vs. Observed Transfers

Figure 3·2 presents the patterns of consumption expenditure and pre-tax income across locations and over time, thereby highlighting the patterns of spatial transfer. In Figure 3·2a, we plot the log of observed consumption expenditure against the log of observed pre-tax income for each state in 1980 (labeled "observed"). As discussed in Section 3.4.2, the regression slope is 0.62 and flatter than one, reflecting the fact that both public and private transfers contribute to the reallocation of income from high- to low-income states. The figure overlays the same relationship under our optimal allocation (labeled "optimal"). The observed and optimal consumption expenditures perfectly coincide in 1980, which is true by construction from our calibration of Pareto weights (Section 3.4.4).

In Figure 3·2b, we present the same relationships in 2000. We find that the slope of the log of consumption expenditure to the log of pre-tax income under the optimal allocation (0.525) is flatter than the same slope under the observed allocation (0.635). These results indicate that optimal transfer entails more redistribution from high- to low-income states in 2000.

In Figure 3·2c, we focus on the changes in consumption expenditure and pre-tax income over time. In the data, a one percent increase in the growth rates of the pre-tax income between 1980 to 2000 is associated with a 0.41 percent increase in the growth rates of consumption expenditure. As documented in Figure 3·1b of Section 3.4.2, these patterns are driven by the fact that both public and private transfers play a role in regional risk sharing. In contrast, under optimal transfers, we find a slope of 0.07. In other words, under our parametrization, more than 90 percent of the differences in the state-level income growth should be redistributed from high-income-growth to low-income-growth states. These patterns indicate that, in our context, the benefits of providing additional consumption insurance outweigh the costs of distorting migration decisions.

Finally, in Figure 3·2d, we present the regression coefficients of log of consumption expenditure on log of pre-tax income period-by-period under the observed and optimal allocations. As already mentioned, these regression coefficients are identical between the observed and optimal transfers in 1980. However, from 1980 onward, the regression coefficients under the observed allocation diverge and stay higher than those under the optimal allocation. In the long run, the coefficient in the observed equilibrium is around 0.55, while the coefficient under the optimal transfer is around 0.45. Therefore, optimal transfer entails more redistribution from high- to low-income states in the long run.

In Table 3.3, we provide a sensitivity analysis of these results under alternative calibration. In Row (1), we report the regression coefficients of the log of consumption expenditure on the log of observed pre-tax income in 2000 and 2075. For 2075, we use our assumption about the evolution of fundamentals and transfers as we described in Section 3.4.4. We also report the regression coefficients of the log difference of consumption expenditure on the log difference of observed pre-tax income from 1980 to 2000.

In Row (2), we report these regression coefficients under optimal transfers prescribed from our baseline specification. Consistent with Figure 3·2, we find that the optimal transfer entails more redistribution from high-income states to low-income states, and particularly more so from high-income-growth states to low-income-growth states.

Figure 3·2: Observed vs Optimal Transfer in the U.S. Over Time

(a) 1980

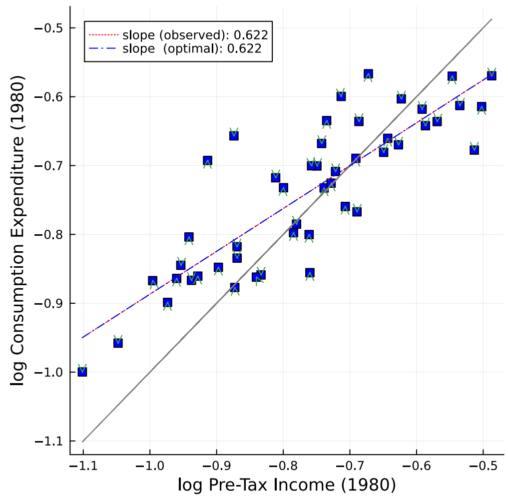

(b) 2000

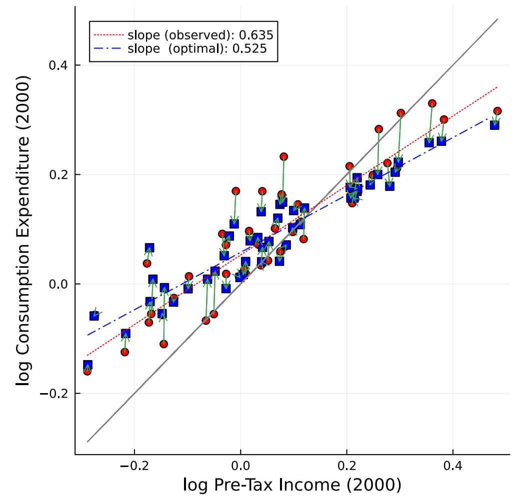

(c) Change from 1980-2000

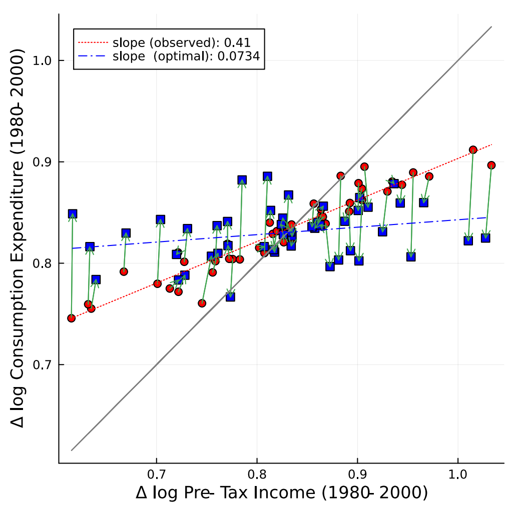

(d) Net Transfer Rates Over Time

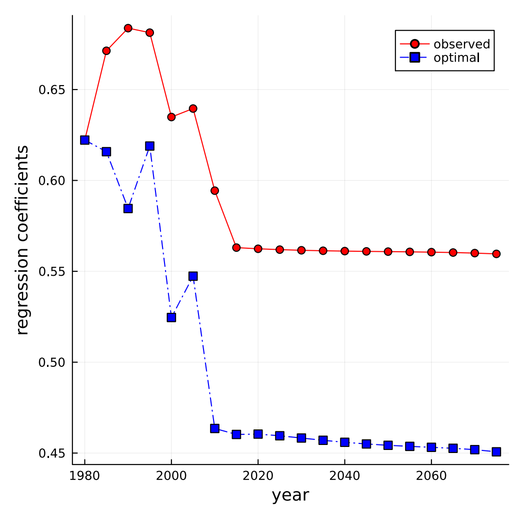

Notes: These figures compare the patterns of the log of consumption expenditure against the log of pre-tax income for each state, under observed allocation (labeled "observed") and under optimal allocation (labeled "optimal"). Figure 3·2d presents the regression coefficients of log of consumption expenditure on log of pre-tax income period-by-period under observed and optimal allocation.

In Rows (3) and (4), we report our results under alternative values for the relative risk aversion parameter ρ instead of the log utility ($\rho = 1$). When we assume risk neutrality ($\rho = 0$; Row 3), we find that the regression coefficients are significantly larger than the observed ones, consistent with the interpretation that the insurance benefit arises because of the concavity of the utility function. It is worth noting that the regression coefficients are less than one even under risk neutrality. This pattern arises because of the spatial differences in the consumer price index. That is, in a location where the price index is low, the marginal value of transfer is higher. In our context, poorer states tend to have lower price indexes (Diamond and Moretti, 2021), which supports positive net transfers toward these areas even under risk neutrality. When we assume the relative risk aversion of $\rho = 3$ (Row 4), we find that the regression coefficients are smaller than our baseline specification. This finding is consistent with the interpretation that the insurance benefit is higher if agents are more risk-averse.

In Rows (5) and (6), we report our results under alternative values for the migration elasticity θ instead of our baseline calibration ($\theta = 0.5$). As we increase the migration elasticity to $\theta = 1$ (Row 5), $\theta = 2$ (Row 6), and $\theta = 5$ (Row 7), the regression coefficients of log of consumption expenditure on log of pre-tax income in 2000 and 2075 increase and exceed the coefficients under observed allocation. This finding is consistent with the interpretation that a higher migration elasticity implies a larger distortion in migration decisions for the same amount of transfer. At the same time, even if we set the migration elasticity to $\theta = 5$ (Row 7), the regression coefficient of log consumption expenditure growth on log pre-tax income growth from 1980 to 2000 is 0.31, which is below the coefficient under the observed allocation (0.41). Therefore, even if we set the migration elasticity to a much higher value than the typical estimates in the existing literature, optimal transfers still entail more redistribution from high-income-growth states to low-income-growth states.

In Rows (8) and (9), we shut down the productivity and amenity spillover, respectively. In both scenarios, we find that the regression coefficients are slightly smaller than the ones in our specification, indicating that the presence of agglomeration externalities favors less redistribution toward low-income states. At the same time, the differences in regression coefficients from our baseline specification are relatively modest, highlighting a larger role of consumption insurance and migration responses in shaping the optimal transfers in our dynamic context.

In Rows (10) and (11), we find that alternative values for the trade elasticities σ have minimal effects on our results. In Rows (12) and (13), we find that the alternative values of the distance elasticity on trade cost used for calibrating trade flows (ϕ) play a minimal role in the regression coefficients. Therefore, conditional on the observed differences of price indices (which are always calibrated at the same value across all specifications), trade frictions across states play a minimal role in governing the shape of optimal transfers.

In Row (13), we find that a higher value of the discount factor ($\beta = 0.98^5$) leads to lower regression coefficients than our baseline specification, especially in the long run (2075). This pattern is consistent with the interpretation that a larger discount factor implies a greater benefit from consumption smoothing.

Finally, in Row (14), we consider a utilitarian Planner (i.e., equal Pareto weights $\varkappa_i = 1$ for all i) instead of imposing the implied Pareto weights to rationalize the observed patterns of transfer in 1980. We find a substantially smaller regression coefficient in 2000 at 0.18, which increases to 0.41 in 2075. These patterns are consistent with the interpretation that a utilitarian Planner will redistribute more toward low-income states starting from the initial period than our baseline specification. However, the gap decreases over time, because the initial location in 1980 does not affect each worker's location in the long run.

Table 3.3: Optimal Net Transfer Rates under Alternative Specifications

	Specification	Slope: log Cons. Exp. on log Pre-Tax-and-Transfer Income		
		2000	2075	Growth 1980-2000
(1)	Observed	0.63	0.56	0.41
(2)	Baseline	0.52	0.45	0.07
(3)	Lower Relative Risk Aversion $\rho = 0$	0.98	0.77	0.82
(4)	Higher Relative Risk Aversion $\rho = 3$	0.48	0.30	-0.02
(5)	Higher Migration Elasticity $\theta = 1$	0.62	0.66	0.15
(6)	Higher Migration Elasticity $\theta = 2$	0.75	0.90	0.25
(7)	Higher Migration Elasticity $\theta = 5$	0.96	1.2	0.31
(8)	No Productivity Spillover $\gamma^Z = 0$	0.52	0.43	0.06
(9)	No Amenity Spillover $\gamma^A = 0$	0.49	0.42	0.05
(10)	Lower Trade Elasticity $\sigma = 3$	0.53	0.45	0.07
(11)	Higher Trade Elasticity $\sigma = 8$	0.52	0.45	0.07
(12)	Lower Dist. Elas. on Trade Cost $\phi = -0.5$	0.52	0.45	0.07
(13)	Higher Dist. Elas. on Trade Cost $\phi = -2$	0.53	0.45	0.07
(14)	Higher Discount Factor $\beta = 0.98^5$	0.52	0.39	0.06
(15)	Utilitarian Planner	0.18	0.41	0.14

Notes: This table reports the regression coefficients of the log of consumption expenditure on the log pre-tax income, in 2000 and 2075, as well as those of the growth rates from 1980 to 2000, under observed allocation (Row 1) and under optimal allocation (Row 2 onwards).

3.5.2 Optimal vs. Observed Population Distribution

Figure 3·3 shows how the optimal spatial transfer changes the spatial distribution of population across the U.S. In Figure 3·3a, we show the log change in population size between 1980 and 2000 against the log change in pre-tax income during the same period, as directly observed in data ("observed") and under optimal spatial transfers ("optimal"). Figure 3·3b presents the same patterns for log changes from 1980 to 2075.

In the data, we find that there is a weak negative correlation between these two variables. In particular, sun-belt states (e.g., Nevada, Arizona, Florida, Texas) have experienced more modest growth in pre-tax income than the average between 1980 and 2000, while increasing their population size at the fastest rates. These patterns are consistent with an interpretation that these sun-belt states experienced a relative decline in housing costs and attracted population, as documented by Glaeser and Gottlieb (2009).

We now compare the observed growth in population size and pre-tax income against those under the optimal policy. In both 2000 and 2075, we find that states with faster growth in nominal income tend to show slower population growth relative to those under the observed spatial transfers. These patterns are consistent with the interpretation that optimal consumption insurance toward regions of slower income growth indeed reallocates the population toward these states.

Figure 3·3: Observed vs Optimal Population Size in the U.S. Over Time

(a) 2000 **(b)** 2075

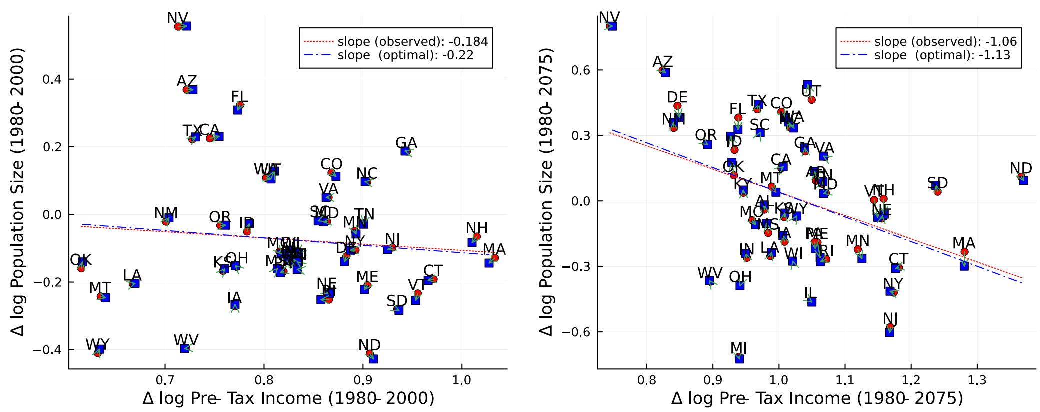

Notes: These figures compare the patterns of the population size changes as observed in the data (labeled "observed") and under the optimal transfer (labeled "optimal").

3.5.3 Welfare Gains from Optimal Transfers

Table 3.4 presents the consumption-equivalent welfare gains by implementing optimal spatial transfer relative to those under the observed allocation. In our baseline calibration (Row 1), we find welfare gains of 0.05 percent by implementing the optimal spatial transfer. While this is a significant gain considering the size of the U.S. economy and population, the welfare gains are smaller by an order of magnitude compared to the welfare gains from optimal policy in a static framework (Fajgelbaum and Gaubert, 2020; Rossi-Hansberg et al., 2019). As we further argue below, this difference arises because migration is frictional and it takes time for the effects and the gain of the policy to materialize.

In Rows (2)-(4), we report the welfare gains under higher values for the migration elasticity θ. We find that the welfare gains increase to 0.076 percent with $\theta = 1$ (Row 2), 0.16 percent with $\theta = 2$ (Row 3), and 0.42 percent with $\theta = 5$ (Row 4). These results are consistent with the interpretation that a larger migration elasticity implies that the population reallocates more quickly as a response to spatial transfers and hence there is significantly less discounting of the welfare gains.

In Rows (5) and (6), we report how our results change by shutting down productivity and amenity spillovers. We find that the welfare gains are 0.05 percent without productivity spillovers and 0.058 percent without amenity spillovers, which are similar to 0.05 percent in our baseline scenario. Therefore, in our context, we find that the welfare gains arising from the agglomeration externality are substantially smaller than those resulting from the benefit of consumption insurance.

In Rows (7) and (8), we find that alternative values for the trade elasticities σ have minimal effects on our results. In Rows (9) and (10), we also find that alternative values for ϕ, the elasticity of trade flows on geographic distance, have modest effects on our results.

In Row (11), we find that a larger discount factor ($\beta = 0.98^5$) leads to 0.055 percent welfare gains, which is slightly higher than our baseline specification. This result is consistent with the interpretation that a higher discount factor leads to more weight in the future after the population responds enough to the policy change.

3.6 Conclusion

We study optimal transfer policy in dynamic spatial equilibrium models with frictional migration and incomplete financial markets. A key policy trade-off is to provide con-

Table 3.4: Optimal Net Transfer Rates under Alternative Specifications

	Specification	Welfare Gains (%)
(1)	Baseline	0.050
(2)	Higher Migration Elasticity $\theta = 1$	0.076
(3)	Higher Migration Elasticity $\theta = 2$	0.16
(4)	Higher Migration Elasticity $\theta = 5$	0.42
(5)	No Productivity Spillover $\gamma^Z = 0$	0.050
(6)	No Amenity Spillover $\gamma^A = 0$	0.058
(7)	Lower Trade Elasticity $\sigma = 3$	0.049
(8)	Higher Trade Elasticity $\sigma = 8$	0.050
(9)	Lower Dist. Elas. on Trade Cost $\phi = -0.5$	0.051
(10)	Higher Dist. Elas. on Trade Cost $\phi = -2$	0.039
(11)	Higher Discount Factor $\beta = 0.98^5$	0.055

Notes: This table reports the consumption-equivalent welfare gains in 1980 under optimal spatial transfer after 1980 relative to those under the observed spatial transfer. We assume Pareto weights calibrated in 1980 to rationalize the observed transfers in 1980.

sumption insurance while minimizing the distortion of migration flows. We derive a recursive formula for optimal spatial transfers that strikes this balance. We calibrate our model to U.S. states and find that the U.S. economy would benefit from increased transfers to low-income-growth states. Welfare gains from optimal transfers are substantial but smaller than in a framework abstracting from slow migration adjustment.

Our framework can be applied to many different applications. In ongoing work, we plan to apply our framework to particular sources of regional shocks, such as import competition shocks from China, and examine optimal policy in the presence and absence of those shocks. Such an exercise will clarify how the government should manage regional shocks dynamically, such as gradually decreasing transfers as a response to negative persistent shocks to provide dynamic incentives for out-migration.

CHAPTER ONE APPENDIX

A.1 Theory Derivations & Proofs

A.1.1 Equilibrium Conditions

To derive relative prices of technology-specific goods, plug the intermediate limit price condition (1.20) into the intermediate demand condition (1.18) to obtain the equilibrium level of output for intermediates:

$$y_{j\iota t} = \frac{(1-\alpha) p_{jt} Y_{jt} \Upsilon a_{j\iota t}}{\gamma w_{\ell t}}. \tag{A.1}$$

Plugging this into the production function for technology-specific goods (1.2) gives us

$$Y_{jt} = \Lambda_{jt} \Big(\frac{(1-\alpha) p_{jt} \Upsilon A_{jt}}{\gamma w_{\ell t}} \Big)^{\frac{1-\alpha}{\alpha}}. \tag{A.2}$$

Plugging Equation (A.2) into the input demand condition (1.17) gives us

$$p_{jt} = \frac{\big(\frac{1}{\alpha}\big)^{\alpha} \big(\frac{1}{1-\alpha}\big)^{1-\alpha} \big(r_j + \omega_j \tau_t\big)^{\alpha} \big(\gamma w_{\ell t}/\Upsilon\big)^{1-\alpha}}{A_{jt}^{1-\alpha}}, \tag{A.3}$$

which implies that relative prices of technology-specific goods follow

$$\frac{p_{jt}}{p_{Jt}} = \Big(\frac{r_j + \omega_j \tau_t}{r_J + \omega_J \tau_t} \Big)^{\alpha} \Big(\frac{A_{jt}}{A_{Jt}} \Big)^{\alpha-1}, \tag{A.4}$$

which gives us Equation (1.22) from the main text.

Turning to the innovation side of the economy, the research firm solves the problem described in Equation (1.23). This yields the condition

$$\xi_{jt} \chi_j \eta s_{jt}^{\eta-1} \phi_{jt} \Pi_{jt} = w_{st}, \tag{A.5}$$

which gives us the research condition (1.25) in the main text.

To establish the existence and uniqueness of the economy's equilibrium path, it suffices to show that there exists a unique equilibrium allocation of scientists $\{s_{jt}\}$ in each period. Consider that the research condition (1.25) and fixed supply of scientists (1.12) defines, at each point in time, the mapping $\mathbb{T}s$ which follows

$$\mathbb{T}s_j = \frac{\big(\chi_j \xi_{jt} \phi_j(\{A_{it-1}\}) \Pi_{jt}(\{s_i, A_{it-1}\}, \tau_t)\big)^{\frac{1}{1-\eta}}}{\sum_i (\chi_i \xi_{it} \phi_{it} \Pi_{it})^{\frac{1}{1-\eta}}} \mathcal{S}. \tag{A.6}$$

Note that innovation profits $\{\Pi_{jt}\}$ are a function of realized technology $\{A_{jt}\}$, which can be written as a function of scientists $\{s_{jt}\}$ and inherited technology $\{A_{jt-1}\}$ via the law of motion for technology (1.8). Now, we need to establish that there exists a unique fixed-point for $\mathbb{T}s$. The fixed supply of scientists (1.12) establishes a compact convex domain set $\mathbb{X}_s$ for $\mathbb{T}s$, and the definition of $\mathbb{T}s$ implies that it maps into $\mathbb{X}_s$ as well. $\mathbb{T}s$ is continuous in $\{s_{jt}\}$ by assumption, so Brouwer's fixed point theorem establishes existence.

For uniqueness, define the function $\mathbb{U}s \equiv s - \mathbb{T}s$, so a fixed-point of $\mathbb{T}s$ is a root of $\mathbb{U}s$. A root of $\mathbb{U}s$ will be unique if the Jacobian of $\mathbb{U}s$ is either a P-matrix for every element of a closed rectangular region of $\mathbb{R}^J$ that contains $\mathbb{X}_s$ or positive (negative) quasi-definite for every element of $\mathbb{X}_s$. Hence, this condition implies that the equilibrium allocation of scientists is unique as well.

To derive demand for technology-specific goods under the nested-CES production structure described in Section 1.2.5, first consider the final producer's demand for sector-level output. The final producer solves

$$\max_{\{E_{\theta t}\}} \mathcal{Y}_t - \sum_\theta p_{\theta t} E_{\theta t}, \tag{A.7}$$

where $p_{\theta t}$ is the price for sector θ. This yields the condition

$$\Omega_t^{\frac{\lambda-1}{\lambda}} \nu_\theta^{\frac{1}{\lambda}} \left(\frac{\mathcal{Y}_t}{E_{\theta t}}\right)^{\frac{1}{\lambda}} = p_{\theta t}. \tag{A.8}$$

Next, producers at the sector level solve

$$\max_{Y_{\theta dt},Y_{\theta ct}} p_{\theta t}E_{\theta t} - p_{\theta ct}Y_{\theta ct} - p_{\theta dt}Y_{\theta dt} \tag{A.9}$$

where $p_{\theta ct}$ and $p_{\theta dt}$ are the price of clean and dirty production in sector θ, respectively. This yields the condition

$$p_{\theta t}\left(\frac{E_{\theta t}}{Y_{\theta et}}\right)^{\frac{1}{\sigma}} = p_{\theta et}, \tag{A.10}$$

where $e \in \{c, d\}$.

A.1.2 Task-Based Microfoundation for CES Shares

Following Acemoglu and Restrepo (2022b), let final output be a CES aggregator over a unit interval of tasks

$$\mathcal{Y}_t = \Omega_t\Big(\int_0^1 q_{xt}^{\frac{\lambda-1}{\lambda}}\,dx\Big)^{\frac{\lambda}{\lambda-1}}, \tag{A.11}$$

where q_{xt} is task-level output.

Task-level output follows

$$q_{xt} = \sum_{\theta} \upsilon_{\theta x} e_{\theta xt}, \tag{A.12}$$

where $e_{\theta xt}$ is the demand for sector θ goods in task x, and $\upsilon_{\theta x}$ is the productivity of sector θ goods in task x. Market clearing requires total demand for sector θ goods across tasks to equal supply:

$$\int_0^1 e_{\theta xt}dx = E_{\theta t}. \tag{A.13}$$

Task-level output is perfect substitutes, so the sector with the lowest marginal cost is used to produce a given task. Thus, task-level prices follow

$$p_{xt} = \min_{\theta} \frac{p_{\theta t}}{\upsilon_{\theta x}}, \tag{A.14}$$

and task-level demand follows

$$\Omega_t^{\frac{\lambda-1}{\lambda}}\left(\frac{\mathcal{Y}_t}{q_{xt}}\right)^{\frac{1}{\lambda}} = p_{xt}. \tag{A.15}$$

Suppose that the unit interval of tasks is partitioned into subsets where only one sector can produce. That is, each sector has a set of tasks

$$\begin{aligned} \mathcal{X}_\theta &= \left\{x \in [0,1] \middle| \upsilon_{\theta x} = 1\right\} \qquad \text{(A.16)} \\ \forall x \notin \mathcal{X}_\theta : \quad \upsilon_{\theta x} &= 0 \\ \cup_\theta \mathcal{X}_\theta &= [0,1] \\ \forall \theta, \theta' : \quad \mathcal{X}_\theta \cap \mathcal{X}_{\theta'} &= \emptyset. \end{aligned}$$

Plugging task-level demand (A.15) into the market clearing condition (A.13) allows us to derive equilibrium sector prices

$$p_{\theta t} = |\mathcal{X}_\theta|^{\frac{1}{\lambda}} \Omega_t^{\frac{\lambda-1}{\lambda}} \left(\frac{\mathcal{Y}_t}{E_{\theta t}}\right)^{\frac{1}{\lambda}}, \tag{A.17}$$

where $|\mathcal{X}_\theta| \equiv \int_{x \in \mathcal{X}_\theta} dx$ is the mass of tasks produced by sector θ. Plugging this into the ideal price index

$$1 = \Omega_t^{-1}\left(\int_0^1 p_{xt}^{1-\lambda} dx\right)^{\frac{1}{1-\lambda}} \tag{A.18}$$

gives us an aggregate representation of final output:

$$\mathcal{Y}_t = \Omega_t \left(\sum_\theta |\mathcal{X}_\theta|^{\frac{1}{\lambda}} E_{\theta t}^{\frac{\lambda-1}{\lambda}}\right)^{\frac{\lambda}{\lambda-1}}. \tag{A.19}$$

Therefore, we can interpret the CES shares ν_θ of Equation (1.28) as the mass of tasks $|\mathcal{X}_\theta|$ produced by sector θ, as claimed in the main text.

A.1.3 Substitution Matrix for Nested-CES Production

I will omit time subscripts for the sake of parsimony. To derive the substitution matrix $\mathbf{\Sigma}$, consider the function $G : \mathbb{R}^{2(J-1)} \to \mathbb{R}^{J-1}$, which follows

$$G_i\Big(\big\{ \ln\big(Y_j/Y_J\big), \ln\big(p_J/p_j\big)\big\}\Big) = \ln\big(p_J/p_i\big) \tag{A.20}$$
$$- \Big[\frac{1}{\lambda}\ln\big(\nu_{\theta(J)}/\nu_{\theta(i)}\big) + \big(\frac{1}{\sigma} - \frac{1}{\lambda}\big)\ln\big(E_{\theta(J)}/E_{\theta(i)}\big) - \frac{1}{\sigma}\ln\big(Y_J/Y_i\big)\Big].$$

That is, G is a function valued at the zero vector when the equilibrium price conditions (A.8) and (A.10) are satisfied. Note that $\theta(i)$ is the sector associated with technology i. To derive demand responses, we can apply the Implicit Function Theorem to G. This gives us

$$\frac{\partial G_i}{\partial \ln\big(Y_j/Y_J\big)} = -\Big[\frac{1}{\sigma}\mathbb{1}(i=j) + \big(\frac{1}{\lambda} - \frac{1}{\sigma}\big)\varepsilon^E_{\theta(i)j}\mathbb{1}(\theta(i)=\theta(j))\Big] \tag{A.21}$$
$$\frac{\partial G_i}{\partial \ln\big(p_J/p_j\big)} = \mathbb{1}(i=j), \tag{A.22}$$

where $\varepsilon^E_{\theta(i)j}$ is the elasticity of sector $\theta(i)$ with respect to good j. Equation (A.21) relies on the fact that $E_{\theta(J)} = Y_J$ because the final sector Θ only has a clean form of production.

Applying the Implicit Function Theorem gives us

$$\mathbf{\Sigma} = -\mathbf{D_y G}^{-1}, \tag{A.23}$$

where $\mathbf{D_y G}$ is the Jacobian of G with respect to log relative quantities. We can ignore the Jacobian of G with respect to log relative prices because it is just the identity matrix $\mathbf{I}$. Using the fact that the Jacobian of G with respect to log relative quantities is block diagonal, we can derive

$$\boldsymbol{\Sigma} = \begin{pmatrix} \tilde{\boldsymbol{\Sigma}}_1 & 0 & \dots & 0 \\ 0 & \tilde{\boldsymbol{\Sigma}}_2 & \dots & \vdots \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & \dots & \tilde{\boldsymbol{\Sigma}}_{\Theta-1} \end{pmatrix}, \tag{A.24}$$

where

$$\tilde{\boldsymbol{\Sigma}}_\theta = \begin{pmatrix} \lambda + (\sigma - \lambda)\varepsilon^E_{\theta d} & (\lambda - \sigma)\varepsilon^E_{\theta d} \\ (\lambda - \sigma)\varepsilon^E_{\theta c} & \lambda + (\sigma - \lambda)\varepsilon^E_{\theta c} \end{pmatrix}, \tag{A.25}$$

where, again, $\varepsilon^E_{\theta d}$ is the elasticity of sector θ with respect to its dirty form of production, and $\varepsilon^E_{\theta c}$ is defined analogously. We can also interpret these elasticities as income shares. Note that in the special case where $\sigma = \lambda$, we have $\boldsymbol{\Sigma} = \sigma\mathbf{I}$ as $\sigma = \lambda$ would imply a standard CES.

A.1.4 Proof of Proposition 1.1

I will omit time subscripts from variables that are time invariant in steady-state. Any balanced growth steady-state requires equal levels of innovation across technologies: $z_i = z_J$. From Equation (1.10), this implies

$$\ln\big(s_i/s_J\big) = -\frac{1}{\eta}\ln\big(\chi_i/\chi_J\big) - \frac{1}{\eta}\ln\big(\phi_i/\phi_J\big). \tag{A.26}$$

Plugging this into the research condition (1.25) and using the no-arbitrage condition (1.24) and relative good prices (1.22), we have

$$(1-\alpha)\ln(\bar{A}_i) = \ln(\Xi_i) + \frac{1}{\eta}\ln\big(\chi_i/\chi_J\big) + \frac{1}{\eta}\ln\big(\phi_i/\phi_J\big) + \alpha\ln(\mathcal{R}_i) + \ln\big(Y_{it}/Y_{Jt}\big), \tag{A.27}$$

where $\Xi_i \equiv \xi_i/\xi_J$ denotes relative innovation subsidies, and $\mathcal{R}_i \equiv (r_i+\omega_i\tau)/(r_J+\omega_J\tau)$ denotes relative input costs inclusive of the carbon price. Note that assuming that production is constant returns implies we can write relative quantities demanded just in terms of relative prices, and assuming spillovers are homogeneous of degree zero implies we can write spillovers just in terms of relative technology. Thus, relative technology in any balanced growth steady-state must satisfy Equation (A.27).

To consider existence and uniqueness of the steady-state, use Equation (A.27) to define the mapping $\mathbb{T}\ln(\bar{A})$ which follows

$$\mathbb{T}\ln(\bar{A}_i) = \frac{1}{1-\alpha}\Big[\ln(\Xi_i) + \frac{1}{\eta}\ln\big(\chi_i/\chi_J\big) + \frac{1}{\eta}\ln\big(\phi_i/\phi_J\big) + \alpha\ln(\mathcal{R}_i) + \ln\big(Y_{it}/Y_{Jt}\big)\Big], \tag{A.28}$$

where the functional dependence on log relative technology comes from relative spillovers and relative quantities. If $\mathbb{T}\ln(\bar{A})$ is a contraction, then we have that there exists a unique steady-state. I also constructively establish the existence of a unique steady-state in Appendix A.1.5 when the production structure is CES and the spillover structure is Cobb-Douglas.

More generally, if there exists a steady-state, a fixed-point of $\mathbb{T}\ln(\bar{A})$ will be unique if the Jacobian of the function $\mathbb{U}\ln(\bar{A}) \equiv \ln(\bar{A}) - \mathbb{T}\ln(\bar{A})$ is either a P-matrix for every element of $\mathbb{R}^{J-1}$ or positive (negative) quasi-definite for every element of $\mathbb{R}^{J-1}$.

To derive the Jacobian of log steady-state relative technology with respect to log relative innovation subsidies $\mathbf{D_\Xi \bar{A}_{ss}}$, we can differentiate

$$(1-\alpha)\frac{\partial\ln(\bar{A}_i)}{\partial\ln(\Xi_j)} = \mathbb{1}(i=j) - \frac{1}{\eta}\sum_q \Phi_{iq}\frac{\partial\ln(\bar{A}_q)}{\partial\ln(\Xi_j)} + (1-\alpha)\sum_q \Sigma_{iq}\frac{\partial\ln(\bar{A}_q)}{\partial\ln(\Xi_j)}, \tag{A.29}$$

which in matrix notation gives us

$$(1-\alpha)\mathbf{D_\Xi\bar{A}_{ss}} = \mathbf{I} - \frac{1}{\eta}\mathbf{\Phi D_\Xi\bar{A}_{ss}} + (1-\alpha)\mathbf{\Sigma D_\Xi\bar{A}_{ss}} \tag{A.30}$$
$$\Rightarrow \mathbf{D_\Xi\bar{A}_{ss}} = \eta\big[\mathbf{\Phi} - \eta(1-\alpha)\big(\mathbf{\Sigma}-\mathbf{I}\big)\big]^{-1}.$$

To derive the Jacobian of log steady-state relative technology with respect to log relative input prices $\mathbf{D_R \bar{A}_{ss}}$, we can differentiate

$$(1-\alpha)\frac{\partial \ln(\bar{A}_i)}{\partial \ln(\mathcal{R}_j)} = \alpha \mathbb{1}(i=j) - \frac{1}{\eta}\sum_q \Phi_{iq}\frac{\partial \ln(\bar{A}_q)}{\partial \ln(\mathcal{R}_j)} + (1-\alpha)\sum_q \Sigma_{iq}\frac{\partial \ln(\bar{A}_q)}{\partial \ln(\mathcal{R}_j)} - \alpha\Sigma_{ij}, \tag{A.31}$$

which in matrix notation gives us

$$(1-\alpha)\mathbf{D_R\bar{A}_{ss}} = \alpha\mathbf{I} - \frac{1}{\eta}\mathbf{\Phi D_R \bar{A}_{ss}} + (1-\alpha)\mathbf{\Sigma D_R\bar{A}_{ss}} - \alpha\mathbf{\Sigma} \tag{A.32}$$

$$\Rightarrow \mathbf{D_R\bar{A}_{ss}} = -\eta\alpha\big[\mathbf{\Phi} - \eta(1-\alpha)\big(\mathbf{\Sigma}-\mathbf{I}\big)\big]^{-1}\big(\mathbf{\Sigma}-\mathbf{I}\big).$$

Combining Equations (A.30) and (A.32), we have

$$d\ln(\bar{A}_{ss}) = \eta\big[\mathbf{\Phi} - \eta(1-\alpha)\big(\mathbf{\Sigma}-\mathbf{I}\big)\big]^{-1}\big[d\ln(\Xi) - \alpha\big(\mathbf{\Sigma}-\mathbf{I}\big)d\ln(\mathcal{R})\big], \tag{A.33}$$

which gives us Equations (1.38) and (1.39) from Proposition 1.1.

□

A.1.5 Steady-State under Parametric Assumptions of Section 1.2.5

Under the parametric assumptions of Section 1.2.5, we can further characterize the steady-state for relative technology. Plugging the equilibrium price conditions (A.8) and (A.10) into Equation (A.27), we have

$$\sum_{j<J}\Phi_{ij}\ln(\bar{A}_j) = \eta\ln(\Xi_i) + \eta(1-\alpha)(\sigma-1)\ln(\bar{A}_i) \tag{A.34}$$

$$+\,\eta(\sigma-\lambda)\ln\big(p_{\theta(i)t}/p_{\theta(J)t}\big) - \eta\alpha(\sigma-1)\ln(\mathcal{R}_i).$$

Denote relative sector prices by $\mathcal{P}_i \equiv p_{\theta(i)t}/p_{\theta(J)t}$. Writing Equation (A.34) in terms of matrix notation, we have

$$\ln(\bar{A}_{ss}) = \eta\big[\mathbf{\Phi} - \eta(1-\alpha)(\sigma-1)\mathbf{I}\big]^{-1}\big[\ln(\Xi) + (\sigma-\lambda)\ln(\mathcal{P}) - \alpha(\sigma-1)\ln(\mathcal{R})\big]. \tag{A.35}$$

Using the sector-level ideal price index, we have that relative sector prices follow

$$\mathcal{P}_i = \frac{p_{\theta(i)t}}{p_{\theta(J)t}} = \left(\frac{p_{\theta ct}^{1-\sigma} + p_{\theta dt}^{1-\sigma}}{p_{Jt}^{1-\sigma}}\right)^{\frac{1}{1-\sigma}} = \left(\mathcal{R}_{\theta c}^{\alpha(1-\sigma)} \bar{A}_{\theta c}^{(1-\alpha)(\sigma-1)} + \mathcal{R}_{\theta d}^{\alpha(1-\sigma)} \bar{A}_{\theta d}^{(1-\alpha)(\sigma-1)}\right)^{\frac{1}{1-\sigma}}, \tag{A.36}$$

where we have used the fact that the final sector Θ only has a clean form of production. Steady-state relative technology is then the solution to a fixed-point problem defined by Equation (A.35). Note that if the production structure is a simple CES, i.e. $\lambda = \sigma$, then Equation (A.35) constructively provides the economy's unique steady-state.

Furthermore, Equation (A.26) implies that

$$\ln(\bar{s}_{ss}) = \ln(\mathcal{V}) + \frac{1}{\eta}\mathbf{\Phi} \ln(\bar{A}_{ss}), \tag{A.37}$$

where $\bar{s}_j \equiv s_j/s_J$ are relative scientists, and $\mathcal{V}_j \equiv \nu_{\theta(j)}/\nu_{\theta(J)}$ are relative CES shares. Thus, we have

$$\ln(\bar{s}_{ss}) = \ln(\mathcal{V}) + \mathbf{\Phi}\big[\mathbf{\Phi} - \eta(1-\alpha)(\sigma-1)\mathbf{I}\big]^{-1}\big[\ln(\Xi) + (\sigma-\lambda)\ln(\mathcal{P}) - \alpha(\sigma-1)\ln(\mathcal{R})\big]. \tag{A.38}$$

We can also characterize corner steady-states in the case where the spillover network is shut down. Then, each sector converges to a single technology, either clean or dirty. First, the technologies that lose out in their sector will converge to zero in relative terms. For the remaining technologies, we can use the same strategy as above. We have

$$\ln(\bar{A}_{ss}) = \frac{1}{(1-\alpha)(1-\lambda)}\ln(\Xi) + \frac{\alpha}{1-\alpha}\ln(\mathcal{R}), \tag{A.39}$$

but only for technologies that persist in the long-run, one for each sector.

A.1.6 Proof of Proposition 1.2

I will describe the equilibrium evolution of technology as a dynamic process with log relative technology $\{\ln(\bar{A}_{jt})\}$ as the state variable and scientists $\{s_{jt}\}$ as the control variable. Thus, we can write the dynamic process in vector notation as

$$\ln(\bar{A}_t) = \mathcal{H}\Big(s_t\big(\ln(\bar{A}_{t-1})\big), \ln(\bar{A}_{t-1})\Big), \tag{A.40}$$

where $s_t\big(\ln(\bar{A}_{t-1})\big)$ is the equilibrium mapping of the control to the previous period's state. By definition, the steady-state $\ln(\bar{A}_{ss})$ is a fixed point of $\mathcal{H}$, so taking a first-order approximation around the steady-state, we have

$$\ln(\bar{A}_t) - \ln(\bar{A}_{ss}) \approx \boldsymbol{\mathcal{J}}\big(\ln(\bar{A}_{t-1}) - \ln(\bar{A}_{ss})\big), \tag{A.41}$$

where $\boldsymbol{\mathcal{J}}$ is the Jacobian of $\mathcal{H}$, evaluated at the steady-state. Thus, because $\mathcal{A}_t$ is defined as the log deviation of relative technology from steady-state, we have Equation (1.41) of Proposition 1.2.

To unpack the Jacobian $\boldsymbol{\mathcal{J}}$, we can describe the dynamic process $\mathcal{H}$ using the law of motion for technology (1.8), the research condition (1.25), the no-arbitrage condition (1.24), and the market clearing condition for scientists (1.12). I will omit time subscripts on policy variables. This yields

$$\begin{aligned} \ln(\bar{A}_{it}) &= \ln(\gamma)\big[\chi_i s_{it}^{\eta}\phi_{it} - \chi_J s_{Jt}^{\eta}\phi_{Jt}\big] + \ln(\bar{A}_{it-1}) \\ (1-\eta)\big[\ln(s_{it}) - \ln(s_{Jt})\big] &= \ln\big(\chi_i/\chi_J\big) + \ln\big(\xi_i/\xi_J\big) + \ln\big(\phi_{it}/\phi_{Jt}\big) \\ &\quad + \ln\big(p_{it}/p_{Jt}\big) + \ln\big(Y_{it}/Y_{Jt}\big) \\ \sum_i s_{it} &= \mathcal{S}. \end{aligned} \tag{A.42}$$

Note that assuming that production is constant returns implies we can write relative quantities demanded just in terms of relative prices, and assuming spillovers are homogeneous of degree zero implies we can write spillovers just in terms of relative technology.

Denote the elasticity of variable x_{it} with respect to $\bar{A}_{jt-1}$ by $\varepsilon^x_{ij} \equiv \frac{\partial \ln(x_{it})}{\partial \ln(\bar{A}_{jt-1})}$, which implies $\varepsilon^{\bar{A}}_{ij} = \mathcal{J}_{ij}$. Differentiating the dynamical system and evaluating at the steady state, we have

$$\begin{aligned} \varepsilon^{\bar{A}}_{ij} &= g\big[\eta(\varepsilon^s_{ij} - \varepsilon^s_{Jj}) - \Phi_{ij}\big] + \mathbb{1}(i=j) \qquad \text{(A.43)} \\ (1-\eta)(\varepsilon^s_{ij} - \varepsilon^s_{Jj}) &= -\Phi_{ij} - (1-\alpha)\epsilon^{\bar{A}}_{ij} + \sum_q \Sigma_{iq}(1-\alpha)\epsilon^{\bar{A}}_{qj} \\ \Rightarrow \big[(1-\eta) + g\eta(1-\alpha)\big]\epsilon^{\bar{A}}_{ij} &= (1-\eta)\mathbb{1}(i=j) - g\Phi_{ij} + g\eta(1-\alpha)\sum_q \Sigma_{iq}\epsilon^{\bar{A}}_{qj}, \end{aligned}$$

where the derivative of relative good prices with respect to relative technology comes from Equation (A.4). Transforming into matrix notation, we have Equation (1.42) from Proposition 1.2:

$$\begin{aligned} \big[(1-\eta)\mathbf{I} - g\eta(1-\alpha)\big(\mathbf{\Sigma} - \mathbf{I}\big)\big]\boldsymbol{\mathcal{J}} &= (1-\eta)\mathbf{I} - g\mathbf{\Phi} \qquad \text{(A.44)} \\ \Rightarrow \boldsymbol{\mathcal{J}} &= \big[(1-\eta)\mathbf{I} - g\eta(1-\alpha)\big(\mathbf{\Sigma} - \mathbf{I}\big)\big]^{-1}\big[(1-\eta)\mathbf{I} - g\mathbf{\Phi}\big]. \end{aligned}$$

□

A.1.7 Proof of Corollary 1.2

I will assume that the transition matrix $\boldsymbol{\mathcal{J}}$ has $J-1$ distinct real eigenvalues. We have

$$\begin{aligned} &\boldsymbol{\mathcal{J}}\boldsymbol{Q} = \boldsymbol{Q}\boldsymbol{D}(\kappa) \qquad \text{(A.45)} \\ \Rightarrow &\big[(1-\eta)\mathbf{I} - g\mathbf{\Phi}\big]\boldsymbol{Q} = \big[(1-\eta)\mathbf{I} - g\eta(1-\alpha)\big(\mathbf{\Sigma} - \mathbf{I}\big)\big]\boldsymbol{Q}\boldsymbol{D}(\kappa) \\ \Rightarrow &\big[(1-\eta)\mathbf{I} - g\eta(1-\alpha)\big(\mathbf{\Sigma} - \mathbf{I}\big)\big]\boldsymbol{Q}\boldsymbol{D}(1-\kappa) = g\big[\mathbf{\Phi} - \eta(1-\alpha)\big(\mathbf{\Sigma} - \mathbf{I}\big)\big]\boldsymbol{Q} \\ \Rightarrow &\boldsymbol{\mathcal{M}} = g\boldsymbol{Q}\boldsymbol{D}(1-\kappa)^{-1}\boldsymbol{Q}^{-1}\big[(1-\eta)\mathbf{I} - g\eta(1-\alpha)\big(\mathbf{\Sigma} - \mathbf{I}\big)\big]^{-1}, \end{aligned}$$

which gives us Equation (1.50) of Corollary 1.2.

□

A.1.8 Proof of Proposition 1.4

As described in the Definition 1.6, the Planner solves

$$\max_{\{c_t,\{\Lambda_{jt},\{\ell_{j\imath t}\},A_{jt},s_{jt}\}\}} \sum_{t\geq 0} \frac{1}{(1+\rho)^t} u(c_t) \;\; s.t. \tag{A.46}$$

$$\mathcal{Y}_t = c_t + \sum_j r_j \Lambda_{jt} : \quad \varkappa_t$$

$$L = \sum_j \int_0^1 \ell_{j\imath t} d\imath : \quad \varpi_{\ell t}$$

$$\ln(A_{jt}) = \ln(\gamma)\chi_j s_{jt}^{\eta} \phi_{jt} + \ln(A_{jt-1}) : \quad \epsilon_{jt}$$

$$\mathcal{S} = \sum_j s_{jt} : \quad \varpi_{st},$$

where next to each constraint is the corresponding multiplier. Note that I have transformed the law of motion for technology to be in terms of logs.

Taking the FOC with respect to c_t, we have

$$\frac{u_t'}{(1+\rho)^t} = \varkappa_t, \tag{A.47}$$

which I will use to solve out $\varkappa_t$ in the remaining FOCs. Taking the FOC with respect to Λ_{jt}, we have

$$\frac{\partial \mathcal{Y}_t}{\partial \Lambda_{jt}} = r_j - \omega_j \sum_{\hat{t}\geq t} \frac{\mathcal{Y}_{\hat{t}}}{(1+\rho)^{\hat{t}-t}} \frac{u_{\hat{t}}'}{u_t'} \frac{\partial \ln(\Omega_{\hat{t}})}{\partial \mathcal{E}_t}. \tag{A.48}$$

Taking the FOC with respect to $\ell_{j\imath t}$, we have

$$\frac{\partial \mathcal{Y}_t}{\partial \ell_{j\imath t}} = \frac{\varpi_{\ell t}}{u_t'/(1+\rho)^t}. \tag{A.49}$$

Taking the FOC with respect to A_{jt}, we have

$$\frac{\epsilon_{jt}}{A_{jt}} = \frac{u'_t}{(1+\rho)^t}\frac{\partial \mathcal{Y}_t}{\partial A_{jt}} + \frac{\epsilon_{jt+1}}{A_{jt}} + \sum_i \epsilon_{it+1} \ln(\gamma)\chi_i s_{it+1}^{\eta}\frac{\partial \phi_{it+1}}{\partial A_{jt}}. \tag{A.50}$$

Taking the FOC with respect to s_{jt}, we have

$$\varpi_{st} = \epsilon_{jt} \ln(\gamma)\chi_j \eta s_{jt}^{\eta-1}\phi_{jt}. \tag{A.51}$$

We can now determine the prices and policy instruments that support this allocation as a competitive equilibrium. First, to close the markup of the intermediate producer, the Planner sets the intermediate subsidy equal to the markup: $\Upsilon = \gamma$. This guarantees that intermediates are produced at marginal cost. Next, the Planner sets the wage for workers to reflect the shadow price of labor: $w_{\ell t} = \frac{\varpi_{\ell t}}{u'_t/(1+\rho)^t}$. Then, if the carbon price is set properly, prices $\{p_{jt}, \{p_{j\iota t}\}\}$ set equal to marginal products/marginal costs will be efficient. Note that, with these prices and policies, we now have average intermediate producer profit $\Pi_{jt} = (\gamma-1)(1-\alpha)p_{jt}Y_{jt}$.

Next, define $R_{t+1} \equiv (1+\rho)u'_t/u'_{t+1}$ as the Planner's intertemporal marginal rate of substitution. To align the input demand condition (1.17) with the Planner's input FOC (A.48), the Planner sets the carbon price according to

$$\tau_t = -\sum_{\hat{t}\geq t}\prod_{\hat{s}=1}^{\hat{t}-t}\frac{1}{R_{t+\hat{s}}}\mathcal{Y}_{\hat{t}}\frac{\partial \ln(\Omega_{\hat{t}})}{\partial \mathcal{E}_t}, \tag{A.52}$$

which is the social cost of carbon. This Pigouvian correction gives us Equation (1.61) of Proposition 1.4.

For innovation subsidies, we can compare the optimality condition of the research problem (A.5) with the Planner's scientist FOC (A.51). We have

$$\chi_j \eta s_{jt}^{\eta-1} \phi_{jt} \xi_{jt} \Pi_{jt} = w_{st} \tag{A.53}$$
$$\chi_j \eta s_{jt}^{\eta-1} \phi_{jt} \epsilon_{jt} = \frac{\varpi_{st}}{\ln(\gamma)}.$$

The Planner can set the wage for scientists according to

$$w_{st} = \frac{\varpi_{st} \;\ln(\gamma)}{u_t'/(1+\rho)^t}. \tag{A.54}$$

Plugging this into the two research conditions (A.53), we can see that the incentives for innovation will be efficient if the Planner sets innovation subsidies according to

$$\xi_{jt} \Pi_{jt} = \frac{\epsilon_{jt}}{u_t'/(1+\rho)^t}. \tag{A.55}$$

Note that any positive multiple of $\{w_{st}, \{\xi_{jt}\}\}$ would achieve the same outcome, but I am focusing on what I view as a natural normalization for these prices and policies. Plugging this into the Planner's technology FOC (A.50), we get the recursion

$$\xi_{jt} \Pi_{jt} = (1-\alpha) S_{jt} \mathcal{Y}_t + \frac{1}{R_{t+1}} \big[\xi_{jt+1} \Pi_{jt+1} + \sum_i \xi_{it+1} \Pi_{it+1} g_{it+1} \varphi_{ijt+1} \big], \tag{A.56}$$

which gives us Equation (1.62) of Proposition 1.4. Finally, the Planner sets the lump-sum tax D_t to balance the government's budget (1.27) in each period.

□

A.1.9 Proof of Corollary 1.3

Manipulating Equation (1.62), we have

$$\tilde{\xi}_{jt} = \frac{S_{jt}}{\gamma - 1} + \frac{\mathcal{Y}_{t+1}/\mathcal{Y}_t}{R_{t+1}} \big[\tilde{\xi}_{jt+1} + \sum_i \tilde{\xi}_{it+1} g_{it+1} \varphi_{ijt+1} \big]. \tag{A.57}$$

I have assumed the economy is in steady-state, so income shares for technologies are constant and the growth rate of output is equal to a constant g_y. For example, if carbon pollution stops and the damage function settles to a constant, then the growth rate of output g_y will be equal to the growth rate of technology g.[1] More generally, we would have $g_y \leq g$ as damages could worsen with continued carbon pollution.

Define $\tilde{R} \equiv R/(1+g_y)$ as the Planner's growth-adjusted intertemporal marginal rate of substitution, evaluated at the steady-state. Then, evaluating Equation (A.57) in matrix notation at the steady-state, we have

$$\tilde{\xi}' = \frac{1}{\gamma - 1} S' \Big[(1 - \tilde{R}^{-1})\mathbf{I} - g\tilde{R}^{-1}\boldsymbol{\varphi} \Big]^{-1}, \tag{A.58}$$

which gives us Equation (1.63) of Corollary 1.3. If all income shares remain positive in steady-state, this is equivalent to

$$\xi' = \frac{1}{\gamma - 1} \vec{1}' \Big[(1 - \tilde{R}^{-1})\mathbf{I} - g\tilde{R}^{-1}\boldsymbol{D}(S)\boldsymbol{\varphi}\boldsymbol{D}(S)^{-1} \Big]^{-1}. \tag{A.59}$$

□

[1] This is the steady-state under the parametric assumptions of Sections 1.2.5 and 1.5.2.

A.1.10 Proof of Proposition 1.5

As described in the Definition 1.7, the Planner solves

$$\max_{\{c_t,\{\Lambda_{jt},\{\ell_{j\iota t}\},A_{jt},s_{jt}\}\}} \sum_{t\geq 0} \frac{1}{(1+\rho)^t} u(c_t) \;\; s.t. \tag{A.60}$$

$$\mathcal{Y}_t = c_t + \sum_j r_j \Lambda_{jt} : \;\; \varkappa_t$$

$$\ln(A_{jt}) = \ln(\gamma)\chi_j s_{jt}^{\eta}\phi_{jt} + \ln(A_{jt-1}) : \;\; \epsilon_{jt}$$

$$\mathcal{S} = \sum_j s_{jt} : \;\; \varpi_{st}$$

$$\{\Lambda_{jt},\{\ell_{j\iota t}\}\} = \text{argmax}\left[\mathcal{Y}_t - \sum_j r_j \Lambda_{jt} - \hat{\tau}_t \mathcal{E}_t\right] \;\; s.t. \;\; L = \sum_j \int_0^1 \ell_{j\iota t} d\iota,$$

where next to each constraint is the corresponding multiplier. As before, I have transformed the law of motion for technology to be in terms of logs. Now the Planner must consider the incentive compatibility constraint for factors of production. Note that I am assuming the intermediate subsidy is set according to $\Upsilon = \gamma$.

Taking the FOC with respect to c_t, we have

$$\frac{u_t'}{(1+\rho)^t} = \varkappa_t, \tag{A.61}$$

which I will use to solve out $\varkappa_t$ in the remaining FOCs. Taking the FOC with respect to s_{jt}, we have

$$\varpi_{st} = \epsilon_{jt} \ln(\gamma)\chi_j \eta s_{jt}^{\eta-1}\phi_{jt}. \tag{A.62}$$

Taking the FOC with respect to A_{jt}, we have

$$\frac{\epsilon_{jt}}{A_{jt}} = \frac{u_t'}{(1+\rho)^t}\Big[\frac{\partial \mathcal{Y}_t}{\partial A_{jt}} + \sum_i \int_0^1 \frac{\partial \mathcal{Y}_t}{\partial \ell_{i\iota t}}\frac{\partial \ell_{i\iota t}}{\partial A_{jt}} d\iota + \sum_i \big(\frac{\partial \mathcal{Y}_t}{\partial \Lambda_{it}} - r_i\big)\frac{\partial \Lambda_{it}}{\partial A_{jt}} - \tau_t \frac{\partial \mathcal{E}_t}{\partial A_{jt}}\Big] \tag{A.63}$$
$$+ \frac{\epsilon_{jt+1}}{A_{jt}} + \sum_i \epsilon_{it+1} \ln(\gamma)\chi_i s_{it+1}^{\eta} \frac{\partial \phi_{it+1}}{\partial A_{jt}},$$

where τ_t is the same social cost of carbon from Equation (1.61). Now consider the producer optimality conditions

$$\begin{aligned} \frac{\partial \mathcal{Y}_t}{\partial \Lambda_{jt}} &= r_j + \omega_j \hat{\tau}_t \\ \frac{\partial \mathcal{Y}_t}{\partial \ell_{j\iota t}} &= w_{\ell t}. \end{aligned} \tag{A.64}$$

Plugging these conditions into the Planner's optimality condition for technology (A.63), and noting that $\sum_j \int_0^1 \frac{\partial \ell_{j\iota t}}{\partial x} d\iota = 0$ for any x due to the fixed supply of labor, we have

$$\frac{\epsilon_{jt}}{A_{jt}} = \frac{u'_t}{(1+\rho)^t}\Big[\frac{\partial \mathcal{Y}_t}{\partial A_{jt}} - (\tau_t - \hat{\tau}_t)\frac{\partial \mathcal{E}_t}{\partial A_{jt}}\Big] + \frac{\epsilon_{jt+1}}{A_{jt}} + \sum_i \epsilon_{it+1} \ln(\gamma)\chi_i s_{it+1}^{\eta} \frac{\partial \phi_{it+1}}{\partial A_{jt}}. \tag{A.65}$$

As before, the Planner can create efficient incentive for innovation by satisfying Equations (A.54) and (A.55). Thus, innovation subsidies follow the recursion

$$\xi_{jt}\Pi_{jt} = (1-\alpha)S_{jt}\mathcal{Y}_t - (\tau_t - \hat{\tau}_t)\mathcal{E}_t \frac{\partial \ln(\mathcal{E}_t)}{\partial \ln(A_{jt})} + \frac{1}{R_{t+1}}\big[\xi_{jt+1}\Pi_{jt+1} + \sum_i \xi_{it+1}\Pi_{it+1} g_{it+1}\varphi_{ijt+1}\big], \tag{A.66}$$

which gives us Equation (1.68) of Proposition 1.5. By construction, prices on the production side of the economy are set in equilibrium. Finally, the Planner sets the lump-sum tax D_t to balance the government's budget (1.27) in each period.

□

A.2 Details on Calibration & Simulation

A.2.1 Details on Numerical Representations

This appendix contains details on how I can represent equilibrium outcomes of my model in terms of fundamentals: technology, policy, and structural parameters. I make the parametric assumptions of Section 1.2.5. First, I will define what I call pseudo prices. These follow

$$\tilde{p}_{jt} \equiv \big(r_j + \omega_j \tau_t\big)^{\alpha} / A_{jt}^{1-\alpha}. \tag{A.67}$$

From Equation (A.3), we have that pseudo prices are proportional to prices and satisfy

$$p_{jt} = \Big(\frac{1}{\alpha}\Big)^{\alpha}\Big(\frac{1}{1-\alpha}\Big)^{1-\alpha}\Big(\frac{\gamma w_{\ell t}}{\Upsilon}\Big)^{1-\alpha}\tilde{p}_{jt}. \tag{A.68}$$

We can define similar pseudo prices at higher levels using ideal price indices

$$\tilde{p}_{\theta t} \equiv \big(\tilde{p}_{\theta ct}^{1-\sigma} + \tilde{p}_{\theta dt}^{1-\sigma}\big)^{\frac{1}{1-\sigma}} \tag{A.69}$$

$$\tilde{P}_t \equiv \Omega_t^{-1}\Big(\sum_{\theta} \nu_\theta \tilde{p}_{\theta t}^{1-\lambda}\Big)^{\frac{1}{1-\lambda}}, \tag{A.70}$$

which must also be proportional to their corresponding price.

Next, combining the intermediate production function (1.4), equilibrium intermediate output (A.1), and labor supply (1.6) gives us

$$\frac{\gamma w_{\ell t}}{\Upsilon} = (1-\alpha)\frac{\mathcal{Y}_t}{L}. \tag{A.71}$$

Plugging the wage (A.71) into the factor of proportionality and combining the final pseudo price (A.70) with the fact that output is the numeraire, we have for output

$$\mathcal{Y}_t = \alpha^{\frac{\alpha}{1-\alpha}} L \tilde{P}_t^{\frac{-1}{1-\alpha}}. \tag{A.72}$$

Next, using the demand conditions for sectors (A.8) and goods within sectors (A.10), we can derive income shares for sectors and goods within sectors. We have

$$S_{\theta t} \equiv \frac{p_{\theta t} E_{\theta t}}{\mathcal{Y}_t} = \frac{\nu_\theta \tilde{p}_{\theta t}^{1-\lambda}}{\Omega_t^{1-\lambda}\tilde{P}_t^{1-\lambda}} \tag{A.73}$$

$$S_{et}^{\theta} \equiv \frac{p_{\theta et} Y_{\theta et}}{p_{\theta t} E_{\theta t}} = \frac{\tilde{p}_{\theta et}^{1-\sigma}}{\tilde{p}_{\theta t}^{1-\sigma}}, \tag{A.74}$$

where $e \in \{c, d\}$. The income share for technology $j = \theta e$ is then $S_{jt} = S_{\theta t} \cdot S^{\theta}_{et}$. The input demand condition (1.17) implies

$$\Lambda_{jt} = \frac{\alpha S_{jt} \mathcal{Y}_t}{r_j + \omega_j \tau_t}, \tag{A.75}$$

which gives us total emissions via $\mathcal{E}_t = \sum_j \omega_j \Lambda_{jt}$ and consumption via $c_t = \mathcal{Y}_t - \sum_j r_j \Lambda_{jt}$.

I define clean quantity shares by sector as $q^{\theta}_{ct} \equiv Y_{\theta ct}/(Y_{\theta ct} + Y_{\theta dt})$. To map between clean income shares and clean quantity shares, note that the demand condition within sectors (A.10) implies that

$$S^{\theta}_{ct} = \frac{p_{\theta ct} Y_{\theta ct}}{p_{\theta ct} Y_{\theta ct} + p_{\theta dt} Y_{\theta dt}} = \frac{1}{1 + \left((1 - q^{\theta}_{ct})/q^{\theta}_{ct}\right)^{\frac{\sigma-1}{\sigma}}}. \tag{A.76}$$

Finally, to map my model of innovation to the empirical evidence that informs my choice of the elasticity of innovation with respect to scientists η, first define technology specific R&D expenditures as $R\&D_{jt} = w_{st} s_{jt}$. To consider the elasticity of patents with respect to R&D expenditure, plug this definition into the innovation production function (1.10) and note

$$\begin{aligned} z_{jt} &= \chi_j \left(\frac{R\&D_{jt}}{w_{st}} \right)^{\eta} \phi_{jt} \\ \Rightarrow \frac{\partial \ln (z_{jt})}{\partial \ln (R\&D_{jt})} &= \eta, \end{aligned} \tag{A.77}$$

where I am taking the mass of innovations z_{jt} as the object in my model analogous to patents. Next, to consider the elasticity of R&D expenditure with respect to the price of research, note that the optimality condition of the research problem (A.5) implies

$$\eta \chi_j \left(\frac{R\&D_{jt}}{w_{st}} \right)^{\eta} \phi_{jt} = R\&D_{jt} \tag{A.78}$$
$$\Rightarrow \frac{\partial \ln (R\&D_{jt})}{\partial \ln (w_{st})} = -\frac{\eta}{1-\eta},$$

which gives a demand elasticity of unity when $\eta = 0.5$.

A.2.2 Details on Selection of US Policy in the 2010s

This appendix provides further details on my procedure for selecting subsidies that describe US climate policy throughout the 2010s. These subsidies are an input into the simulations of Section 1.5.4. To determine the clean input subsidy for the transportation sector $\bar{\xi}_{car,c}$, consider public spending on the subsidy as a proportion of total spending in the transportation sector. We have

$$\frac{\bar{\xi}_{car,c} r_{car,c} \Lambda_{car,ct}}{p_{car,t} E_{car,t}} = \frac{\bar{\xi}_{car,c}}{1-\bar{\xi}_{car,c}} \frac{\alpha p_{car,ct} Y_{car,ct}}{p_{car,t} E_{car,t}} = \frac{\bar{\xi}_{car,c}}{1-\bar{\xi}_{car,c}} \alpha S_{ct}^{car}, \tag{A.79}$$

where the first equality comes from the input demand condition (1.17).

Let $\hat{cred}_{car,ct}$ denote data on federal spending on the plug-in electric vehicle tax credit as a proportion of total spending in the transportation sector. The numerator comes from US Congressional Research Service Report IF11017, which contains federal outlays on the tax credit for the years 2011 to 2018. This report was published in 2019 and contains projections out to 2022, but I only include spending that occurred before the report was published. The denominator comes from the series on motor vehicle output discussed in Section 1.5.3. Next, I take estimates of clean income shares in the transportation sector $\hat{S}_{ct}^{car}$ from my data on clean quantity shares using the mapping from Equation (A.76). Plugging these data series into Equation (A.79), we have

$$\frac{\bar{\xi}_{car,c}}{1-\bar{\xi}_{car,c}} \sum_{t=2011}^{2018} \alpha \hat{S}_{ct}^{car} = \sum_{t=2011}^{2018} \hat{cred}_{car,ct}, \tag{A.80}$$

which allows me to back out the value $\bar{\xi}_{car,c} = 0.012$.

To determine innovation subsidies $\{\xi_j\}$, consider public spending on each subsidy as a proportion of total spending on R&D. Using the optimality condition of the research problem (A.5) and scientist supply (1.12), we have

$$\frac{(\xi_j - 1) z_{jt} \Pi_{jt}}{w_{st} \sum_j s_{jt}} = \frac{\xi_j - 1}{\xi_j} \frac{s_{jt}}{\eta \mathcal{S}}. \tag{A.81}$$

As discussed in the main text, I normalize the innovation subsidy on the general technology to one as only relative innovation subsidies influence the composition of scientists across technologies.

Let $\hat{pubrd}_{jt}$ denote data on public spending on R&D in technology j as a proportion of total spending on R&D. The numerator comes from IEA data on public spending on R&D by technology. This series goes until 2015. Table A.1 contains information on the assignment of spending types to the technologies in my model. The denominator comes from the BEA. I then set innovation subsidies to match this series according to

$$\frac{\xi_j - 1}{\xi_j} \sum_{t=2010}^{2015} \frac{s_{jt}}{\eta \mathcal{S}} = \sum_{t=2010}^{2015} \hat{pubrd}_{jt}. \tag{A.82}$$

Unlike clean input subsidies, which can be set directly from the data, finding the innovation subsidies that match the data requires simulating the model to solve Equation (A.82), so I calibrate separate innovation subsidies for the versions of the model with and without spillovers. In the model with spillovers this yields $(\xi_{car,c}, \xi_{car,d}, \xi_{elec,c}, \xi_{elec,d}) = (1.006, 1.014, 1.057, 1.013)$, and in the model without spillovers this yields $(\xi_{car,c}, \xi_{car,d}, \xi_{elec,c}, \xi_{elec,d}) = (1.212, 1.005, 1.293, 1.016)$. The clean innovation subsidies are higher in the model without spillovers because clean technologies receive less R&D overall, so the same public R&D spending requires higher subsidy rates.

Table A.1: Assignment of Public R&D Spending to Technologies

Transportation		Electricity Generation	
Description	Codes	Description	Codes
Clean		Clean	
Vehicle Batteries/Storage Technologies	1311	Renewable Energy Sources	3
Advanced EV/HEV/FCV Systems	1312	–Excluding Biofuels	34
Electric Vehicle Infrastructure	1314	Nuclear	4
Dirty		Dirty	
Advanced Combustion Engines	1313	Coal	22
Oil & Gas (1/2)	21	Oil & Gas (1/2)	21

Notes: Clean electricity generation is assigned all of the spending in the renewables category except biofuels. I assign spending in the oil and gas category equally between dirty transportation and electricity generation because these fuels are used in both sectors. Data on public R&D spending by technology comes from the IEA.

A.2.3 First-Best Simulation Solution Method

To simulate the first-best policy path, I solve a root-finding problem. I make the parametric assumptions of Sections 1.2.5 and 1.5.2. Using the results described in Appendix A.2.1, my economy can be reduced to the sequence $\{\tau_{1t}, \tau_{2t}, \mathcal{C}_{1t}, \mathcal{C}_{2t}, \{\tilde{\xi}_{jt}, A_{jt}\}\}$.

For the carbon price, combining atmospheric carbon concentrations (1.70), the damage function (1.73), and the social cost of carbon (1.61) gives us

$$\tau_t = \varrho \sum_{\hat{t} \geq t} \prod_{\hat{s}=1}^{\hat{t}-t} \frac{1}{R_{t+\hat{s}}} \mathcal{Y}_{\hat{t}} \big(\psi_p + (1-\psi_p)\psi_0 \psi^{\hat{t}-t}\big). \tag{A.83}$$

This allows us to write the carbon price τ_t as a two-dimensional recursion $\{\tau_{1t}, \tau_{2t}\}$, which provides separate Pigouvian corrections for the permanent and transitory component of carbon pollution. These follow

$$\tau_{1t} = \varrho \psi_p \mathcal{Y}_t + \frac{1}{R_{t+1}} \tau_{1t+1} \tag{A.84}$$

$$\tau_{2t} = \varrho (1-\psi_p) \psi_0 \mathcal{Y}_t + \frac{1}{R_{t+1}} \psi \tau_{2t+1}. \tag{A.85}$$

The sum of these two components gives us the carbon price: $\tau_t = \tau_{1t} + \tau_{2t}$. The recursive formulation for $\{\mathcal{C}_{1t}, \mathcal{C}_{2t}\}$ is given by Equations (1.71) and (1.72), and these together sum to atmospheric carbon concentrations: $\mathcal{C}_t = \mathcal{C}_{1t} + \mathcal{C}_{2t}$.

Innovation subsidies times income shares $\{\tilde{\xi}_{jt}\}$ follow the recursive formula described in Equation (A.57), and technology $\{A_{jt}\}$ follows the law of motion (1.8). Note that the research condition (1.25) and scientist supply (1.12) allow us to write $\{s_{jt}\}$ as a function of $\{\tilde{\xi}_{jt}\}$ and $\{A_{jt-1}\}$. That is, we have

$$(1-\eta)\ln(\bar{s}_t) = \ln(\tilde{\Xi}_t) - \eta\ln(\mathcal{V}) - \mathbf{\Phi}\ln(\bar{A}_{t-1}), \tag{A.86}$$

where $\bar{s}_{jt} \equiv s_{jt}/s_{Jt}$ are relative scientists, $\tilde{\Xi}_t \equiv \tilde{\xi}_{jt}/\tilde{\xi}_{Jt}$ are relative innovation subsidies times income shares, and $\mathcal{V}_j \equiv \nu_{\theta(j)}/\nu_{\theta(J)}$ are relative CES shares. Scientist supply (1.12) then pins down the scale of the scientist allocation. Furthermore, as I describe in Appendix A.2.1, all of the other endogenous outcomes of the economy, such as output $\mathcal{Y}_t$, consumption c_t, income shares S_{jt}, and emissions $\mathcal{E}_t$, can be written as functions of τ_t and $\{A_{jt}\}$.

In summary, the root system follows

$$\begin{aligned}
\tau_{1t} &= \varrho\psi_p\mathcal{Y}_t + \frac{1}{R_{t+1}}\tau_{1t+1} \\
\tau_{2t} &= \varrho(1-\psi_p)\psi_0\mathcal{Y}_t + \frac{1}{R_{t+1}}\psi\tau_{2t+1} \\
\mathcal{C}_{1t} &= \psi_p\mathcal{E}_t + \mathcal{C}_{1t-1} \\
\mathcal{C}_{2t} &= (1-\psi_p)\psi_0\mathcal{E}_t + \psi\mathcal{C}_{2t-1} \\
\tilde{\xi}_{jt} &= \frac{S_{jt}}{\gamma-1} + \frac{\mathcal{Y}_{t+1}/\mathcal{Y}_t}{R_{t+1}}\big[\tilde{\xi}_{jt+1} + \sum_i \tilde{\xi}_{it+1}g_{it+1}\varphi_{ijt+1}\big] \\
\ln(A_{jt}) &= \ln(\gamma)\chi\Big(\frac{s_{jt}}{\nu_{\theta(j)}}\Big)^{\eta}\phi_{jt} + \ln(A_{jt-1}).
\end{aligned} \tag{A.87}$$

The state variables $\{\mathcal{C}_{1t}, \mathcal{C}_{2t}, \{A_{jt}\}\}$ are backward-looking, so I specify initial conditions using the strategies described in Sections 1.5.2 and 1.5.3. The policy variables $\{\tau_{1t}, \tau_{2t}, \{\tilde{\xi}_{jt}\}\}$ are forward-looking, so I specify a terminal condition by assuming the economy is in steady-state beyond my final period T. My simulation extends for 500 periods; long enough for the economy to be near its steady-state. Denote carbon prices relative to output by $\tilde{\tau}_t \equiv \tau_t/\mathcal{Y}_t$. These achieve a steady-state which follows

$$\tilde{\tau}_1 = \frac{\varrho\psi_p}{1-\tilde{R}^{-1}} \tag{A.88}$$

$$\tilde{\tau}_2 = \frac{\varrho(1-\psi_p)\psi_0}{1-\psi\tilde{R}^{-1}}, \tag{A.89}$$

where, as before, $\tilde{R} \equiv R/(1+g)$ is the Planner's growth-adjusted intertemporal marginal rate of substitution, evaluated at the steady-state. From this, we can see that carbon prices asymptote to infinity as the economy continues to grow, so emissions and the income shares of dirty technologies must go to zero. Hence, the growth rate of output g_y goes to the growth rate of technology g because climate damages asymptote to a constant as the transitory component of carbon concentrations decays to zero. Therefore, I assume that both consumption and output grow at rate g in the periods beyond my simulation. Terminal carbon prices then follow $\tau_{1T+1} = (1+g)\mathcal{Y}_T\tilde{\tau}_1$ and $\tau_{2T+1} = (1+g)\mathcal{Y}_T\tilde{\tau}_2$.

By Corollary 1.3, the steady-state for $\{\tilde{\xi}_j\}$ follows

$$\tilde{\xi}' = \frac{1}{\gamma-1}S'\Big[(1-\tilde{R}^{-1})\mathbf{I} - g\tilde{R}^{-1}\boldsymbol{\varphi}\Big]^{-1}. \tag{A.90}$$

To select the steady-state income shares S and growth rate g, we can derive steady-state relative technology by combining Equations (A.37) and (A.86) to achieve

$$\ln(\bar{A}_{ss}) = \eta\boldsymbol{\Phi}^{-1}\big[\ln(\tilde{\Xi}) - \ln(\mathcal{V})\big], \tag{A.91}$$

which, combined with the requirement that dirty income shares are zero, defines a fixed-point problem to solve for steady-state $\{\tilde{\xi}_j\}$. Note that the growth rate g comes from Equation (1.9).

Finding the steady-state when there are no cross-technology spillovers $\boldsymbol{\varphi} = \mathbf{0}$ is a special case because the inverse of $\mathbf{\Phi}$ does not exist. In that case, we have $\tilde{\xi}_j = S_j/(\gamma - 1)(1 - \tilde{R}^{-1})$, which implies $\tilde{\xi}_j = 0$ for dirty technologies because they have zero income share. Thus, scientists for dirty technologies must also be zero. The steady-state scientist condition (A.37) and scientist supply (1.12) can be satisfied by setting $s_j = \nu_{\theta(j)}$ for all clean technologies. Plugging this into Equation (A.86) gives us that income shares for clean technologies also follow $S_j = \nu_{\theta(j)}$. From Equation (A.73), this implies that $\bar{A}_j = 1$ for clean technologies and $\bar{A}_j = 0$ for dirty technologies.

A.2.4 Second-Best Simulation Solution Method

To simulate the second-best policy path, I use a similar strategy as in the first-best, reducing the economy to the sequence $\{\tau_{1t}, \tau_{2t}, \mathcal{C}_{1t}, \mathcal{C}_{2t}, \{\tilde{\xi}_{jt}, A_{jt}\}\}$.

Carbon concentrations $\{\mathcal{C}_{1t}, \mathcal{C}_{2t}\}$ and technology $\{A_{jt}\}$ follow the same laws of motion as before. The two components of the social cost of carbon $\{\tau_{1t}, \tau_{2t}\}$ follow the same recursion as before, except now it is $\hat{\tau}_t$ that influences production decisions.

Manipulating Equation (1.68), we have that second-best innovation subsidies times income shares $\{\tilde{\xi}_{jt}\}$ follow the recursive formula

$$\tilde{\xi}_{jt} = \frac{S_{jt}}{\gamma - 1} - \frac{\tau_t - \hat{\tau}_t}{(\gamma - 1)(1 - \alpha)} \frac{\mathcal{E}_t}{\mathcal{Y}_t} \frac{\partial \ln(\mathcal{E}_t)}{\partial \ln(A_{jt})} + \frac{\mathcal{Y}_{t+1}/\mathcal{Y}_t}{R_{t+1}} \big[\tilde{\xi}_{jt+1} + \sum_i \tilde{\xi}_{it+1} g_{it+1} \varphi_{ijt+1}\big]. \tag{A.92}$$

As before, we can write $\{s_{jt}\}$ as a function of $\{\tilde{\xi}_{jt}\}$ and $\{A_{jt-1}\}$ using Equation (A.86), but now that innovation subsidies can be negative, we will set $s_{jt} = 0$ whenever $\tilde{\xi}_{jt} \leq 0$. For the effect of innovation on equilibrium emissions, we can use the results of Appendix A.2.1 to derive

$$\frac{\mathcal{E}_t}{\mathcal{Y}_t}\frac{\partial \ln(\mathcal{E}_t)}{\partial \ln(A_{jt})} = \alpha \sum_\theta \frac{\omega_{\theta d} S_{\theta dt}}{r_{\theta d} + \omega_{\theta d}\hat{\tau}_t}\big[(1 + (1-\alpha)(1-\lambda))S_{jt} + \mathbb{1}(\theta = \theta(j))(1-\alpha)(\lambda - \sigma)S_{jt}^\theta + \mathbb{1}(\theta d = j)(1-\alpha)(\sigma - 1)\big]. \tag{A.93}$$

In summary, the root system follows

$$\begin{aligned}
\tau_{1t} &= \varrho\psi_p\mathcal{Y}_t + \frac{1}{R_{t+1}}\tau_{1t+1} \qquad (A.94)\\
\tau_{2t} &= \varrho(1-\psi_p)\psi_0\mathcal{Y}_t + \frac{1}{R_{t+1}}\psi\tau_{2t+1}\\
\mathcal{C}_{1t} &= \psi_p\mathcal{E}_t + \mathcal{C}_{1t-1}\\
\mathcal{C}_{2t} &= (1-\psi_p)\psi_0\mathcal{E}_t + \psi\mathcal{C}_{2t-1}\\
\tilde{\xi}_{jt} &= \frac{S_{jt}}{\gamma - 1} - \frac{\tau_t - \hat{\tau}_t}{(\gamma-1)(1-\alpha)}\frac{\mathcal{E}_t}{\mathcal{Y}_t}\frac{\partial \ln(\mathcal{E}_t)}{\partial \ln(A_{jt})} + \frac{\mathcal{Y}_{t+1}/\mathcal{Y}_t}{R_{t+1}}\big[\tilde{\xi}_{jt+1} + \sum_i \tilde{\xi}_{it+1} g_{it+1}\varphi_{ijt+1}\big]\\
\ln(A_{jt}) &= \ln(\gamma)\chi\Big(\frac{s_{jt}}{\nu_{\theta(j)}}\Big)^\eta\phi_{jt} + \ln(A_{jt-1}),
\end{aligned}$$

where endogenous outcomes, such as emissions $\mathcal{E}_t$, are determined by $\hat{\tau}_t$, rather than the social cost of carbon. The state variables $\{\mathcal{C}_{1t}, \mathcal{C}_{2t}, \{A_{jt}\}\}$ are backward-looking, so I specify initial conditions as before. The policy variables $\{\tau_{1t}, \tau_{2t}, \{\tilde{\xi}_{jt}\}\}$ are forward-looking, so I again assume the economy is in steady-state beyond my final period T. The steady-states for the two components of the social cost of carbon relative to output again follow Equations (A.88) and (A.89).

Denote by $\mathcal{T}_{jt} \equiv \frac{\tau_t - \hat{\tau}_t}{1-\alpha}\frac{\mathcal{E}_t}{\mathcal{Y}_t}\frac{\partial \ln(\mathcal{E}_t)}{\partial \ln(A_{jt})}$ the component of the innovation subsidy that reflects the pollution distortion, excluding $(\gamma - 1)$. Similar to Corollary 1.3, the steady-state for $\{\tilde{\xi}_j\}$ follows

$$\tilde{\xi}' = \frac{1}{\gamma - 1}\big(S' - \mathcal{T}'\big)\Big[(1 - \tilde{R}^{-1})\mathbf{I} - g\tilde{R}^{-1}\boldsymbol{\varphi}\Big]^{-1}. \tag{A.95}$$

The question, then, is the steady-state behavior of $\mathcal{T}$. I will assume that $\hat{\tau}_t$ either converges to a weakly positive constant or grows to infinity. If it grows to infinity, I will assume it remains below the social cost of carbon, so the asymptote of $\hat{\tau}_t/\mathcal{Y}_t$ is below $\tilde{\tau}_1 + \tilde{\tau}_2$. We have

$$\begin{aligned}\mathcal{T}_j = (\tilde{\tau}_1 + \tilde{\tau}_2 - \hat{\tau}_t/\mathcal{Y}_t)\frac{\alpha}{1-\alpha}\mathcal{Y}_t \sum_\theta \frac{\omega_{\theta d} S_{\theta d}}{r_{\theta d} + \omega_{\theta d}\hat{\tau}_t}\big[&(1 + (1-\alpha)(1-\lambda))S_j \\ &+ \mathbb{1}(\theta = \theta(j))(1-\alpha)(\lambda - \sigma)S_j^\theta \\ &+ \mathbb{1}(\theta d = j)(1-\alpha)(\sigma - 1)\big].\end{aligned} \quad \text{(A.96)}$$

Thus, there are two cases. If $\hat{\tau}_t$ grows over time, then dirty income shares must be zero in steady-state and we will have $\mathcal{T} = \vec{0}$. If $\hat{\tau}_t$ converges to a weakly positive constant, then $\mathcal{T}$ will diverge to infinity unless steady-state dirty income shares are zero. Without a growing carbon price, this requires steady-state relative dirty technologies to be zero as well. In either case, $\mathcal{T} = \vec{0}$ and steady-state $\{\tilde{\xi}_j\}$ is the solution to the fixed-point problem defined by Equation (A.95), but in the latter case, we have the added restriction that steady-state relative dirty technologies must be zero.

A.3 Additional Tables & Figures

Figure A·1: Tesla Mule 1

Notes: Tesla engineers modifying a gas-powered Lotus Elise to make a prototype electric vehicle.

Figure A·2: Bell Labs Solar Cell

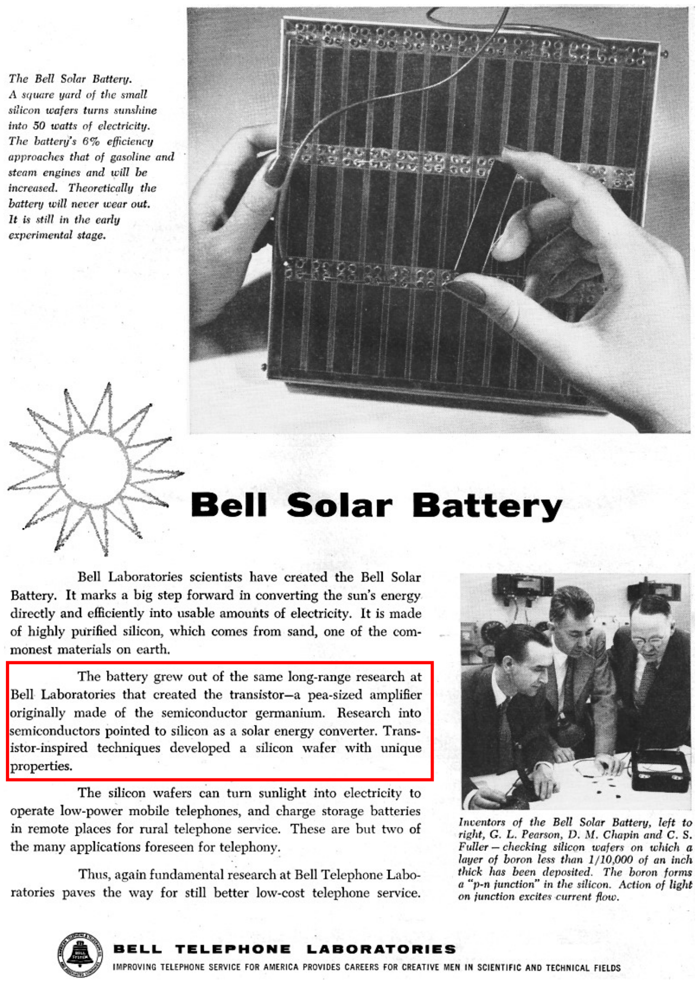

Notes: Bell Telephone Laboratories July 1954 Ad. Red box added to emphasize the description of knowledge spillovers.

Figure A·3: Linearized Transition Path Accurately Approximates Full Simulation

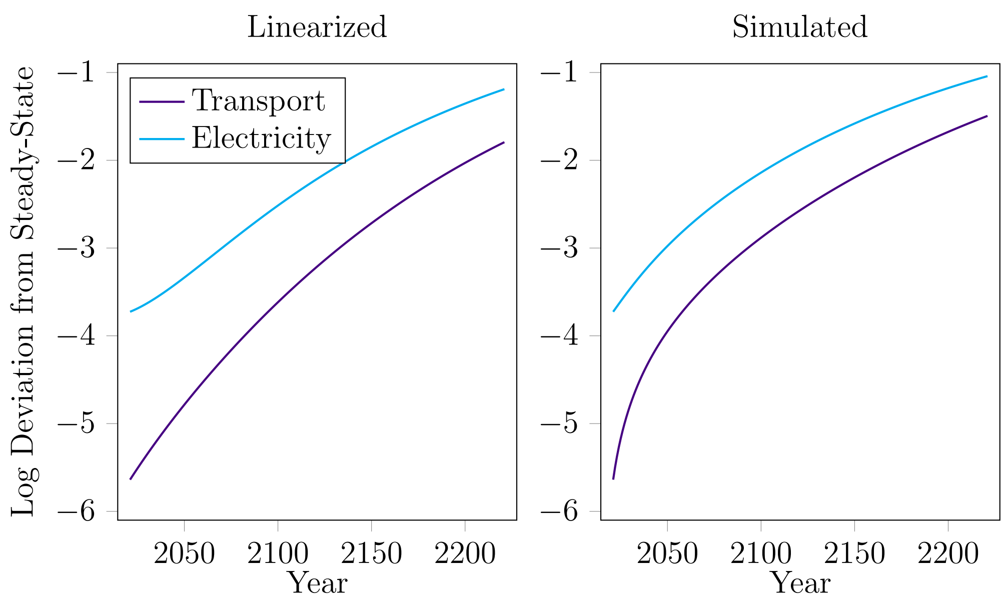

Notes: Technology's transition path following the policy reform of Section 1.6.1. The left panel displays the linearized path using Proposition 1.2, while the right panel displays the fully simulated path.

Figure A·4: Clean Income Share Path

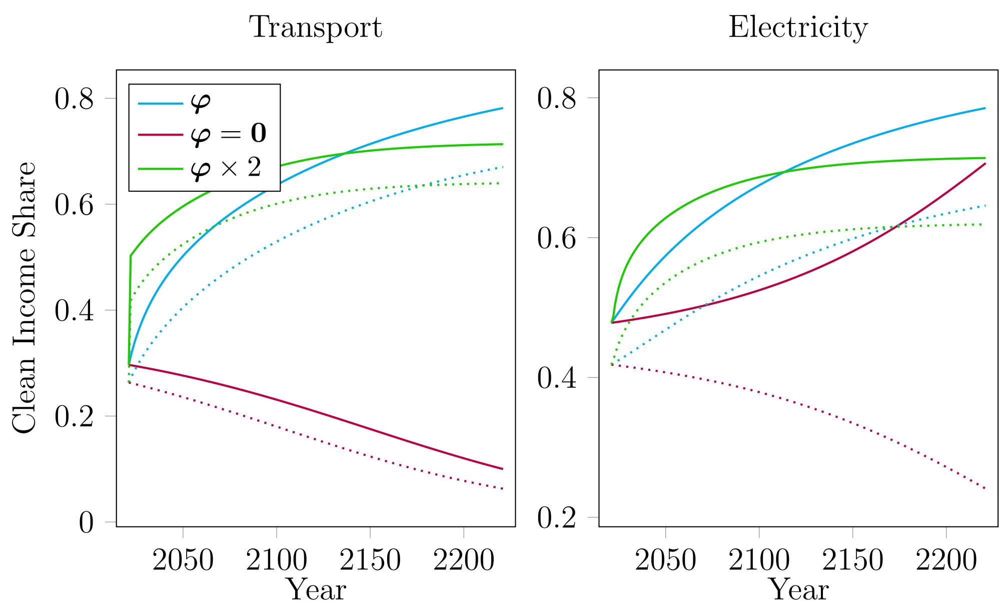

Notes: Impact of introducing a carbon price at the Biden Administration's estimate of the SCC ($51) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Dotted lines indicate laissez-faire paths.

Figure A·5: Determinants of Steady-State Policy Impact

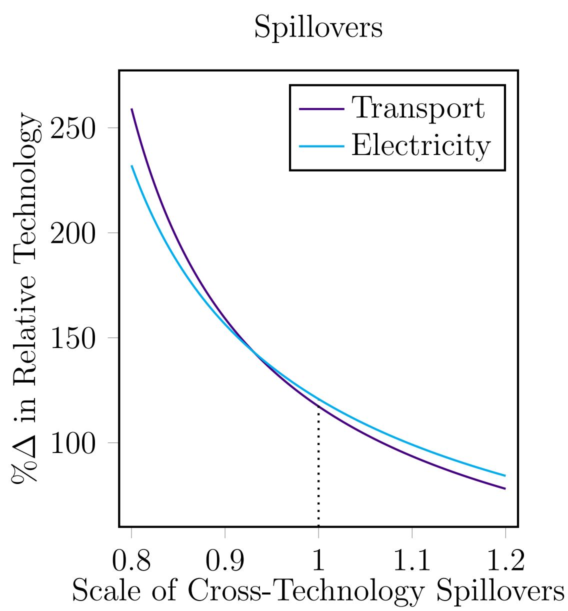

Notes: Impact of the policy reform of Section 1.6.1 as a function of the level of cross-technology knowledge spillovers and the elasticity of substitution between clean and dirty goods. Vertical dotted lines represent the benchmark calibration.

Figure A·6: Determinants of Increasing Returns to Innovation

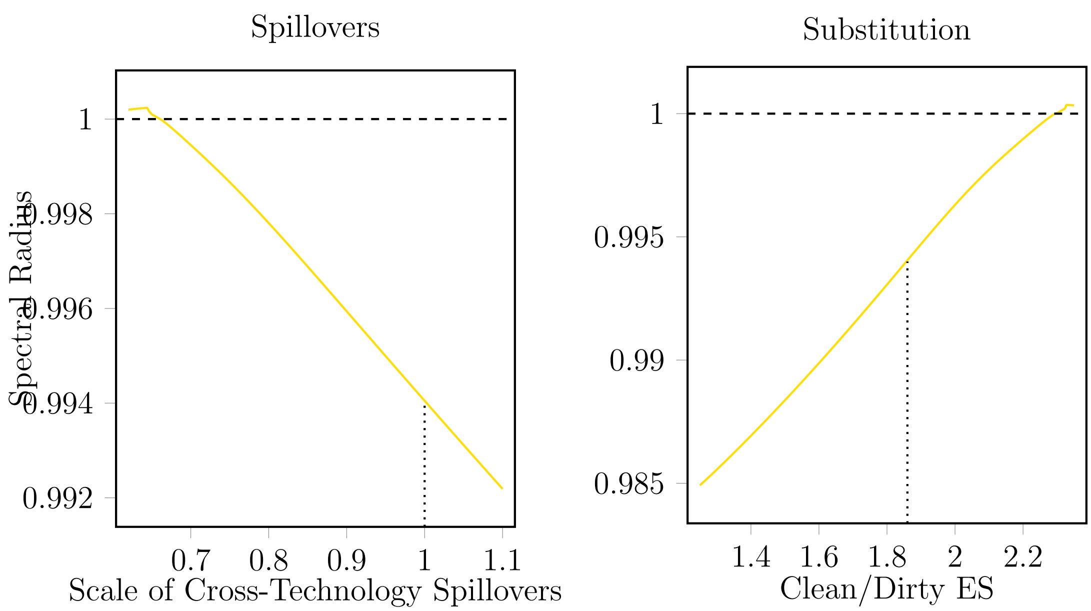

Notes: Spectral radius of the transition matrix as a function of the level of cross-technology knowledge spillovers and the elasticity of substitution between clean and dirty goods. The spectral radius passes one when spillovers reduce to 66.1% of their calibrated level or the elasticity of substitution increases to 2.29. Vertical dotted lines represent the benchmark calibration.

Table A.2: Impact of Policy Reform ($190 Carbon Price)

	No Spillovers	**Calibrated Spillovers**	**Double Spillovers**
Long-Run Impacts			
Relative Clean Technology by Sector			
$\%\Delta\bar{B}_{car}$	0%	+219.99%	+56.93%
$\%\Delta\bar{B}_{elec}$	$+\infty$%	+216.84%	+63.79%
Clean Income Shares by Sector			
ΔS_c^{car}	0 pp	+17.95 pp	+13.95 pp
ΔS_c^{elec}	+100 pp	+21.28 pp	+16.98 pp
Emissions Intensity			
$\%\Delta\bar{\omega}$	-87.48%	-93.6%	-84.48%
Transitional Impacts			
Half-Lives of Convergence by Sector			
$t_{car}^{(1/2)}$	–	151 years	21 years
$t_{elec}^{(1/2)}$	–	147 years	25 years
Carbon Emissions by Year			
$\%\Delta\mathcal{E}_{2035}$	-73.93%	-79.86%	-81.51%
$\%\Delta\mathcal{E}_{2060}$	-76.36%	-83.84%	-82.78%
Degree of Increasing Returns to Innovation			
Spectral Radius			
$\max\lvert\kappa_j\rvert$	1.01	0.994	0.978

Notes: Impact of introducing a carbon price at the EPA's proposed SCC ($190) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Changes in relative technology are listed in log points. For path dependent economies, long-run impacts refer to corner, rather than interior, steady-states.

Figure A·7: Technology Path ($190 Carbon Price)

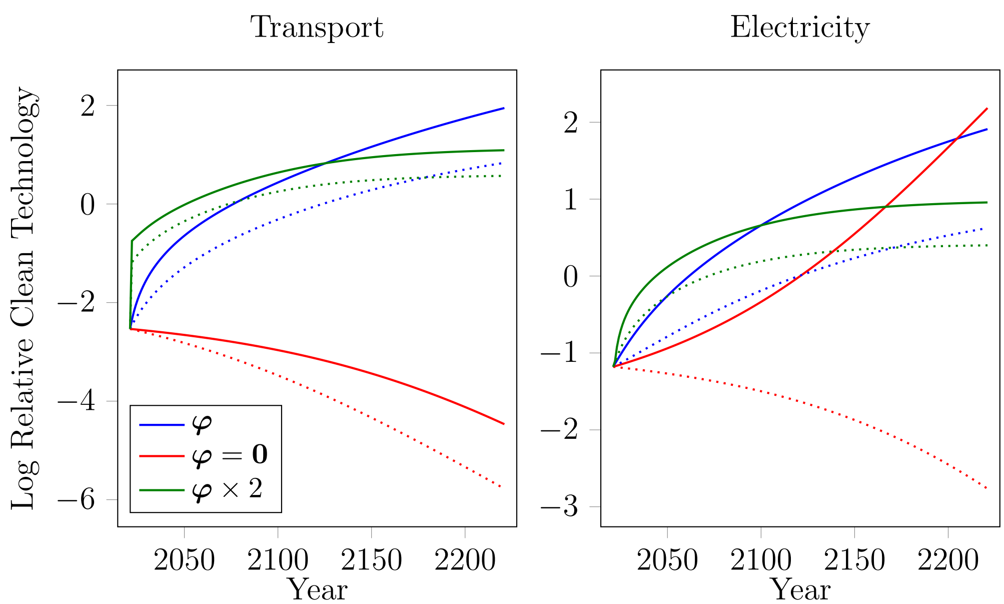

Notes: Impact of introducing a carbon price at the EPA's proposed SCC ($190) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Dotted lines indicate laissez-faire paths.

Figure A·8: Clean Income Share Path ($190 Carbon Price)

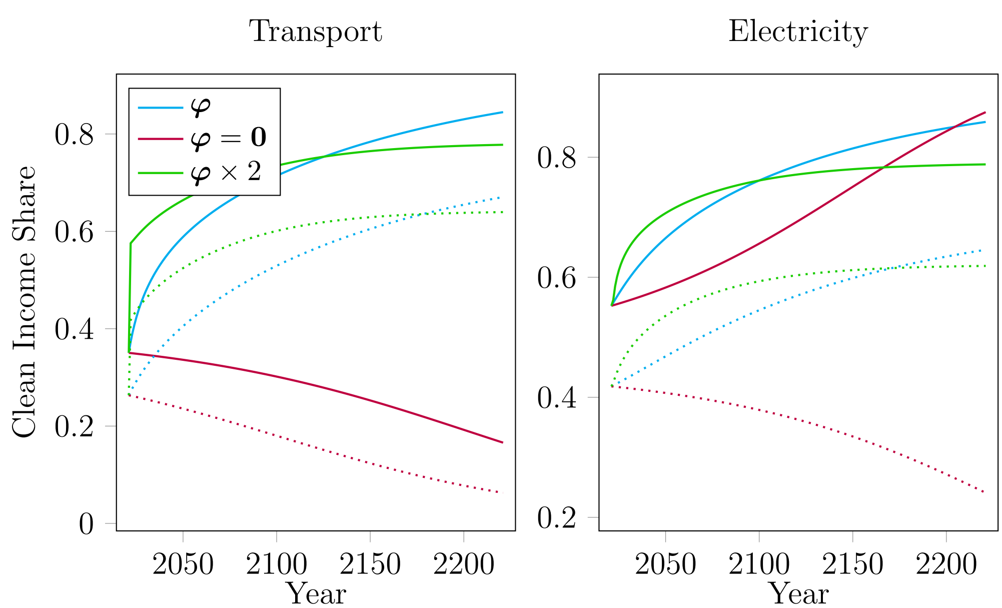

Notes: Impact of introducing a carbon price at the EPA's proposed SCC ($190) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Dotted lines indicate laissez-faire paths.

Figure A·9: Pollution Path ($190 Carbon Price)

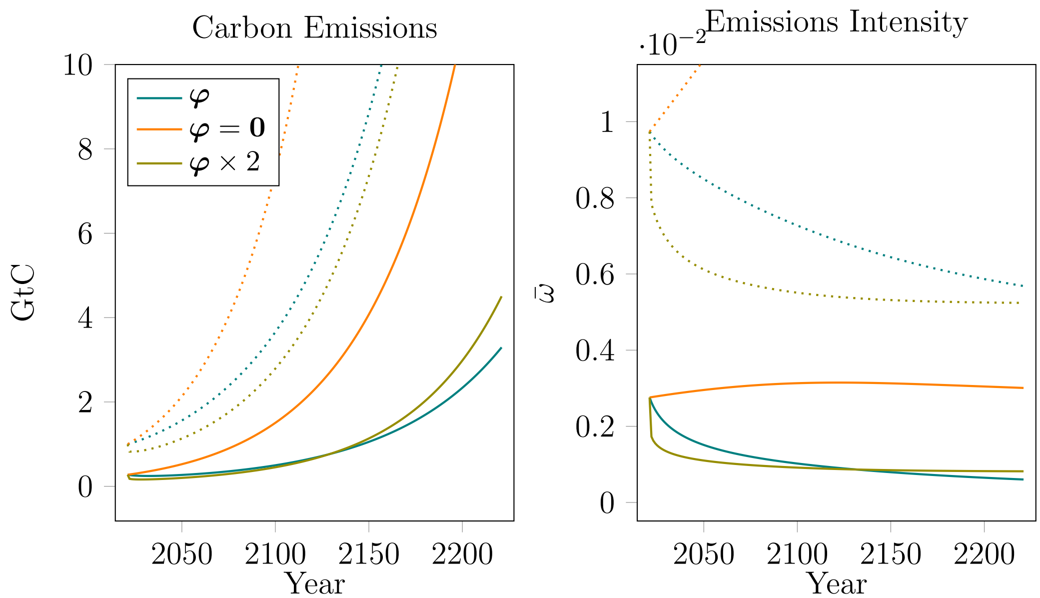

Notes: Impact of introducing a carbon price at the EPA's proposed SCC ($190) and clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Dotted lines indicate laissez-faire paths.

Figure A·10: Carbon Price Necessary for Clean Growth (No Spillovers)

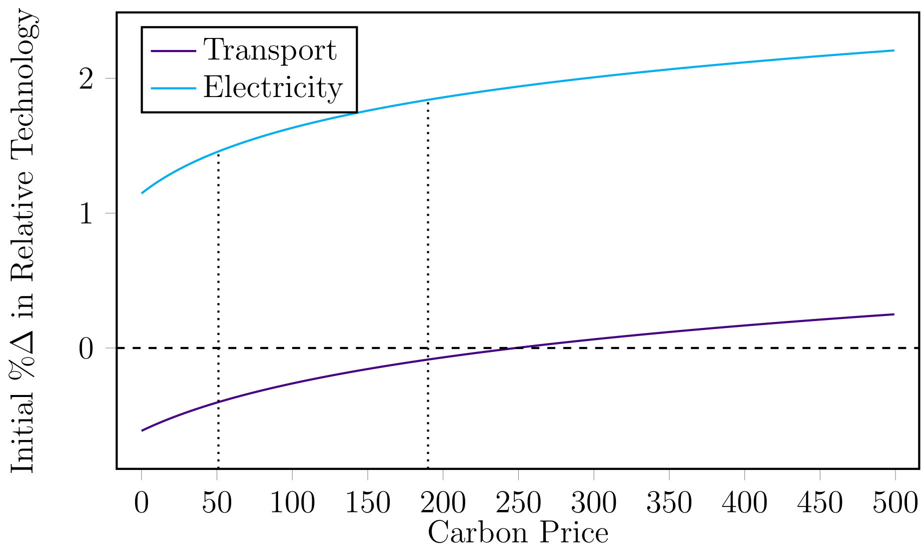

Notes: Direction of innovation as a function of the carbon price when the spillover network has been shut down. Values above zero indicate a sector is in its clean basin of attraction. Carbon prices are in addition to a clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Vertical dotted lines reference the Biden Administration's SCC ($51) and the EPA's proposed SCC ($190). Transportation switches to clean growth with a carbon price of about $248, while electricity generation switches without a carbon price.

Figure A·11: Clean Innovation Subsidy Necessary for Clean Growth (No Spillovers)

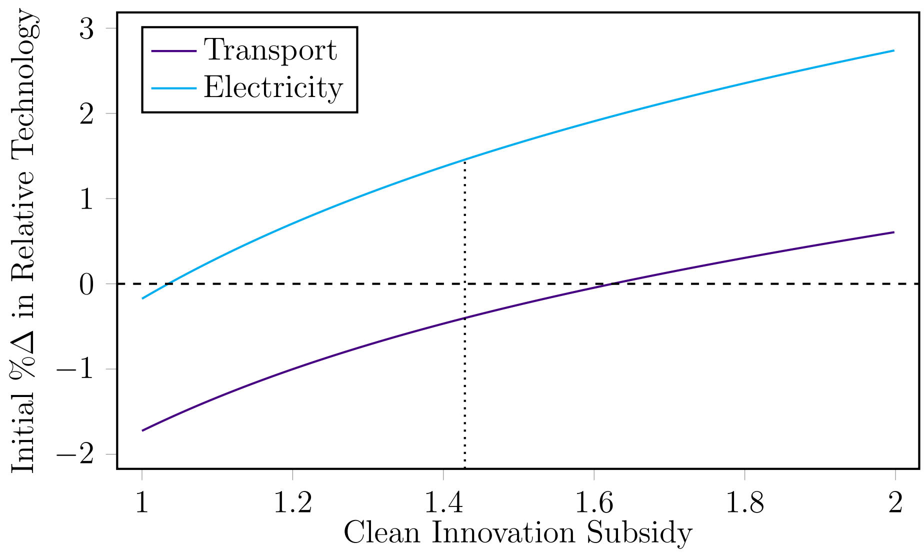

Notes: Direction of innovation as a function of the clean innovation subsidy when the spillover network has been shut down. Values above zero indicate a sector is in its clean basin of attraction. Clean innovation subsidies are in addition to a carbon price at the Biden Administration's SCC ($51). Vertical dotted line references clean innovation subsidy equivalent to a 30% tax credit ($\xi_c = 1.43$). Transportation switches to clean growth with a subsidy of 1.62, while electricity generation switches with a subsidy of 1.03.

Figure A·12: Half-Lives of Convergence (Double Spillovers)

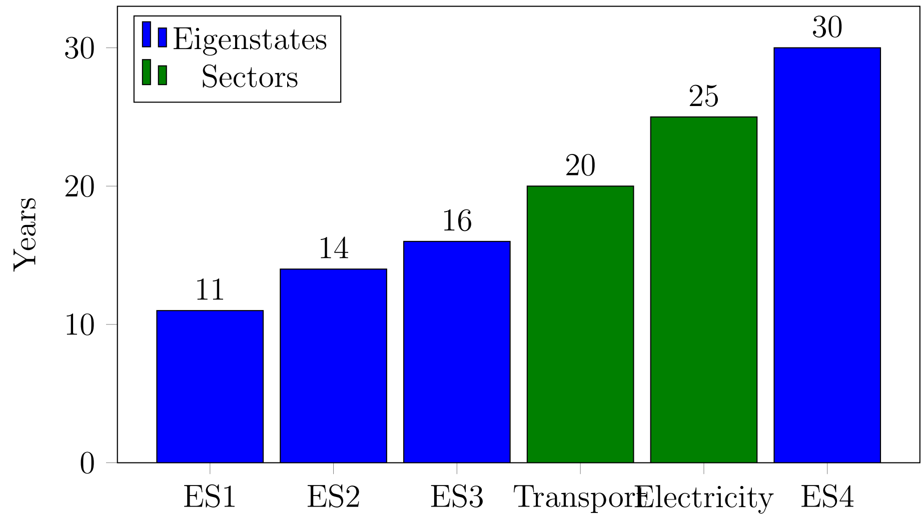

Notes: Transition speeds following the policy reform of Section 1.6.1. Cross-technology spillovers have been doubled.

Figure A·13: Half-Lives of Convergence (No Within-Sector Spillovers)

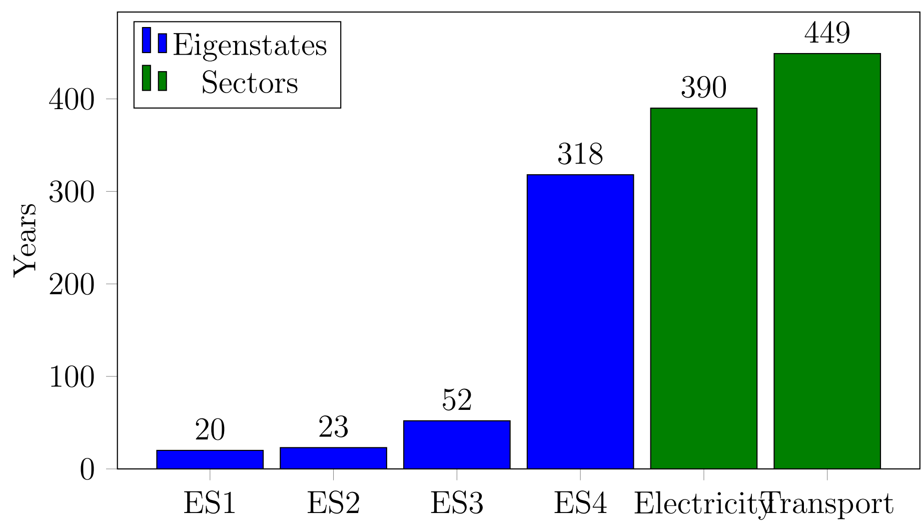

Notes: Transition speeds following the policy reform of Section 1.6.1. Within-sector spillovers are shut down by transferring them to the diagonal of the gross spillover network.

Figure A·14: Degree of Increasing Returns Doesn't Depend on Policy Reform

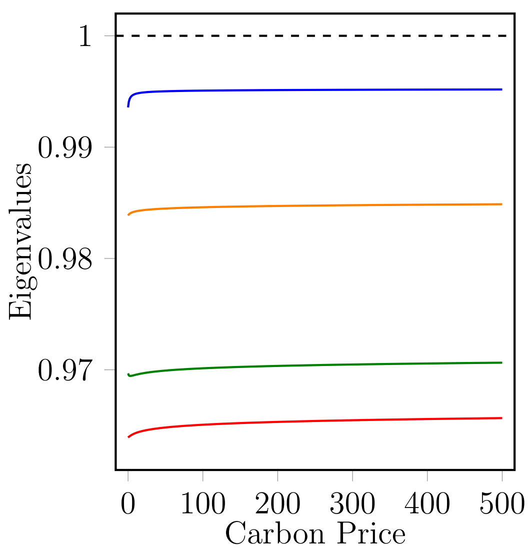

Notes: Eigenvalues of the transition matrix as a function of carbon prices and clean innovation subsidies. In each case, the other policy instrument is kept at its value from the policy reform of Section 1.6.1.

Figure A·15: First-Best Clean Income Share Path

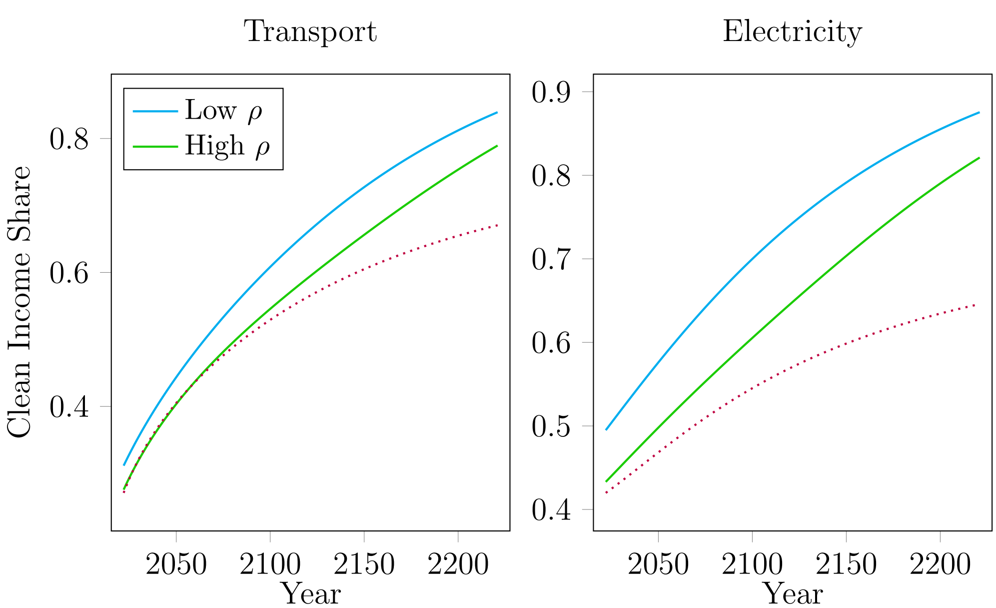

Notes: Dotted lines indicate laissez-faire paths.

Figure A·16: First-Best Temperature Path

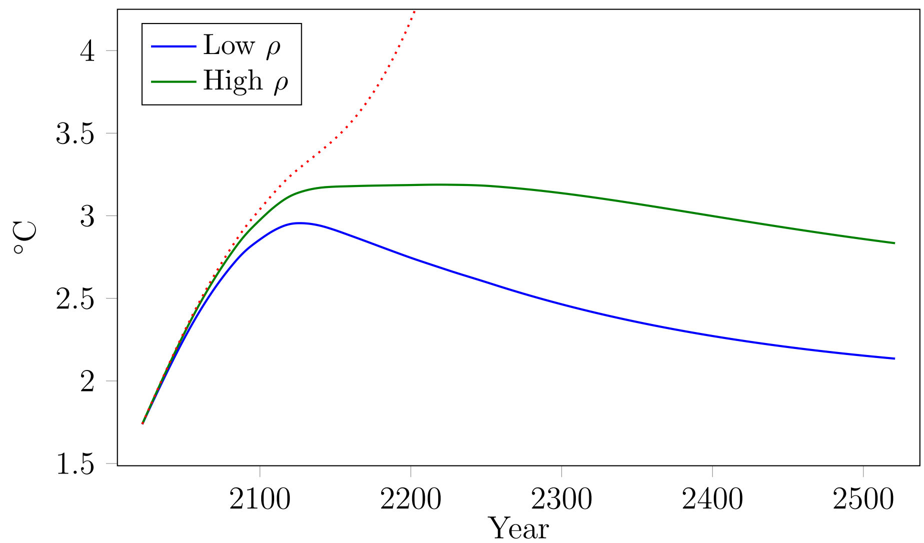

Notes: Temperature increases follow from the equation $\Delta T_t = \Gamma \ln(\mathcal{C}_t/\bar{\mathcal{C}})/\ln(2)$, with $\Gamma = 3$. See Footnote 36 for further discussion. Outside emissions come from the 2010 RICE model. The dotted line indicates the laissez-faire path.

Figure A·17: First-Best Policy Path (High Damages)

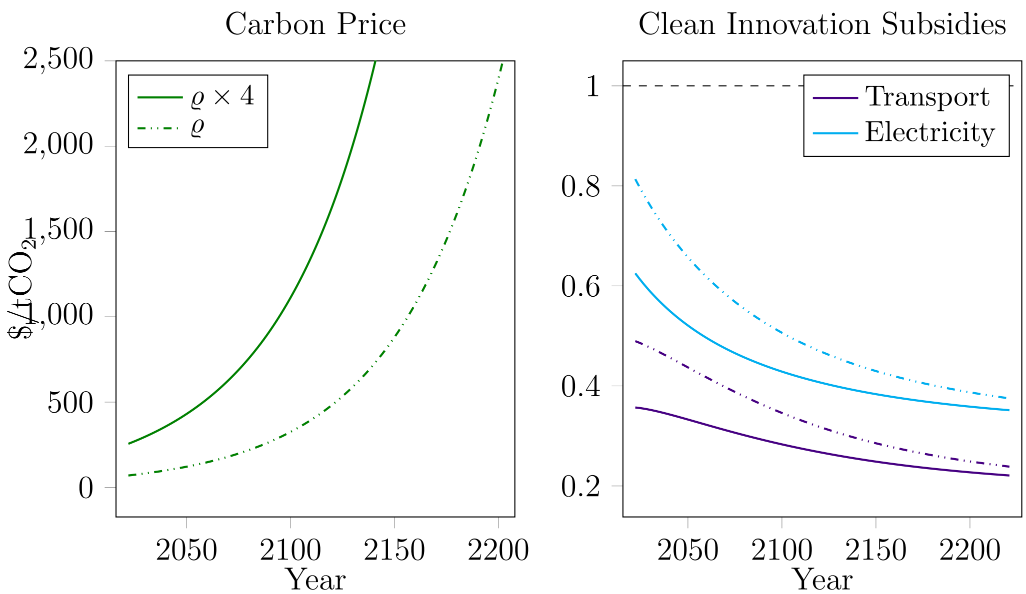

Notes: Optimal policy when climate damages are quadrupled. Dashed-dotted lines represent the benchmark calibration. Both cases use the low discount rate. Innovation subsidies are listed as a fraction of the baseline innovation wedge.

Figure A·18: Second-Best Clean Income Share Path

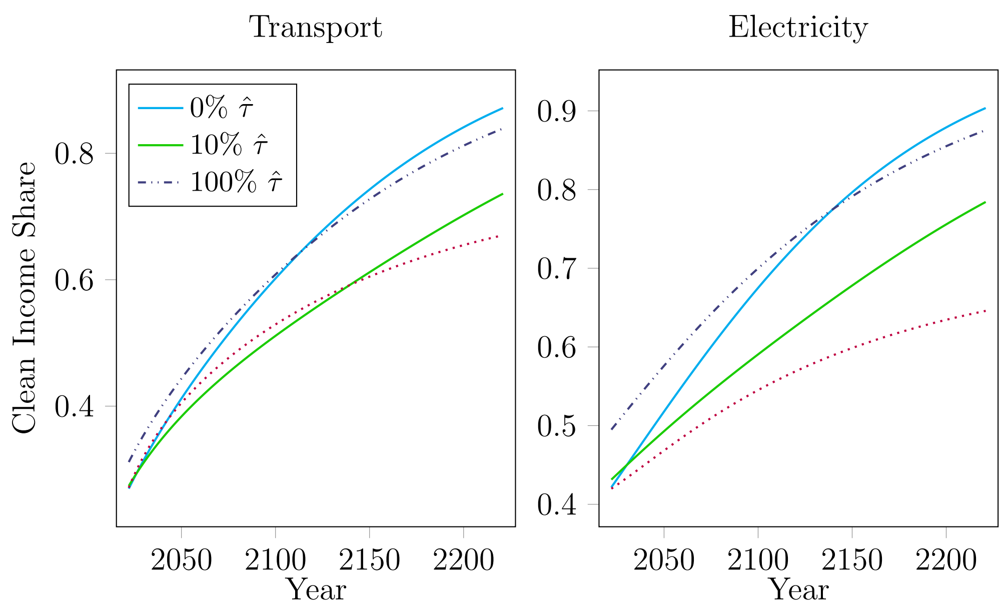

Notes: Policy paths use the low discount rate. The external carbon price $\hat{\tau}$ is a proportion of the social cost of carbon, so 100% is the first-best. Dotted lines indicate laissez-faire paths.

Figure A·19: Second-Best Temperature Path

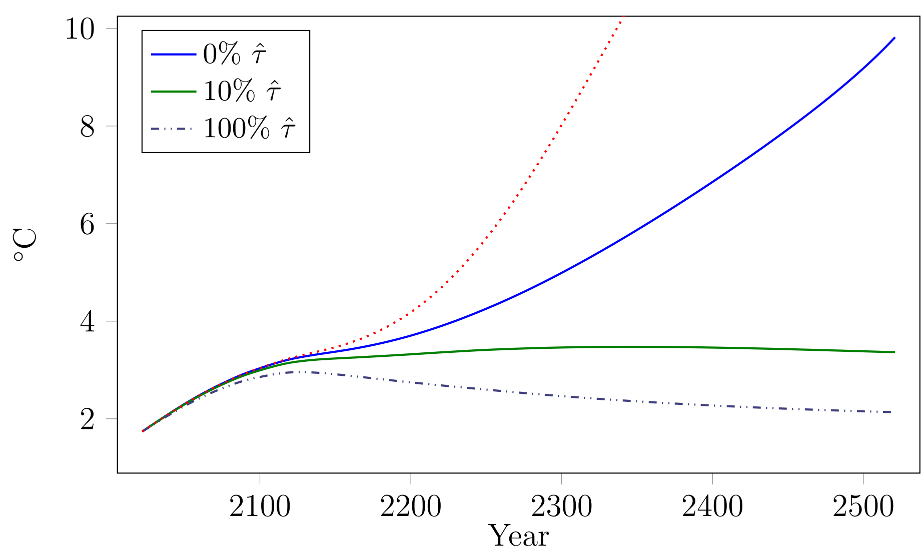

Notes: Temperature increases follow from the equation $\Delta T_t = \Gamma \ln(\mathcal{C}_t/\bar{\mathcal{C}})/\ln(2)$, with $\Gamma = 3$. See Footnote 36 for further discussion. Outside emissions come from the 2010 RICE model. Policy paths use the low discount rate. The external carbon price $\hat{\tau}$ is a proportion of the social cost of carbon, so 100% is the first-best. The dotted line indicates the laissez-faire path.

Figure A·20: Second-Best Policy Path (High Discounting)

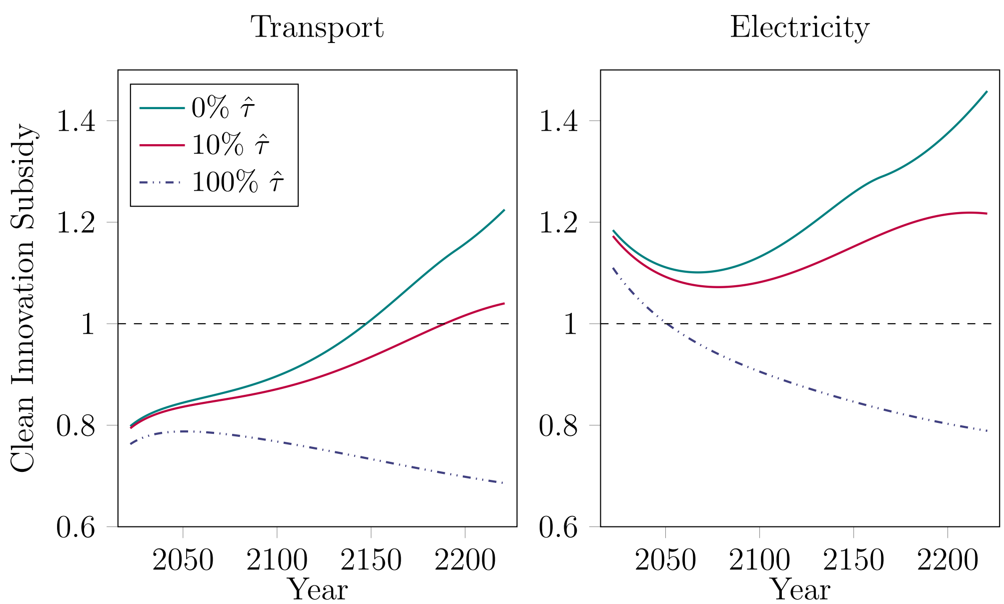

Notes: Optimal second-best policy with the high discount rate. The external carbon price $\hat{\tau}$ is a proportion of the social cost of carbon, so 100% is the first-best. Innovation subsidies are listed as a fraction of the baseline innovation wedge.

CHAPTER TWO APPENDIX

B.1 Derivations and Proofs for Static Model

B.1.1 The Aggregate Factor Production Function

The distribution of factors across tasks is determined by the cost-minimization problem:

$$\min_{y(.)} \int_{\Omega_K} \frac{R}{A_K \alpha_K(x)} y(x) dx + \sum_i \int_{\Omega_{L_i}} \frac{W_i}{A_{L_i} \alpha_{L_i}(x)} y(x) dx \;\; s.t. \;\; Y = F$$

where, as described in the text, the tasks are allocated to factors according to which factor has the lowest marginal cost.

With the ideal price index equal to one, this gives us the familiar condition

$$y(x) = Y/p(x)^{\sigma}$$

If we then consider market clearing for factor J, we can see

$$J = \int_{\Omega_J} \ell_J(x) dx = \int_{\Omega_J} \frac{y(x)}{\alpha_J(x) A_J} dx = \int_{\Omega_J} \frac{Y\big(\alpha_J(x) A_J\big)^{\sigma-1}}{W_J^{\sigma}}$$

$$\Rightarrow W_J = \Big(\int_{\Omega_J} \alpha_J(x)^{\sigma-1} dx\Big)^{\frac{1}{\sigma}} A_J^{\frac{\sigma-1}{\sigma}} \Big(\frac{Y}{J}\Big)^{\frac{1}{\sigma}}$$

Using this formula, one can plug marginal costs into the ideal price index equation and solve out for the desired result.

$$F\big(X, L, K\big) = \Big(\Lambda_K^{\frac{1}{\sigma}} \big(A_K K\big)^{\frac{\sigma-1}{\sigma}} + \sum_i \Lambda_{L_i}^{\frac{1}{\sigma}} \big(A_{L_i} L_i\big)^{\frac{\sigma-1}{\sigma}}\Big)^{\frac{\sigma}{\sigma-1}}$$

B.1.2 Derivation of the Jacobian of the Task Allocation

We have the level set

$$\Big\{(X,L,K,\theta)\in\mathbb{R}^{2N+2}\Big|\ \frac{F_K(1+\theta)}{A_K\alpha_K(x_i)}=\frac{F_{L_i}}{A_{L_i}\alpha_{L_i}(x_i)}\ \ \forall i\Big\}$$

Thus, define

$$G\equiv\begin{pmatrix}\frac{F_K(1+\theta)}{A_K\alpha_K(x_1)}-\frac{F_{L_1}}{A_{L_1}\alpha_{L_1}(x_1)}\\ \vdots\\ \frac{F_K(1+\theta)}{A_K\alpha_K(x_N)}-\frac{F_{L_N}}{A_{L_N}\alpha_{L_N}(x_N)}\end{pmatrix}$$

which, by construction, equals a vector of zeros when evaluated on the above level set.

Therefore, implicitly defining the task thresholds as a function of (L,K,θ), we have by the Implicit Function Theorem

$$\begin{pmatrix}\frac{\partial x_1}{\partial\theta}&\frac{\partial x_1}{\partial K}&\frac{\partial x_1}{\partial L_1}&\cdots&\frac{\partial x_1}{\partial L_N}\\ \vdots&\vdots&\vdots&\ddots&\\ \frac{\partial x_N}{\partial\theta}&\frac{\partial x_N}{\partial K}&\frac{\partial x_N}{\partial L_1}&\cdots&\frac{\partial x_N}{\partial L_N}\end{pmatrix}=-\begin{pmatrix}\frac{\partial G_1}{\partial x_1}&\cdots&\frac{\partial G_1}{\partial x_N}\\ \vdots&\ddots&\\ \frac{\partial G_N}{\partial x_1}&\cdots&\frac{\partial G_N}{\partial x_N}\end{pmatrix}^{-1}\begin{pmatrix}\frac{\partial G_1}{\partial\theta}&\frac{\partial G_1}{\partial K}&\frac{\partial G_1}{\partial L_1}&\cdots&\frac{\partial G_1}{\partial L_N}\\ \vdots&\vdots&\vdots&\ddots&\\ \frac{\partial G_N}{\partial\theta}&\frac{\partial G_N}{\partial K}&\frac{\partial G_N}{\partial L_1}&\cdots&\frac{\partial G_N}{\partial L_N}\end{pmatrix}\tag{B.1}$$

Let's consider the various parts of this matrix:

$$\frac{\partial G_i}{\partial x_i}=\frac{F_K(1+\theta)}{A_K\alpha_K(x_i)}\Big(\frac{\partial F_K/\partial x_i}{F_K}-\frac{\alpha_K'(x_i)}{\alpha_K(x_i)}\Big)-\frac{F_{L_i}}{A_{L_i}\alpha_{L_i}(x_i)}\Big(\frac{\partial F_{L_i}/\partial x_i}{F_{L_i}}-\frac{\alpha_{L_i}'(x_i)}{\alpha_{L_i}(x_i)}\Big)$$

Define $\tilde{P}_i\equiv\frac{F_K(1+\theta)}{A_K\alpha_K(x_i)}=\frac{F_{L_i}}{A_{L_i}\alpha_{L_i}(x_i)}$ and note that

$$\frac{\partial F_K}{\partial x_i}=\frac{F_K}{\sigma}\Big(\frac{\alpha_K(x_i)^{\sigma-1}}{\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx}+\frac{F_{x_i}}{F}\Big)$$

and

$$\frac{\partial F_{L_i}}{\partial x_i}=\frac{F_{L_i}}{\sigma}\Big(\frac{-\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_{L_i}}\alpha_{L_i}(x)^{\sigma-1}dx}+\frac{F_{x_i}}{F}\Big)$$

so we have

$$\frac{\partial G_i}{\partial x_i}=\tilde{P}_i\Big(\frac{1}{\sigma}\Big[\frac{\alpha_K(x_i)^{\sigma-1}}{\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx}+\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_{L_i}}\alpha_{L_i}(x)^{\sigma-1}dx}\Big]+\Big[\frac{\alpha_{L_i}'(x_i)}{\alpha_{L_i}(x_i)}-\frac{\alpha_K'(x_i)}{\alpha_K(x_i)}\Big]\Big)\ \ \forall i$$

Finally,

$$\left.\frac{\partial G_i}{\partial x_j}\right|_{i\neq j} = \frac{\tilde{P}_i}{\sigma}\frac{\alpha_K(x_j)^{\sigma-1}}{\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx}$$

Thus, we have

$$\frac{\partial G}{\partial x} = \begin{pmatrix} \frac{\tilde{P}_1}{\sigma}\left(\frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_{L_1}}\alpha_{L_1}(x)^{\sigma-1}dx} + \sigma\Delta\alpha_1\right. & & 0 \\ & \ddots & \\ 0 & & \frac{\tilde{P}_N}{\sigma}\left(\frac{\alpha_{L_N}(x_i)^{\sigma-1}}{\int_{\Omega_{L_N}}\alpha_{L_N}(x)^{\sigma-1}dx} + \sigma\Delta\alpha_N\right) \end{pmatrix}$$

$$+ \frac{1}{\sigma\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx}\begin{pmatrix}\tilde{P}_1 \\ \vdots \\ \tilde{P}_N\end{pmatrix}\begin{pmatrix}\alpha_K(x_1)^{\sigma-1} & \dots & \alpha_K(x_N)^{\sigma-1}\end{pmatrix},$$

where $\Delta\alpha_j \equiv \frac{\alpha'_{L_j}(x_j)}{\alpha_{L_j}(x_j)} - \frac{\alpha'_K(x_j)}{\alpha_K(x_j)}$. Therefore, we have a matrix of the form $A + uv^T$, which means, to find its inverse, we can apply the Sherman-Morrison formula where

$$\left(A + uv^T\right)^{-1} = A^{-1} - \frac{A^{-1}uv^TA^{-1}}{1 + v^TA^{-1}u}$$

Luckily, because A is diagonal, we have that A^{-1} is just the multiplicative inverse along the diagonal. Therefore, we have

$$A^{-1}uv^TA^{-1} = \left(\frac{u_iv_j}{A_iA_j}\right)_{i,j}$$

$$1 + v^TA^{-1}u = 1 + \sum_p \frac{u_pv_p}{A_p}$$

Thus, we have

$$\Big(\frac{\partial G}{\partial x}\Big)^{-1} = \frac{1}{1+\sum_p \frac{u_p v_p}{A_p}}\Bigg[\begin{pmatrix} \frac{1}{A_1}+\sum_p \frac{u_p v_p}{A_p A_1} & & 0 \\ & \ddots & \\ 0 & & \frac{1}{A_N}+\sum_p \frac{u_p v_p}{A_p A_N}\end{pmatrix} - \Big(\frac{u_i v_j}{A_i A_j}\Big)_{i,j}\Bigg]$$
$$= c\begin{pmatrix} \frac{1}{A_1}+\sum_{p\neq 1}\frac{u_p v_p}{A_p A_1} & \cdots & -\frac{u_1 v_N}{A_1 A_N} \\ \vdots & \ddots & \\ -\frac{u_N v_1}{A_1 A_N} & \cdots & \frac{1}{A_1}+\sum_{p\neq N}\frac{u_p v_p}{A_p A_1}\end{pmatrix}$$

where

$$c \equiv \frac{1}{1+\sum_p \frac{u_p v_p}{A_p}} \tag{B.2}$$

Now by defining

$$\hat{\alpha}_i \equiv \Big(\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_{L_i}}\alpha_{L_i}(x)^{\sigma-1}dx} + \sigma\Big[\frac{\alpha'_{L_i}(x_i)}{\alpha_{L_i}(x_i)} - \frac{\alpha'_K(x_i)}{\alpha_K(x_i)}\Big]\Big) \tag{B.3}$$

we have

$$A_i = \frac{\tilde{P}_i}{\sigma}\Big(\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_{L_i}}\alpha_{L_i}(x)^{\sigma-1}dx} + \sigma\Big[\frac{\alpha'_{L_i}(x_i)}{\alpha_{L_i}(x_i)} - \frac{\alpha'_K(x_i)}{\alpha_K(x_i)}\Big]\Big) = \frac{\tilde{P}_i\hat{\alpha}_i}{\sigma}$$

$$u_i = \frac{\tilde{P}_i}{\sigma\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx}$$
$$v_i = \alpha_K(x_i)^{\sigma-1}$$

Therefore,

$$\frac{u_i v_j}{A_i A_j} = \frac{\sigma}{\tilde{P}_j\hat{\alpha}_i\hat{\alpha}_j}\frac{\alpha_K(x_j)^{\sigma-1}}{\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx}$$

Furthermore,

$$\frac{u_p v_p}{A_p A_i} = \frac{\sigma}{\tilde{P}_i\hat{\alpha}_p\hat{\alpha}_i}\frac{\alpha_K(x_p)^{\sigma-1}}{\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx}$$

Finally,

$$\Big(\frac{\partial G}{\partial x}\Big)^{-1} = \sigma c\begin{pmatrix} \frac{1}{\tilde{P}_1\hat{\alpha}_1}+\sum_{p\neq 1}\frac{1}{\tilde{P}_1\hat{\alpha}_p\hat{\alpha}_1}\frac{\alpha_K(x_p)^{\sigma-1}}{\Lambda_K} & \cdots & \frac{-1}{\tilde{P}_N\hat{\alpha}_1\hat{\alpha}_N}\frac{\alpha_K(x_N)^{\sigma-1}}{\Lambda_K} \\ \vdots & \ddots & \\ \frac{-1}{\tilde{P}_1\hat{\alpha}_N\hat{\alpha}_1}\frac{\alpha_K(x_1)^{\sigma-1}}{\Lambda_K} & \cdots & \frac{1}{\tilde{P}_N\hat{\alpha}_N}+\sum_{p\neq N}\frac{1}{\tilde{P}_N\hat{\alpha}_p\hat{\alpha}_N}\frac{\alpha_K(x_p)^{\sigma-1}}{\Lambda_K}\end{pmatrix} \tag{B.4}$$

Now note that

$$\begin{pmatrix} \frac{\partial G_1}{\partial \theta} & \frac{\partial G_1}{\partial K} & \frac{\partial G_1}{\partial L_1} & \cdots & \frac{\partial G_1}{\partial L_N} \\ \vdots & \vdots & \vdots & \ddots & \\ \frac{\partial G_N}{\partial \theta} & \frac{\partial G_N}{\partial K} & \frac{\partial G_N}{\partial L_1} & \cdots & \frac{\partial G_N}{\partial L_N} \end{pmatrix} = \begin{pmatrix} \frac{\tilde{P}_1}{1+\theta} & -\frac{\tilde{P}_1}{\sigma}\frac{1}{K} & \frac{\tilde{P}_1}{\sigma}\frac{1}{L_1} & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \\ \frac{\tilde{P}_N}{1+\theta} & -\frac{\tilde{P}_N}{\sigma}\frac{1}{K} & 0 & \cdots & \frac{\tilde{P}_N}{\sigma}\frac{1}{L_N} \end{pmatrix} \tag{B.5}$$

Combining equations (B.4) and (B.5) as in (B.3.2) we have

$$\begin{pmatrix} \frac{\partial x_1}{\partial \theta} & \frac{\partial x_1}{\partial K} & \frac{\partial x_1}{\partial L_1} & \cdots & \frac{\partial x_1}{\partial L_N} \\ \vdots & \vdots & \vdots & \ddots & \\ \frac{\partial x_N}{\partial \theta} & \frac{\partial x_N}{\partial K} & \frac{\partial x_N}{\partial L_1} & \cdots & \frac{\partial x_N}{\partial L_N} \end{pmatrix} = \tag{B.6}$$

$$c \begin{pmatrix} -\sigma\frac{\hat{\alpha}_1^{-1}}{1+\theta} & \frac{\hat{\alpha}_1^{-1}}{K} & -\frac{1}{L_1}\tilde{\alpha}_1 & \cdots & \frac{\hat{\alpha}_1^{-1}\hat{\alpha}_N^{-1}}{L_N}\frac{\alpha_K(x_N)^{\sigma-1}}{\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx} \\ \vdots & \vdots & \vdots & \ddots & \\ -\sigma\frac{\hat{\alpha}_N^{-1}}{1+\theta} & \frac{\hat{\alpha}_N^{-1}}{K} & \frac{\hat{\alpha}_N^{-1}\hat{\alpha}_1^{-1}}{L_1}\frac{\alpha_K(x_1)^{\sigma-1}}{\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx} & \cdots & -\frac{1}{L_N}\tilde{\alpha}_N \end{pmatrix}$$

where $\tilde{\alpha}_j \equiv \hat{\alpha}_j^{-1} + \sum_{p\neq j}\hat{\alpha}_p^{-1}\hat{\alpha}_j^{-1}\frac{\alpha_K(x_p)^{\sigma-1}}{\int_{\Omega_K}\alpha_K(x)^{\sigma-1}dx})$. Hence, it is easy to verify that $\frac{\partial X}{\partial K} \propto \frac{\partial X}{\partial \theta}$.

B.1.3 Deriving the Static Lagrangian

First, we will modify the government constraint by using the CRTS property of the production function. Note that

$$G+\sum_i n_i D = \sum_i \tau_L F_{L_i} L_i + \tau_K F_K K \iff G+\sum_i n_i D = \tilde{F} - \sum_i (1-\tau_L) F_{L_i} L_i - (1-\tau_K) F_K K \tag{B.7}$$

Next, one can verify by inspection that the conjunction of propositions

$$(1-\tau_L)F_{L_i} = -u_{i,l}/u_{i,c}\ \forall i$$

$$(1-\tau_K)F_K = \phi'(K)$$

$$G + \sum_i n_i D = \sum_i \tau_L F_{L_i} L_i + \tau_K F_K K$$

$$C_i = (1-\tau_L)F_{L_i}L_i + n_i D + \omega_i\big((1-\tau_K)F_K K - \phi(K)\big)\ \forall i$$

is equivalent to

$$C_i = \frac{-u_{i,l}}{u_{i,c}} L_i + n_i D + \omega_i\big(\phi'(K)K - \phi(K)\big)\ \forall i$$

$$G + \sum_i n_i D = \tilde{F} - \sum_i \frac{-u_{i,l}}{u_{i,c}} Ł_i - \phi'(K)K$$

$$\frac{-u_{1,l}/u_{1,c}}{-u_{i,l}/u_{i,c}} = \frac{F_{L_1}}{F_{L_i}}\ \forall i \neq 1$$

Namely, select taxes to satisfy the factor optimality conditions and the rest follows.

Therefore, we can solve the Ramsey problem with the following Lagrangian:

$$\begin{aligned}\mathcal{L} = W(u_1, u_2, \dots, u_N) &+ \sum_i \lambda_i\big(\frac{-u_{i,l}}{u_{i,c}} L_i + n_i D + \omega_i\big[\phi'(K)K - \phi(K)\big] - C_i\big) \quad \text{(B.8)}\\ &+ \mu\big(\tilde{F} - \sum_i \frac{-u_{i,l}}{u_{i,c}} Ł_i - \phi'(K)K - G - \sum_i n_i D\big)\\ &+ \sum_{i\neq 1} \eta_i\big(\frac{-u_{1,l}/u_{1,c}}{-u_{i,l}/u_{i,c}} - \frac{F_{L_1}}{F_{L_i}}\big)\end{aligned}$$

B.1.4 Proof of Proposition 2.1

The FOC for capital is

$$\begin{aligned}&\sum_i \lambda_i \omega_i(\phi''(K)K) + \mu\big(F_K - \phi'(K) - \phi''(K)K\big)\\ &\qquad + \Big(\mu \sum_j \frac{\partial F}{\partial x_j}\frac{\partial x_j}{\partial K} - \sum_j \sum_{i\neq 1} \eta_i \frac{\partial\big(\frac{F_{L_1}}{F_{L_i}}\big)}{\partial x_j}\frac{\partial x_j}{\partial K}\Big)\\ &\qquad = \sum_{i\neq 1} \eta_i \frac{1}{F_{L_i}^2}\big(F_{L_1 K} F_{L_i} - F_{L_1} F_{L_i K}\big)\end{aligned}$$

With a CES aggregator, we have $\forall i$

$$F_{L_i} = \Big(\int_{\Omega_{L_i}} \alpha_{L_i}(j)^{\sigma-1} dj\Big)^{\frac{1}{\sigma}} A_{L_i}^{\frac{\sigma-1}{\sigma}} \Big(\frac{F}{L_i}\Big)^{\frac{1}{\sigma}}$$

Thus, $\forall i$

$$F_{L_1K}F_{L_i} - F_{L_1}F_{L_iK} = 0$$

Furthermore, note that $\frac{\partial x_i}{\partial K} = \frac{-(1+\theta)}{\sigma K}\frac{\partial x_i}{\partial \theta}$, so by the automation FOC

$$\Big(\mu \sum_j \frac{\partial F}{\partial x_j}\frac{\partial x_j}{\partial K} - \sum_j \sum_{i \neq 1} \eta_i \frac{\partial\big(\frac{F_{L_1}}{F_{L_i}}\big)}{\partial x_j}\frac{\partial x_j}{\partial K}\Big) = 0$$

Therefore, we have

$$\sum_i \lambda_i \omega_i (\phi''(K)K) + \mu\big(F_K - \phi'(K) - \phi''(K)K\big) = 0$$

so with the definition of the elasticity of capital and a few additional algebraic steps, we can derive equation (2.17).

$$\frac{\tau_K}{1-\tau_K} = \frac{\sum_i \omega_i(\mu - \lambda_i)}{\mu}\frac{1}{\varepsilon_K}$$

Equation (2.18) follows straightforwardly from differentiating the Lagrangian of equation (B.8) with respect to D.

$$\frac{\sum_i \lambda_i n_i}{\sum_i n_i} = \mu$$

One final point worthy of note is that if we combine equations (2.17) and (2.18) and assume equal per-capita wealth holdings, then the optimal tax on capital will be zero. This is because a tax on equal wealth holdings would be lump-sum and thus superfluous with a UBI.

□

B.1.5 Proof of Proposition 2.2

The FOC for θ is:

$$\mu \sum_j \frac{\partial F}{\partial x_j} \frac{\partial x_j}{\partial \theta} = \sum_j \sum_{i \neq 1} \eta_i \frac{\partial \left(\frac{F_{L_1}}{F_{L_i}} \right)}{\partial x_j} \frac{\partial x_j}{\partial \theta}$$

Note that

$$\left. \frac{\partial \left(\frac{F_{L_1}}{F_{L_i}} \right)}{\partial x_j} \right|_{j \neq i,1} = \frac{\frac{\partial F_{L_1}}{\partial x_j} F_{L_i} - F_{L_1} \frac{\partial F_{L_i}}{\partial x_j}}{F_{L_i}^2} = \frac{F_{L_1}}{F_{L_i}} \frac{\frac{\partial F / \partial x_j}{F} - \frac{\partial F / \partial x_j}{F}}{\sigma} = 0$$

Thus, we have

$$\mu \sum_j \frac{\partial F}{\partial x_j} \frac{\partial x_j}{\partial \theta} = \sum_{i \neq 1} \eta_i \Big(\frac{\partial \left(\frac{F_{L_1}}{F_{L_i}} \right)}{\partial x_1} \frac{\partial x_1}{\partial \theta} + \frac{\partial \left(\frac{F_{L_1}}{F_{L_i}} \right)}{\partial x_i} \frac{\partial x_i}{\partial \theta} \Big)$$

Note that

$$\frac{\partial \left(\frac{F_{L_1}}{F_{L_i}} \right)}{\partial x_1} = \frac{\frac{\partial F_{L_1}}{\partial x_1} F_{L_i} - F_{L_1} \frac{\partial F_{L_i}}{\partial x_1}}{F_{L_i}^2} = \frac{-1}{\sigma} \frac{F_{L_1}}{F_{L_i}} \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_1} \alpha_{L_1}(x)^{\sigma-1} dx}$$

and

$$\frac{\partial \left(\frac{F_{L_1}}{F_{L_i}} \right)}{\partial x_i} = \frac{\frac{\partial F_{L_1}}{\partial x_i} F_{L_i} - F_{L_1} \frac{\partial F_{L_i}}{\partial x_i}}{F_{L_i}^2} = \frac{1}{\sigma} \frac{F_{L_1}}{F_{L_i}} \frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_i} \alpha_{L_i}(x)^{\sigma-1} dx}$$

Finally, by cancelling the constant on $\frac{\partial x}{\partial \theta}$, we have equation (2.19).

$$\mu \sum_i \frac{\partial F}{\partial x_i} \hat{\alpha}_i^{-1} = \sum_{i \neq 1} \frac{\eta_i}{\sigma} \frac{F_{L_1}}{F_{L_i}} \Big(\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_i} \alpha_{L_i}(x)^{\sigma-1} dx} \hat{\alpha}_i^{-1} - \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_1} \alpha_{L_1}(x)^{\sigma-1} dx} \hat{\alpha}_1^{-1} \Big)$$

B.1.6 Derivation of the Optimality Conditions for the Labor Tax

Consider the FOC of labor for person $i \neq 1$:

$$W_i u_{i,l} + \lambda_i\big(\frac{\partial MRS_i}{\partial L_i} L_i + MRS_i\big) + \mu\big(F_{L_i} - MRS_i - \frac{\partial MRS_i}{\partial L_i} L_i\big)$$
$$+ \Big(\mu \sum_j \frac{\partial F}{\partial x_j}\frac{\partial x_j}{\partial L_i} - \sum_j \sum_{p\neq 1} \eta_p \frac{\partial\big(\frac{F_{L_1}}{F_{L_p}}\big)}{\partial x_j}\frac{\partial x_j}{\partial L_i}\Big)$$
$$= \eta_i \frac{MRS_1}{MRS_i^2}\frac{\partial MRS_i}{\partial L_i} + \sum_{j\neq 1} \eta_j \frac{\partial F_{L_1}/F_{L_j}}{\partial L_i}$$

First, as noted in the case of capital, only the relative wage of person i will change with the labor supply of person i. We may also eliminate many of the changes in relative wages from changes in thresholds, as noted in the automation FOC.

Next, note that we can divide by i's MRS and substitute for their implementation constraint and pseudo-Frisch elasticity: $\hat{\varepsilon}_{L_i} \equiv \frac{MRS_i}{L_i \cdot \partial MRS_i/\partial L_i}$. Note that $\hat{\varepsilon}_{L_i}$ is the Frisch elasticity in the case of additively separable utility. Then, we have:

$$\frac{\lambda_i - \mu}{\hat{\varepsilon}_{L_i}} + \mu\frac{\tau_L}{1-\tau_L} = W_i u_{i,c} - \lambda_i + \eta_i \frac{MRS_1}{MRS_i}\frac{1}{MRS_i L_i}\frac{1}{\hat{\varepsilon}_{L_i}} + \eta_i \frac{F_{L_1}}{F_{L_i}}\frac{1}{MRS_i L_i}\frac{1}{\sigma} \quad \text{(B.9)}$$
$$- \frac{1}{MRS_i}\Bigg(\mu \sum_j \frac{\partial F}{\partial x_j}\frac{\partial x_j}{\partial L_i}$$
$$- \sum_{j\neq 1} \frac{\eta_j}{\sigma}\frac{F_{L_1}}{F_{L_j}}\Big[\frac{\alpha_{L_j}(x_j)^{\sigma-1}}{\int_{\Omega_j}\alpha_{L_j}(x)^{\sigma-1}dx}\frac{\partial x_j}{\partial L_i} - \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_1}\alpha_{L_1}(x)^{\sigma-1}dx}\frac{\partial x_1}{\partial L_i}\Big]\Bigg)$$

Define $I_i \equiv \frac{-u_{i,l}}{u_{i,c}} L_i$ as the labor income of type i. Now consider the FOC of consumption for person $i \neq 1$ and note that $\frac{MRS_1}{MRS_i} = \frac{F_{L_1}}{F_{L_i}}$:

$$W_i u_{i,c} - \lambda_i = \big(\mu - \lambda_i\big)\frac{\partial MRS_i}{\partial C_i} L_i + \eta_i \frac{F_{L_1}}{F_{L_i}}\frac{1}{I_i}\frac{\partial MRS_i}{\partial C_i} L_i \quad \text{(B.10)}$$

Denote by $\varepsilon^h_{L_i}$ the Hicksian labor supply elasticity of type i. Note that by implicitly differentiating the constraint on the expenditure minimization problem, $u(c, l) = \bar{u}$, we can derive

$$\frac{\partial C}{\partial L} = \frac{-u_{i,l}}{u_{i,c}}$$

Now, differentiating the FOC of the expenditure minimization problem, $MRS = W$, with respect to the wage gives us

$$\frac{\partial MRS}{\partial C}\frac{\partial C}{\partial L}\frac{\partial L}{\partial W} + \frac{\partial MRS}{\partial L}\frac{\partial L}{\partial W} = 1$$

$$\Rightarrow \frac{\partial MRS}{\partial C}L + \frac{\frac{\partial MRS}{\partial L}L}{MRS} = \frac{1}{\varepsilon_L^h} \tag{B.11}$$

Therefore, combining equations (B.9), (B.10), and (B.11), we have

$$\begin{aligned} \frac{\tau_L}{1-\tau_L} = \frac{\mu - \lambda_i}{\mu}\frac{1}{\varepsilon^h_{L_i}} + \frac{\eta_i}{\mu}\frac{F_{L_1}}{F_{L_i}}\frac{1}{I_i}\Big(\frac{1}{\varepsilon^h_{L_i}} + \frac{1}{\sigma}\Big) \qquad \text{(B.12)} \\ - \frac{1}{\mu MRS_i}\Bigg(\mu\sum_j \frac{\partial F}{\partial x_j}\frac{\partial x_j}{\partial L_i} \\ - \sum_{j\neq 1}\frac{\eta_j}{\sigma}\frac{F_{L_1}}{F_{L_j}}\Big[\frac{\alpha_{L_j}(x_j)^{\sigma-1}}{\int_{\Omega_j}\alpha_{L_j}(x)^{\sigma-1}dx}\frac{\partial x_j}{\partial L_i} - \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_1}\alpha_{L_1}(x)^{\sigma-1}dx}\frac{\partial x_1}{\partial L_i}\Big]\Bigg) \quad \forall i \neq 1 \end{aligned}$$

This term has two main components. The first line has behavioral responses (efficiency terms) multiplied by social values (equity terms). The second line is the automation effect from changes in labor supply. It is like that of the FOC for the threshold rule, except with new weights which represent the change in thresholds from labor rather than the threshold rule. Thus, as these new weights differ from the old weights, we will have the labor tax change as well.

B.1.7 Proof of Proposition 2.3

The capital FOC is as before:

$$\sum_i \lambda_i\omega_i(\phi''(K)K) + \mu\big(F_K - \phi'(K) - \phi''(K)K\big)$$

$$+\Big(\mu\sum_j \frac{\partial F}{\partial x_j}\frac{\partial x_j}{\partial K} - \sum_j\sum_{i\neq 1}\eta_i\frac{\partial\big(\frac{F_{L_1}}{F_{L_i}}\big)}{\partial x_j}\frac{\partial x_j}{\partial K}\Big)$$

$$= \sum_{i\neq 1}\eta_i\frac{1}{F_{L_i}^2}\big(F_{L_1K}F_{L_i} - F_{L_1}F_{L_iK}\big)$$

except now $\frac{\partial F}{\partial X} = 0$, so we have

$$\sum_i \lambda_i\omega_i(\phi''(K)K)+\mu\big(F_K - \phi'(K) - \phi''(K)K\big) =$$

$$\sum_{i\neq 1}\frac{\eta_i}{\sigma}\frac{F_{L_1}}{F_{L_i}}\Big(\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\int_{\Omega_i}\alpha_{L_i}(x)^{\sigma-1}dx}\frac{\partial x_i}{\partial K} - \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_1}\alpha_{L_1}(x)^{\sigma-1}dx}\frac{\partial x_1}{\partial K}\Big)$$

Equation (2.23) follows from the definition of capital supply elasticity and some additional algebra:

$$\frac{\tau_K}{1-\tau_K} = \frac{\sum_i\omega_i(\mu-\lambda_i)}{\mu}\frac{1}{\varepsilon_K} + \frac{c}{\mu}\frac{1}{\phi'(K)K}\sum_{i\neq 1}\frac{\eta_i}{\sigma}\frac{F_{L_1}}{F_{L_i}}\Big(\frac{\alpha_{L_i}(x_i)^{\sigma-1}}{\Lambda_{L_i}}\hat{\alpha}_i^{-1} - \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\Lambda_{L_1}}\hat{\alpha}_1^{-1}\Big)$$

The labor tax derivation is the same as in B.1.6, but the effect of threshold changes on production is now zero. Therefore, we have

$$\frac{\tau_L}{1-\tau_L} = \frac{\mu-\lambda_i}{\mu}\frac{1}{\varepsilon^h_{L_i}} + \frac{\eta_i}{\mu}\frac{F_{L_1}}{F_{L_i}}\frac{1}{I_i}\Big(\frac{1}{\varepsilon^h_{L_i}} + \frac{1}{\sigma}\Big) \tag{B.13}$$

$$+\frac{1}{\mu MRS_i}\sum_{j\neq 1}\frac{\eta_j}{\sigma}\frac{F_{L_1}}{F_{L_j}}\Big[\frac{\alpha_{L_j}(x_j)^{\sigma-1}}{\int_{\Omega_j}\alpha_{L_j}(x)^{\sigma-1}dx}\frac{\partial x_j}{\partial L_i} - \frac{\alpha_{L_1}(x_1)^{\sigma-1}}{\int_{\Omega_1}\alpha_{L_1}(x)^{\sigma-1}dx}\frac{\partial x_1}{\partial L_i}\Big] \quad \forall i \neq 1$$

□

B.2 Derivations and Proofs for Dynamic Model

B.2.1 The Dynamic Ramsey Problem

We want to solve the problem

$$\max_{\{C_{g,t}\},\{L_{g,t}\},\{a_{g,t}\},\{D_t\},\{\theta_t\},\{B_t\},\{\tau_{L,t}\},\{\tau_{K,t}\}} W\big(\{V_g\}\big) \;\; s.t.$$

$$\big[\text{IC1}\big]: \; (1-\tau_{L,t})F_{L_i,t} = \frac{-V_{g,l_t}}{V_{g,c_t}} \;\; \forall i,t;\; \forall g \in \mathcal{G}_{i,t}$$

$$\big[\text{IC2}\big]: \; V_{g,c_{t-1}} = \big(1+(1-\tau_{K,t})F_{K,t}-\delta\big)V_{g,c_t} \;\; \forall g \in \mathcal{G};\; \forall t \in \{t_g+1,\dots t_g+T_g\}$$

$$\big[\text{BC1}\big]: \; a_{g,t+1}+C_{g,t} = \big(1+(1-\tau_{K,t})F_{K,t}-\delta\big)a_{g,t} + (1-\tau_{L,t})F_{L_i,t}L_{g,t} + D_t$$

$$\big[\text{BC2}\big]: \; a_{g,t+1}+C_{g,t} = (1-\tau_{L,t})F_{L_i,t}L_{g,t} + D_t \;\; \forall i,t;\; \forall g \in \mathcal{G}^B_{i,t} \setminus \mathcal{G}^D_{i,t}$$

$$\big[\text{BC3}\big]: \; C_{g,t} = \big(1+(1-\tau_{K,t})F_{K,t}-\delta\big)a_{g,t} + (1-\tau_{L,t})F_{L_i,t}L_{g,t} + D_t \;\; \forall i,t;\; \forall g \in \mathcal{G}^D_{i,t} \setminus \mathcal{G}^B_{i,t}$$

$$\big[\text{BC4}\big]: \; C_{g,t} = (1-\tau_{L,t})F_{L_i,t}L_{g,t} + D_t \;\; \forall i,t;\; \forall g \in (\mathcal{G}^B_{i,t} \cap \mathcal{G}^D_{i,t})$$

$$\big[\text{GBC}\big]: \; G_t + D_t|\mathcal{G}_t| + \big(1+F_{K,t}-\delta\big)B_t = B_{t+1} + \sum_{g\in\mathcal{G}_t} \tau_{K,t}F_{K,t}a_{g,t} + \sum_i \sum_{g\in\mathcal{G}_{i,t}} \tau_{L,t}F_{L_i,t}L_{g,t}$$

Note that we have four budget constraint cases here. This is from the permutations generated by the possibility of having been just born or being about to die.

Furthermore, using $K_t + B_t = \sum a_{g,t}$, the government budget constraint becomes

$$G_t + D_t|\mathcal{G}_t| + \big(1+(1-\tau_{K,t})F_{K,t}-\delta\big)B_t = B_{t+1} + \tau_{K,t}F_{K,t}K_t + \sum_i \sum_{g\in\mathcal{G}_{i,t}} \tau_{L,t}F_{L_i,t}L_{g,t}$$

One could also get this result by imagining that government bonds are tax exempt, so by no arbitrage, the coupon payments are $(1-\tau_{K,t})F_{K,t}-\delta$, rather than $F_{K,t}-\delta$ with a tax liability of $\tau_{K,t}F_{K,t}$. The two setups are equivalent.

Note that the satisfaction of the various budget constraints in each period implies the satisfaction of the resource constraint in each period.

Using the linear homogeneity of the production function,

$$
\begin{aligned}
& G_t + D_t|\mathcal{G}_t| + \big(1 + (1-\tau_{K,t})F_{K,t} - \delta\big)B_t = B_{t+1} + \tilde{F}_t \\
& \quad - (1-\tau_{K,t})F_{K,t}K_t - \sum_i \sum_{g\in\mathcal{G}_{i,t}} (1-\tau_{L,t})F_{L_i,t}L_{g,t} \\
\Rightarrow\; & G_t + D_t|\mathcal{G}_t| + (1-\delta)B_t = B_{t+1} + \tilde{F}_t \\
& \quad - \sum_{g\in\mathcal{G}_t\setminus\mathcal{G}_t^B} (1-\tau_{K,t})F_{K,t}a_{g,t} - \sum_i \sum_{g\in\mathcal{G}_{i,t}} (1-\tau_{L,t})F_{L_i,t}L_{g,t}
\end{aligned}
$$

Plugging the implementation constraints into the budget constraints, we can transform the problem as follows:

$$
\begin{aligned}
\mathcal{L} &= W\big(\{V_g\}\big) \\
&+ \sum_{g\in\mathcal{G}_0\setminus(\mathcal{G}_0^B\cup\mathcal{G}_0^D)} \lambda_{g,0}\big((1+(1-\tau_{K,0})F_{K,0}-\delta)a_{g,0} + \frac{-V_{g,l_0}}{V_{g,c_0}}L_{g,0} + D_0 - a_{g,1} - C_{g,0}\big) \\
&+ \sum_{t\neq 0}\sum_{g\in\mathcal{G}_t\setminus(\mathcal{G}_t^B\cup\mathcal{G}_t^D)} \lambda_{g,t}\big(\frac{V_{g,c_{t-1}}}{V_{g,c_t}}a_{g,t} + \frac{-V_{g,l_t}}{V_{g,c_t}}L_{g,t} + D_t - a_{g,t+1} - C_{g,t}\big) \\
&+ \sum_{g\in\mathcal{G}_0^D\setminus\mathcal{G}_0^B} \lambda_{g,0}\big((1+(1-\tau_{K,0})F_{K,0}-\delta)a_{g,0} + \frac{-V_{g,l_0}}{V_{g,c_0}}L_{g,0} + D_0 - C_{g,0}\big) \\
&+ \sum_{t\neq 0}\sum_{g\in\mathcal{G}_t^D\setminus\mathcal{G}_t^B} \lambda_{g,t}\big(\frac{V_{g,c_{t-1}}}{V_{g,c_t}}a_{g,t} + \frac{-V_{g,l_t}}{V_{g,c_t}}L_{g,t} + D_t - C_{g,t}\big) \\
&+ \sum_{t}\sum_{g\in\mathcal{G}_t^B\setminus\mathcal{G}_t^D} \lambda_{g,t}\big(\frac{-V_{g,l_t}}{V_{g,c_t}}L_{g,t} + D_t - a_{g,t+1} - C_{g,t}\big) \\
&+ \sum_{t}\sum_{g\in(\mathcal{G}_t^B\cap\mathcal{G}_t^D)} \lambda_{g,t}\big(\frac{-V_{g,l_t}}{V_{g,c_t}}L_{g,t} + D_t - C_{g,t}\big) \\
&+ \mu_0\big(B_1 + \tilde{F}_0 - \sum_{g\in\mathcal{G}_0\setminus\mathcal{G}_0^B}(1-\tau_{K,0})F_{K,0}a_{g,0} \\
&- \sum_{g\in\mathcal{G}_0}\frac{-V_{g,l_0}}{V_{g,c_0}}L_{g,0} - G_0 - D_0|\mathcal{G}_0| - (1-\delta)B_0\big) \\
&+ \sum_{t\neq 0}\mu_t\big(B_{t+1} + \tilde{F}_t - \sum_{g\in\mathcal{G}_t\setminus\mathcal{G}_t^B}\big(\frac{V_{g,c_{t-1}}}{V_{g,c_t}} + \delta - 1\big)a_{g,t} \\
&- \sum_{g\in\mathcal{G}_t}\frac{-V_{g,l_t}}{V_{g,c_t}}L_{g,t} - G_t - D_t|\mathcal{G}_t| - (1-\delta)B_t\big) \\
&+ \sum_{t}\sum_{i}\sum_{g\in\mathcal{G}_{i,t}\setminus\{g'_t\}} \eta_{g,t}\big(\frac{-V_{g'_t,l_t}/V_{g'_t,c_t}}{-V_{g,l_t}/V_{g,c_t}} - \frac{F_{L_1,t}}{F_{L_i,t}}\big) \\
&+ \sum_{t}\sum_{g\in\mathcal{G}_t\setminus(\mathcal{G}_t^D\cup\{g''_t\})} \gamma_{g,t}\big(\frac{V_{g''_t,c_t}}{V_{g''_t,c_{t+1}}} - \frac{V_{g,c_t}}{V_{g,c_{t+1}}}\big)
\end{aligned}
\tag{B.14}
$$

There are six terms for household budget constraints in this Lagrangian. We have the four permutations discussed earlier and two additional terms because the initial tax on capital can be directly chosen by the Planner. This is because the Euler

Equation substitution need not actually hold for the initial period as initial wealth is exogenous. The initial capital tax can, therefore, be freely chosen without respect for any implementation constraints. I am assuming here that the tax schedules are announced at the beginning of $t = 0$, so the previous period's savings are unaffected by $\tau_{K,0}$.

We then have two government budget constraint terms, with the first reflecting the direct choice of capital taxation in the initial period. The final two terms ensure that the tax rates are the same for all households. g'_t is some household picked from labor group one in each period, while g''_t is some household picked from the households that are not old in each period. Note this choice can differ across periods; the point is just the necessity of making some choices each period to maintain a uniform tax schedule across households.

B.2.2 Proof of Proposition 2.4

The FOC for $a_{g,t+1}|_{g\in\mathcal{G}_t\setminus\mathcal{G}_t^D}$ is then:

$$
\begin{aligned}
&-\lambda_{g,t} + \lambda_{g,t+1}\frac{V_{g,c_t}}{V_{g,c_{t+1}}} + \mu_{t+1}F_{K,t+1} - \mu_{t+1}\Big(\frac{V_{g,c_t}}{V_{g,c_{t+1}}} + \delta - 1\Big) \\
&+ \Big(\mu_{t+1}\sum_i \frac{\partial F_{t+1}}{\partial x_{i,t+1}}\frac{\partial x_{i,t+1}}{\partial K_{t+1}} - \sum_i \sum_{g\in\mathcal{G}_{i,t+1}\setminus\{g'_{t+1}\}} \eta_{g,t+1}\sum_j \frac{\partial\big(\frac{F_{L_1,t+1}}{F_{L_i,t+1}}\big)}{\partial x_{j,t+1}}\frac{\partial x_{j,t+1}}{\partial K_{t+1}}\Big) \\
&= \sum_i \sum_{g\in\mathcal{G}_{i,t+1}\setminus\{g'_{t+1}\}} \eta_{g,t+1}\frac{\partial\big(\frac{F_{L_1,t+1}}{F_{L_i,t+1}}\big)}{\partial K_{t+1}}
\end{aligned}
$$

As noted before, relative wages are not affected by the capital accumulation, so the last term is zero. Furthermore, the effect of capital on the automation thresholds is proportional to that of the threshold rule, so the middle term is zero due to the threshold rule FOC (2.34).

Therefore, we have equation (2.32)

$$\lambda_{g,t+1}\big(1+(1-\tau_{K,t+1})F_{K,t+1}-\delta\big)+\mu_{t+1}\tau_{K,t+1}F_{K,t+1}=\lambda_{g,t}$$

The derivation of equation (2.33) just requires differentiating the Lagrangian of (B.14) with respect to D_t to yield

$$\frac{\sum_{g\in\mathcal{G}_t}\lambda_{g,t}}{|\mathcal{G}_t|}=\mu_t$$

□

B.2.3 Derivation of Additional Optimal Policies in Dynamic Environment

To derive optimal debt policy, consider the FOC for B_{t+1} in the Lagrangian of equation (B.14):

$$\begin{aligned}
&\mu_t-\mu_{t+1}\big(1+F_{K,t+1}-\delta\big)\\
&-\mu_{t+1}\sum_i\frac{\partial F_{t+1}}{\partial x_{i,t+1}}\frac{\partial x_{i,t+1}}{\partial K_{t+1}}-\sum_{i,g\in\mathcal{G}_{i,t+1}\setminus\{g'_{t+1}\}}\eta_{g,t+1}\sum_j\frac{\partial\big(\frac{F_{L_1,t+1}}{F_{L_i,t+1}}\big)}{\partial x_{j,t+1}}\frac{\partial x_{j,t+1}}{\partial K_{t+1}}\\
&=-\sum_i\sum_{g\in\mathcal{G}_{i,t+1}\setminus\{g'_{t+1}\}}\eta_{g,t+1}\frac{\partial\big(\frac{F_{L_1,t+1}}{F_{L_i,t+1}}\big)}{\partial K_{t+1}}
\end{aligned}$$

There is then a clear parallel between debt policy and capital taxation. Government debt *could* be a tool for the control of automation as crowding out reduces the capital stock, but the automation term above is zero when the threshold rule FOC of equation (2.34) is satisfied. Therefore, debt is an unnecessary policy tool in the context of an optimally chosen threshold rule. Furthermore, as before, the final term is zero as relative wages are unaffected by capital with a CES production function. We will then have

$$\mu_t=\mu_{t+1}\big(1+F_{K,t+1}-\delta\big) \tag{B.15}$$

This equation equalizes the marginal benefit from revenue today with the marginal cost from tomorrow's debt burden and forgone capital.

If control of the extensive margin of automation is infeasible, we have (moving back the time index):

$$\mu_{t-1} + \frac{c}{K_t} \sum_{i \neq 1} \sum_{g \in \mathcal{G}_{i,t} \setminus \{g'_t\}} \frac{\eta_{g,t}}{\sigma} \frac{F_{L_1,t}}{F_{L_i,t}} \Big(\frac{\alpha_{L_i,t}(x_{i,t})^{\sigma-1}}{\Lambda_{L_i,t}} \hat{\alpha}_{i,t}^{-1} - \frac{\alpha_{L_1,t}(x_{1,t})^{\sigma-1}}{\Lambda_{L_1,t}} \hat{\alpha}_{1,t}^{-1} \Big)$$
$$= \mu_t \big(1 + F_{K,t} - \delta \big) \tag{B.16}$$

We again have the same marginal costs and benefits from equation (B.15) but with the addition of tomorrow's equity effect of automation on the benefit side. The equity effect is now a benefit because of crowding out.

For the choice of $\tau_{K,0}$, we have

$$\frac{\partial \mathcal{L}}{\partial \tau_{K,0}} = F_{K,0} \sum_{g \in \mathcal{G}_0 \setminus \mathcal{G}_0^B} \big(\mu_0 - \lambda_{g,0} \big) a_{g,0} \tag{B.17}$$

Therefore, welfare will be increasing in the initial tax on capital income so long as society remains unequal: where those with low private value of funds relative to the social value of funds own much of the wealth.

If we impose the natural upper bound for $\tau_{K,0}$ of one, it seems plausible that the above derivative will remain positive right up until capital income is fully expropriated.

B.3 Additional Information on Numerical Model

B.3.1 Construction of Occupational Data

The variables n_i, L_i, and W_i will all be taken from the ACS 2010 5yr on IPUMS where the occupation coding is occ1990. I will only include those ages 18 to 65 who worked last year and drop those who reported no earnings. I will use two sets of crosswalks, one from Webb (2019) and another from Acemoglu and Autor (2011), to go from occ1990 to a ten occupation partition used in Acemoglu and Autor (2011).

Denote by p_j^i the person weight of person j in occupation i. Likewise, w_j^i the number of weeks worked last year, u_j^i the usual hours worked in a week, and I_j^i the pre-tax wage and salary income. For the number of people in each occupation, I will compute

$$n_i = \frac{\sum_j p_j^i}{\sum_l \sum_j p_j^l/10} \tag{B.18}$$

That is, sum person weight across all the members of each occupation and divide by the average of summed person weight. This normalizes the average number of people in each occupation to one. Note that these will constitute welfare weights as well as determining the per-person allocation of the household.

For labor supply, I will compute

$$L_i = \sum_j p_j^i \cdot w_j^i \cdot u_j^i \tag{B.19}$$

That is, sum by occupation the person weight times the weeks worked last year times the usual hours worked in a week. This gives a measure of the total hours worked in an occupation.

For wages, I will compute

$$W_i = \frac{\sum_j p_j^i \frac{I_j^i}{w_j^i \cdot u_j^i}}{\sum_j p_j^i} \tag{B.20}$$

That is, an average of hourly 2010 wages, weighted by person weight.

Occupational income shares are computed as

$$S_{L_i} = (1 - S_K)\frac{W_i L_i}{\sum_j W_j L_j} \tag{B.21}$$

For exposure by occupation, I will use the Webb (2019) percentile data. This data is in occ1990dd levels, so I will take a weighted average to aggregate to the ten occupations I am using. The weighting will be the labor supply weight used for the construction of the exposure percentiles in Webb (2019).

B.3.2 Aggregate Elasticity of Substitution

Karabarbounis and Neiman (2014) estimate a long-run elasticity of substitution of 1.33 from the equation $\frac{\partial \ln S_K}{\partial \ln R} = 1-\Sigma$. To derive the long-run elasticity of substitution in the model, we will consider the following level-set:

$$\Bigg\{ \big(\{L_i\}, K, \{W_i\}, R, \xi\big) \in \mathbb{R}^{2N+3} \Bigg| \big(1-\tau_L\big) W_i = \hat{n}_i L_i^{1/\varepsilon_L} C_i^{\gamma},$$

$$\big(1-\tau_K\big) R = \xi K^{1/\varepsilon_K},$$

$$\frac{W_i}{R} = \Big(\frac{\Lambda_i}{\Lambda_K}\Big)^{\frac{1}{\sigma}} \Big(\frac{A_{L_i}}{A_K}\Big)^{\frac{\sigma-1}{\sigma}} \Big(\frac{K}{L_i}\Big)^{\frac{1}{\sigma}},$$

$$1 = \Big(\frac{R}{A_K}\Big)^{1-\sigma} \Lambda_K + \sum_i \Big(\frac{W_i}{A_{L_i}}\Big)^{1-\sigma} \Lambda_i \Bigg\}$$

Where $\Lambda_i \equiv \int_{\Omega_i} \alpha_i(x)^{\sigma-1} dx$ and $\hat{n}_i \equiv \frac{\big(1/n_i\big)^{1+1/\varepsilon_L}}{\big(1/n_i\big)^{1-\gamma}}$. We have $x_i = \Big(\frac{W_i}{R}\frac{A_K}{A_{L_i}}\Big)^{\frac{1}{\zeta_i}}$, so

$$\Lambda_i = \int_{x_i}^{\overline{x}_i} x^{\zeta_j \nu_j (\sigma-1)} dx = \Bigg[\overline{x}_i^{1+\zeta_i\nu_i(\sigma-1)} - \Big(\frac{W_i}{R}\frac{A_K}{A_{L_i}}\Big)^{\frac{1+\zeta_i\nu_i(\sigma-1)}{\zeta_i}}\Bigg] \Big/ \big(1+\zeta_i\nu_i(\sigma-1)\big)$$

and

$$\Lambda_K = \sum_j \int_0^{x_j} x^{\zeta_j(\nu_j-1)(\sigma-1)} dx = \sum_j \Big(\frac{W_j}{R}\frac{A_K}{A_{L_j}}\Big)^{\frac{1+\zeta_j(\nu_j-1)(\sigma-1)}{\zeta_j}} \Big/ \big(1+\zeta_j(\nu_j-1)(\sigma-1)\big)$$

Also, we have

$$C_i = (1-\tau_L) W_i L_i + \omega_i \Big((1-\tau_K) R K - \frac{\xi}{1+1/\varepsilon_K} K^{1+1/\varepsilon_K}\Big)$$

Note that the equations defining the above level set are just the FOCs of the labor and capital markets, relative wages, and the ideal price index as the numeraire.

We will use the Implicit Function Theorem to derive log equilibrium variables as functions of log capital costs. One can easily verify that there are $2N+2$ endogenous variables and $2N+2$ level-set equations. Thus, if we define $G : \mathbb{R}^{2N+3} \to \mathbb{R}^{2N+2}$ using the log versions of the level-sets above, we can derive

$$\begin{pmatrix} \frac{\partial \ln L_1}{\partial \ln \xi} \\ \vdots \\ \frac{\partial \ln K}{\partial \ln \xi} \\ \frac{\partial \ln W_1}{\partial \ln \xi} \\ \vdots \\ \frac{\partial \ln R}{\partial \ln \xi} \end{pmatrix} = - \begin{pmatrix} \frac{\partial G_1}{\partial \ln L_1} & \cdots & \frac{\partial G_1}{\partial \ln K} & \frac{\partial G_1}{\partial \ln W_1} & \cdots & \frac{\partial G_1}{\partial \ln R} \\ \vdots & \cdots & \cdots & \ddots & & \\ \frac{\partial G_{2N+2}}{\partial \ln L_1} & \cdots & \frac{\partial G_{2N+2}}{\partial \ln K} & \frac{\partial G_{2N+2}}{\partial \ln W_1} & \cdots & \frac{\partial G_{2N+2}}{\partial \ln R} \end{pmatrix}^{-1} \begin{pmatrix} \frac{\partial G_1}{\partial \ln \xi} \\ \vdots \\ \frac{\partial G_{2N+2}}{\partial \ln \xi} \end{pmatrix}$$

Note that, during calibration, we will have $x_i = \min \left\{ \overline{x}_i, \left(\frac{W_i}{R} \frac{A_K}{A_{L_i}} \right)^{\frac{1}{\zeta_i}} \right\}$. Thus, the derivatives of prices that are mediated through x_i will be zero when $\left(\frac{W_i}{R} \frac{A_K}{A_{L_i}} \right)^{\frac{1}{\zeta_i}} > \overline{x}_i$.

To relate the above objects to an estimate of the long-run elasticity of substitution, we will consider the ratio of log changes in the capital share and user cost of capital from a perturbation of the production cost of capital:

$$1 - \Sigma = \frac{\partial \ln S_K / \partial \ln \xi}{\partial \ln R / \partial \ln \xi}$$

B.3.3 Estimated Effect of Exposure on Log Wages

I will interpret regression estimates of the effect of exposure as entirely reflecting an automation effect induced by changes in the user cost of capital.

If we just consider the effect of prices on task allocations, we will have

$$\mathrm{d} \ln W_i = \frac{1}{\sigma} \frac{\partial \ln \Lambda_i}{\partial \ln W_i} \mathrm{d} \ln W_i + \frac{1}{\sigma} \frac{\partial \ln \Lambda_i}{\partial \ln R} \mathrm{d} \ln R$$

$$\Rightarrow \mathrm{d}\ln W_i = \frac{\frac{1}{\sigma}\frac{\partial \ln \Lambda_i}{\partial \ln R}}{1 - \frac{1}{\sigma}\frac{\partial \ln \Lambda_i}{\partial \ln W_i}} \mathrm{d}\ln R \equiv \lambda_i(e^r, e^s)\mathrm{d}\ln R$$

Now consider a regression for change in log wages of the form

$$100\Delta \ln W_{ij} = \alpha_j + \beta_r e_i^r + \beta_s e_i^s + \varepsilon_{ij}$$

Where i indexes occupations and j indexes industries. We can interpret the RHS of this regression equation as, in part, a linear approximation of the automation effect derived above. This gives us the calibration equation

$$\sum_i \frac{w_i}{\sum_i w_i} \frac{\partial \lambda_i}{\partial e^p} 100\Delta \ln R = \mathbb{E}\left(\frac{\partial \lambda_i}{\partial e^p}\right) 100\Delta \ln R \approx \beta_p \quad p \in \{r, s\} \tag{B.22}$$

where w_i are the occupational weights used in the regressions for $\beta_{r,s}$.

B.3.4 Additional Tables and Figures

Figure B·1: Wealth Shape Parameter

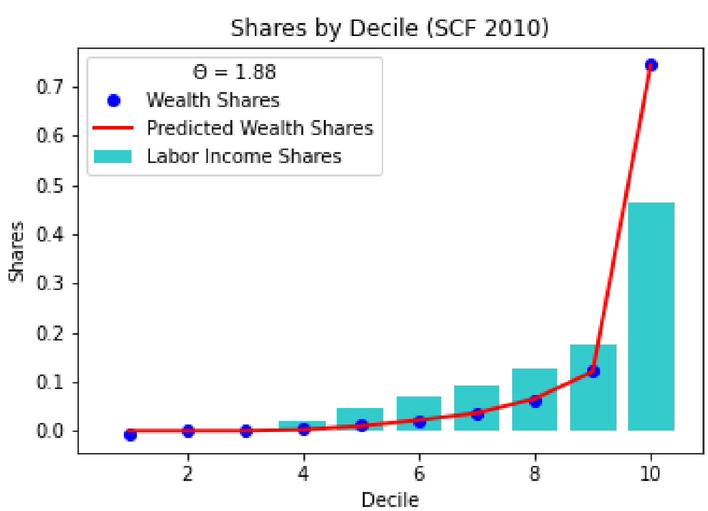

Figure B·2: Occupational Exposure

Table B.1: Sensitivity Analysis of Capital Parameters

	$\Theta = 1$ $\varepsilon_K = 2$ (1)	$\Theta = 1.88$ $\varepsilon_K = 2$ (2)	$\Theta = 1$ $\varepsilon_K = 1$ (3)	$\Theta = 1.88$ $\varepsilon_K = 1$ (4)
Policy Choices				
τ_L	0.343	0.348	0.323	0.327
τ_K	0.119	0.205	0.202	0.329
θ	0.03	0.034	0.029	0.034
UBI	0.089	0.135	0.108	0.168
Comparative Statics				
Ψ	1.007	1.023	1.010	1.046
$\%\Delta Y$	-0.0524	-0.1062	-0.0674	-0.1241
$\%\Delta L$	-0.0604	-0.0824	-0.0555	-0.0851

Notes: This table presents the policy choices from local maxima of the Planner's problem. Comparative statics are in reference to the 2010 American status quo, which is assumed to have no threshold rule or lump-sum transfer. Note that the status quo changes alongside the parameters governing the wealth distribution and elasticity of capital supply.

CHAPTER THREE APPENDIX

C.1 Theory Derivations and Proofs

C.1.1 Proof of Lemma 3.1

Economy with Heterogeneous Preferences. Consider a problem of households living in location j at time t, deciding where to migrate. We index each individual by $\omega \in [0,1]$, and $\{\epsilon_{jk}(\omega)\}_k$ denote the preference draw of individual ω. Each individual solves the following problem:

$$\begin{aligned} v_{jt}(\omega) = u(C_{jt}, a_{jt}) + \max_{\{\mathbb{I}_{jkt}(\omega)\}_k} \sum_k \mathbb{I}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] \\ \text{s.t.} \quad \sum_k \mathbb{I}_{jkt}(\omega) = 1, \end{aligned} \tag{C.1}$$

where $\mathbb{I}_{jkt}(\omega) \in \{0,1\}$ is an indicator function for location choice of individual ω. The fraction of individuals migrating from j to k is given by

$$\mu_{jkt} = \int_0^1 \mathbb{I}_{jkt}(\omega) d\omega. \tag{C.2}$$

Economy with Representative Agent. Define the following function:

$$\begin{aligned} \psi_{jkt}(\{\mu_{jkt}\}_k) = -\max_{\{\mathbb{I}_{jkt}(\omega)\}_{\omega,k}} \int_0^1 \sum_k (\chi_{jkt+1} + \epsilon_{jk}(\omega))\mathbb{I}_{jkt}(\omega) d\omega \\ \text{s.t.} \quad \int_0^1 \mathbb{I}_{jkt}(\omega) d\omega = \mu_{jkt} \\ \sum_k \mathbb{I}_{jkt}(\omega) = 1. \end{aligned} \tag{C.3}$$

The representative agent solves

$$V_{jt} = u(C_{jt}, a_{jt}) + \max_{\{\mu_{jkt}\}_k} \sum_k \mu_{jkt} \beta V_{kt+1} - \psi_{jkt}(\{\mu_{jkt}\}_k) \tag{C.4}$$

Equivalence Result. We formally restate the equivalence result of Lemma 3.1 as follows.

Lemma C.1. *Suppose* $\{\mathbb{I}_{jkt}(\omega)\}_k$ *solves* (C.1) *for all* ω. *Then,* $\{\mu_{jkt}\}_k$*, given by* (C.2)*, solves* (C.4)*. Conversely, suppose* $\{\mu_{jkt}\}_k$ *solves* (C.4)*. Then* $\{\mathbb{I}_{jkt}(\omega)\}_{\omega,k}$*, given by the solution to* (C.3) *associated with* $\{\mu_{jkt}\}_k$*, solves* (C.1) *for almost all* ω*. Moreover, the expected utility in the economy with heterogeneous preferences equals the utility of the representative agent:*

$$\int_0^1 v_{jt}(\omega) d\omega = V_{jt}$$

Proof. We prove the first part. Suppose to the contrary, there exists $\{\tilde{\mu}_{jkt}\}_k$ such that

$$\sum_k \tilde{\mu}_{jkt} \beta V_{kt+1} - \psi_{jkt}(\{\tilde{\mu}_{jkt}\}_k) > \sum_k \mu_{jkt} V_{kt+1} - \psi_{jk}(\{\mu_{jkt}\}_k). \tag{C.5}$$

Let $\{\tilde{\mathbb{I}}_{jkt}(\omega)\}_{\omega,k}$ denote the solution to (C.3) associated with $\{\tilde{\mu}_{jkt}\}_k$. Plugging into (C.5),

$$\int_0^1 \sum_k \tilde{\mathbb{I}}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega) d\omega\right.$$
$$> \int_0^1 \sum_k \mathbb{I}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] d\omega, \tag{C.6}$$

where $\sum_k \tilde{\mathbb{I}}_{jkt}(\omega) = 1$ and $\sum_k \mathbb{I}_{jkt}(\omega) = 1$ for all ω. However, this is a contradiction because by our presumption, for any ω,

$$\sum_k \mathbb{I}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] \geq \sum_k \tilde{\mathbb{I}}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right]$$

for all $\tilde{\mathbb{I}}_{jkt}(\omega)$, which would imply

$$\int_0^1 \sum_k \tilde{\mathbb{I}}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega) d\omega \right.$$
$$\leq \int_0^1 \sum_k \mathbb{I}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] d\omega. \quad \text{(C.7)}$$

Now we prove the converse. Suppose to the contrary, there exists $\{\tilde{\mathbb{I}}_{jkt}(\omega)\}_k$ such that

$$\sum_k \tilde{\mathbb{I}}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] > \sum_k \mathbb{I}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] \quad \text{(C.8)}$$

and $\sum_k \tilde{\mathbb{I}}_{jkt}(\omega) = 1$ hold for all $\omega \in \Omega$, where $\Omega \subset [0, 1]$ and $|\Omega| > 0$. Define

$$\tilde{\mu}_{jkt} = \int_0^1 \tilde{\mathbb{I}}_{jkt}(\omega) d\omega. \quad \text{(C.9)}$$

Then

$$\begin{aligned} \sum_k \mu_{jkt} \beta V_{kt+1} - \psi_{jkt}(\{\mu_{jkt}\}_k) &= \int_0^1 \sum_k \mathbb{I}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] d\omega \\ &< \int_0^1 \sum_k \tilde{\mathbb{I}}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] d\omega \\ &\leq \sum_k \tilde{\mu}_{jkt} \beta V_{kt+1} - \psi_{jkt}(\{\tilde{\mu}_{jkt}\}_k). \end{aligned}$$

This is a contradiction that $\{\mu_{jkt}\}_k$ is a solution to (C.4).

We need to show that the expected utility coincides with each other in the two economies. This immediately follows given the above result. Let $\{\mathbb{I}_{jkt}(\omega)\}_{\omega,k}$ be the solution to (C.1) for all ω, and let $\{\mu_{jkt}\}_k$ denote the solution to (C.4). Then

$$\int_0^1 \sum_k \mathbb{I}_{jkt}(\omega) \left[\beta V_{kt+1} + \chi_{jkt+1} + \epsilon_{jk}(\omega)\right] d\omega = \sum_k \mu_{jkt} \beta V_{kt+1} - \psi_{jkt}(\{\mu_{jkt}\}_k). \quad \text{(C.10)}$$

□

C.1.2 Proof of Proposition 3.1

Instead of solving the original planning problem, we consider the following relaxed version of the problem:

$$\max_{\left\{\{q_{kjt},\mu_{jkt}\},L_{jt+1},V_{jt}\right\}_{t\geq 0}} \sum_i \varkappa_i L_{i0} V_{i0} \tag{C.11}$$

$$\text{s.t. } Z_{it}(L_{it})L_{it} = \sum_j \kappa_{ijt} q_{ijt} L_{jt}$$

$$L_{jt+1} = \sum_i L_{it}\mu_{ijt}$$

$$V_{jt} = u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1}$$

$$\{\mu_{jkt}\} \in \arg\max_{\{\mu_{jkt}\}:\sum_k \mu_{jkt}=1} u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1}.$$

Here, we dropped constraints (3.1), (3.3), (3.2), (3.8), (3.11), and (3.12). We later show that the solution to the above problem can be implemented as an equilibrium that satisfies all these dropped constraints. Let $\tilde{\Lambda}_{it}$ denote the Lagrange multiplier on the resource constraint. Let $\hat{\mu}_{jkt}(\{V_{it+1}\})$ denote the solution to the household's migration problem. Then the dual of the above problem can be written as

$$\max_{\left\{\{q_{kjt},\mu_{jkt}\},L_{jt+1},V_{jt}\right\}_{t\geq 0}} \sum_{t=0}^{\infty} \prod_{s\leq t} \frac{1}{R_s} \sum_i \Lambda_{it}\Big(Z_{it}(L_{it})L_{it} - \sum_j \kappa_{ijt} q_{ijt} L_{jt}\Big) \tag{C.12}$$

$$\text{s.t.} \quad V_{j0} \geq \underline{V}_{j0}$$

$$L_{jt+1} = \sum_i L_{it}\mu_{ijt}$$

$$V_{jt} = u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1}$$

$$\mu_{jkt} = \hat{\mu}_{jkt}(\{V_{it+1}\}),$$

where we consider the normalized multiplier $\Lambda_{it} \equiv \tilde{\Lambda}_{it}/\tilde{\Lambda}_{1t}$ for all t and denote $R_t \equiv \tilde{\Lambda}_{1t-1}/\tilde{\Lambda}_{1t}$.[1] Note that this problem has been normalized by $1/\tilde{\Lambda}_{1,0}$.

We solve the recursive form of the above problem. The dual of the optimal transfer problem can be recursively rewritten as

$$\begin{aligned}
\mathcal{S}_t\big(\{L_{jt}, V_{jt}\}\big) = & \max_{\{\{q_{ijt},\mu_{ijt}\},L_{jt+1},V_{jt+1}\}} \sum_i \Lambda_{it}\Big(Z_{it}(L_{it})L_{it} - \sum_j \kappa_{ijt}q_{ijt}L_{jt}\Big) \qquad \text{(C.13)} \\
& + \frac{1}{R_{t+1}}\mathcal{S}_{t+1}\big(\{L_{jt+1}, V_{jt+1}\}\big) \\
\text{s.t.} \quad L_{jt+1} = & \sum_i L_{it}\mu_{ijt} : \quad \tilde{\Delta}_{jt+1} \\
V_{jt} \leq & \; u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta\sum_k \mu_{jkt}V_{kt+1} : \quad \Xi_{jt} \\
\mu_{ijt} = & \; \hat{\mu}_{ijt}\big(\{V_{kt+1}\}\big) : \quad \Omega_{ijt},
\end{aligned}$$

where next to each constraint is the corresponding multiplier.

First, taking the FOC with respect to q_{ijt}, we have

$$\Xi_{jt}\partial_c u_{jt}\frac{\partial c_{jt}}{\partial q_{ijt}} = \Lambda_{it}\kappa_{ijt}L_{jt}. \tag{C.14}$$

Denote by $\tilde{P}_{jt}$ the Planner's ideal price index, defined with respect to shadow prices, for consumption in location j, [2] then we have

$$\Xi_{jt} = \frac{\tilde{P}_{jt}L_{jt}}{\partial_c u_{jt}}. \tag{C.15}$$

Next, taking the FOC with respect to L_{jt+1}, we have

$$\frac{1}{R_{t+1}}\frac{\partial S_{t+1}}{\partial L_{jt+1}} = \tilde{\Delta}_{jt+1} \equiv \frac{1}{R_{t+1}}\Delta_{jt+1}. \tag{C.16}$$

[1] We also take the normalization that $R_0 = 1$.

[2] That is, $\tilde{P}_{jt} = \min_{\{q_{ijt}\}}\left[\sum_i \Lambda_{it}\kappa_{ijt}q_{ijt}\right]$ *s.t.* $c_{jt}\big(\{q_{ijt}\}\big) \geq 1$.

Taking the FOC with respect to V_{jt+1}, we have

$$\sum_i \sum_k \Omega_{ikt} \frac{\partial \hat{\mu}_{ikt}}{\partial V_{jt+1}} + \beta \sum_i \Xi_{it} \mu_{ijt} = -\frac{1}{R_{t+1}} \frac{\partial S_{t+1}}{\partial V_{jt+1}}. \tag{C.17}$$

Taking the FOC with respect to μ_{ijt}, we have

$$\begin{aligned} \Omega_{ijt} &= \frac{1}{R_{t+1}} \Delta_{jt+1} L_{it} + \Xi_{it} \Big(\beta V_{jt+1} - \frac{\partial \psi_{it}}{\partial \mu_{ijt}} \Big) \\ &= \frac{1}{R_{t+1}} \Delta_{jt+1} L_{it} + \Xi_{it} \zeta_{it}, \end{aligned} \tag{C.18}$$

where the second line comes from household optimization Equation (C.26).

The envelope conditions are

$$\frac{\partial \mathcal{S}_t}{\partial L_{jt}} = \Lambda_{jt} \big(Z_{jt} + \frac{dZ_{jt}}{dL_{jt}} L_{jt} \big) - \sum_i \Lambda_{it} \kappa_{ijt} q_{ijt} + \Xi_{jt} \partial_a u_{jt} \frac{dA_{jt}}{dL_{jt}} + \frac{1}{R_{t+1}} \sum_k \mu_{jkt} \Delta_{kt+1} \tag{C.19}$$

$$\frac{\partial \mathcal{S}_t}{\partial V_{jt}} = -\Xi_{jt}. \tag{C.20}$$

Combining Equations (C.15), (C.17), (C.18), and (C.20), as well as noting that $\sum_k \frac{\partial \hat{\mu}_{ikt}}{\partial V_{jt}} = 0$, we have

$$\begin{aligned} \frac{1}{R_{t+1}} \sum_i \sum_k \Delta_{kt+1} L_{it} \frac{\partial \hat{\mu}_{ikt}}{\partial V_{jt+1}} + \beta \sum_i \frac{P_{it} L_{it}}{\partial_c u_{it}} \mu_{ijt} &= \frac{1}{R_{t+1}} \frac{P_{jt+1} L_{jt+1}}{\partial_c u_{jt+1}} \\ \Leftrightarrow \quad \sum_i \sum_k \Delta_{kt+1} L_{it} \frac{\partial \hat{\mu}_{ikt}}{\partial V_{jt+1}} &= \sum_i L_{it} \mu_{ijt} \Big(\frac{P_{jt+1}}{\partial_c u_{jt+1}} - \frac{\beta R_{t+1} P_{it}}{\partial_c u_{it}} \Big) \\ \Leftrightarrow \quad -\sum_i L_{it} \sum_k \Delta_{kt+1} \frac{1}{P_{jt+1}} \frac{\partial \hat{\mu}_{ikt}}{\partial c_{jt+1}} &= \sum_i L_{it} \mu_{ijt} \Bigg(\frac{\beta R_{t+1} \partial_c u_{jt+1} / P_{jt+1}}{\partial_c u_{it} / P_{it}} - 1 \Bigg), \end{aligned} \tag{C.21}$$

where $\frac{\partial \hat{\mu}_{ikt}}{\partial c_{jt+1}} \equiv \frac{\partial \hat{\mu}_{ikt}}{\partial V_{jt+1}} \partial_c u_{jt}$. This gives us Equation (3.14) of Proposition 3.1. We remove the hat notation for movement flows in Proposition 3.1 to simplify notation.

Now we show that the above solution can be supported as an equilibrium with transfers $\tau_{jt} = P_{jt} c_{jt} - w_{jt} - s_{jt}$. Supporting wages and prices are set according to

$w_{it} = \Lambda_{it} Z_{it}$ and $p_{ijt} = \kappa_{ijt} w_{it}/Z_{it}$, so the producer price of location 1, w_{1t}/Z_{1t}, is set as the numeraire in each period. Thus, the Planner and household ideal price indices coincide, and since (C.14) is equivalent to the household's consumption FOC, the household's equilibrium consumption allocation coincides with that of the Planner. Therefore, (3.1), (3.2), and (3.8) are satisfied. Since all of the resource constraints and household budget constraints are satisfied by construction, the government budget balances by Walras' law, implying (3.11) and (3.12) are satisfied. We can then interpret R_t as the Planner's marginal rate of substitution with respect to the numeraire good across time.

To derive the formula for Δ_{jt}, simply combine Equations (C.15), (C.16), and (C.19) and evaluate at equilibrium prices to achieve

$$\Delta_{jt} = w_{jt} - P_{jt} c_{jt} + w_{jt} \gamma^z_{jt} + \frac{\partial_a u_{jt}}{\partial_c u_{jt}} P_{jt} A_{jt} \gamma^a_{jt} + \frac{1}{R_{t+1}} \sum_k \mu_{jkt} \Delta_{kt+1}. \tag{C.22}$$

□

C.1.3 First-Best Allocation

In this section, we characterize the first-best planning problem, in which the Planner can directly control consumption and migration flows. For the sake of exposition, we will consider the case of a utilitarian Planner. Utilitarianism allows us to economize on notation while still illustrating the basic features of the first-best allocation.[3] Unlike in the planning problem of Section 3.3, direct control over the allocation allows the Planner to separate migration flows from the consumption offered in each location. While the information requirements of the first-best are not realistic, it provides us a useful benchmark with an important lesson: when the Planner can choose con-

[3]In the more general case, where we sketch the entire initial Pareto frontier, the Planner may allocate different consumption and migration choices on the basis of initial location to otherwise identical members of the population. Thus, this requires the Planner to keep track of everyone's initial location, which substantially complicates the notation of the planning problem.

sumption separately from migration flows, they equalize the marginal resource cost of consumption across locations. When this separation is impossible, as in Section 3.3, this introduces a trade-off between equalizing marginal utility and efficient migration.

Definition C.1 (Utilitarian First-Best Planning Problem)**.** *Starting with initial state variables $\{L_{j0}\}$, the Planner solves*

$$\max_{\{\{q_{kjt},\mu_{jkt}\},L_{jt+1},V_{jt}\}_{t\geq 0}} \sum_j L_{j0}V_{j0} \tag{C.23}$$

$$\begin{aligned} s.t. \qquad & \sum_j \kappa_{ijt}q_{ijt}L_{jt} = Z_{it}(L_{it})L_{it} \\ & L_{jt} = \sum_i L_{it-1}\mu_{ijt-1} \\ & V_{jt} = u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt}V_{kt+1} \\ & \sum_j \mu_{ijt} = 1. \end{aligned}$$

The solution to this problem has the following characteristics.

Proposition C.1. *The utilitarian first-best allocation satisfies:*

(i) Perfect insurance. That is, the marginal utility of consumption in time t is set according to

$$\frac{\partial_c u_{jt}}{\tilde{P}_{jt}} = \frac{\partial_c u_{kt}}{\tilde{P}_{kt}} \quad \forall j,k, \tag{C.24}$$

where $\tilde{P}_{jt}$ is the Planner's ideal shadow price index for consumption in location j.

(ii) The movement flow condition for sending location i is set according to

$$\tilde{\beta}_{t+1}\Big(V_{jt+1} - \frac{1}{\beta}\frac{\partial\psi_{it}}{\partial\mu_{ijt}}\Big) + \Delta_{jt+1} = \tilde{\beta}_{t+1}\Big(V_{kt+1} - \frac{1}{\beta}\frac{\partial\psi_{it}}{\partial\mu_{ikt}}\Big) + \Delta_{kt+1} \quad \forall j,k, \tag{C.25}$$

where $\tilde{\beta}_{t+1}$ is a discount term defined below, and Δ_{jt} is given recursively by Equation (3.13)*, defined with respect to the Planner's shadow prices.*

The proof of Proposition C.1 is below, but we have established the most important feature of the first-best allocation: perfect insurance. The Planner's ability to smooth consumption comes from the ability to control migration flows separately from consumption. Once migration flows have taken place, the distribution of consumption across locations is simply a static problem to maximize a concave objective. In that setting, the objective's margins will be equalized up to real differences in marginal cost, as a failure to do so is akin to misallocation in the social objective.

For migration flows, the Planner equalizes, for a given sending location, the marginal benefit of additional population across prospective locations. This is a similar condition to that of the household (C.26) where the marginal mover is indifferent between prospective locations. However, the Planner's condition includes the marginal fiscal externality of additional population Δ_{jt}. Thus, the household decision has a wedge in that they do not internalize the fiscal externality of their location choices.

It is also clear that this first-best allocation cannot be generally achieved under the spatial transfer policy. To see this, note that the first-order condition of the migration decision (3.7), which the Planning problem in Definition 3.2 takes as an incentive compatibility constraint, is given by

$$\beta V_{kt+1} - \frac{\partial \psi_{it}}{\partial \mu_{ikt}} = \beta V_{jt+1} - \frac{\partial \psi_{it}}{\partial \mu_{ijt}}. \tag{C.26}$$

This equation is not compatible with Equation (C.25) unless $\Delta_{jt} = \Delta_{kt}$ for all j, k.

Finally, a couple of comments are in order. In principle, the first-best choice set includes the possibility of consumption and movement flows that depend on the entire history of idiosyncratic shocks for each member of the household. However, given our assumption that preference shocks are separable from consumption and iid, consumption only needs to condition on contemporaneous location, and migration decisions only need to condition on contemporaneous preference shocks. Next, because the first-best describes an allocation, not an equilibrium, the conditions of Proposition C.1 are written in terms of the Planner's shadow prices. However, if the first-best allocation were to be implemented as an equilibrium, these shadow prices would suffice as supporting prices, just as in Proposition 3.1.

Proof. We follow the same strategy as in Appendix C.1.2 and solve the Planner's recursive dual problem, except now we can drop the incentive compatibility constraint and replace it with an accounting constraint. The Planner solves

$$\begin{aligned}
\mathcal{S}_t\big(\{L_{jt}, V_{jt}\}\big) = & \max_{\{\{q_{ijt},\mu_{ijt}\},L_{jt+1},V_{jt+1}\}} \sum_i \Lambda_{it}\Big(Z_{it}(L_{it})L_{it} - \sum_j \kappa_{ijt} q_{ijt} L_{jt}\Big) \\
& + \frac{1}{R_{t+1}} \mathcal{S}_{t+1}\big(\{L_{jt+1}, V_{jt+1}\}\big) \\
\text{s.t.} \quad L_{jt+1} = & \sum_i L_{it}\mu_{ijt} : \quad \tilde{\Delta}_{jt+1} \\
V_{jt} \leq & \; u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1} : \quad \Xi_{jt} \\
& \sum_j \mu_{ijt} = 1 : \quad \Omega_{it}
\end{aligned} \tag{C.27}$$

where next to each constraint is the corresponding multiplier.

As before, taking the FOCs with respect to q_{ijt}, L_{jt+1}, V_{jt+1}, and μ_{ijt}, we have

$$\begin{aligned}
\Xi_{jt} &= \frac{\tilde{P}_{jt} L_{jt}}{\partial_c u_{jt}} \\
\frac{1}{R_{t+1}} \frac{\partial S_{t+1}}{\partial L_{jt+1}} &= \tilde{\Delta}_{jt+1} \equiv \frac{1}{R_{t+1}} \Delta_{jt+1} \\
\beta \sum_i \Xi_{it}\mu_{ijt} &= -\frac{1}{R_{t+1}} \frac{\partial S_{t+1}}{\partial V_{jt+1}} \\
\Omega_{it} &= \frac{1}{R_{t+1}} \Delta_{jt+1} L_{it} + \Xi_{it}\big(\beta V_{jt+1} - \frac{\partial \psi_{it}}{\partial \mu_{ijt}}\big),
\end{aligned}$$

where $\tilde{P}_{jt}$ is the Planner's ideal shadow price index for consumption of location j. The envelope conditions again follow Equations (C.19) and (C.20).

Combining the FOC and envelope condition for V_{jt+1} with the law of motion for population, as well as iterating backward, gives us

$$\Xi_{jt} = \Big[\prod_{s\leq t} R_s\Big]\beta^t \frac{L_{jt}}{\tilde{\Lambda}_{10}}. \tag{C.28}$$

Here we have used the result from the primal problem that $\Xi_{j0} = L_{j0}/\tilde{\Lambda}_{10}$, where the factor of proportionality $\tilde{\Lambda}_{10}$ comes from the normalization applied to the dual problem (C.12). Plugging this into the FOC for consumption gives us Equation (C.24) of Proposition C.1.

If we define $\tilde{\beta}_{t+1} \equiv \big[\prod_{s\leq t+1} R_s\big]\beta^{t+1}/\tilde{\Lambda}_{10}$, then plugging equation (C.28) into the FOC for migration flows gives us Equation (C.25) of Proposition C.1.

□

C.1.4 Complete Financial Markets and Proof of Corollary 3.1

In this section, we study the case where the regional risk-sharing mechanism is complete. There exists a representative, competitive financial intermediary that offers state-contingent contracts to households. Crucially, the financial intermediary has a perfect enforcement technology and can observe the entire history of households' idiosyncratic preference shocks, so contracts are complete. These contracts specify a sequence of consumption and migration decisions which conditions on households' history of preference shocks. In exchange, households are guaranteed an expected utility level of $\underline{V}$.

We assume that contracting takes place before the realization of any preference shocks and before households learn their initial location. There is a known distribution of initial locations, but households do not know their own initial location. This guarantees our assumption that payments $\{s_{jt}\}$ do not vary within location as all households undertake contracts from the same position. Finally, given our assumption that preference shocks are separable from consumption and iid, consumption only needs to condition on contemporaneous location, and migration decisions only need to condition on contemporaneous preference shocks. The financial intermediary then faces the same movement cost function as the representative household from Lemma 3.1 as they select movers from the population so as to minimize the overall opportunity cost of moving for a given array of movement flows.

The optimal contract solves

$$\max_{\{\{\mu_{ijt}\},c_{jt},L_{jt+1},V_{jt}\}_{t\geq 0}} \sum_{t=0}^{\infty} \prod_{s\leq t} \frac{1}{R_s} \sum_j \big(w_{jt} + \tau_{jt} - P_{jt}c_{jt}\big) L_{jt} \tag{C.29}$$

$$\begin{aligned}
\text{s.t.} \quad & \sum_j L_{j0} V_{j0} \geq \underline{V} : \quad \hat{\Psi} \\
& L_{jt+1} = \sum_i L_{it}\mu_{ijt} : \quad \Big[\prod_{s\leq t} \frac{1}{R_s}\Big] \hat{\Delta}_{jt+1} \\
& V_{jt} = u(c_{jt}, a_{jt}) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta \sum_k \mu_{jkt} V_{kt+1} : \quad \Big[\prod_{s\leq t} \frac{1}{R_s}\Big] \hat{\Xi}_{jt} \\
& \sum_j \mu_{ijt} = 1 : \quad \Big[\prod_{s\leq t} \frac{1}{R_s}\Big] \hat{\Omega}_{it},
\end{aligned}$$

where next to each constraint is the corresponding multiplier with any normalization factor in square brackets. Hat notation signifies that the multiplier corresponds to the financial intermediary's problem, as opposed to the Planner's problem.

First, taking the FOC with respect to c_{jt}, we have

$$P_{jt} L_{jt} = \hat{\Xi}_{jt} \partial_c u_{jt}. \tag{C.30}$$

Next, taking the FOC with respect to L_{jt}, we have

$$\hat{\Delta}_{jt} = w_{jt} + \tau_{jt} - P_{jt}c_{jt} + \frac{1}{R_{t+1}} \sum_k \mu_{jkt} \hat{\Delta}_{kt+1}. \tag{C.31}$$

Taking the FOC with respect to V_{jt}, we have

$$\frac{1}{R_t} \hat{\Xi}_{jt} = \beta \sum_i \hat{\Xi}_{it-1} \mu_{ijt-1} \tag{C.32}$$

$$\hat{\Xi}_{j0} = L_{j0} \hat{\Psi}. \tag{C.33}$$

Taking the FOC with respect to μ_{ijt}, we have

$$\hat{\Omega}_{it} = \frac{1}{R_{t+1}}\hat{\Delta}_{jt+1}L_{it} + \hat{\Xi}_{it}\Big(\beta V_{jt+1} - \frac{\partial \psi_{it}}{\partial \mu_{ijt}}\Big). \tag{C.34}$$

Combining Equations (C.32) and (C.33) with the law of motion for population gives us

$$\hat{\Xi}_{jt} = \Big[\prod_{s\leq t} R_s\Big]\beta^t\hat{\Psi}L_{jt}. \tag{C.35}$$

Plugging Equation (C.35) into Equation (C.30) gives us complete risk sharing:

$$\frac{\partial_c u_{jt}}{P_{jt}} = \frac{1}{\big[\prod_{s\leq t} R_s\big]\beta^t}\frac{1}{\hat{\Psi}}. \tag{C.36}$$

To determine optimal transfers in the context of complete markets, the Planner must satisfy an incentive compatibility constraint that reflects complete risk sharing. However, we will solve a relaxed problem that drops this IC constraint and show that the Planner's solution still satisfies Equation (C.36).

As in Appendix C.1.2, the Planner solves the problem (C.81). The FOCs for q_{ijt}, L_{jt+1}, and V_{jt+1} again follow Equations (C.15), (C.14), (C.16), and (C.17). The envelope conditions again follow Equations (C.19) and (C.20). Comparing the Planner's consumption FOC (C.14) and the financial intermediary's consumption FOC (C.30), we have that $\Xi_{jt} = \hat{\Xi}_{jt}$ if supporting prices are set such that private price indices equal the Planner's shadow prices.

We will guess that $\tau_{jt} = w_{jt}\gamma^z_{jt} + \frac{\partial_a u_{jt}}{\partial_c u_{jt}}P_{jt}A_{jt}\gamma^a_{jt} - K_t$, as in Corollary 3.1, and verify that this provides a solution to the Planner's problem.[4] Comparing the Planner's population FOC (C.16) and the financial intermediary's population FOC (C.31), we have that $\hat{\Delta}_{jt} = \Delta_{jt} - \mathcal{K}_t$, where $\mathcal{K}_t \equiv K_t + \sum_{\hat{t}\geq 1}\prod_{s=1}^{\hat{t}}\frac{1}{R_{t+s}}K_{t+\hat{t}}$ is the present value

[4]By government budget balance (3.11), the lump-sum tax is equal to the average externality: $K_t = \sum_k L_{kt}\big(w_{kt}\gamma^z_{kt} + \frac{\partial_a u_{kt}}{\partial_c u_{kt}}P_{kt}A_{kt}\gamma^a_{kt}\big)$.

of lump-sum taxes. Taking the FOC with respect to μ_{ijt}, we have

$$\begin{aligned}\Omega_{ijt} &= \frac{1}{R_{t+1}}\Delta_{jt+1}L_{it} + \Xi_{it}\Big(\beta V_{jt+1} - \frac{\partial\psi_{it}}{\partial\mu_{ijt}}\Big) \qquad \text{(C.37)}\\ &= \frac{1}{R_{t+1}}\Delta_{jt+1}L_{it} + \hat{\Xi}_{it}\Big(\beta V_{jt+1} - \frac{\partial\psi_{it}}{\partial\mu_{ijt}}\Big)\\ &= \hat{\Omega}_{it} + \frac{1}{R_{t+1}}\mathcal{K}_{t+1}L_{it}, \qquad \text{(C.38)}\end{aligned}$$

where the last line comes from the financial intermediary's optimization Equation (C.34).

Combining Equations (C.15), (C.17), (C.37), and (C.20), as well as noting that $\sum_k \frac{\partial\hat{\mu}_{ikt}}{\partial V_{jt}} = 0$, we have

$$\sum_i L_{it}\mu_{ijt}\Big(\frac{\beta R_{t+1}\partial_c u_{jt+1}/P_{jt+1}}{\partial_c u_{it}/P_{it}} - 1\Big) = 0, \qquad \text{(C.39)}$$

which will be satisfied if the Planner selects consumption according to complete risk sharing (C.36). Hence, the requirement that consumption provides complete risk sharing is slack, so our relaxation was valid.

Therefore, the transfers described in Corollary 3.1 solve the Planner's problem with supporting wages and prices set according to $w_{it} = \Lambda_{it}Z_{it}$ and $p_{ijt} = \kappa_{ijt}w_{it}/Z_{it}$. Furthermore, this transfer policy implements the utilitarian first-best allocation as a competitive equilibrium. That is, the utilitarian Planner's first-best optimality conditions (C.24) and (C.25) coincide with those of the financial intermediary when transfers correct externalities.

Risk Neutrality and Equalization of Consumption Prices An alternative setup for the complete market is to assume risk neutrality (constant marginal utility $\partial_c u_{jt}$ across locations and time) and equalized consumption prices across locations ($P_{it} = P_{jt}$ for all $i, j \in J$). As noted in the main text, linear utility functions and equalized price indices across locations imply that the RHS of Equation (3.14) is equal to zero. Policy should then equalize the fiscal externality across locations as this guarantees that the LHS of Equation (3.14) is equal to zero.

Equalization of the fiscal externality at all points in time implies that

$$\tau_{jt} = \tau_{1t} + w_{jt}\gamma^z_{jt} - w_{1t}\gamma^z_{1t} + \frac{\partial_a u_{jt}}{\partial_c u_{jt}} P_{jt} A_{jt} \gamma^a_{jt} - \frac{\partial_a u_{1t}}{\partial_c u_{1t}} P_{1t} A_{1t} \gamma^a_{1t}. \tag{C.40}$$

Plugging this into the government budget constraint (3.11), we have

$$\tau_{jt} = w_{jt}\gamma^z_{jt} + \frac{\partial_a u_{jt}}{\partial_c u_{jt}} P_{jt} A_{jt} \gamma^a_{jt} - \sum_k L_{kt}\big(w_{kt}\gamma^z_{kt} + \frac{\partial_a u_{kt}}{\partial_c u_{kt}} P_{kt} A_{kt} \gamma^a_{kt}\big), \tag{C.41}$$

which gives us Equation (3.16) of Corollary 3.1. One can see that the budget-balancing transfer K_t is equal to the average externality in each period.

To see that this transfer policy implements the utilitarian first-best allocation as a competitive equilibrium, consider the utilitarian Planner's first-best optimality conditions (C.24) and (C.25). Equation (C.24) is satisfied by assumption. Plugging the condition for household movement flows (C.26) into (C.25), we have equalization of the fiscal externality $\Delta_{jt} = \Delta_t \; \forall t$. This is the exact condition the transfers of Equation (C.41) are designed to achieve.

□

C.2 Theory Extensions

C.2.1 Dynamic and Cross-Regional Agglomeration Externalities

We follow the same strategy as in Appendix C.1.2 and solve a relaxation of the Planner's recursive dual problem. Again, one can back out supporting prices that implement the solution as an equilibrium. The Planner solves

$$
\begin{aligned}
\mathcal{S}_t\big(\{L_{jt}, V_{jt}, \bar{Z}_{jt}\}\big) = & \max_{\{\{q_{ijt},\mu_{ijt}\},L_{jt+1},V_{jt+1}\}} \sum_i \Lambda_{it}\Big(Z_{it}L_{it} - \sum_j \kappa_{ijt}q_{ijt}L_{jt}\Big) \\
& + \frac{1}{R_{t+1}}\mathcal{S}_{t+1}\big(\{L_{jt+1}, V_{jt+1}, \bar{Z}_{jt+1}\}\big) \qquad \text{(C.42)} \\
\text{s.t.} \quad L_{jt+1} = & \sum_i L_{it}\mu_{ijt} : \quad \tilde{\Delta}_{jt+1} \\
V_{jt} \leq & \; u\big(c_{jt}(\{q_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu_{jkt}\}\big) + \beta\sum_k \mu_{jkt}V_{kt+1} : \quad \Xi_{jt} \\
\mu_{ijt} = & \; \hat{\mu}_{ijt}\big(\{V_{kt+1}\}\big) : \quad \Omega_{ijt},
\end{aligned}
$$

where labor productivity $Z_{jt} = g_{jt}\big(\bar{Z}_{jt}, \{L_{kt}\}\big)$ and fundamental productivity $\bar{Z}_{jt} = h_{jt}\big(\{\bar{Z}_{kt-1}, L_{kt-1}\}\big)$ allow for general externalities. Again, next to each constraint is the corresponding multiplier.

As before, taking the FOCs with respect to q_{ijt}, L_{jt+1}, V_{jt+1}, and μ_{ijt}, we have

$$
\begin{aligned}
\Xi_{jt} &= \frac{\tilde{P}_{jt}L_{jt}}{\partial_c u_{jt}} \\
\frac{1}{R_{t+1}}\frac{\partial S_{t+1}}{\partial L_{jt+1}} &= \tilde{\Delta}_{jt+1} \equiv \frac{1}{R_{t+1}}\Delta_{jt+1} \\
\sum_i\sum_k \Omega_{ikt}\frac{\partial\hat{\mu}_{ikt}}{\partial V_{jt+1}} + \beta\sum_i \Xi_{it}\mu_{ijt} &= -\frac{1}{R_{t+1}}\frac{\partial S_{t+1}}{\partial V_{jt+1}} \\
\Omega_{ijt} &= \frac{1}{R_{t+1}}\Delta_{jt+1}L_{it} + \Xi_{it}\zeta_{it},
\end{aligned}
$$

where $\tilde{P}_{jt}$ is the Planner's ideal shadow price index for consumption of location j and ζ_{it} is the household's multiplier on their accounting constraint for movement flows.

For the envelope conditions, we have

$$\frac{\partial \mathcal{S}_t}{\partial L_{jt}} = \Lambda_{jt} Z_{jt} - \sum_i \Lambda_{it}\kappa_{ijt} q_{ijt} + \sum_k \Lambda_{kt}\frac{\partial Z_{kt}}{\partial L_{jt}} L_{kt} + \Xi_{jt}\partial_a u_{jt}\frac{\partial A_{jt}}{\partial L_{jt}}$$
$$+ \frac{1}{R_{t+1}}\sum_k \left(\frac{\partial \mathcal{S}_{t+1}}{\partial \bar{Z}_{kt+1}}\frac{\partial \bar{Z}_{kt+1}}{\partial L_{jt}} + \mu_{jkt}\Delta_{kt+1}\right)$$
$$\frac{\partial \mathcal{S}_t}{\partial V_{jt}} = -\Xi_{jt}$$
$$\frac{\partial \mathcal{S}_t}{\partial \bar{Z}_{jt}} = \Lambda_{jt}\frac{\partial Z_{jt}}{\partial \bar{Z}_{jt}} L_{jt} + \frac{1}{R_{t+1}}\sum_k \frac{\partial \mathcal{S}_{t+1}}{\partial \bar{Z}_{kt+1}}\frac{\partial \bar{Z}_{kt+1}}{\partial \bar{Z}_{jt}}.$$

Denote by $\varsigma_{jt} \equiv \frac{\partial \mathcal{S}_t}{\partial \ln(\bar{Z}_{jt})}$, $\gamma_{jt}^{Z,k} \equiv \frac{\partial \ln(z_{jt})}{\partial \ln(L_{kt})}$, $\gamma_{jt}^{Z\bar{Z}} \equiv \frac{\partial \ln(z_{jt})}{\partial \ln(\bar{Z}_{jt})}$, $\gamma_{jt}^{\bar{Z},k} \equiv \frac{\partial \ln(\bar{Z}_{jt})}{\partial \ln(L_{kt-1})}$, and $\gamma_{jt}^{\bar{Z}\bar{Z},k} \equiv \frac{\partial \ln(\bar{Z}_{jt})}{\partial \ln(\bar{Z}_{kt-1})}$. Combining equations in a similar manner as that used to derive Equation (C.21), we have

$$-\sum_i L_{it}\sum_k \Delta_{kt+1}\frac{1}{P_{jt+1}}\frac{\partial \mu_{ikt}}{\partial c_{jt+1}} = \sum_i L_{it}\mu_{ijt}\left(\frac{\beta R_{t+1}\partial_c u_{jt+1}/P_{jt+1}}{\partial_c u_{it}/P_{it}} - 1\right) \qquad \text{(C.43)}$$
$$\Delta_{jt} = w_{jt} - P_{jt}c_{jt} + \sum_k \frac{L_{kt}}{L_{jt}} w_{kt}\gamma_{kt}^{Z,j} + \frac{\partial_a u_{jt}}{\partial_c u_{jt}} P_{jt} a_{jt}\gamma_{jt}^{A}$$
$$+ \frac{1}{R_{t+1}}\sum_k \left(\varsigma_{kt+1}\frac{\gamma_{kt+1}^{\bar{Z},j}}{L_{jt}} + \mu_{jkt}\Delta_{kt+1}\right)$$
$$\varsigma_{jt} = w_{jt}L_{jt}\gamma_{jt}^{Z\bar{Z}} + \frac{1}{R_{t+1}}\sum_k \varsigma_{kt+1}\gamma_{kt+1}^{\bar{Z}\bar{Z},j}.$$

Just as in Appendix C.1.2, setting supporting transfers, wages, and prices according to $\tau_{jt} = P_{jt}c_{jt} - w_{jt} - s_{jt}$, $w_{it} = \Lambda_{it}Z_{it}$, and $p_{ijt} = \kappa_{ijt}w_{it}/Z_{it}$ implements the Planner's allocation as an equilibrium with the producer price of location 1 set as the numeraire.

As argued in the main text, one can see that the logic of Proposition 3.1 extends to a setting with a more general structure of externalities. The optimality condition for this more general setting, Equation (C.43), has the same interpretation as before; the only substantive difference comes from the additional recursive terms in the fiscal externality which reflect the dynamic nature of the fundamental productivity process.

C.2.2 Heterogeneous Household Types

We consider an extension to the baseline model by allowing for the presence of ex-ante heterogeneous household types. We assume each household is indexed by type $\theta \in \Theta$.

Competitive Equilibrium

The value function of each household, which is analogous to (3.1), is now given by

$$V_{jt}^{\theta} = \max_{\{\mu_{jkt}^{\theta}\}} u^{\theta}(c_{jt}^{\theta}, a_{jt}^{\theta}) - \psi_{jt}^{\theta}\big(\{\mu_{jkt}^{\theta}\}\big) + \beta^{\theta} \sum_k \mu_{jkt}^{\theta} V_{kt+1}^{\theta}$$
$$\text{s.t.} \quad \sum_k \mu_{jkt}^{\theta} = 1.$$

Therefore, the optimal migration choices solve

$$\{\mu_{jkt}^{\theta}\} \in \arg\max_{\{\mu_{jkt}^{\theta}\}} u^{\theta}(c_{jt}^{\theta}, a_{jt}^{\theta}) - \psi_{jt}^{\theta}\big(\{\mu_{jkt}^{\theta}\}\big) + \beta^{\theta} \sum_k \mu_{jkt}^{\theta} V_{kt+1}^{\theta}. \tag{C.44}$$
$$\text{s.t.} \quad \sum_k \mu_{jkt}^{\theta} = 1.$$

The household's consumption basket is $c_{jt}^{\theta} = c_{jt}^{\theta}\big(\{q_{ijt}^{\theta}\}\big)$, and we assume $c_{jt}^{\theta}\big(\{q_{ijt}^{\theta}\}\big)$ is homogeneous of degree one. The households solve

$$\begin{aligned} \{q_{ijt}^{\theta}\}_i \in \arg\max_{\{q_{ijt}^{\theta}\}} c_{jt}^{\theta}\big(\{q_{ijt}^{\theta}\}\big) \\ \text{s.t.} \quad \sum_i p_{ijt} q_{ijt}^{\theta} \leq w_{jt}^{\theta} + \tau_{jt}^{\theta} + s_{jt}^{\theta}. \end{aligned} \tag{C.45}$$

The associated price index is

$$P_{jt}^{\theta} = \min_{\{q_{ijt}^{\theta}\}} \big[\sum_i p_{ijt} q_{ijt}^{\theta}\big] \quad \text{s.t.} \quad c_{jt}^{\theta}\big(\{q_{ijt}^{\theta}\}\big) \geq 1. \tag{C.46}$$

We allow the spatial transfer of the government to depend on the household's type,

τ_{jt}^{θ}. The budget constraint of each household is

$$c_{jt}^{\theta} = \frac{w_{jt}^{\theta} + \tau_{jt}^{\theta} + s_{jt}^{\theta}}{P_{jt}^{\theta}}. \tag{C.47}$$

On the production side, we postulate the production function of each location i as $F_{it}\big(\{Z_{it}^{\theta} L_{it}^{\theta}\}_{\theta}\big)$, where F_{it} is constant returns to scale. The firms choose the labor inputs to maximize profits:

$$\max_{L_{it}^{\theta}} p_{it} F_{it}(\{Z_{it}^{\theta} L_{it}^{\theta}\}_{\theta}) - \sum_{\theta} w_{it}^{\theta} L_{it}^{\theta},$$

where p_{it} denotes the price of locally produced goods. Under perfect competition and constant returns to scale, the price is equal to the unit cost of production:

$$p_{it} = \min_{\{L_{it}^{\theta}\}_{\theta}} \sum_{\theta} w_{it}^{\theta} L_{it}^{\theta} \quad \text{s.t.} \quad F_{it}\big(\{Z_{it}^{\theta} L_{it}^{\theta}\}_{\theta}\big) \geq 1, \tag{C.48}$$

and the optimality conditions are

$$p_{it} \frac{\partial F_{it}\big(\{Z_{it}^{\theta} L_{it}^{\theta}\}_{\theta}\big)}{\partial (Z_{it}^{\theta} L_{it}^{\theta})} Z_{it}^{\theta} = w_{it}^{\theta}. \tag{C.49}$$

As before, we assume that the presence of iceberg trade costs, κ_{ijt}, when shipped from location i to j at time t. The price of the goods produced in location i faced by the consumers in location j is, therefore,

$$p_{ijt} = p_{it} \kappa_{ijt}. \tag{C.50}$$

Following Fajgelbaum and Gaubert (2020), we allow a general form of within-location spillovers both for amenity and productivity:

$$a_{it}^{\theta} = A_{it}^{\theta}\big(\{L_{it}^{\theta'}\}_{\theta'}\big), \tag{C.51}$$

$$Z_{it}^{\theta} = Z_{it}^{\theta}\big(\{L_{it}^{\theta'}\}_{\theta'}\big). \tag{C.52}$$

The market clearing conditions are

$$\sum_j \sum_\theta \kappa_{ijt} q^\theta_{ijt} L^\theta_{jt} = F_{it}\big(\{Z^\theta_{it} L^\theta_{it}\}_\theta\big). \tag{C.53}$$

The population distribution evolves according to

$$L^\theta_{jt+1} = \sum_i \mu^\theta_{ijt} L^\theta_{it}. \tag{C.54}$$

The government budget constraint is

$$\sum_j \sum_\theta \tau^\theta_{jt} L^\theta_{jt} = 0. \tag{C.55}$$

Likewise, the private transfer satisfies

$$\sum_j \sum_\theta s^\theta_{jt} L^\theta_{jt} = 0. \tag{C.56}$$

Given public and private transfers, $\{\tau^\theta_{jt}, s^\theta_{jt}\}$ that satisfy (C.55) and (C.56) Competitive equilibrium consists of

- prices: $\{\{w^\theta_{jt}, P^\theta_{jt}\}_\theta, \{p_{ijt}\}_i, p_{jt}, \}$,
- quantities: $\{\{\mu^\theta_{ijt}, q^\theta_{ijt}\}_i, L^\theta_{jt}, a^\theta_{jt}, Z^\theta_{jt}, c^\theta_{jt}\}$,

such that equations (C.44)-(C.56) hold.

Planning Problem

The planner's problem in this economy is analogous to Definition 3.2:

Definition C.2 (Planning Problem with Heterogeneous Households)**.** *Given Pareto weights $\{\varkappa^\theta_i\}$ attached to individuals of type θ in location i in period* 0*, the Planner* *solves*

$$\max_{\left\{\{q^{\theta}_{kjt},\mu^{\theta}_{jkt},p_{kjt}p_{kt}\},c^{\theta}_{jt},L^{\theta}_{jt+1},V^{\theta}_{jt},P_{jt},a^{\theta}_{jt},w^{\theta}_{jt},\tau^{\theta}_{jt},s^{\theta}_{jt}\right\}_{t\geq 0}} \sum_{i}\sum_{\theta} \varkappa^{\theta}_{i} L^{\theta}_{i0} V^{\theta}_{i0}$$

$$s.t. \quad V^{\theta}_{jt} = u^{\theta}(c^{\theta}_{jt}, a^{\theta}_{jt}) - \psi^{\theta}_{jt}\big(\{\mu^{\theta}_{jkt}\}\big) + \beta^{\theta} \sum_{k} \mu^{\theta}_{jkt} V^{\theta}_{kt+1} \tag{PK}$$

$$\{\mu^{\theta}_{jkt}\} \in \arg \max_{\{\mu^{\theta}_{jkt}\}:\sum_k \mu^{\theta}_{jkt}=1} u^{\theta}(c^{\theta}_{jt}, a^{\theta}_{jt}) - \psi^{\theta}_{jt}\big(\{\mu^{\theta}_{jkt}\}\big) + \beta^{\theta} \sum_{k} \mu^{\theta}_{jkt} V^{\theta}_{kt+1} \tag{IC}$$

(C.45) - (C.56).

The set of solutions to this problem given arbitrary Pareto weights $\{\varkappa^{\theta}_{i}\}$ *defines the Pareto frontier.*

As in the baseline, we consider the following relaxed dual problem in tackling the above problem:

$$\max_{\left\{\{q^{\theta}_{kjt},\mu^{\theta}_{jkt}\},L^{\theta}_{jt+1},V^{\theta}_{jt}\right\}_{t\geq 0}} \sum_{t=0}^{\infty} \prod_{s\leq t} \frac{1}{R_s} \sum_{i} \Lambda_{it}\Big(F_{it}(\{Z^{\theta}_{it}(\{L^{\theta'}_{it}\}_{\theta'})L^{\theta}_{it}\}_{\theta}) - \sum_{j}\sum_{\theta} \kappa_{ijt} q^{\theta}_{ijt} L^{\theta}_{jt}\Big) \tag{C.57}$$

s.t.

$$V_{j0} \geq \underline{V}_{j0}$$

$$L^{\theta}_{jt+1} = \sum_{i,\theta} L^{\theta}_{it} \mu^{\theta}_{ijt}$$

$$V^{\theta}_{jt} = u\big(c^{\theta}_{jt}(\{q^{\theta}_{ijt}\}), A^{\theta}_{jt}(L^{\theta}_{jt})\big) - \psi_{jt}\big(\{\mu^{\theta}_{jkt}\}\big) + \beta^{\theta} \sum_{k} \mu^{\theta}_{jkt} V^{\theta}_{kt+1}$$

$$\mu^{\theta}_{jkt} = \hat{\mu}^{\theta}_{jkt}(\{V^{\theta}_{it+1}\}),$$

Recursive formulation is

$$\begin{aligned}
\mathcal{S}_t\big(\{L^\theta_{jt}, V^\theta_{jt}\}\big) = &\max_{\{\{q^\theta_{ijt},\mu^\theta_{ijt}\},L^\theta_{jt+1},V^\theta_{jt+1}\}} \sum_i \Lambda_{it}\Big(F_{it}(\{Z^\theta_{it}(\{L^{\theta'}_{it}\}_{\theta'})L^\theta_{it}\}_\theta) - \sum_{j,\theta}\kappa_{ijt}q^\theta_{ijt}L^\theta_{jt}\Big) \\
&+ \frac{1}{R_{t+1}}\mathcal{S}_{t+1}\big(\{L^\theta_{jt+1}, V^\theta_{jt+1}\}\big) \qquad \text{(C.58)}\\
\text{s.t.} \quad & L^\theta_{jt+1} = \sum_i L^\theta_{it}\mu^\theta_{ijt} : \quad \tilde{\Delta}^\theta_{jt+1} \\
& V^\theta_{jt} \le u^\theta\big(c^\theta_{jt}(\{q^\theta_{ijt}\}), A^\theta_{jt}(\{L^{\theta'}_{jt}\}_{\theta'})\big) - \psi^\theta_{jt}\big(\{\mu^\theta_{jkt}\}\big) + \beta^\theta \sum_k \mu^\theta_{jkt}V^\theta_{kt+1} : \quad \Xi^\theta_{jt} \\
& \mu^\theta_{ijt} = \hat{\mu}^\theta_{ijt}\big(\{V^\theta_{kt+1}\}\big) : \quad \Omega^\theta_{ijt},
\end{aligned}$$

The first-order conditions are

$$\Lambda_{it}\kappa_{ijt} = \frac{\partial u^\theta(c^\theta_{jt}, a^\theta_{jt})}{\partial c^\theta_{jt}} \frac{\partial c^\theta_{jt}(\{q^\theta_{ijt}\})}{\partial q^\theta_{ijt}} \Xi^\theta_{jt}/L^\theta_{jt} \qquad \text{(C.59)}$$

$$\tilde{\Delta}^\theta_{jt+1}L^\theta_{it} = \Omega^\theta_{ijt} \qquad \text{(C.60)}$$

$$\tilde{\Delta}^\theta_{jt+1} = \frac{1}{R_t}\frac{\partial \mathcal{S}_{t+1}}{\partial L^\theta_{jt+1}} \qquad \text{(C.61)}$$

$$\frac{1}{R_t}\frac{\partial \mathcal{S}_{t+1}}{\partial V^\theta_{jt+1}} + \sum_k \sum_l \Omega^\theta_{klt+1}\frac{\partial \hat{\mu}^\theta_{klt}}{\partial V^\theta_{jt+1}} + \beta^\theta \sum_k \Xi^\theta_{kt}\mu^\theta_{kjt} = 0 \qquad \text{(C.62)}$$

Define the planner's price index for location j for type θ at time t as

$$P^\theta_{jt}c^\theta_{jt} = \min_{\{q^\theta_{ijt}\}} \sum_i \Lambda_{it}\kappa_{ijt}q^\theta_{ijt} \quad \text{s.t.} \quad c^\theta_{jt}(\{q^\theta_{ijt}\}) \ge c^\theta_{jt} \qquad \text{(C.63)}$$

Given this definition, we can multiply both sides by q^θ_{ijt} and sum over i to rewrite equation (C.59) as

$$1 = \frac{\partial u^\theta(c^\theta_{jt}, a^\theta_{jt})}{\partial c^\theta_{jt}} \frac{1}{P^\theta_{jt}} \Xi^\theta_{jt}/L^\theta_{jt}, \qquad \text{(C.64)}$$

where we used the fact that c^θ_{jt} is homogeneous of degree one.

The envelope conditions are

$$\frac{\partial \mathcal{S}_t}{\partial L^\theta_{jt}} = \Lambda_{jt}\left[\frac{\partial F_{jt}(\{Z^\theta_{jt}(\{L^{\theta'}_{jt}\}_{\theta'})L^\theta_{jt}\}_\theta)}{\partial(Z^\theta_{jt}L^\theta_{jt})}Z^\theta_{jt}\right.$$
$$\left. + \sum_{\tilde\theta}\frac{\partial F_{jt}(\{Z^\theta_{jt}(\{L^{\theta'}_{jt}\}_{\theta'})L^\theta_{jt}\}_\theta)}{\partial(Z^{\tilde\theta}_{jt}L^{\tilde\theta}_{jt})}L^{\tilde\theta}_{jt}\frac{\partial Z^\theta_{jt}(\{L^{\theta'}_{jt}\}_{\theta'})}{\partial L^{\tilde\theta}_{jt}}\right.$$
$$- P^\theta_{jt}c^\theta_{jt} + \sum_{\tilde\theta}\Xi^{\tilde\theta}_{jt}\frac{\partial u^{\tilde\theta}(c^{\tilde\theta}_{jt}, a^{\tilde\theta}_t)}{\partial a^{\tilde\theta}}\frac{\partial a^{\tilde\theta}(\{L^{\theta'}_{jt}\}_{\theta'})}{\partial L^\theta_{jt}} + \sum_k \mu^\theta_{jkt}\tilde\Delta^\theta_{jt+1} = 0 \tag{C.65}$$
$$\frac{\partial \mathcal{S}_t}{\partial V^\theta_{jt}} = -\Xi^\theta_{jt} \tag{C.66}$$

Substituting (C.60), (C.64), and (C.66) into (C.62), and defining $\Delta_{jt} \equiv \frac{\partial \mathcal{S}_t}{\partial L^\theta_{jt}}$, we obtain

$$-\sum_k\sum_l \Delta^\theta_{lt+1}L^\theta_{kt}\frac{\partial\hat\mu^\theta_{klt}}{\partial c^\theta_{jt+1}}\frac{1}{P^\theta_{jt+1}} = \sum_k L^\theta_{kt}\mu^\theta_{kjt}\left(\beta^\theta R_t\frac{\frac{\partial u^\theta(c^\theta_{jt+1},a^\theta_{jt+1})}{\partial c^\theta_{jt+1}}\frac{1}{P^\theta_{jt+1}}}{\frac{\partial u^\theta(c^\theta_{kt},a^\theta_{kt})}{\partial c^\theta_{kt}}\frac{1}{P^\theta_{kt}}} - 1\right) \tag{C.67}$$

Plugging (C.59) and (C.61) into (C.65) and rearranging,

$$\Delta^\theta_{jt} = \Lambda_{jt}\left[\frac{\partial F_{jt}(\{Z^\theta_{jt}(\{L^{\theta'}_{jt}\}_{\theta'})L^\theta_{jt}\}_\theta)}{\partial(Z^\theta_{jt}L^\theta_{jt})}Z^\theta_{jt} + \sum_{\tilde\theta}\frac{\partial F_{jt}(\{Z^\theta_{jt}(\{L^{\theta'}_{jt}\}_{\theta'})L^\theta_{jt}\}_\theta)}{\partial(Z^{\tilde\theta}_{jt}L^{\tilde\theta}_{jt})}Z^{\tilde\theta}_{jt}\gamma^{z,\tilde\theta,\theta}_{jt}\right]$$
$$-P^\theta_{jt}c^\theta_{jt} + \sum_{\tilde\theta}\frac{\frac{\partial u^{\tilde\theta}(c^{\tilde\theta}_{jt},a^{\tilde\theta}_t)}{\partial a^{\tilde\theta}}}{\frac{\partial u^{\tilde\theta}(c^{\tilde\theta}_{jt},a^{\tilde\theta}_{jt})}{\partial c^{\tilde\theta}_{jt}}}P^{\tilde\theta}_{jt}A^{\tilde\theta}_{jt}\gamma^{a,\tilde\theta,\theta}_{jt} + \frac{1}{R_t}\sum_k \mu^\theta_{jkt}\Delta^\theta_{jt+1}, \tag{C.68}$$

where

$$\gamma^{z,\tilde\theta,\theta}_{jt} \equiv \frac{\partial \ln Z^{\tilde\theta}_{jt}(\{L^{\theta'}_{jt}\}_{\theta'})}{\partial \ln L^\theta_{jt}} \tag{C.69}$$
$$\gamma^{a,\tilde\theta,\theta}_{jt} \equiv \frac{\partial \ln A^{\tilde\theta}_{jt}(\{L^{\theta'}_{jt}\}_{\theta'})}{\partial \ln L^\theta_{jt}}. \tag{C.70}$$

Define the wage of type θ in location j at time t in the planner's solution to be

$$w_{jt}^{\theta} = \Lambda_{jt} \frac{\partial F_{jt}(\{Z_{jt}^{\theta}(\{L_{jt}^{\theta'}\}_{\theta'}) L_{jt}^{\theta}\}_{\theta})}{\partial (Z_{jt}^{\theta} L_{jt}^{\theta})} Z_{jt}^{\theta}. \tag{C.71}$$

Then (C.68) can be rewritten as

$$\Delta_{jt}^{\theta} = w_{jt}^{\theta} + \sum_{\tilde{\theta}} w_{jt}^{\tilde{\theta}} \gamma_{jt}^{z,\tilde{\theta},\theta} - P_{jt}^{\theta} c_{jt}^{\theta} + \sum_{\tilde{\theta}} \frac{\frac{\partial u^{\tilde{\theta}}(c_{jt}^{\tilde{\theta}}, a_t^{\tilde{\theta}})}{\partial a^{\tilde{\theta}}}}{\frac{\partial u^{\tilde{\theta}}(c_{jt}^{\tilde{\theta}}, a_{jt}^{\tilde{\theta}})}{\partial c_{jt}^{\tilde{\theta}}}} P_{jt}^{\tilde{\theta}} A_{jt}^{\tilde{\theta}} \gamma_{jt}^{a,\tilde{\theta},\theta} + \frac{1}{R_t} \sum_{k} \mu_{jkt}^{\theta} \Delta_{jt+1}^{\theta}. \tag{C.72}$$

One can easily verify, as in the baseline model, that the above allocation can be supported as a competitive equilibrium with the following transfer policies:

$$\tau_{jt}^{\theta} = P_{jt}^{\theta} c_{jt}^{\theta} - w_{jt}^{\theta} - s_{jt}^{\theta}. \tag{C.73}$$

C.2.3 Origin-Destination Specific Transfer

In the main text, we restricted ourselves to location-specific transfers. Such policy is arguably realistic, it is natural to consider richer policy tools. Here, we consider a transfer policy that can depend on the households' origin and destination location. We denote τ_{jt}^{o} as the transfer to households living in location j at time t and in location o at time $t-1$. With such a policy, the consumption of households with past location o and current location j at time t is

$$c_{jt}^{o} = \frac{w_{jt} + \tau_{jt}^{o} + s_{jt}^{o}}{P_{jt}}, \tag{C.74}$$

which replaces equation (3.1). Here, we allow the private transfer to also depend on origin and destination.

Competitive Equilibrium

The value function of a household currently living in location j with past location o is

$$V_{jt}^{o} = \max_{\{\mu_{jkt}^{o}\}} u(c_{jt}^{o}, a_{jt}) - \psi_{jt}\big(\{\mu_{jkt}^{o}\}\big) + \beta \sum_{k} \mu_{jkt}^{o} V_{kt+1}^{j}$$
$$\text{s.t.} \quad \sum_{k} \mu_{jkt}^{o} = 1.$$

We denote the solution to the above problem as

$$\hat{\mu}_{ijt}^{o}\big(\{V_{kt+1}^{\theta}\}\big) \equiv \arg\max_{\{\mu_{jkt}^{o}\}} u(c_{jt}^{o}, a_{jt}) - \psi_{jt}\big(\{\mu_{jkt}^{o}\}\big) + \beta \sum_{k} \mu_{jkt}^{o} V_{kt+1}^{j} \tag{C.75}$$
$$\text{s.t.} \quad \sum_{k} \mu_{jkt}^{o} = 1. \tag{C.76}$$

The consumption choices for each location's goods and the associated price indices remain the same as in the baseline model for each o:

$$\begin{aligned} \{q_{ijt}^{o}\}_i &\in \arg\max_{\{q_{ijt}^{o}\}} c_{jt}\big(\{q_{ijt}^{o}\}\big) \\ \text{s.t.} \quad \sum_{i} p_{ijt} q_{ijt}^{o} &\leq w_{jt} + \tau_{jt}^{o} + s_{jt}^{o}. \end{aligned} \tag{C.77}$$

The associated price index remains the same as in (3.2).

The production side of the economy remains identical to the main text with the population in each location given by

$$L_{jt} = \sum_{o} L_{jt}^{o},$$

and L_{jt}^{o} evolves according to

$$L_{jt+1}^{o} = \sum_{i} \mu_{ijt}^{o} L_{it}^{o}. \tag{C.78}$$

The government budget constraint is

$$\sum_j \sum_o \tau^o_{jt} L^o_{jt} = 0. \tag{C.79}$$

Likewise, the private transfer satisfies

$$\sum_j \sum_o s^\theta_{jt} L^o_{jt} = 0. \tag{C.80}$$

Planning Problem

The planning problem can be recursively written as

$$\begin{aligned}
\mathcal{S}_t\big(\{L^o_{jt}, V^o_{jt}\}\big) = & \max_{\{\{q^o_{ijt}, \mu^o_{ijt}\}\}_{i,j,o}, \{L^i_{jt+1}, V^i_{jt+1}\}_{i,j}} \sum_i \Lambda_{it}\Big(Z_{it}(L_{it})L_{it} - \sum_{j,o} \kappa_{ijt} q^o_{ijt} L^o_{jt}\Big) \\
& + \frac{1}{R_{t+1}} \mathcal{S}_{t+1}\big(\{L^i_{jt+1}, V^i_{jt+1}\}\big) \qquad \text{(C.81)} \\
\text{s.t.} \quad L_{it} &= \sum_o L^o_{it} \\
L^i_{jt+1} &= \sum_o \mu^o_{ijt} L^o_{it}. : \quad \tilde{\Delta}^i_{jt+1} \\
V^o_{jt} &\leq u\big(c_{jt}(\{q^o_{ijt}\}), A_{jt}(L_{jt})\big) - \psi_{jt}\big(\{\mu^o_{jkt}\}\big) + \beta \sum_k \mu^o_{jkt} V^j_{kt+1} : \quad \Xi^o_{jt} \\
\mu^o_{ijt} &= \hat{\mu}^o_{ijt}\big(\{V^i_{kt+1}\}\big) : \quad \Omega^o_{ijt},
\end{aligned}$$

The first-order conditions are

$$1 = \frac{\partial u(c^o_{jt}, A_{jt})}{\partial c^o_{jt}} \frac{1}{P_{jt}} \Xi^o_{jt} / L^o_{jt} \tag{C.82}$$

$$\tilde{\Delta}^i_{jt+1} L^o_{it} = \Omega^o_{ijt} \tag{C.83}$$

$$\tilde{\Delta}^i_{jt+1} = \frac{1}{R_t} \frac{\partial \mathcal{S}_{t+1}}{\partial L^i_{jt+1}} \tag{C.84}$$

$$\frac{1}{R_t} \frac{\partial \mathcal{S}_{t+1}}{\partial V^i_{jt+1}} + \sum_o \sum_l \Omega^o_{ilt} \frac{\partial \hat{\mu}^o_{ilt}}{\partial V^i_{jt+1}} + \beta \sum_o \Xi^o_{it} \mu^o_{ijt} = 0 \tag{C.85}$$

where

$$P_{jt} = \min_{\{q^o_{ijt}\}} [\sum_i \Lambda_{it} \kappa_{ijt} q^o_{ijt}] \quad \text{s.t.} \quad c_{jt}(\{q^o_{ijt}\}) \geq 1. \tag{C.86}$$

The envelope conditions are

$$\frac{\partial \mathcal{S}_t\big(\{L^o_{jt}, V^o_{jt}\}\big)}{\partial L^o_{jt}} = \Lambda_{jt} Z_{jt}(1+\gamma^z_{jt}) - P_{jt} c^o_{jt} + \Xi^o_{jt} \frac{\partial u(c^o_{jt}, A_{jt})}{\partial A_{jt}} A'_{jt}(L_{jt}) + \sum_k \mu^o_{jkt} \tilde{\Delta}^j_{kt+1} \tag{C.87}$$

$$\frac{\partial \mathcal{S}_t\big(\{L^o_{jt}, V^o_{jt}\}\big)}{\partial V^o_{jt}} = -\Xi^o_{jt} \tag{C.88}$$

Defining $\Delta^o_{jt} \equiv \frac{\partial S_t}{\partial L^o_{jt}}$, we can combine above equations to obtain

$$-L_{it} \sum_l \Delta^i_{lt+1} \frac{\partial \hat{\mu}_{ilt}}{\partial c^i_{jt+1}} \frac{1}{P_{jt+1}} = \beta R_t \sum_o L^o_{it} \mu^o_{ijt} \left(\frac{\frac{\partial u(c^i_{jt+1}, A_{jt+1})}{\partial c^i_{jt+1}} \frac{1}{P_{jt+1}}}{\frac{\partial u(c^o_{it}, A_{it})}{\partial c^o_{it}} \frac{1}{P_{it}}} - 1 \right) \tag{C.89}$$

$$\Delta^o_{jt} = w_{jt}(1+\gamma^z_{jt} \frac{L^o_{jt}}{L_{jt}}) - P_{jt} c^o_{jt} + \frac{\frac{\partial u(c^o_{jt}, A_{jt})}{\partial A_{jt}}}{\frac{\partial u(c^o_{jt}, A_{jt})}{\partial c_{jt}}} P_{jt} A_{jt} \gamma^A_{jt} \frac{L^o_{jt}}{L_{jt}} + \frac{1}{R_t} \sum_k \mu^o_{jkt} \Delta^j_{kt+1} \tag{C.90}$$

where $w_{jt} = \Lambda_{jt} Z_{jt}$. One can verify that the above allocation can be supported as a competitive equilibrium with the following transfer policies:

$$\tau^o_{jt} = P_{jt} c^o_{jt} - w_{jt} - s^o_{jt}. \tag{C.91}$$

C.3 Numerical Illustrations

In this section, we consider the setting with only two locations, $\{1, 2\}$. We set the structural parameters $\{\beta, \rho, \theta, \sigma, \gamma^A, \gamma^Z\}$ using our calibration for the U.S. economy in Section 3.4, as summarized by Table 3.2. For location fundamentals, we set the following values: $\bar{Z}_{jt} = \bar{A}_{jt} = 1$ for all j, t; $\kappa_{ijt} = 1.2$ for $i \neq j$ and $\kappa_{ijt} = 1$ for $i = j$; $\chi_{ijt} = 0.1$ for $i \neq j$ and $\chi_{ij} = 1$ for $i = j$. We always choose the consumption basket in location 1 in each period as numeraire.

Permanent and Transient Shocks In Figures C·1 and C·2, we show the impulse responses to the permanent and transitory negative shock of 30% to location 1, respectively. To simulate these impulse responses, we assume that the economy is in a steady state up to $t = 0$. At $t = 0$, an unanticipated permanent and temporary change of productivity in location 1 materializes, $\overline{Z}_{1t}$, whose process is given by the top left panel ("Productivity") of Figures C·1 and C·2, respectively. The solid lines in each panel show the equilibrium responses of the employment, consumption, real wages, optimal transfers, and consumer prices for locations 1 and 2, under the optimal transfer characterized by Proposition 3.1. To benchmark our results, we also simulate our model under an alternative assumption that transfer is set to zero in the dashed lines.

Right after the negative productivity shocks to location 1, net transfer sharply increases in location 1 to mitigate the differences in consumption between the two locations. Interestingly, even if the shock is permanent, net transfer decreases over time until it reaches the steady-state level. This back-loading of incentives is a natural consequence of the insurance-incentive trade-off highlighted in Proposition 3.1. On the one hand, the Planner wants to provide consumption insurance to location 1. On the other hand, the Planner wants to incentivize out-migration from location

1. The most effective way of ensuring incentives for the Planner is to commit to cut transfer in the future. Reflecting this back-loaded pattern of transfer, the consumption in location 1 is substantially higher under optimal transfer than the no-transfer benchmark, while the gap decreases over time. Lastly, given the transfer, there is less migration than no transfer benchmark, and a larger size of the population stays in location 1 under the optimal transfer.

Figure C·1: Permanent Shock

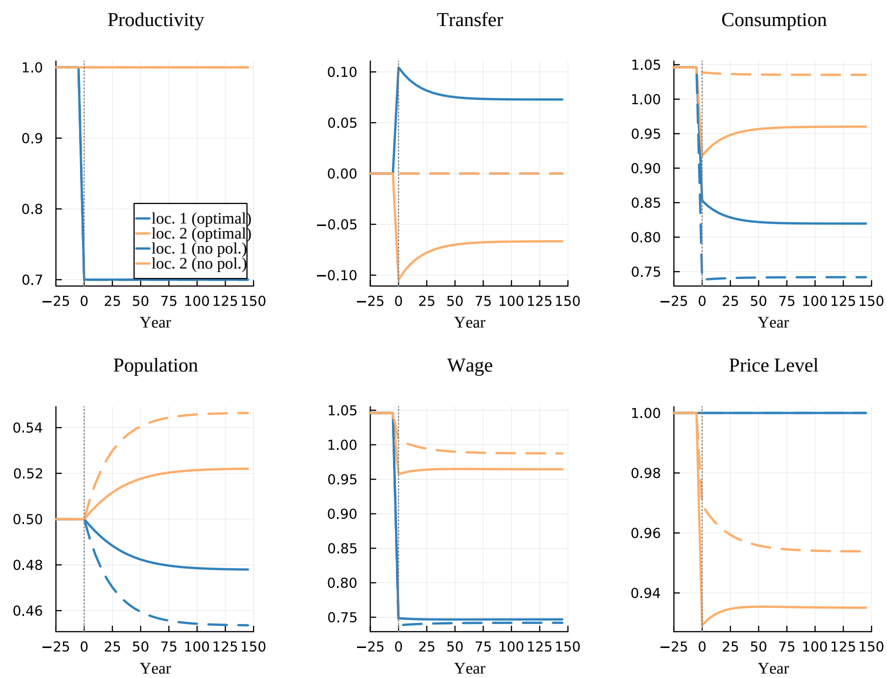

Responses to Anticipated Shocks In Figure C·3, we alternatively consider the impulse responses to anticipated permanent and transitory shocks in the future, respectively. To simulate these impulse responses, we assume that the economy is in a steady state up to $t = 0$. At $t = 0$, agents learn that the productivity in location 1, $\overline{Z}_{1t}$, will drop by 30 percent at $t = 10$ and for all subsequent period, as shown in the top left panel ("Productivity") of Figures C·1.

In both cases, net transfer in location 1 is positive and increasing up to $t = 10$,

Figure C·2: Transitory Shock

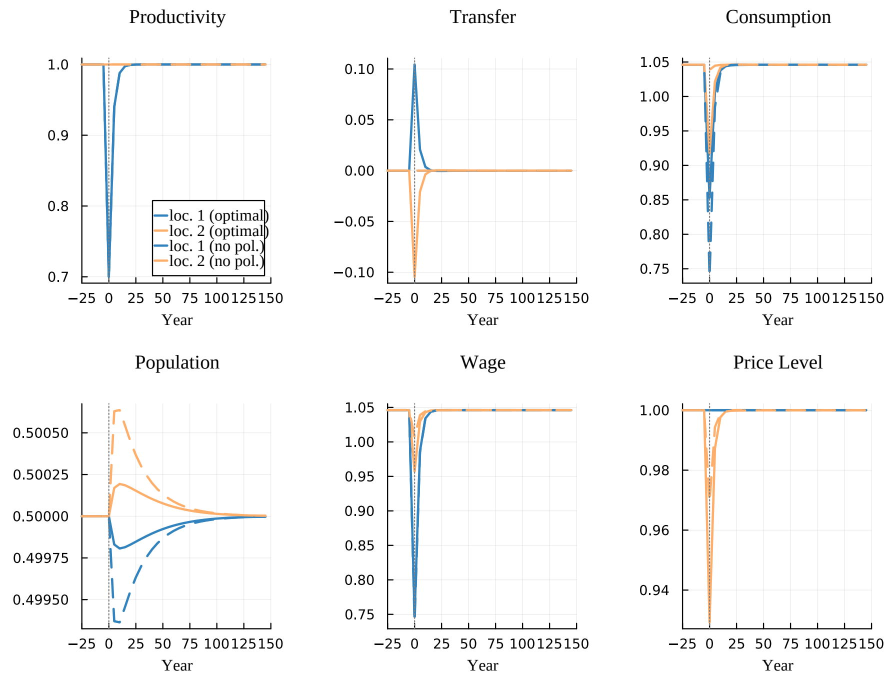

and turns negative after $t = 10$. Before the negative shock materializes, the Planner wants to encourage outmigration from location 1. However, after the shock hits, the Planner has the incentive to provide consumption insurance to mitigate the negative shock. At the same time, for the same reasons as the unanticipated permanent shock in Figure C·1, the transfer gradually decreases over time to mitigate the moral-hazard problem of migration decisions.

Figure C·3: Permanent Anticipated Shock

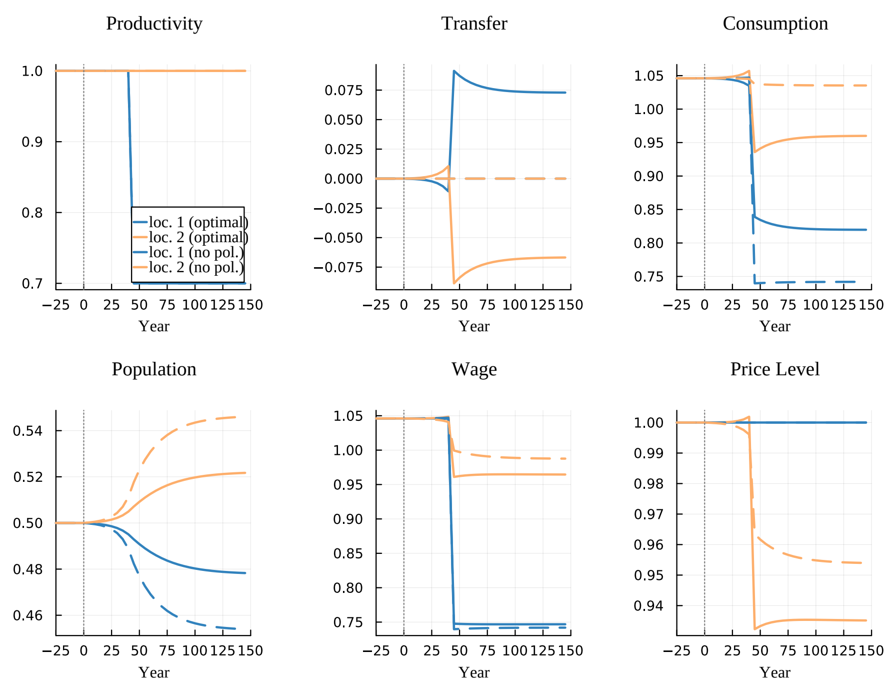

C.4 Appendix for Parametric Model

Given productivities $\{Z_{jt}(\cdot)\}$, amenities $\{A_{jt}(\cdot)\}$, migration frictions $\{\psi_{ijt}(\cdot)\}$, trade frictions $\{\kappa_{ijt}\}$, private and public transfers $\{s_{jt}, \tau_{jt}\}$, and initial population distribution $\{l_{j0}\}$,

- Consumption $\{c_{jt}, q_{jt}\}$:

$$c_{jt} = \frac{w_{jt} + \tau_{jt} + s_{jt}}{P_{jt}} \tag{C.92}$$

  and

$$q_{ijt} = \frac{P_{jt} c_{jt} \Psi_{ijt}}{p_{it}}, \qquad \Psi_{ijt} = \frac{p_{ijt}^{1-\sigma}}{\sum p_{ijt}^{1-\sigma}} \tag{C.93}$$

- Migration decisions $\{\mu_{ijt}, V_{jt}\}$:

$$\mu_{jk} = \frac{\chi_{jk} \exp(\theta V_k)}{\sum_l \chi_{jl} \exp(\theta V_l)} \tag{C.94}$$

$$V_{jt} = u\big(c_{jt}, a_{jt}\big) + \beta \frac{1}{\theta} \log \sum_k \chi_{jkt} \exp(\theta V_{kt+1}) \tag{C.95}$$

- Population evolution $\{l_{jt}\}$:

$$L_{jt} = \sum_i \mu_{ijt} L_{it-1} \tag{C.96}$$

- Prices $\{p_{jt}, P_{jt}\}$:

$$P_{jt} = \left(\sum_l (\kappa_{ljt} w_{lt} / Z_{lt})^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{C.97}$$

$$p_{ijt} = \kappa_{ijt} w_{it} / z_{it} \tag{C.98}$$

  where we take location 1's price index as numeraire such that $P_{1t} = 1$ for all t.

- Market clearing $\{w_{jt}\}$:[5]

$$w_i L_i = \sum_j \Psi_{ijt} P_{jt} c_{jt} L_{jt} \tag{C.99}$$

- Amenity and productivity $\{a_{jt}, z_{jt}\}$:

$$z_{jt} = \bar{Z}_{jt} L_{jt}^{\gamma^Z}, \quad a_{jt} = \bar{A}_{jt} L_{jt}^{\gamma^A} \tag{C.100}$$

Optimal Allocation. Under optimal allocation, we replace the consumption equation (C.92) with the one in Proposition 3.1 for $t > 0$:

$$-\sum_i L_{it} \sum_k \Delta_{kt+1} \frac{1}{P_{jt+1}} \frac{\partial \mu_{ikt}}{\partial c_{jt+1}} = \sum_i L_{it} \mu_{ijt} \left(\frac{\beta R_{t+1} \partial_c u_{jt+1} / P_{jt+1}}{\partial_c u_{it} / P_{it}} - 1 \right) \tag{C.101}$$

where for the i.i.d. logit case,

$$\frac{\partial \mu_{ikt}}{\partial c_{jt+1}} = \partial_c u_{jt+1} \times \begin{cases} -\theta \mu_{jit} & \text{if } k \neq i \\ \theta \left(1 - \mu_{jit}\right) & \text{if } k = i \end{cases}$$

and $\{\Delta_{jt}\}$ is determined by

$$\Delta_{jt} = w_{jt} - P_{jt} c_{jt} + w_{jt} \gamma_{jt}^z + \frac{\partial_a u_{jt}}{\partial_c u_{jt}} P_{jt} A_{jt} \gamma_{jt}^a + \frac{1}{R_{t+1}} \sum_k \mu_{jkt} \Delta_{kt+1} \tag{C.102}$$

and R_t is determined so that the government budget constraint is satisfied, i.e.,

$$\sum_i P_{it} c_{it} L_{it} = \sum_i w_{it} L_{it} \tag{C.103}$$

[5] We combine equation (C.98) and goods market clearing ($\sum_j \kappa_{ijt} q_{ijt} L_{jt} = z_{it} L_{it}$) yields the above expression.

For $t = 0$, the level of consumption is $\{c_{j0}\}$ is determined by

$$\frac{\varkappa_j \partial_c u_{j0}}{P_{j0}} = \frac{\varkappa_i \partial_c u_{i0}}{P_{i0}} \tag{C.104}$$

where $\{\varkappa_j\}$ are the Pareto weights.

C.5 Data and Calibration Appendix

C.5.1 Calibration of Trade Flows

This appendix provides a detail of the calibration of trade flows. We we follow the approach of Fajgelbaum and Gaubert (2020) to use the gravity equation of trade flows to find the unique estimates of trade flows between states that are consistent with the actual distance between states, observed pre-tax income, post-tax-and-transfer income, under the assumption of trade frictions with respect to distance.

Under the parametric model (CES goods demand), trade flows from i to j in period t follow the gravity structure:

$$\Psi_{ijt} = \frac{\delta_{it}^{O} \kappa_{ijt}}{\sum \delta_{jt}^{O} \kappa_{ijt}},$$

where δ_{jt}^{O} and δ_{it}^{D} are endogenous variables that depend on equilibrium variables such as wages and prices. Furthermore, from market clearing condition, we have

$$w_i L_i = \sum_j \Psi_{ijt} P_{jt} c_{jt} L_{jt}.$$

Finally, we impose a parametric restriction of the trade cost so that it is an iso-elastic function of geodesic distance between i and j such that

$$\kappa_{ijt} = \text{dist}_{ij}^{\psi},$$

where ψ is a parameter. Given the assumption of ψ and the data $\{\text{dist}_{ij}, w_{it}, P_{it}, c_{it}, L_{it}\}$, we can use the above equations to identify $\left\{\delta_{it}^{O}\right\}$ (up to scale), and hence trade flows $\{\Psi_{ijt}\}$.

C.5.2 Calibration of Bilateral Migration Frictions

We first back out the migration frictions, $\{\chi_{ijT}\}$, for the terminal period in the data $t = T$. To do so, we follow the following steps:

1. We guess $\{V_{jT+1}\}_j$. Using the migration flow data at T, we recover

$$\chi_{ijT} = \frac{\mu_{ijT}}{\mu_{iiT}} \exp(\theta(V_{iT+1} - V_{jT+1})),$$

   where recall our normalize $\chi_{iiT} = 1$.

2. Using the future fundamentals and assumptions of public and transfers for $t > T$, including the assumption that $\chi_{ijt} = \chi_{ijT}$ for $t > T$, we simulate the model forward starting at $t = T$ all the way toward $t = \infty$. This simulation yields the equilibrium paths of continuation value, including those at $t = T + 1$, $\{V^*_{jT+1}\}_j$.

3. We repeat the step 1 and 2 by iteratively updating $\{V_{jT+1}\}_j$ with $\{V^*_{jT+1}\}_j$.

4. Stop if this procedure converges.

Once we obtain $\{\chi_{ijT}\}$, we sequentially back out $\{\chi_{ijt}\}$ for $t < T$ backward starting from $t = T - 1$ using the following procedure:

1. Using $\{V_{jt+1}\}_j$, we obtain $\{V_{jt}\}_j$ using the Bellman equation:

$$V_{jt} = u\big(c_{jt}, a_{jt}\big) + \beta \frac{1}{\theta} \log \sum_k \chi_{jkt} \exp(\theta V_{kt+1})$$

2. Using $\{V_{jt}\}_j$, we back out $\{\chi_{ijt-1}\}$

$$\chi_{ijt-1} = \frac{\mu_{ijt-1}}{\mu_{iit-1}} \exp(\theta(V_{it} - V_{jt}))$$

REFERENCES

Acemoglu, D. (1998). Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality. *The Quarterly Journal of Economics*, 113(4):1055–1089.

Acemoglu, D. (2002). Directed Technical Change. *The Review of Economic Studies*, 69(4):781–809.

Acemoglu, D. (2023). Distorted Innovation: Does the Market Get the Direction of Technology Right? *AEA Papers and Proceedings*, 113:1–28.

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. Working Paper 31657, National Bureau of Economic Research.

Acemoglu, D., Aghion, P., Bursztyn, L., and Hémous, D. (2012). The Environment and Directed Technical Change. *American Economic Review*, 102(1):131–66.

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. (2018). Innovation, Reallocation, and Growth. *American Economic Review*, 108(11):3450–91.

Acemoglu, D., Akcigit, U., Hanley, D., and Kerr, W. (2016). Transition to Clean Technology. *Journal of Political Economy*, 124(1):52–104.

Acemoglu, D. and Autor, D. (2011). Skills, Tasks, and Technologies: Implications for Employment and Earnings. In *Handbook of Labor Economics*, volume 4, pages 1043–1171. Elsevier.

Acemoglu, D. and Finkelstein, A. (2008). Input and Technology Choices in Regulated Industries: Evidence from the Health Care Sector. *Journal of Political Economy*, 116(5):837–880.

Acemoglu, D. and Linn, J. (2004). Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry. *The Quarterly Journal of Economics*, 119(3):1049–1090.

Acemoglu, D., Manera, A., and Restrepo, P. (2020). Does the US Tax Code Favor Automation? *Brookings Papers on Economic Activity*.

Acemoglu, D. and Restrepo, P. (2018). The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment. *American Economic Review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6):2188–2244.

Acemoglu, D. and Restrepo, P. (2022a). Demographics and Automation. *The Review of Economic Studies*, 89(1):1–44.

Acemoglu, D. and Restrepo, P. (2022b). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, 90(5):1973–2016.

Adao, R., Costinot, A., and Donaldson, D. (2017). Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade. *American Economic Review*, 107(3):633–689.

Aghion, P., Bénabou, R., Martin, R., and Roulet, A. (2023). Environmental Preferences and Technological Choices: Is Market Competition Clean or Dirty? *American Economic Review: Insights*, 5(1):1–20.

Aghion, P., Dechezleprêtre, A., Hémous, D., Martin, R., and Van Reenen, J. (2016). Carbon Taxes, Path Dependency, and Directed Technical Change: Evidence from the Auto Industry. *Journal of Political Economy*, 124(1):1–51.

Aghion, P., Hepburn, C., Teytelboym, A., and Zenghelis, D. (2019). Path Dependence, Innovation, and the Economics of Climate Change. In *Handbook on Green Growth*, pages 67–83. Edward Elgar Publishing.

Aghion, P. and Howitt, P. (1992). A Model of Growth Through Creative Destruction. *Econometrica*, 60(2):323–351.

Akcigit, U., Hanley, D., and Serrano-Velarde, N. (2020). Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth. *The Review of Economic Studies*, 88(1):1–43.

Akcigit, U. and Kerr, W. (2018). Growth through Heterogeneous Innovations. *Journal of Political Economy*, 126(4):1374–1443.

Allen, T. and Donaldson, D. (2020). Persistence and Path Dependence in the Spatial Economy. Working Paper 28059, National Bureau of Economic Research.

Anderson, S., De Palma, A., and Thisse, J.-F. (1992). *Discrete Choice Theory of Product Differentiation*. MIT Press.

Archer, D. (2005). Fate of Fossil Fuel CO2 in Geologic Time. *Journal of Geophysical Research: Oceans*, 110(C9).

Archer, D., Eby, M., Brovkin, V., Ridgwell, A., Cao, L., Mikolajewicz, U., Caldeira, K., Matsumoto, K., Munhoven, G., Montenegro, A., and Tokos, K. (2009). Atmospheric Lifetime of Fossil Fuel Carbon Dioxide. *Annual Review of Earth and Planetary Sciences*, 37:117–134.

Arkolakis, C. and Walsh, C. (2023). Clean Growth. Working Paper 31615, National Bureau of Economic Research.

Arrow, K. (1962). Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity*, pages 609–626. Princeton University Press.

Artuç, E., Chaudhuri, S., and McLaren, J. (2010). Trade Shocks and Labor Adjustment: A Structural Empirical Approach. *American Economic Review*, 100(3):1008–1045.

Asdrubali, P., Sørensen, B., and Yosha, O. (1996). Channels of Interstate Risk Sharing: United States 1963–1990. *The Quarterly Journal of Economics*, 111(4):1081–1110.

Atkinson, A. and Stiglitz, J. (1976). The Design of Tax Structure: Direct Versus Indirect Taxation. *Journal of Public Economics*, 6(1-2):55–75.

Autor, D., Dorn, D., and Hanson, G. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6):2121–2168.

Bailey, M. and Danziger, S. (2013). *Legacies of the War on Poverty*. Russell Sage Foundation.

Baily, M. (1978). Some Aspects of Optimal Unemployment Insurance. *Journal of Public Economics*, 10(3):379–402.

Balboni, C. (2019). In Harm's Way? Infrastructure Investments and the Persistence of Coastal Cities. *Working Paper*.

Balke, N. and Lamadon, T. (2022). Productivity Shocks, Long-Term Contracts, and Earnings Dynamics. *American Economic Review*, 112(7):2139–77.

Barrage, L. (2018). Be Careful What You Calibrate For: Social Discounting in General Equilibrium. *Journal of Public Economics*, 160:33–49.

Barrage, L. (2020). Optimal Dynamic Carbon Taxes in a Climate–Economy Model with Distortionary Fiscal Policy. *The Review of Economic Studies*, 87(1):1–39.

Barrage, L. and Nordhaus, W. (2024). Policies, Projections, and the Social Cost of Carbon: Results from the DICE-2023 Model. *Proceedings of the National Academy of Sciences*, 121(13).

Bartik, T. (2020). Using Place-Based Jobs Policies to Help Distressed Communities. *Journal of Economic Perspectives*, 34(3):99–127.

Bhagwati, J. (1958). Immiserizing Growth: A Geometrical Note. *The Review of Economic Studies*, 25(3):201–205.

Bhagwati, J. (1968). Distortions and Immiserizing Growth: A Generalization. *The Review of Economic Studies*, 35(4):481–485.

Bilal, A. and Rossi-Hansberg, E. (2023). Anticipating Climate Change Across the United States. Working Paper 31323, National Bureau of Economic Research.

Bloom, N., Griffith, R., and Van Reenen, J. (2002). Do R&D Tax Credits Work? Evidence from a Panel of Countries 1979–1997. *Journal of Public Economics*, 85(1):1–31.

Bloom, N., Jones, C., Van Reenen, J., and Webb, M. (2020). Are Ideas Getting Harder to Find? *American Economic Review*, 110(4):1104–1144.

Bloom, N., Romer, P., Terry, S., and Van Reenen, J. (2021). Trapped Factors and China's Impact on Global Growth. *The Economic Journal*, 131(633):156–191.

Bloom, N., Schankerman, M., and Van Reenen, J. (2013). Identifying Technology Spillovers and Product Market Rivalry. *Econometrica*, 81(4):1347–1393.

Bloom, N., Van Reenen, J., and Williams, H. (2019). A Toolkit of Policies to Promote Innovation. *Journal of Economic Perspectives*, 33(3):163–84.

Blundell, R., Griffith, R., and Windmeijer, F. (2002). Individual Effects and Dynamics in Count Data Models. *Journal of Econometrics*, 108(1):113–131.

Böhringer, C. and Rutherford, T. (2009). Integrated Assessment of Energy Policies: Decomposing Top-Down and Bottom-Up. *Journal of Economic Dynamics and Control*, 33(9):1648–1661.

BP (2022). BP Statistical Review of World Energy, 71th Edition. Technical report.

Bryan, K. and Williams, H. (2021). Innovation: Market Failures and Public Policies. In *Handbook of Industrial Organization*, volume 5, pages 281–388. Elsevier.

Budish, E., Roin, B., and Williams, H. (2015). Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials. *American Economic Review*, 105(7):2044–85.

Burdett, K. and Coles, M. (2003). Equilibrium Wage-Tenure Contracts. *Econometrica*, 71(5):1377–1404.

Cai, S., Caliendo, L., Parro, F., and Xiang, W. (2022). Mechanics of Spatial Growth. Working Paper 30579, National Bureau of Economic Research.

Cai, Y. and Lontzek, T. (2019). The Social Cost of Carbon with Economic and Climate Risks. *Journal of Political Economy*, 127(6):2684–2734.

Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and Labor Market Dynamics: General Equilibrium Analysis of the China Trade Shock. *Econometrica*, 87(3):741–835.

Cameron, S., Chaudhuri, S., and McLaren, J. (2007). Trade Shocks and Labor Adjustment: Theory. Working Paper 13463, National Bureau of Economic Research.

Chamley, C. (1986). Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives. *Econometrica*, 54(3):607–622.

Chari, V. and Kehoe, P. (1990). Sustainable Plans. *Journal of Political Economy*, 98(4):783–802.

Chetty, R. (2006). A General Formula for the Optimal Level of Social Insurance. *Journal of Public Economics*, 90(10-11):1879–1901.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review*, 101(3):471–475.

Costinot, A. and Werning, I. (2023). Robots, Trade, and Luddism: A Sufficient Statistic Approach to Optimal Technology Regulation. *The Review of Economic Studies*, 90(5):2261–2291.

Cruz, J.-L. (2021). Global Warming and Labor Market Reallocation. *Working Paper*.

Cruz, J.-L. and Rossi-Hansberg, E. (2023). The Economic Geography of Global Warming. *The Review of Economic Studies*, 91(2):899–939.

Dasgupta, P. (2008). Discounting Climate Change. *Journal of Risk and Uncertainty*, 37:141–169.

Davis, M. and Gregory, J. (2021). Place-Based Redistribution in Location Choice Models. Working Paper 29045, National Bureau of Economic Research.

Davis, S. and Boundy, R. (2022). Transportation Energy Data Book: Edition 40. Technical report, U.S. Department of Energy, Energy Vehicle Technologies Office, Oak Ridge National Laboratory.

Desmet, K., Nagy, D. K., and Rossi-Hansberg, E. (2018). The Geography of Development. *Journal of Political Economy*, 126(3):903–983.

Diamond, P. and Mirrlees, J. (1971a). Optimal Taxation and Public Production I: Production Efficiency. *The American Economic Review*, 61(1):8–27.

Diamond, P. and Mirrlees, J. (1971b). Optimal Taxation and Public Production II: Tax Rules. *The American Economic Review*, 61(3):261–278.

Diamond, R. and Moretti, E. (2021). Where is Standard of Living the Highest? Local Prices and the Geography of Consumption. Working Paper 29533, National Bureau of Economic Research.

Domar, E. (1961). On the Measurement of Technological Change. *The Economic Journal*, 71(284):709–729.

Dubé, J.-P., Joo, J., and Kim, K. (2022). Discrete-Choice Models and Representative Consumer Theory. Working Paper 30130, National Bureau of Economic Research.

Dvorkin, M. (2023). Heterogeneous Agents Dynamic Spatial General Equilibrium. *Working Paper*.

Fajgelbaum, P. and Gaubert, C. (2020). Optimal Spatial Policies, Geography, and Sorting. *The Quarterly Journal of Economics*, 135(2):959–1036.

Farhi, E., Sleet, C., Werning, I., and Yeltekin, S. (2012). Non-Linear Capital Taxation Without Commitment. *The Review of Economic Studies*, 79(4):1469–1493.

Farhi, E. and Werning, I. (2007). Inequality and Social Discounting. *Journal of Political Economy*, 115(3):365–402.

Frey, C. B. (2019). *The Technology Trap: Capital, Labor, and Power in the Age of Automation*. Princeton University Press.

Fried, S. (2018). Climate Policy and Innovation: A Quantitative Macroeconomic Analysis. *American Economic Journal: Macroeconomics*, 10(1):90–118.

Galor, O. (2007). *Discrete Dynamical Systems*. Springer.

Gans, J. (2012). Innovation and Climate Change Policy. *American Economic Journal: Economic Policy*, 4(4):125–145.

Gaubert, C., Kline, P., Vergara, D., and Yagan, D. (2021a). Trends in US Spatial Inequality: Concentrating Affluence and a Democratization of Poverty. *AEA Papers and Proceedings*, 111:520–525.

Gaubert, C., Kline, P., and Yagan, D. (2021b). Place-Based Redistribution. Working Paper 28337, National Bureau of Economic Research.

Giannone, E., Li, Q., Paixão, N., and Pang, X. (2023a). Unpacking Moving. *Working Paper*.

Giannone, E., Miyauchi, Y., Paixão, N., Pang, X., and Suzuki, Y. (2023b). Living in a Ghost Town: The Geography of Depopulation and Aging. *Working Paper*.

Gillingham, K., Nordhaus, W., Anthoff, D., Blanford, G., Bosetti, V., Christensen, P., McJeon, H., and Reilly, J. (2018). Modeling Uncertainty in Integrated Assessment of Climate Change: A Multimodel Comparison. *Journal of the Association of Environmental and Resource Economists*, 5(4):791–826.

Glaeser, E. and Gottlieb, J. (2008). The Economics of Place-Making Policies. Working Paper 14373, National Bureau of Economic Research.

Glaeser, E. and Gottlieb, J. (2009). The Wealth of Cities: Agglomeration Economies and Spatial Equilibrium in the United States. *Journal of Economic Literature*, 47(4):983–1028.

Glaeser, E. and Luttmer, E. (2003). The Misallocation of Housing Under Rent Control. *American Economic Review*, 93(4):1027–1046.

Golosov, M., Hassler, J., Krusell, P., and Tsyvinski, A. (2014). Optimal Taxes on Fossil Fuel in General Equilibrium. *Econometrica*, 82(1):41–88.

Greaney, B. (2023). Homeownership and the Distributional Effects of Uneven Regional Growth. *Working Paper*.

Griliches, Z. (1990). Patent Statistics as Economic Indicators: A Survey. *Journal of Economic Literature*, 28(4):1661–1707.

Grossman, G. and Helpman, E. (1991). Quality Ladders in the Theory of Growth. *The Review of Economic Studies*, 58(1):43–61.

Guerreiro, J., Rebelo, S., and Teles, P. (2022). Should Robots be Taxed? *The Review of Economic Studies*, 89(1):279–311.

Hall, B. (1993). R&D Tax Policy during the 1980s: Success or Failure? *Tax Policy and the Economy*, 7:1–35.

Hall, B., Jaffe, A., and Trajtenberg, M. (2005). Market Value and Patent Citations. *The RAND Journal of Economics*, 36(1):16–38.

Hall, B. and Van Reenen, J. (2000). How Effective are Fiscal Incentives for R&D? A Review of the Evidence. *Research policy*, 29(4-5):449–469.

Hall, B. and Ziedonis, R. (2001). The Patent Paradox Revisited: An Empirical Study of Patenting in the U.S. Semiconductor Industry, 1979-1995. *The RAND Journal of Economics*, 32(1):101–128.

Hanlon, W. W. (2015). Necessity Is the Mother of Invention: Input Supplies and Directed Technical Change. *Econometrica*, 83(1):67–100.

Hanson, G. (2023). Local Labor Market Impacts of the Energy Transition: Prospects and Policies. Working Paper 30871, National Bureau of Economic Research.

Harberger, A. (1954). Monopoly and Resource Allocation. *The American Economic Review*, 44(2):77–87.

Hassler, J., Krusell, P., and Olovsson, C. (2021). Directed Technical Change as a Response to Natural Resource Scarcity. *Journal of Political Economy*, 129(11):3039–3072.

Hazell, J., Herreño, J., Nakamura, E., and Steinsson, J. (2022). The Slope of the Phillips Curve: Evidence from U.S. States. *The Quarterly Journal of Economics*, 137(3):1299–1344.

Hicks, J. (1932). *The Theory of Wages*. London: Macmillan Company & Co.

Howard, G. and Shao, H. (2022). Internal Migration and the Microfoundations of Gravity. *Working Paper*.

Hubmer, J. (2023). The Race Between Preferences and Technology. *Econometrica*, 91(1):227–261.

Hulten, C. (1978). Growth Accounting with Intermediate Inputs. *The Review of Economic Studies*, 45(3):511–518.

Jaffe, A., Trajtenberg, M., and Henderson, R. (1993). Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations. *The Quarterly Journal of Economics*, 108(3):577–598.

Jakobsen, K., Jakobsen, K., Kleven, H., and Zucman, G. (2020). Wealth taxation and Wealth Accumulation: Theory and Evidence from Denmark. *The Quarterly Journal of Economics*, 135(1):329–388.

Jones, B. and Summers, L. (2021). A Calculation of the Social Returns to Innovation. In *Innovation and Public Policy*, pages 13–60. University of Chicago Press.

Jones, C. (1995). R&D-Based Models of Economic Growth. *Journal of Political Economy*, 103(4):759–784.

Jones, C. (2022). The Past and Future of Economic Growth: A Semi-Endogenous Perspective. *Annual Review of Economics*, 14:125–152.

Jones, C. and Vollrath, D. (2013). *Introduction to Economic Growth*. W.W. Norton.

Jones, C. and Williams, J. (1998). Measuring the Social Return to R&D. *The Quarterly Journal of Economics*, 113(4):1119–1135.

Karabarbounis, L. and Neiman, B. (2014). The Global Decline of the Labor Share. *The Quarterly Journal of Economics*, 129(1):61–103.

Kennan, J. and Walker, J. (2011). The Effect of Expected Income on Individual Migration Decisions. *Econometrica*, 79(1):211–251.

Klein, P., Krusell, P., and Rios-Rull, J.-V. (2008). Time-Consistent Public Policy. *The Review of Economic Studies*, 75(3):789–808.

Kleinman, B., Liu, E., and Redding, S. (2023). Dynamic Spatial General Equilibrium. *Econometrica*, 91(2):385–424.

Kline, P. and Moretti, E. (2014a). Local Economic Development, Agglomeration Economies, and the Big Push: 100 Years of Evidence from the Tennessee Valley Authority. *The Quarterly Journal of Economics*, 129(1):275–331.

Kline, P. and Moretti, E. (2014b). People, Places, and Public Policy: Some Simple Welfare Economics of Local Economic Development Programs. *Annual Review of Economics*, 6(1):629–662.

Krusell, P. and Smith, A. (2022). Climate Change Around the World. Working Paper 30338, National Bureau of Economic Research.

Kydland, F. and Prescott, E. (1977). Rules Rather than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy*, 85(3):473–491.

Känzig, D. (2023). The Unequal Economic Consequences of Carbon Pricing. Working Paper 31221, National Bureau of Economic Research.

Lanzi, E. and Sue Wing, I. (2011). Directed Technical Change in the Energy Sector: An Empirical Test of Induced Directed Innovation. *Working Paper*.

Lanzi, E., Verdolini, E., and Haščič, I. (2011). Efficiency-Improving Fossil Fuel Technologies for Electricity Generation: Data Selection and Trends. *Energy Policy*, 39(11):7000–7014.

Lemoine, D. and Traeger, C. (2014). Watch Your Step: Optimal Policy in a Tipping Climate. *American Economic Journal: Economic Policy*, 6(1):137–166.

Liu, E. and Ma, S. (2021). Innovation Networks and R&D Allocation. Working Paper 29607, National Bureau of Economic Research.

Lucas, R. (2003). Macroeconomic Priorities. *American Economic Review*, 93(1):1–14.

Melo, P., Graham, D., and Noland, R. (2009). A Meta-Analysis of Estimates of Urban Agglomeration Economies. *Regional Science and Urban Economics*, 39(3):332–342.

Mirrlees, J. (1972). The Optimum Town. *The Swedish Journal of Economics*, pages 114–135.

Moretti, E. (2012). *The New Geography of Jobs*. Houghton Mifflin Harcourt.

Moscona, J. and Sastry, K. (2022). Inappropriate Technology: Evidence from Global Agriculture. *Working Paper*.

Myers, K. and Lanahan, L. (2022). Estimating Spillovers from Publicly Funded R&D: Evidence from the US Department of Energy. *American Economic Review*, 112(7):2393–2423.

Naito, H. (1999). Re-Examination of Uniform Commodity Taxes under a Non-Linear Income Tax System and its Implication for Production Efficiency. *Journal of Public Economics*, 71(2):165–188.

Newell, R., Jaffe, A., and Stavins, R. (1999). The Induced Innovation Hypothesis and Energy-Saving Technological Change. *The Quarterly Journal of Economics*, 114(3):941–975.

Nordhaus, W. (1992). An Optimal Transition Path for Controlling Greenhouse Gases. *Science*, 258(5086):1315–1319.

Nordhaus, W. (2002). Modeling Induced Innovation in Climate-Change Policy. In *Technological Change and the Environment*, pages 182–209. Routledge.

Nordhaus, W. (2007). A Review of the Stern Review on the Economics of Climate Change. *Journal of Economic Literature*, 45(3):686–702.

Nordhaus, W. (2010). Economic Aspects of Global Warming in a Post-Copenhagen Environment. *Proceedings of the National Academy of Sciences*, 107(26):11721–11726.

Nordhaus, W. (2017). Revisiting the Social Cost of Carbon. *Proceedings of the National Academy of Sciences*, 114(7):1518–1523.

Oberfield, E. and Raval, D. (2021). Micro Data and Macro Technology. *Econometrica*, 89(2):703–732.

O'Hara, F. (1990). *Glossary: Carbon Dioxide and Climate*. Oak Ridge National Laboratory.

Ord, T. (2020). *The Precipice: Existential Risk and the Future of Humanity*. Hachette Books.

Pang, X. and Sun, P. (2023). Moving into Risky Floodplains: The Spatial Implications of Flood Relief Policies. *Working Paper*.

Papageorgiou, C., Saam, M., and Schulte, P. (2017). Substitution Between Clean and Dirty Energy Inputs: A Macroeconomic Perspective. *The Review of Economics and Statistics*, 99(2):281–290.

Parfit, D. (1984). *Reasons and Persons*. Oxford University Press.

Peters, M. (2022). Market Size and Spatial Growth—Evidence From Germany's Post-War Population Expulsions. *Econometrica*, 90(5):2357–2396.

Popp, D. (2002). Induced Innovation and Energy Prices. *American Economic Review*, 92(1):160–180.

Popp, D. (2004). ENTICE: Endogenous Technological Change in the DICE Model of Global Warming. *Journal of Environmental Economics and Management*, 48(1):742–768.

Romer, P. (1990). Endogenous Technological Change. *Journal of Political Economy*, 98(5, Part 2):S71–S102.

Roser, M. (2020). Why Did Renewables Become so Cheap so Fast? Our World in Data.

Rossi-Hansberg, E., Sarte, P.-D., and Schwartzman, F. (2019). Cognitive Hubs and Spatial Redistribution. Working Paper 26267, National Bureau of Economic Research.

Saez, E. and Zucman, G. (2016). Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data. *The Quarterly Journal of Economics*, 131(2):519–578.

Stern, N. (2007). *The Economics of Climate Change: The Stern Review*. Cambridge University Press.

Straub, L. and Werning, I. (2020). Positive Long-Run Capital Taxation: Chamley-Judd Revisited. *American Economic Review*, 110(1):86–119.

Thuemmel, U. (2022). Optimal Taxation of Robots. *Journal of the European Economic Association*.

Van der Werf, E. (2008). Production Functions for Climate Policy Modeling: An Empirical Analysis. *Energy Economics*, 30(6):2964–2979.

Wagner, G. and Weitzman, M. (2016). *Climate Shock: The Economic Consequences of a Hotter Planet*. Princeton University Press.

Webb, M. (2019). The Impact of Artificial Intelligence on the Labor Market. *Working Paper*.

Weitzman, M. (2009). On Modeling and Interpreting the Economics of Catastrophic Climate Change. *The Review of Economics and Statistics*, 91(1):1–19.

Wilson, D. (2009). Beggar Thy Neighbor? The In-State, Out-of-State, and Aggregate Effects of R&D Tax Credits. *The Review of Economics and Statistics*, 91(2):431–436.

CURRICULUM VITAE

[redacted]

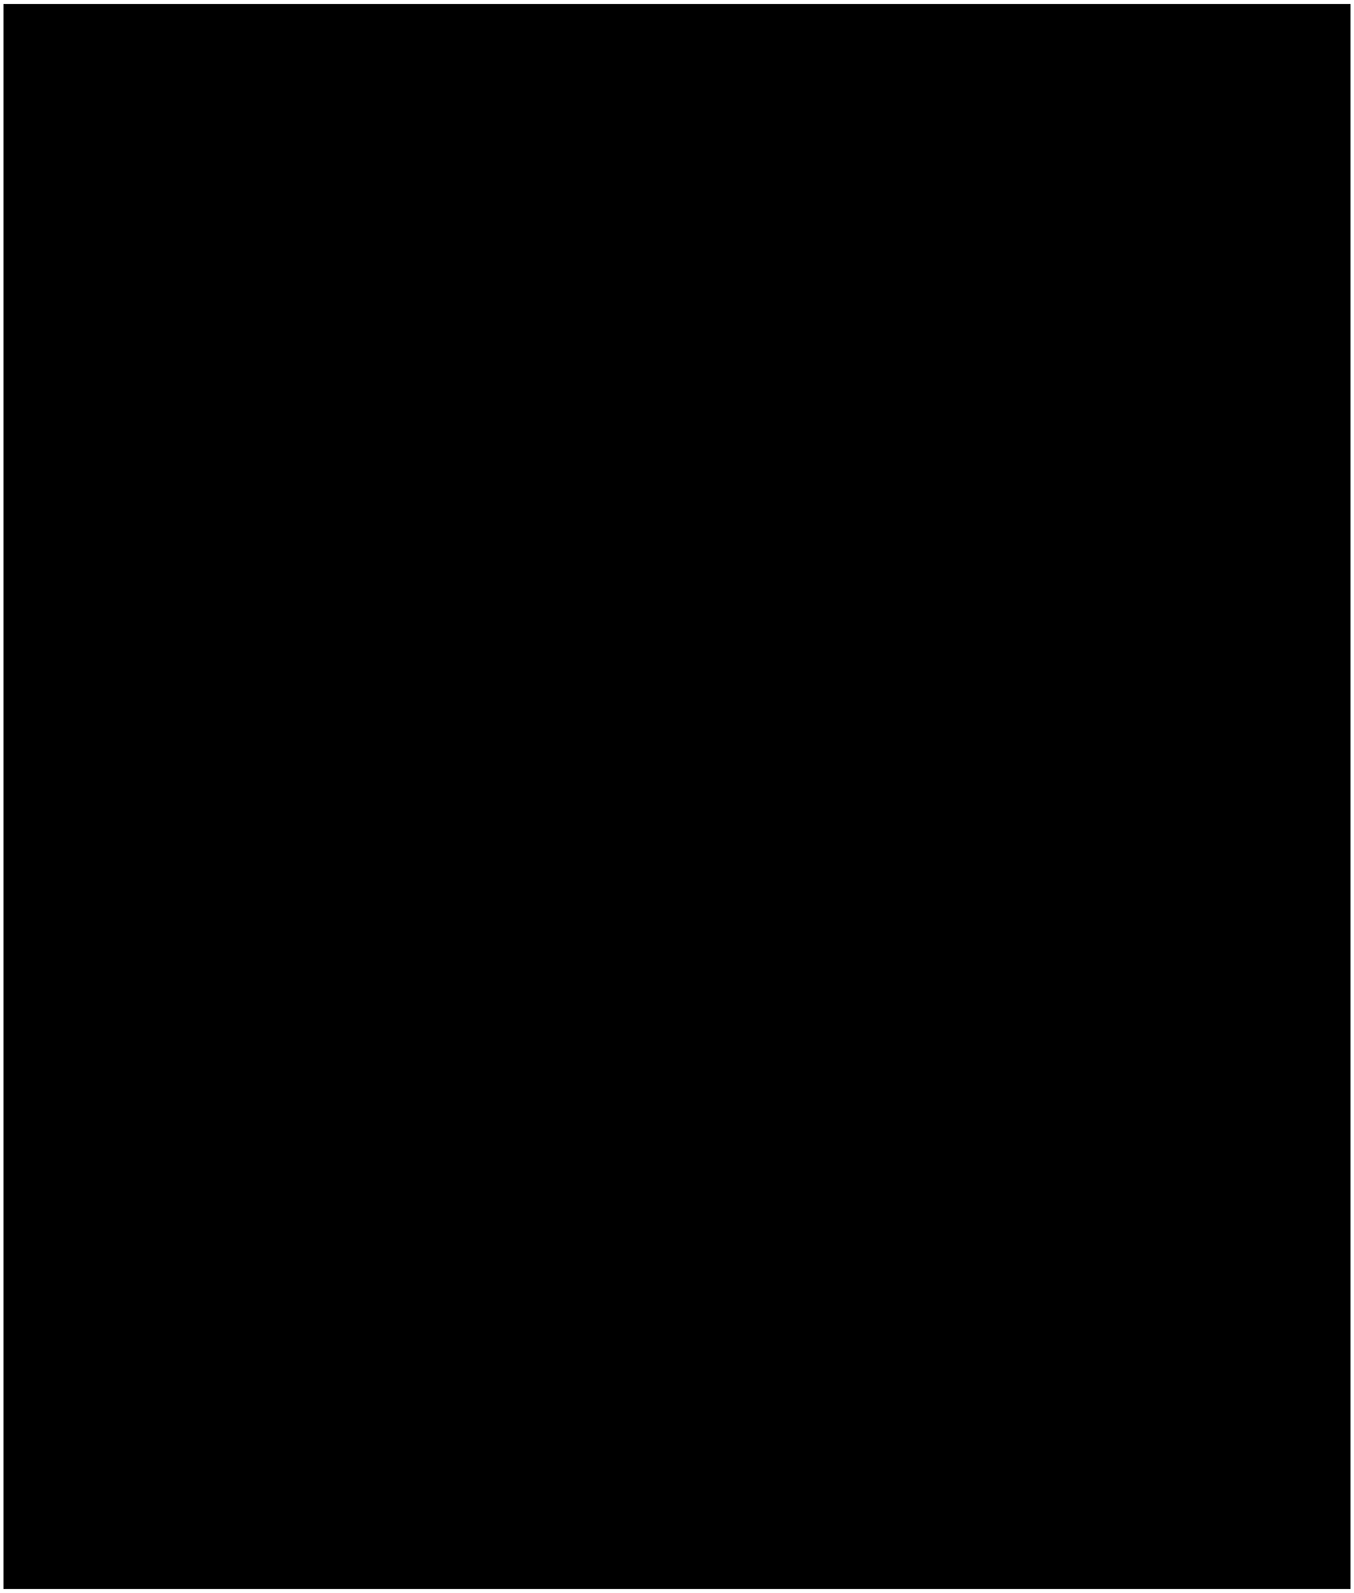

[redacted]