Green Transitions Into Alternate Histories

Konstantin Kucheryavyy*
Baruch College

Jeffrey E. Sun†
University of Toronto

May 2, 2025

Preliminary and incomplete.

Abstract

We study green transitions in settings where "greener" technologies already exist but have not been optimally adopted due to path-dependence and adjustment costs. We study this through the lens of dynamic models with multiple equilibria and adjustment costs, and investigate the design of a "transition policy" designed to shift the economy from one equilibrium to another. We illustrate using the example of rail and automobiles in the United States. The marginal cost of new rail projects often outweighs their marginal benefit, due to greater historical investment in, or "agglomeration into," automobile infrastructure. We model agglomeration into sectors using an approach similar to spatial models of agglomeration into locations. A hypothetical policy designed to shift the economy from the automobile equilibrium to the rail equilibrium represents an optimal policy problem in multiple time-varying instruments within a dynamic model. We propose a geometric approach to this class of problems: the state space is equipped with a Riemannian metric that captures the costs of transitioning between states, such that the optimal policy arises as a geodesic. The solution algorithm is surprisingly simple and is robust to different assumptions about household expectations, including fully myopic or forward-looking and in between.

*konstantin.kucheryavyy@baruch.cuny.edu
†jesun@princeton.edu

The ruler is like a mirror, reflecting light, doing nothing, and yet, beauty and ugliness present themselves; a scale establishing equilibrium.

–Shen Buhai, ∼351 BCE

1 Introduction

The recent expansion of interest in policies aimed at encouraging or inducing "green transitions" suggests that policymakers believe that these transitions will not arise adequately in the competitive equilibrium. The reason for this is potentially more subtle than the mere existence of externalities. Rather, it suggests that policies such as subsidies may not even have positive net *marginal* benefit, though a sufficiently large, sustained policy may be able to enact a shift to a new equilibrium in which further green investments are marginally beneficial in some sense.

Examples abound in public policy and the economics literature. Network effects, learning by doing, and increasing returns to scale can be seen as reasons why large and unprofitable investments in the present can "crowd-in" profitable investments from a competitive private sector. The recent large-scale investment in electric vehicle charging stations in Quebec, for example, arguably induced many households to adopt electric vehicles, which has been followed by an increase in the construction of new charging stations by the private sector. Economically, these policies are intriguing in that they suggest a role for moving policy analysis "away from the margin," whereby a policy with negative marginal benefit can, if it is large enough, have positive overall benefit.

In this paper, we therefore study that aspect of green transition policy which is orthogonal to technological change *per se*. That is, we focus on the setting where "greener" alternatives to existing technologies already exist, but have not been optimally adopted. Our illustrative example is that of rail versus automobile transportation in the United States. These two sectors are substitutable to some degree, but the marginal benefits of new investments depend strongly on the stock of existing infrastructure.

We provide two historical arguments for path-dependence in the rail and automobile sectors. First, rail technology is well-developed and in many countries much more highly utilized than in the United States. In Russia, for instance, passenger rail usage is 1,220 passenger-kilometers per capita, compared with 80 for the United States. Second, passenger rail was once highly utilized in the United States, reaching a high of 1,046 passenger-kilometers per capita in 1945, and its decline is often attributed,[1] at least in part, to lobbying efforts by the automobile industry, leading to an antitrust conviction in 1949[2] and another congressional

[1] See, *inter alia*, https://web.archive.org/web/20071120111335/http://www.worldcarfree.net/resources/freesources/American.htm

[2] https://casetext.com/case/united-states-v-national-city-lines-4

antitrust hearing in 1974.[3] Bringing the model to this example, we seek to examine the counterfactual path of passenger rail utilization in the absence of these lobbying efforts; the potential costs and benefits of effecting, in the present, a transition to a potential rail equilibrium; the optimal design of such a policy; and how the public's expectations affect the optimal policy design.

To study this type of policy, we use a model in which external returns to scale, or "agglomeration effects," within industries lead to multiple equilibria, which together with adjustment costs generate path-dependence. The model is a relatively simple combination of external economics-of-scale production, costly reallocation of factors between sectors, and nested CES aggregation over sectors. The main methodological contribution is in the focus on the optimal policy problem of setting optimal *time-varying* sector-level subsidies and the introduction of a novel solution method for this problem.

The search space for this "optimal policy path" is large. Each policy instrument, a sector-level subsidy, is not a static value but a function of time. As time is continuous, the search space is infinite-dimensional. Furthermore, there are multiple policy instruments to be set jointly. Finally, the joint policy must take into account intertemporal equilibrium effects, as forward-looking agents may choose their actions in the present in anticipation of future policy changes. That is, the optimal policy, a many-valued function of time, is defined with respect to an infinite-dimensional operator mapping the policy function to the resulting equilibrium transition path.

The problem of planning optimal transition paths between different states of the world involves trading off (1) between the speed of the transition and the costs of adjustment; as well as possibly (2) between the "shortest path" between states and more roundabout paths which pass through more favorable intermediate states. The optimal policy path must take into account equilibrium effects, expectations, and both of these trade-offs. Our solution method achieves this without recourse to functional analysis and reduces the dimension of the search space from infinity to two.

Three transformations reduce the complexity of the problem. First, we isolate a set of sufficient statistics for the social welfare function, which allows us to search over the path of these sufficient statistics directly, only later backing out the policy path necessary to induce it. This is analogous to solving for the first-best allocation in a static model. (In the case where the first-best allocation is not feasible, this first transformation is not possible, but the following steps are still feasible.) Our model admits a two-dimensional sufficient statistic, allowing us to effectively deal with a two-dimensional state space. This allows us to separate the problem of finding the optimal transition from the problem of finding the policy which induces this transition.

Second, we equip the state space with a Riemannian metric which "reverse-engineers" the social planner's objective—an integral over time of within-period social welfare or "felicity"—in the sense that the length of a path under this metric is equal to the social planner's objective function. This allows us to separate the

[3] https://libraryarchives.metro.net/DPGTL/testimony/1974_statement_bradford_c_snell_s1167.pdf

problem of finding the optimal set of intermediate states from the problem of finding the optimal *speed* with which to traverse those states.

Third, we use the property of Riemanniann manifolds that the shortest path between two points is a geodesic, or a straight line with respect to the metric. That is, in order to find this geodesic, it suffices to find the initial velocity of the path. The geodesic equations are then used to extend the path forward "in a straight line," with respect to the metric. The goal is to find such a line which terminates in the desired state, for which a simple shooting algorithm is sufficient and standard. In a two-dimensional state space, this initial velocity is a two-dimensional vector, reducing the dimension of the search space to two.

The baseline model has only four key elasticities: the strength of agglomeration effects within a sector, the elasticity of substitution between the rail and automobile sectors, the elasticity of substitution between the aggregate "transportation sector" and the rest of the economy, and the adjustment cost of reallocating factors between sectors. In addition to these, the model can be generalized in various directions with the addition of only one parameter, for example, the degree to which households are forward-looking, the degree to which adjustment costs are external, and the degree to which agglomeration effects are external. In this way, the model can be thought of as a complement to empirical work—specifying and precisely defining key elasticities which can yield clean results, and a way of interpreting those elasticities in general, without imposing a one-size-fits-all estimation or identification strategy.

This paper proceeds as follows. In Section 3, we describe a growth model with agglomeration affects, adjustment costs, and possibly multiple equilibria and local maxima. In Section 4, we apply the model to the example of rail and automobile infrastructure in the United States. In Section 5, we present the solution method, decomposing the problem into primal and Pigouvian problems and recasting the primal as a geodesic problem in a specially-designed Riemannian manifold. In Section 6, we present the results of the model and the optimal policy path. Section 7 concludes.

2 Related Literature

Our paper contributes to the literature on economic policies designed to guide the economy toward a desired state. Much of the existing research in this field focuses on static environments. For instance, the literature on international- trade and industrial policies often examines how a social planner can achieve the economy's social optimum (see Caliendo and Parro, 2022, for a survey). Theoretical derivations of government policies under simplifying assumptions dominate this body of work (e.g., Bartelme et al., forthcoming; Lashkaripour and Lugovskyy, 2023). Exceptions include Ossa [2014], which employs mathematical programming with equilibrium constraints (as developed by Su and Judd, 2012) to numerically solve for optimal policies, and Wang et al. [2024], which leverages machine learning methods to solve for optimal trade and industrial policies.

A key limitation of this literature is its focus on economies with a unique equilibrium. In contrast, the tools we develop in this paper are applicable to environments with either unique or multiple equilibria. For example, our primary case study — investment in cars versus railways — features multiple equilibria, where the planner seeks to transition from a “bad” car equilibrium to a “good” railways equilibrium.

Our paper also relates to the literature on optimal government policies in dynamic environments. Many studies in this area aim to identify policy-induced trajectories of economic variables that deviate the least from a “desired” path (e.g., Benigno and Woodford, 2012; McKay and Wolf, 2023). Unlike these approaches, we focus on identifying an optimal policy path that explicitly transitions the economy to a desired state. Additionally, while much of this literature approximates the problem using linear-quadratic programming, we consider a more general non-linear framework.

This paper also relates to the macroeconomic literature that employs the shooting algorithm (e.g., see Chapter 11 in Ljungqvist and Sargent, 2018) — the primary numerical method we use to solve for the geodesic that determines the optimal transition path. In this literature, government policies are typically exogenously specified, and the task is to identify the economy’s transition path to a new steady state. By contrast, our focus is on determining the optimal transition path that actively guides the economy to a desired state.

Regarding the economic impact of railroads in the United States, our study is closely related to the literature that examines the historical and economic significance of rail infrastructure. For example, Donaldson and Hornbeck [2016] explore the transformative economic effects of railroads on the 19th-century U.S. economy.

On the historical development of U.S. railroads, Gallamore and Meyer [2014] provide an insightful account of 20th-century American railroads. They highlight how government policy critically shaped the trajectory of the rail industry, discussing the detrimental effects of excessive regulation and the subsequent renaissance spurred by deregulation in the 1970s.

3 Model

We describe a dynamic growth model of an economy with three sectors which must split between them a stock of scarce factors. The total stock of scarce factors grows at a fixed and exogenous rate and reallocating factors across sectors incurs adjustment costs. Each sector features increasing returns to scale.

For some parameter values, the model features multiple equilibrium balanced growth paths (BGPs) and/or multiple local maxima of the social welfare function. This generates path-dependence: depending on the initial state of the economy, the BGP to which the economy converges, either in the absence of policy or under a short-sighted social planner, may depend on the initial state.

The goal of this exercise is to study how revenue-neutral time-varying taxes and subsidies can induce a transition from one BGP to another.

In Section 4, we will apply this model to the case of automobile and rail infrastructure in the United States. The motivating feature of this historical experience is the possibility that the automobile-dominant BGP arose due to a policy preference for automobiles, despite this equilibrium being (1) potentially not unique as a BGP, and (2) potentially not statically preferable to a hypothetical rail-dominant equilibrium, in light of our current understanding of the costs of carbon emissions.

3.1 Environment

Consider a closed economy with $K = 3$ sectors in continuous time. In Section 4 we will label sectors 1, 2, and 3 as "rail," "auto," and "everything else," respectively.

At each time t, the economy has access to a finite stock of scarce factors whose total supply grows at an exogenous rate $g > 0$. We will refer to this stock of scarce factors as "capital,"[4] normalize its total supply to 1, and denote each sector's share of this stock as $K_{k,t}$, with,

$$\sum_k K_{k,t} = 1 \quad \forall t.$$

3.2 Production

Each sector k is perfectly competitive and admits a representative firm which produces using this "capital" alone, subject to purely external economies of scale and adjustment costs,[5]

$$\begin{aligned} Q_{k,t} &= \min\left\{A_{k,t}K_{k,t},\ A_{k,t}K_{k,t} + I_{k,t} - \delta_k \dot{K}_{k,t}^2\right\} \\ &= A_{k,t}K_{k,t} - \max\left\{0,\ \delta_k \dot{K}_{k,t}^2 - I_{k,t}\right\}, \end{aligned}$$

where $K_{k,t}$ is the share of capital used by Sector k at time t, $A_{k,t}$ is productivity in sector k, $\dot{K}_{k,t} \equiv \frac{dK_{k,t}}{dt}$ is the rate of change of $K_{k,t}$, and $I_{k,t}$ is the convex cost of adjustment at rate $\dot{K}_{k,t}$.

Sector productivity $A_{k,t}$ is taken as given by the representative firm, but is an endogenous object given by,

$$A_{k,t} = \overline{A}_{k,t} K_{k,t}^{\gamma_k},$$

where $\overline{A}_{k,t}$ is an exogenous productivity shifter, and γ_k governs the strength of economies of scale.

[4]This model's production technology can be thought of as representing production using "equipped capital," and is isomorphic to a model with Cobb-Douglas production in capital and homogeneous, mobile labor. Or, similarly, production using land only, with homogeneous, mobile labor and capital.

[5]The assumptions of perfect competition and external economies of scale are not necessary to solve the model, and we are currently working to provide a solution when weakening this assumption.

The final good is produced by first combining the products of Sectors 1 and 2 using a CES technology, then combining the result with the product of Sector 3 using another CES technology, according to,

$$Q_t = \left[Q_{12,t}^{\frac{\eta-1}{\eta}} + b_{3,t}^{\frac{1}{\eta}} Q_{3,t}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

$$Q_{12,t} = \left[b_{1,t}^{\frac{1}{\sigma}} Q_{1,t}^{\frac{\sigma-1}{\sigma}} + b_{2,t}^{\frac{1}{\sigma}} Q_{2,t}^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}.$$

Here $b_{k,t}$ are exogenous shifters.

In Section 4, we will have $\sigma > 1$ and $\eta < 1$, such that the rail and auto sectors are imperfect substitutes and the combined "transportation sector" is a complement to the rest of the economy.

3.3 Households

The population admits a representative household which is hand-to-mouth in consumption,[6] owns all capital in the economy, and has linear felicity[7] over consumption,

$$U_t = C_t,$$

subject to the budget constraint,

$$P_t C_t = X_t \equiv \sum_k R_{k,t} K_{k,t},$$

where P_t is the price of the final good, X_t is the household's expenditure, and $R_{k,t}$ is the rental rate of capital in sector k. The rental rate of capital can differ across sectors due to capital adjustment costs.

3.4 Capital and Goods Markets

Denote by $P_{k,t}$ the price of sector-k good received by producers in this sector. All producers are perfectly competitive, and their profit-maximization implies

$$P_{k,t} = \frac{R_{k,t}}{\overline{A}_{k,t} K_{k,t}^{\gamma_k}}.$$

Consumer prices are related to producer prices by $\tilde{P}_{k,t} = (1 + \tau_{k,t}) P_{k,t}$, where $\tau_{k,t}$ is the government consumption tax (if $\tau_{k,t} > 0$) or subsidy (if $\tau_{k,t} \in (-1, 0)$). The overall system of taxes and subsidies is

[6]This assumption allows us to remain agnostic to household expectations, but is not necessary for tractability. We are currently working on an extension with saving in which we solve the model under various assumptions about foresight.

[7]Linear utility is currently necessary for tractability, specifically in order for time discounting and economic growth to cancel out as described in Section 3.8.

revenue neutral. Demand is given by

$$C_{1,t} = b_{1,t} \left(\frac{\tilde{P}_{1,t}}{\tilde{P}_{12,t}} \right)^{1-\sigma} \left(\frac{\tilde{P}_{12,t}}{\tilde{P}_t} \right)^{1-\eta} \frac{X_t}{\tilde{P}_{1,t}},$$
$$C_{2,t} = b_{2,t} \left(\frac{\tilde{P}_{2,t}}{\tilde{P}_{12,t}} \right)^{1-\sigma} \left(\frac{\tilde{P}_{12,t}}{\tilde{P}_t} \right)^{1-\eta} \frac{X_t}{\tilde{P}_{2,t}},$$
$$C_{3,t} = b_{3,t} \left(\frac{\tilde{P}_{3,t}}{\tilde{P}_t} \right)^{1-\eta} \frac{X_t}{\tilde{P}_{3,t}};$$

with price indices given by

$$\tilde{P}_{12,t} = \left[b_{1,t} \tilde{P}_{1,t}^{1-\sigma} + b_{2,t} \tilde{P}_{2,t}^{1-\sigma} \right]^{\frac{1}{1-\sigma}},$$
$$\tilde{P}_t = \left[\tilde{P}_{12,t}^{1-\eta} + b_{3,t} \tilde{P}_{3,t}^{1-\eta} \right]^{\frac{1}{1-\eta}}.$$

3.5 Equilibrium

In equilibrium, the following hold.

- Goods market clearing: $C_t + \sum_k I_{k,t} = Q_t$.
- Capital market clearing: $\sum_k K_{k,t} = 1$.
- Government budget: $\sum_k \tau_{k,t} P_{k,t} C_{k,t} = 0$.
- Demand and prices are given by the expressions above, consumer's budget constraint holds.
- Producer's price is given by $P_{k,t} = \dfrac{R_{k,t}}{\overline{A}_{k,t} K_{k,t}^{\gamma_k}}$.

3.6 Terminology

Let $K_t \equiv \{K_{k,t}\}_k$ and $\dot{K}_t \equiv \{\dot{K}_{k,t}\}_k$. In this economy, the household's utility at any time t depends on both K_t and $\dot{K}_t$. If we impose that adjustment costs are fully paid for,[8] $I_{k,t} = \delta_k \dot{K}_{k,t} \forall k$, then these dependencies can be separated as follows,

$$U_t(K_t, \dot{K}_t) = Q_t(K_t) - \sum_k \delta_k \dot{K}_{k,t}^2,$$

where Q_t does not depend on $\dot{K}_t$.

If we denote by $\widetilde{D}$ the matrix with,

$$\widetilde{D}_{ij} = \begin{cases} \delta_i & \text{if } i = j, \\ 0 & \text{otherwise,} \end{cases},$$

[8] It would be good to figure out how to ensure this in equilibrium rather than impose it externally.

the household's utility can be written as,

$$U_t(K_t, \dot{K}_t) = \overline{U}_t(K_t) - \dot{K}_t' \widetilde{D} \dot{K}_t,$$

where $\overline{U}_t(K_t) \equiv Q_t(K_t)$.

We refer to $\overline{U}_t(K_t)$ as the "static felicity" at time t. It is equal to total output at time t and also equal to the utility the household would receive under capital allocations K_t if capital were not moving: $\dot{K}_{k,t} = 0\ \forall k$.

3.7 Problem of Short-Sighted Social Planner

Consider a short-sighted social planner who seeks to maximize welfare in the short term. That is, as we formalize more carefully in Section 5, at each time t the short-sighted social planner takes K_t as given and chooses $\dot{\tau}_{k,t}$ to maximize $\dot{U}_t$, the rate of change of household utility.

Such a social planner, who might be viewed as an economist who makes decisions purely based on marginal costs and benefits relative to the current state, will in effect implement gradient ascent on the static felicity function $\overline{U}_t(K_t)$.

Consequently, an economy led by this planner may converge to a BGP which is locally optimal, a local maximum of $\overline{U}_t(K_t)$, but not statically preferred to another maximum. Furthermore, for some assumptions about time discounting, this local maximum may not be preferable to a transition which incurs adjustment costs and passes through less favorable states, but ultimately leads to a more favorable BGP.

3.8 Problem of Long-Term Social Planner

We now construct the problem of one such long-term social planner. From an initial period which we denote by $t = 0$, a long-term social planner chooses a path of revenue-neutral tax (and subsidy) schedules,

$$\tau_t \equiv \{\tau_{k,t}\}_{k,t},$$

to maximize the net present value of present and future household utility,[9]

$$\{\tau_t^*\}_t \equiv \arg\max_{\{\tau_t\}_t} \int_0^\infty e^{(g-\beta)t} \left(\overline{U}_t\left(K_t\right) - \dot{K}_t' \widetilde{D} \dot{K}_t \right) dt,$$

such that $\{\tau_t\}_t$ is revenue-neutral for all t, and $\{K_s, \dot{K}_s\}$ is the unique transition path resulting from an initial allocation K_{t_0} and the path of tax schedules $\{\tau_s\}_s$. Conditions for this uniqueness are given in Section 5 and are weaker than the conditions necessary for a static tax schedule τ_t to correspond to a unique steady-state allocation K_T without conditioning on initial state K_{t_0}.

[9] We abuse notation by writing the limit of the integral as infinity even when the value of the integral is potentially infinite. The more careful construction of this limit is given in Section 5.

4 Calibration

While the calibration remains a work in progress, we present illustrative results using the following parameter values:

Parameter	Value
$\overline{A}_k$	$[1.2, 1, 1]$
γ_k	$[1.9, 1.9, 0]$
η	0.01
σ	3
δ_k	$[1, 1, 1]$

Table 1: Parameter Values

Figure 1 shows the static felicity $\overline{U}$ as a function of the allocations $\{K_t\}_t$, where implicitly, $K_3 = 1 - K_1 - K_2$. Figure 2 illustrates the trade-off between a "direct" transition and other welfare (static felicity) along intermediate states. Point C lies exactly between points A and B, such that a path passing through Point C would minimize adjustment costs at the cost of static felicity. Point D represents the point of maximal static felicity along the line $K_1 = K_2$. Passing through this point would maximize the felicity of intermediate states at the cost of incurring more adjustment costs.

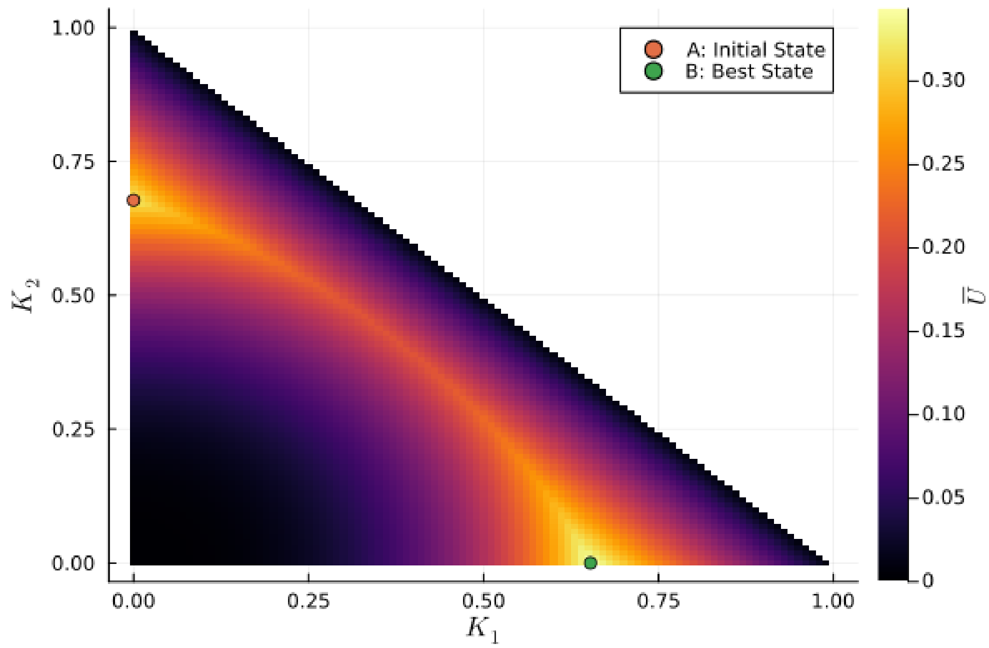

Figure 1: Static felicity as a function of the share of all scarce factors allocated to rail (K_1) and automobile (K_2) infrastructure.

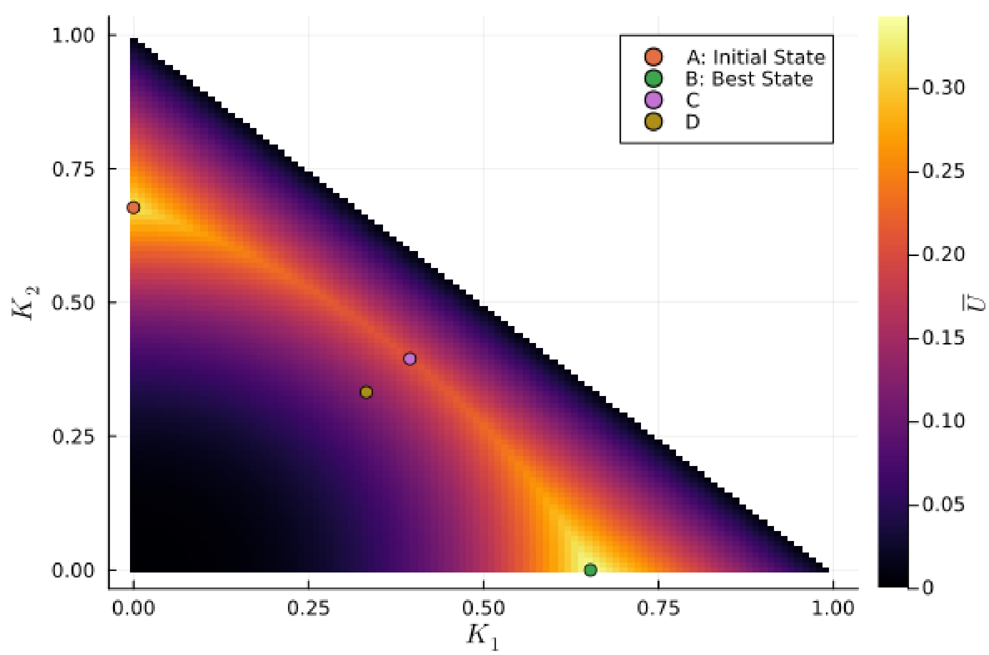

Figure 2: The "true" midpoint C of the Initial and Best States, as well as the point D with maximal static felicity along the line $K_1 = K_2$.

5 Solution Method

This section describes the solution method for the problem of the long-term social planner. This requires several steps, essentially decomposing the problem into several pieces, transforming and/or solving each piece, then recombining the pieces.

The methodological novelty is the discovery that this problem has a "core" which is geometrical in nature and which can be solved using the tools of Riemannian geometry. After properly defining the problem, we can "drill down" to this core through a sequence of transformations which try to isolate the "hard part" of the problem (the high-dimensional nonlinear optimization problem) from the low-dimensional "easy part." This core problem is then revealed to be equivalent to a geodesic problem which admits a vast dimension reduction, allowing it to be solved. The solved geodesic problem can then be lifted back up through the transformations to yield the solution to the original problem.

The structure of the solution method is as follows:

1. Properly define the infinite-horizon social planner's problem.
2. Decompose this into a "primal problem" and a "Pigouvian problem," where the "hard part" is the primal problem.
3. Decompose the primal problem into a "rate problem" (how quickly to adjust) and a "path problem" (in which direction to adjust), where the "hard part" is the path problem.
4. Transform the path problem into a geodesic problem on a certain Riemannian manifold
5. Transform the geodesic problem into an initial-velocity problem using a highly general result from differential geometry.
6. Solve the initial-velocity problem.
7. Use the solution to the initial-velocity problem to solve the geodesic problem.
8. Use the solution to the geodesic problem to solve the primal problem.
9. Use the solution to the primal problem to solve the infinite-horizon social planner's problem.

5.1 The Problem of the Infinite-Horizon Social Planner

In Section 3, we defined the long-term social planner's problem as,

$$\{\tau_t^*\}_t \equiv \arg\max_{\{\tau_t\}_t} \int_0^\infty e^{(g-\beta)t} \left(\overline{U}\left(K_t\right) - \dot{K}_t' \widetilde{D} \dot{K}_t\right) dt,$$

However, the current iteration of the solution method applies specifically to the case where $g - \beta = 0$. This causes the above integral from 0 to ∞ to diverge in general, and so we must take this limit slightly more carefully in order to ensure that the problem remains well-defined.

In particular, consider a social planner with time horizon T. We define the problem of this social planner as,

$$\tau^T \equiv \arg\max_{\{\tau_t\}_t} \int_0^T e^{(g-\beta)t} \left(\overline{U}\left(K_t\right) - \dot{K}_t' \widetilde{D} \dot{K}_t \right) dt.$$

Even though the value of the integral might diverge as $T \to \infty$, this does not necessarily imply that the τ^T diverge. In particular, though the domains of the τ^T differ, we can still consider the pointwise limit,

$$\tau_t^* = \lim_{T \to \infty} \tau_{T,t}^*.$$

Furthermore, if $\overline{U}(K)$ is bounded, then we can examine separately the case $g < \beta$ and $g \geq \beta$. If $g < \beta$, the integral is bounded as $T \to \infty$, and if $g \geq \beta$, the optimal path of allocations $\{K_t\}_t$ must reach the global optimum in finite time and stay there forever,

$$\exists K^*, \overline{T} \text{ s.t } \forall T > \overline{T}, K_T(\tau^T) = K^*,$$
$$\text{where} \quad K^* \equiv \arg\max_K \overline{U}(K).$$

Thus, the following problem is well-defined and has the same solution as the original problem,

$$\{\tau_t^*\}_t \equiv \arg\max_{\{\tau_t\}_t} \int_0^\infty e^{(g-\beta)t} \left(\overline{U}\left(K_t\right) - \dot{K}_t' \widetilde{D} \dot{K}_t - \overline{U}(K^*) \right) dt.$$

To see that it is well-defined, not that $\forall t > T$, the value of the integrand is zero, so that the value of the integral is bounded. To see that it has the same solution as the original problem, note that in finite horizon, this problem is identical to the original problem with the addition of a constant term.

For the remainder of this section, we will consider the case with $g = \beta$,

$$\{\tau_t^*\}_t \equiv \arg\max_{\{\tau_t\}_t} \int_0^\infty \overline{U}\left(K_t\right) - \dot{K}_t' \widetilde{D} \dot{K}_t - \overline{U}(K^*) dt,$$

such that $\{K_t\}_t$ is the equilibrium path of allocations induced by $\{\tau_t\}_t$ and initial allocation K_0.

5.2 Decomposition of Problem into Primal and Pigouvian Problems

Note that the social planner's objective depends only on the path of capital allocations, $\{K_t\}_t$. Thus, the social planner's problem is equivalent to searching over all paths $\{K_t\}_t$ that can be induced by some tax schedule path $\{\tau_t\}_t$, given an initial allocation K_0. Let us denote such capital allocation paths $\{K_t\}_t$ as "feasible."

In general, we require the following assumption:

Assumption 1 All paths that pass through "feasible states" are feasible, where a "feasible state" is a capital allocation K_t that lies in some feasible path. That is, the feasibility of a path is a purely local property.

Proposition 1 If Assumption 1 holds, the social planner's problem can be decomposed into two problems: the primal problem and the Pigouvian problem. Denoting the set of feasible states by $\mathcal{K}$, the primal problem is,

$$\{K_t^*\}_t = \max_{\{K_t\}_t \in C^2(\mathbb{R}^+,\mathcal{K})} \int_0^\infty \overline{U}\left(K_t\right) - \dot{K}_t' D \dot{K}_t - \overline{U}(K^*) dt,$$

and the Pigouvian problem is,

$$\text{Find } \{\tau_t\}_t \text{ s.t. } \{K_t^*\}_t \text{ is induced by } \{\tau_t\}_t.$$

5.2.1 Decomposition for Our Specific Model

In the case of the particular model under consideration, Assumption 1 holds and the set of all feasible states is,

$$\mathcal{K} = \{\{K_k\}_k \mid \sum_k K_k = 1 \wedge K_k > 0 \forall k\}.$$

We can also parameterize this usefully by $[0,1]^2$. Denote an element of $\mathcal{K}$ by X or X_t, and parameterize it thus:

$$X(K_1, K_2) = [K_1, K_2, 1 - K_1 - K_2].$$

Then, if we define,

$$D = \begin{bmatrix} \delta_1 + \delta_3 & \delta_3 \\ \delta_3 & \delta_2 + \delta_3 \end{bmatrix},$$

the social planner's problem under this change of variables becomes,

$$\{X_t^*\}_t \equiv \arg\max_{\{X_t\}_t} \int_0^\infty \overline{U}\left(X_t\right) - \dot{X}_t' D \dot{X}_t - \overline{U}(X^*) dt.$$

5.2.2 Additional Properties

Our particular model also has two more useful properties which are not necessary for tractability but do have useful consequences. They are as follows.

Assumption 2 (Separability of Expectations.) The first-best path of allocations $\{X_t^*\}_t$ is feasible regardless of whether households are forward-looking.

Proposition 2 If Assumption 2 holds, the first-best path $\{X_t^*\}_t$ does not depend on whether the household is forward-looking or not. Thus, even though assumptions about whether the household is forward-looking may affect the optimal tax schedules $\{\tau_t^*\}_t$, these assumptions can be imposed after the primal problem is solved.

Assumption 3 (Strong Decomposability.) For any path of allocations $\{X_t\}_t$, the local response of the economy to a change in the tax schedule $\{\tau_t\}_t$ is fully determined by a function $\psi : (X, \tau, \dot{\tau}) \mapsto \dot{X}$.

Proposition 3 If Assumption 3 holds, then there exists a function $\phi : (X, \dot{X}, \tau) \mapsto \dot{\tau}$ such that $\psi(X, \tau, \phi(X, \dot{X}, \tau)) = \dot{X}$ for all $X, \dot{X}$, and τ consistent with $X, \dot{X}$. This allows the Pigouvian problem to be solved inductively given X_0, τ_0.

5.3 Decomposition of Primal Problem into Rate Problem and Path Problem

Our current problem is as follows,

$$\{X_t^*\}_t \equiv \arg\max_{\{X_t\}_t} \int_0^\infty \overline{U}\left(X_t\right) - \dot{X}_t' D \dot{X}_t - \overline{U}(X^*) dt.$$

Note that $\{X_t^*\}_t$ can be thought of as a function γ,

$$\gamma : \mathbb{R}^+ \to \mathcal{K},$$

and that γ, in a sense, contains two pieces of information: the rate at which the economy should adjust, and the path along which it should adjust. We can decompose this by specifying a function ρ whose rate is fixed, and an increasing "rate function" π as follows,

$$\begin{aligned} \gamma &= \rho \circ \pi, \\ \rho &: \mathbb{R}^+ \to \mathcal{K}, \\ \text{with} \quad &|\rho'(s)| = 1 \ \forall s \in \mathbb{R}^+. \\ \text{and} \quad &\rho(0) = X_0, \\ \pi &: \mathbb{R}^+ \to \mathbb{R}^+, \\ \text{with} \quad &\pi'(t) > 0 \ \forall t \in \mathbb{R}^+, \\ \text{and} \quad &\pi(0) = 0. \end{aligned}$$

where $|\cdot|$ is some norm on values $\rho'(s)$, which we will discuss in more detail soon. Intuitively, ρ carries the information about which states the economy should pass through, and π carries the information about how quickly the economy should pass through them.

Let P be the set of all constant-speed functions $\mathbb{R}^+ \to \mathcal{K}$ as above, and similarly let Π be the set of all continuous increasing functions $\mathbb{R}^+ \to \mathbb{R}^+$ as above.

Note that ρ can be solved for conditional on π. That is, our problem can be expressed as,

$$\max_{\rho \in P} \max_{\pi \in \Pi} \int_0^\infty \overline{U}\left(\rho(\pi(t))\right) - (\rho'(\pi(t))\pi'(t))' D(\rho'(\pi(t))\pi'(t)) - \overline{U}(X^*) dt,$$

using the identities,

$$\begin{aligned} X_t &= \rho(\pi(t)), \\ \dot{X}_t &= \rho'(\pi(t))\pi'(t) \text{ by the chain rule.} \end{aligned}$$

For our particular model, π can be solved in closed form given ρ:

Proposition 4 In this model, conditional on ρ, the optimal path π satisfies,

$$\pi'(t) = \frac{\sqrt{\rho'(\pi(t))' D \rho'(\pi(t))}}{\sqrt{\overline{U}\left(\rho(\pi(s))\right) - \overline{U}(X^*)}},$$

Under a change of variables $s \equiv \pi(t)$ and $\tilde{\pi} \equiv \pi^{-1}$, this becomes,

$$\tilde{\pi}'(s) = \frac{\sqrt{\rho'(s)' D \rho'(s)}}{\sqrt{\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)}} = \pi'(t)^{-1},$$

and the problem becomes,

$$\begin{aligned} \rho^* &= \arg\max_{\rho \in P} \int_0^\infty \left[\overline{U}\left(\rho(s)\right) - (\rho'(s)/\tilde{\pi}'(s))' D(\rho'(s)/\tilde{\pi}'(s)) - \overline{U}(X^*)\right] \tilde{\pi}'(s) ds \\ &= \arg\max_{\rho \in P} \int_0^\infty \left[\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)\right] \tilde{\pi}'(s) - \left[\rho'(s)' D \rho'(s)\right] / \tilde{\pi}'(s) ds \\ &= \arg\max_{\rho \in P} \int_0^\infty \sqrt{\rho'(s)' D \rho'(s)} \sqrt{\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)} + \sqrt{\rho'(s)' D \rho'(s)} \sqrt{\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)} ds \\ &= \arg\max_{\rho \in P} \int_0^\infty 2\sqrt{\rho'(s)' D \rho'(s)} \sqrt{\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)} ds \\ &= \arg\max_{\rho \in P} \int_0^\infty \sqrt{\rho'(s)' D \rho'(s)} \sqrt{\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)} ds \end{aligned}$$

This might seem like a great deal of unmotivated rewriting, but we have made important progress in one crucial respect: we have constrained the search space to the set of constant-speed paths $\mathbb{R} \to P$.

This is important not because we specifically care about constant-speed paths, but because ρ now contains only information about which intermediate states the economy passes through, and *not* how quickly it passes through them.

In particular, the integrand,

$$\sqrt{\rho'(s)' D \rho'(s)} \sqrt{\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)},$$

can be understood as capturing the "cost" of transitioning from state $X \equiv \rho(s)$ to state $X + \Delta X$, if ΔX is small, if the speed of adjustment is chosen *optimally*.

But "cost" here reflects not only the adjustment cost, but the entire contribution of the transition to the integral, including the costs of remaining in a suboptimal state. Indeed, the above optimization problem looks less like a utility-maximization problem and more like a distance-minimization problem. The next section makes this explicit.

5.4 Transformation of Primal Problem into Geodesic Problem

The problem as it currently stands is to solve,

$$\arg\max_{\rho\in P} \int_0^\infty \sqrt{\rho'(s)' D\rho'(s)}\sqrt{\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)}ds.$$

To proceed, note first that for each ρ there exists some $S \in \mathbb{R}$ at which the path reaches X^*, causing the integrand to vanish for all $s > S$. Thus, the problem is equivalent to solving,

$$\arg\max_{\rho\in P} \int_0^{S(\rho)} \sqrt{\rho'(s)' D\rho'(s)}\sqrt{\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)}ds.$$

Second, note that we never defined $|\cdot|$, the norm on $\rho'(s)$. Let us now define it as follows,

$$|\rho'(s)| = \left(\rho'(s)' D\rho'(s)\right)\left(\overline{U}\left(\rho(s)\right) - \overline{U}(X^*)\right).$$

Everything then drops out of our integral,

$$\begin{aligned} &\arg\max_{\rho\in P} \int_0^{S(\rho)} ds \\ =&\arg\max_{\rho\in P} S(\rho). \end{aligned}$$

That is, the problem becomes a minimum-time problem over a certain set of paths ρ.

Finally, note that $|\rho'(s)|$ is a norm on $\rho'(s)$, imposing some notion that ρ is a path of constant speed. Thus, the minimum-time problem is really a minimum-distance problem with respect to the norm $|\cdot|$. Indeed, $|\cdot|$ can be written,

$$\begin{aligned} |\rho'(s)|_X &= \rho'(s)' M(X)\rho'(s), \\ \text{where} \quad M(X) &\equiv \left(\overline{U}(X) - \overline{U}(X^*)\right) D. \end{aligned}$$

In the language of Riemannian geometry, $M(X)$ is a metric tensor on $\mathcal{K}$, and $|\cdot|_X$ is the norm induced by this metric.

In the following, we refer to M as the "Transition Cost Metric."

5.5 Solving The Geodesic Problem

The problem is now to find the shortest path from X_0 to X^* with respect to the Riemannian metric with metric tensor $M(X)$. Two results make this problem tractable.

Theorem 1 The shortest path between two points in a Riemannian manifold is a geodesic. It is locally straight with respect to the Riemannian metric.

Theorem 2 A geodesic ρ on a Riemannian manifold is fully determined by its initial velocity $\rho'(0)$.

The problem can therefore be solved by searching over initial velocities $\rho'(0)$. For each guess, a geodesic can be constructed following the geodesic equations. The following figure illustrates this process.

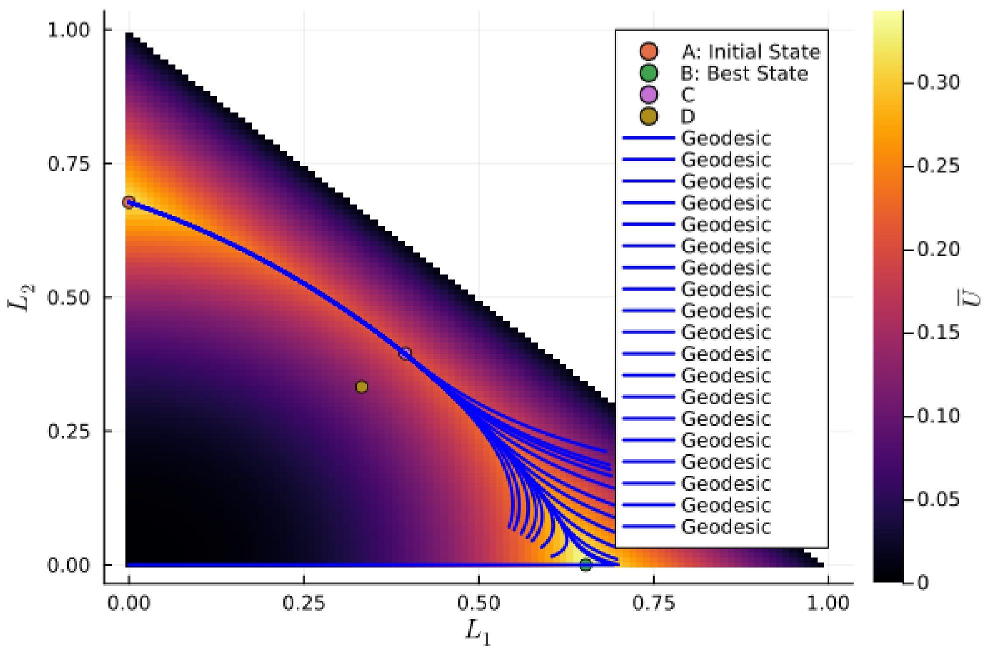

Figure 3: Illustration of the solution search process, displaying "repulsion" by good states.

5.6 Solving the Pigouvian Problem

Finally, it remains to solve the Pigouvian problem. That is, having obtained the optimal path of allocations $\{K_t^*\}_t$, it remains,

$$\text{Find } \{\tau_t\}_t \text{ s.t. } \{K_t^*\}_t \text{ is induced by } \{\tau_t\}_t.$$

In a free market equilibrium, capital is freely mobile across industries, and so $R_k = R$ for all k. Then we have

$$\frac{\tilde{P}_k}{\tilde{P}_1} = \frac{(1+\tau_k)\,P_k}{(1+\tau_1)\,P_1} = \frac{(1+\tau_k)}{(1+\tau_1)} \cdot \frac{\overline{A}_1 K_1^{\gamma}}{\overline{A}_k K_k^{\gamma_k}}.$$

Next, $C_k = Q_k = \overline{A}_k K_k^{1+\gamma_k}$, then condition $\sum_k \tau_k P_k C_k = 0$ can be written as

$$\sum_k \frac{\tau_k}{1+\tau_k} \frac{\tilde{P}_k}{\tilde{P}_1} \overline{A}_k K_k^{1+\gamma_k} = 0,$$

which gives

$$\sum_k \tau_k K_k = 0.$$

We can then solve

$$1 + \tau_k = \frac{\tilde{P}_k Q_k}{\tilde{P}_1 Q_1 + \tilde{P}_2 Q_2 + \tilde{P}_3 Q_3} \cdot \frac{1}{K_k}.$$

6 Results

Our results are presented in Figures 4–6. Figure 4 depicts the computed geodesic between the initial auto-dominated state (State A) and the desired rail-dominated state (State B). Notably, the optimal path in the (K_1, K_2) space is not a straight line; instead, it is convex toward the origin. This curvature arises because, during the transition from State A to State B, the social planner reallocates some capital from the rest of the economy to supplement the capital transferred from the auto industry to the railways, which is illustrated in Figure 5.

By slowing down the withdrawal of capital from the auto industry, the social planner mitigates the productivity loss in the auto industry caused by external economies of scale. Alternatively (and equivalently), this capital reallocation can be interpreted as a strategy to enhance productivity in the railways industry along the transition path. This approach proves optimal because, in this example, the rest of the economy lacks economies of scale, whereas both the auto and railways industries exhibit strong economies of scale. We find that the optimal transition path generates 17.8% more welfare compared to a straight-line transition.

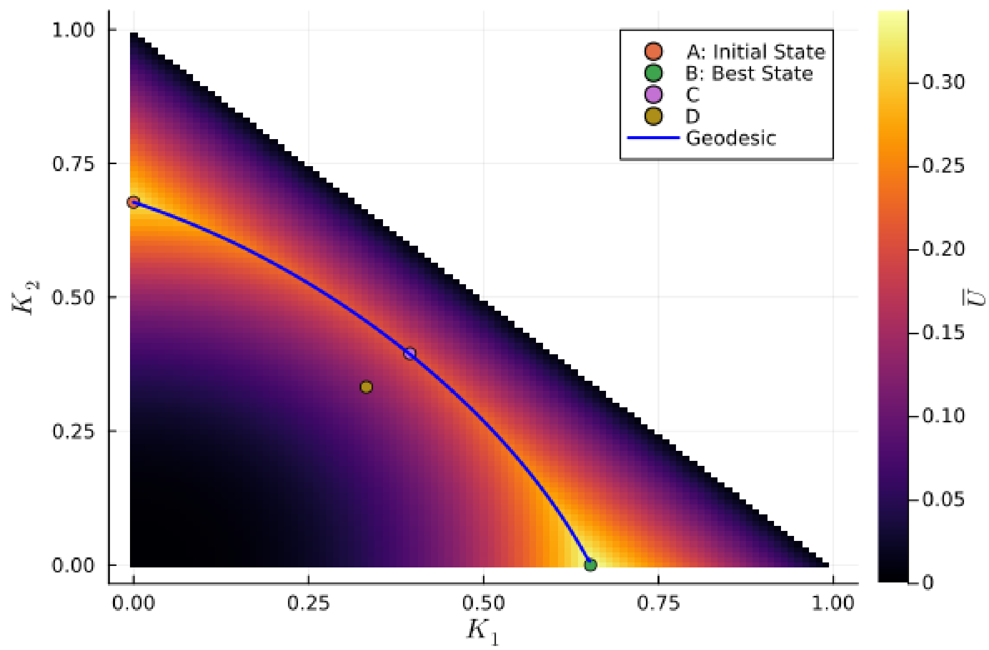

Figure 4: Optimal path from the Initial State to the Best State. A geodesic with respect to the Transition Cost Metric.

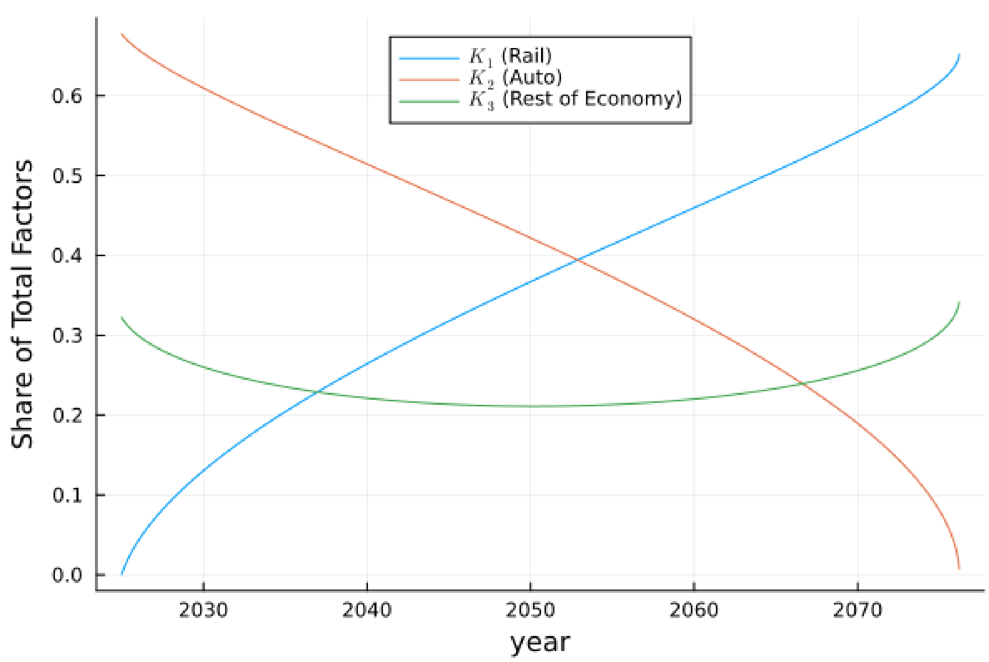

Figure 5: Optimal allocations of scarce factors along the optimal path. The solution to the "primal" problem.

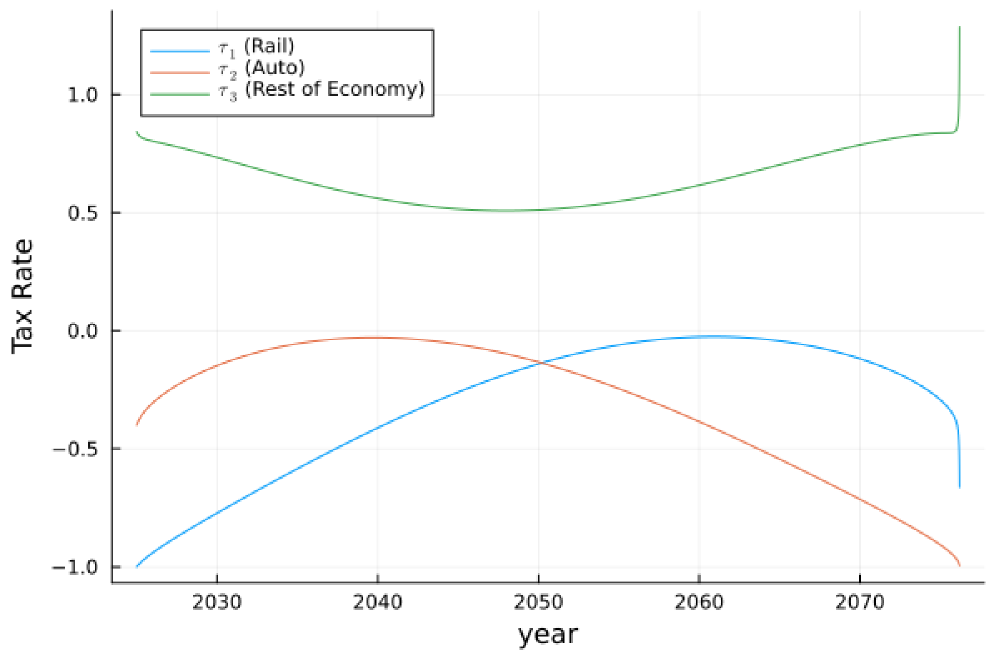

Figure 6: Path of taxes/subsidies which induce the optimal path of allocations. The solution to the Pigouvian problem.

Figure 5 also shows that the optimal transition spans approximately 53 years, from 2025 to 2078. Figure 6 illustrates the optimal Pigouvian taxes supporting this transition. Initially, the social planner heavily subsidizes the rail industry (negative τ_1). The auto industry also receives subsidies, but to a lesser extent than the railways ($|\tau_2| < |\tau_1|$). These subsidies are funded by taxes imposed on the rest of the economy ($\tau_3 > 0$).

In the final state, the social planner continues to subsidize both the auto and railways industries. This is necessary because households and firms do not internalize the productivity-enhancing externalities. Toward the end of the transition, the social planner subsidizes the "dying" car industry more heavily than the railways to prevent the auto industry from collapsing too rapidly, thereby ensuring a smoother transition.

7 Conclusion

In this paper we study green transition policy in settings where "greener" alternatives to existing technologies already exist, but have not been optimally adopted. Our illustrative example is that of rail versus automobile transportation in the United States. These two sectors are substitutable to some degree, but the marginal benefits of new investments depend strongly on the stock of existing infrastructure.

To study this type of policy, we use a model in which external returns to scale, or "agglomeration effects," within industries lead to multiple equilibria, which together with adjustment costs generate path-dependence.

The model is a relatively simple combination of external economics-of-scale production, costly reallocation of factors between sectors, and nested CES aggregation over sectors. The main methodological contribution is in the focus on the optimal policy problem of setting optimal *time-varying* sector-level subsidies and the introduction of a novel solution method for this problem.

A hypothetical policy designed to shift the economy from the automobile equilibrium to the rail equilibrium represents an optimal policy problem in multiple time-varying instruments within a dynamic model. We propose a geometric approach to this class of problems: the state space is equipped with a Riemannian metric that captures the costs of transitioning between states, such that the optimal policy arises as a geodesic. The solution algorithm is surprisingly simple and is robust to different assumptions about household expectations, including fully myopic or forward-looking and in between.

By slowing down the withdrawal of capital from the auto industry, the social planner mitigates the productivity loss in the auto industry caused by external economies of scale. Alternatively (and equivalently), this capital reallocation can be interpreted as a strategy to enhance productivity in the railways industry along the transition path. This approach proves optimal because, in this example, the rest of the economy lacks economies of scale, whereas both the auto and railways industries exhibit strong economies of scale. We find that the optimal transition path generates 17.8% more welfare compared to a straight-line transition.

The baseline model has only four key elasticities: the strength of agglomeration effects within a sector, the elasticity of substitution between the rail and automobile sectors, the elasticity of substitution between the aggregate "transportation sector" and the rest of the economy, and the adjustment cost of reallocating factors between sectors. In addition to these, the model can be generalized in various directions with the addition of only one parameter, for example, the degree to which households are forward-looking, the degree to which adjustment costs are external, and the degree to which agglomeration effects are external. In this way, the model can be thought of as a complement to empirical work—specifying and precisely defining key elasticities which can yield clean results, and a way of interpreting those elasticities in general, without imposing a one-size-fits-all estimation or identification strategy.

While the approach to estimation is somewhat stylized, this exercise represents a first step in applying macroeconomic tools to climate policy analysis away from the margin.

References

Dominick Bartelme, Arnaud Costinot, Dave J. Donaldson, and Andres Rodriguez-Clare. The textbook case for industrial policy: Theory meets data. *Journal of Political Economy*, forthcoming.

Pierpaolo Benigno and Michael Woodford. Linear-quadratic approximation of optimal policy problems. *Journal of Economic Theory*, 147(1):1–42, 2012. ISSN 0022-0531.

Lorenzo Caliendo and Fernando Parro. Chapter 4 - trade policy. In Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, editors, *Handbook of International Economics: International Trade, Volume 5*, volume 5 of *Handbook of International Economics*, pages 219–295. Elsevier, 2022. doi: https://doi.org/10.1016/bs.hesint.2022.02.004. URL https://www.sciencedirect.com/science/article/pii/S1573440422000041.

Dave Donaldson and Richard Hornbeck. Railroads and american economic growth: A "market access" approach. *The Quarterly Journal of Economics*, 131(2):799–858, 02 2016. ISSN 0033-5533. doi: 10.1093/qje/qjw002. URL https://doi.org/10.1093/qje/qjw002.

R.E. Gallamore and J.R. Meyer. *American Railroads: Decline and Renaissance in the Twentieth Century*. Harvard University Press, 2014. ISBN 9780674725645. URL https://books.google.com/books?id=Iov5AwAAQBAJ.

Ahmad Lashkaripour and Volodymyr Lugovskyy. Profits, scale economies, and the gains from trade and industrial policy. *American Economic Review*, 113(10):2759–2808, October 2023. doi: 10.1257/aer.20210419. URL https://www.aeaweb.org/articles?id=10.1257/aer.20210419.

L. Ljungqvist and T.J. Sargent. *Recursive Macroeconomic Theory, fourth edition*. MIT Press, 2018. ISBN 9780262038669. URL https://books.google.com/books?id=IG1qDwAAQBAJ.

Alisdair McKay and Christian K. Wolf. Optimal Policy Rules in HANK. Mimeo, 2023.

Ralph Ossa. Trade wars and trade talks with data. *American Economic Review*, 104(12):4104–46, December 2014. doi: 10.1257/aer.104.12.4104. URL https://www.aeaweb.org/articles?id=10.1257/aer.104.12.4104.

Che-Lin Su and Kenneth L. Judd. Constrained optimization approaches to estimation of structural models. *Econometrica*, 80(5):2213–2230, 2012. doi: https://doi.org/10.3982/ECTA7925. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA7925.

Zi Wang, Xingcheng Xu, Yanqing Yang, and Xiaodong Zhu. Optimal trade and industrial policies in the global economy: A deep learning framework, 2024. URL https://arxiv.org/abs/2407.17731.