How to Fix a Coordination Failure: A "Super-Pigouvian" Approach

John Sturm*

MIT

January 17, 2023

Most recent version available here.

A central concern in industrial policy discussions is that sector-specific external economies of scale may create multiple equilibria—and therefore the potential for coordination failure. Pigouvian policies that address market failures on the margin do not remove the risk of mis-coordination globally. I propose a new "super-Pigouvian" (SP) policy that retains the decentralized spirit of Pigouvian policy—regulating prices rather than quantities—but also prevents coordination failure. The main idea behind SP is to subsidize market behavior, both on and off the equilibrium path, according to the population's willingness to pay for the welfare gains that those behaviors generate (a) directly, like Pigou, and also (b) indirectly, by affecting other households' choices. After demonstrating SP's welfare properties theoretically, I quantify them in a dynamic model of structural transformation calibrated to South Korea's heavy and chemical industry drive in the 1970s. SP modestly improves welfare compared to the worst equilibrium under Pigouvian policy.

*MIT Deptartment of Economics, 50 Memorial Drive, Cambridge, MA 02142. Email: jsturm@mit.edu

I am deeply grateful to my advisors, Daron Acemoglu, Iván Werning, and Arnaud Costinot, for their guidance and encouragement. I would also like to specially thank Nathan Lane for sharing data. For their comments and suggestions, I thank David Atkin, Dave Donaldson, Joel Flynn, Stephen Morris, Daniel O'Connor, Victor Orestes, André Sztutman, and participants of the MIT Macro Seminar.

1 Introduction

Industrial policy is back in fashion. Even relatively free-market countries like the United States are promoting investment in geographically concentrated "innovation zones," subsidizing R&D in the strategic semiconductor industry, and backing nascent green technologies in the hopes they become profitable.

The standard economic rationale for industrial policy is to correct market failures introduced by externalities among firms or between firms and households. These market failures come in two distinct varieties. First, externalities cause individual economic actors to under- or over-prioritize actions relative to their marginal social value. For example, a firm may be unwilling to invest in a green production method if it bears the full cost while others share in the benefit of lower pollution. Second, externalities can create multiple, Pareto-ranked equilibria. If all economic actors believe that the economy will fall into an inefficient equilibrium, then this belief can be self-fulfilling—a *coordination failure*. For example, suppose that a green technology becomes cheaper as more firms adopt it. Even if it is efficient for all firms to adopt the technology, each individual firm may choose not to if no one else does.

The dominant industrial policy paradigm for correcting these externalities is Pigouvian taxation and subsidization. A Pigouvian policymaker taxes or subsidizes an activity—such as a firm's investment in green technology—until its private marginal cost equals its social marginal cost [Pigou, 1924]. This logic underlies policies as wide-ranging as carbon taxes, R&D subsidies, and federal grants for basic science research [Nordhaus, 2006, Görg and Strobl, 2007, Mandt et al., 2020]. Even non-Pigouvian policies, like infant industry protection, are often justified as second-best versions of the Pigouvian ideal [Harrison and Rodríguez-Clare, 2010].

Despite its popularity with academics and policymakers alike, Pigouvian policy has a key limitation: it does not resolve coordination failures. To see why, consider a simple economy with a traditional sector and an industrial sector. In the traditional sector, workers have fixed productivity. The industrial sector features external economies of scale, so that each worker's productivity grows with the number of other industrial workers. Pigouvian policy subsidizes wages in the industrial sector to compensate workers for these externalities on the margin. However, these externalities—and so, Pigouvian subsidies—are zero if all workers choose the traditional sector because, in this case, no one benefits from the increase in industrial productivity caused by any one worker's decision to industrialize. Full non-industrialization is therefore an equilibrium under Pigouvian policy as long as traditional productivity exceeds industrial productivity under full non-industrialization. However, it is efficient only if traditional productivity exceeds industrial productivity under full industrialization.

In this paper, I propose a new, "super-Pigouvian" policy that retains the spirit of Pigouvian policy while also preventing coordination failure. Like Pigouvian policy, this new policy steers clear of central planning, instead simply augmenting the prices faced by households with taxes and subsidies. Moreover, both policies adjust prices in order to align households' private incentives with their marginal contributions to social welfare.

However, super-Pigouvian policy differs from Pigouvian policy in two essential ways. First, although both policies ensure households are paid for the welfare impact of their actions, inclusive of externalities, super-Pigouvian taxation measures this impact in a more comprehensive way. Namely, it seeks to reward households for actions that trigger large shifts in aggregate behavior—away from "bad equilibria" and toward the efficient allocation. Standard Pigouvian policy does not incentivize these actions because it measures a household's impact on welfare *while holding fixed the behavior of all other households.* Super-Pigouvian policy instead compensates each household for not only these direct effects but also the indirect effects of its actions through their influence on other households' behavior.

The second essential feature of super-Pigouvian policy helps to address a question posed by the first: How does any one household's behavior affect another's? So that this notion is well defined, I work in an alternative version of the traditional general equilibrium setup wherein (a) households act sequentially and (b) the economy must follow a competitive equilibrium starting from any history of past behavior. Within this setting, I require that super-Pigouvian policy compensates households for their actions' direct and indirect welfare effects *starting from any history.* This condition—the second key aspect of its definition—shapes behavior off of the equilibrium path and therefore disciplines the effects of one household's action on another's. This differs from Pigouvian taxation, which need only be specified along an equilibrium path.

My main theoretical result is a first welfare theorem for super-Pigouvian policy. Namely, I show that any equilibrium in which policy is set according to the super-Pigouvian rule is Pareto efficient. Three ideas underlie this result. First, I study a dynamic model in which households act sequentially rather than simultaneously. This narrows the scope of coordination failure by allowing each actor to observe actions that have come before them, rather than guessing what they might be. Second, my approach compensates households for their welfare impacts *accounting for effects on other, future behavior*, which ensures households choose the action that maximizes welfare given how others will respond. Third, the fact that super-Pigouvian policy is set off the equilibrium path ensures that a household who deviates away from an inefficient equilibrium knows that those who act afterwards will continue the efficient transition it began. The combination of these three factors implies efficiency by a logic akin to the principle of optimality in dynamic programming.

The idea that an alternative policy can improve on Pigouvian taxation may appear to conflict with standard results in public finance. It is well known that any taxes or subsidies that implement an efficient allocation must be Pigouvian locally to that allocation. Consistent with this fact, I show that, in any strategy profile that can be implemented under super-Pigouvian policy—and so, according to my first welfare theorem, is efficient—policy is Pigouvian locally to the equilibrium path. Where super-Pigouvian taxes and subsidies differ from the Pigouvian prescription is in inefficient strategy profiles. There, super-Pigouvian taxes and subsidies guarantee that the profile cannot be individually rational for some market participant and so is not an equilibrium. I moreover confirm that super-Pigouvian policy is not overly selective by providing a second welfare theorem: as with standard Pigouvian policy, any efficient allocation can be implemented in a super-Pigouvian equilibrium.

A concern with some approaches to industrial policy—such as regulating quantities or "picking winners"—is that they require the social planner to be unrealistically well-informed. One of the advantages of standard Pigouvian policy is that it can be implemented using only information on the extent of externalities on the margin of the equilibrium path—or equivalently, households' willingness to pay (WTP) for one another's actions, holding fixed all future behavior. I show that, analogously to the Pigouvian case, a planner can implement super-Pigouvian policy using only household WTPs for one another's actions, provided that these WTPs account for actions' indirect effects through their influence on future behavior. Importantly, households can form these WTPs using only information about prices and transfers, without any individual knowing the efficient allocation.

After establishing these welfare properties in a general environment, I quantify the welfare gains from super-Pigouvian policy in a dynamic model of structural transformation that I calibrate using South Korea's heavy and chemical industry (HCI) drive in the 1970s. This model captures the idea that there may be strong complementaries between firms—for example steel manufacturers and industrial machine producers—who use one another's output as an input [Rodrik, 1996, Buera et al., 2021]. When these complementarities are strong enough, entrepreneurs' optimism or pessimism about industrialization can be self-fulfilling.

I take this model to the data using estimates of the HCI drive's effects on industrialization from Lane [2022] and use the calibrated model to compute welfare under a range of policies. I find that Pigouvian taxation can support multiple equilibria at different levels of self-fulfilling optimism about investment into new industrial technologies. Relative to the worst Pigouvian equiilbrium, super-Pigouvian policy generates modest welfare gains equivalent to a $\sim$2% increase in HCI value added. I also examine the mechanism by which super-Pigouvian policy incentivizes households to "escape" inefficient Pigouvian equilibria. Super-Pigouvian policy offers large incentives for investment in new technology concentrated in states where firms are nearly indifferent to investment.

Literature: This paper is not the first to propose a policy that can resolve coordination failures without requiring a social planner to know the efficient equilibrium path. Notably, Sandholm [2002, 2005, 2007] studies whether a social planner may implement the efficient allocation by simply committing to engage in Pigouvian taxation at any history the economy may reach. He shows that when—rather than acting with perfect foresight—households update their behavior myopically and with a small enough amount of noise, only welfare-maximizing strategy profiles are played in a non-negligible fraction of periods in the long run. Fujishima [2013] provides a similar result in the case of perfect foresight, showing that in the limit of small frictions in updating behavior, (a) there always exists an equilibrium path from the current state to the efficient steady state and (b) starting from sufficiently close to that state, all equilibria converges to it.

Like the policy these papers propose, super-Pigouvian policy requires taxes to be set appropriately both on and off the equilibrium path. However, super-Pigouvian policy sets these taxes in a non-Pigouvian fashion. I show that this difference leads to stronger welfare properties: super-Pigouvian policy ensures efficiency along the entire path, not just at a steady state, and its efficiency properties do not rely on agents acting myopically or with vanishing frictions. This welfare guarantee is analogous to that of Atkeson et al. [2010], who propose a monetary policy rule that eliminates indeterminacy in the New-Keynesian model. The key difference between our respective approaches to unique implementation is that mine does not require the social planner to (even approximately) know the efficient allocation.

Also related is a literature on dynamic mechanism design. One way to understand the welfare properties of super-Pigouvian policy is by analogy to a dynamic game of common interest with staggered actions. Super-Pigouvian policy ensures that households (a) act in the common interest and (b) even though they are price takers, act as if they can influence the equilibrium path. My first welfare theorem is then analogous to the efficiency of subgame perfect equilibria in sequential-action, discounted, dynamic games of common interest [Lagunoff and Matsui, 1997].[1] The super-Pigouvian idea that households should be compensated for their full contribution to welfare is also familiar from the well-known mechanism of Vickrey [1961], Clarke [1971], Groves [1973] and the more recently-proposed dynamic pivot mechanism [Bergemann and Välimäki, 2010]. The essential distinction between the approach of these papers and my own is that their mechanisms require households to share all private information with a central planner; this makes it easy to avoid coordination failure, as there

[1] Methodologically, the one-shot deviation principle in discounted dynamic games is closely tied to the dynamic programming logic on which my proof relies.

is a single decision-maker.[2] By contrast, I emphasize how super-Pigouvian policy can prevent coordination failure without requiring the central planner to know the efficient allocation.

While the idea of coordination failure applies to many topics, it has received particular attention in the industrial policy sphere [Rosenstein-Rodan, 1943, Murphy et al., 1989, Rodriguez-Clare, 1996, Ciccone, 2002, Rodrik, 2004]. From a modelling perspective, my main contribution to this literature is to nest an empirically-motivated static model of industrialization within the dynamic setup of Frankel and Pauzner [2000]. On the estimation side, I calibrate my model using empirical moments from actual policy variation, rather than steady-state sufficient statistics as in Buera et al. [2021]. Finally, I assess the gains from a novel industrial policy.

Although I focus on industrial policy, my theoretical results apply to numerous other settings where positive externalities can create multiple equilibria. One important case is economic geography, where a long literature has considered whether local spillovers may lead to agglomeration in inefficient locations [Krugman, 1991, Davis and Weinstein, 2002, Bleakley and Lin, 2012, Allen and Donaldson, 2020]. Another application is to directed technological change, where the fact that successive innovations "build on the shoulders" of earlier inventions can create multiple equilibria [Acemoglu et al., 2012]. The theory also applies to business cycle models, in which recent papers have ruled out multiple "sunspot" equilibria with modelling assumptions, using tools from the global games literature [Angeletos and Lian, 2016]. In each of these contexts, many authors either make modelling assumptions that rule out multiplicity or, if they allow for multiplicity, study policies that are *necessary* for efficiency but not typically *sufficient* [Guimaraes and Machado, 2018, Schaal and Taschereau-Dumouchel, 2015, Fajgelbaum and Gaubert, 2020]. Super-Pigouvian taxation is a policy that one can propose while neither foreclosing multiplicity through modelling choices nor narrowing the definition of policy success. The methodological contribution that underlies these applications is my first welfare theorem's reliance on dynamic programming rather than convexity assumptions.

Outline: The paper is organized as follows. In Section 2, I provide a motivating example that introduces the main ideas of the paper in a simplified setting. Section 3 presents the general model. Sections 4 and 5 characterize the welfare properties of Pigouvian and super-Pigouvian properties, respectively, as well as the information required for a planner to implement them. Finally Section 6 quantifies the welfare gains from super-Pigouvian policy

[2] I also differ from this literature—as well as the one on subgame perfect implementation more generally, e.g. Moore and Repullo [1988]—by implicitly assuming that a social planner can elicit household willingnesses to pay. Here I am motivated by the application to large economies in which any one agent's preferences may have a negligible effect on policy.

in a specialized model of industrialization, calibrated using the HCI drive in 1970s South Korea. Section 7 concludes.

2 Motivating example

Before expositing the formal model, I introduce the main ideas behind super-Pigouvian policy in an accessible example.

2.1 A simple model of industrialization

Two households $i = 1, 2$ sequentially decide whether to work in the traditional sector or the industrial sector. After both have chosen a sector, they each inelastically supply one unit of labor and production occurs.

The traditional and industrial sectors are competitive and produce a common final good. In the traditional sector, each household has productivity $A > 0$. The industrial sector features a productivity externality, with each household having productivity $(L^I)^\alpha$, where L^I is the number of households in the industrial sector and $\alpha > 0$.

A government taxes labor in each sector and can condition the second household's taxes on the sectoral choice of the first.

Figure 1 depicts the timing of household sectoral choice and its implications for output.

Remark. *Economists often study Pigouvian taxation in contexts where households' decisions are sensitive to the incentives they face on the margin. It may therefore surprise the reader that this model's equilibria are corner solutions to the household problem, implying that household behavior is locally insensitive to marginal incentives. Nonetheless, Pigouvian taxation plays an important role in determining which of these corners can be supported.*

2.2 Pigouvian policy and coordination failure

We begin by considering what can happen when the government sets taxes in the standard Pigouvian fashion. A Pigouvian equilibrium is a path for the wages faced by firms, the wages faced by households, and households' labor supply and consumption decisions such that:

- The wages faced by firms are their marginal products of labor, i.e. A in the traditional sector and $(L^I)^\alpha$ in the industrial sector.
- The wages faced by households are those set by firms plus the value of the marginal externalities they generate on output, local to the equilibrium path, i.e. A in the

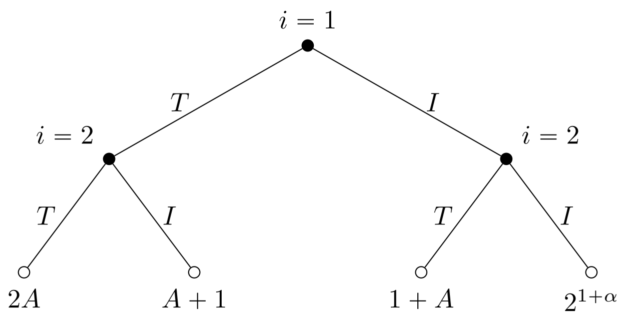

Figure 1: Timing and output in a simple model of industrialization. The structure of the tree represents the timing of household sectoral choice. T is the traditional sector and I is the industrial sector. Terminal nodes show the value of aggregate output at each allocation of households to sectors.

traditional sector and $(L^I)^\alpha + \frac{d}{d\overline{L}}\Big|_{\overline{L}=L^I}(\overline{L})^\alpha L^I = (1+\alpha)(L^I)^\alpha$ in the industrial sector. In other words, the post-tax wage in each sector equals the total, marginal output contribution of labor in that sector, through both the direct and externality channel.

- Taking these paths of wages as given, households make the labor supply decisions that maximize the NPV of their earnings, and they consume up to their budget constraints.

Suppose full industrialization is Pareto efficient, i.e. $2A < 2^{1+\alpha}$. Can Pigouvian taxation support this outcome? Can it support other, inefficient, outcomes?

Pigouvian taxation can indeed support industrialization. On an equilibrium path in which both households industrialize, the post-tax industrial wage is $(1+\alpha)2^\alpha$. The fact that industrialization is efficient implies that this wage exceeds that offered in the traditional sector, rationalizing households' equilibrium behavior:

$$2A \; < \; 2^{1+\alpha} \qquad \Longrightarrow \qquad A \; < \; (1+\alpha)2^\alpha. \tag{1}$$

However, Pigouvian taxation can also support the "development trap" outcome where both households choose the traditional sector. This is simply because, on the equilibrium path where no one works in the industrial sector, (a) productivity in the industrial sector is zero and (b) there is no one to benefit from externalities that boost industrial productivity. So the post-tax industrial wage is zero and both households prefer the traditional sector. The left panel of Figure 2 illustrates households incentives under these post-tax wages. If either household deviates from its equilibrium choice of the traditional sector then, holding the other household's behavior fixed in the standard general equilibrium fashion, it ends up with a post-tax wage of zero instead of A. The dynamic structure of the diagram is unimportant for this point but will become relevant in the next section.

In other words, Pigouvian taxation cannot rule out the *coordination failure* outcome in which each household—expecting that no one will join it, should it migrate to the industrial sector—stays in the traditional sector. This self-fulfilling prophecy arises despite the fact that taxes are set so as to correct for externalities on the margin of the equilibrium path where no one works in industry. The core problem is that non-industrialization corresponds to a *local maximum* of the production function, away from which neither household can move without coordinating with the other.[3]

As Pigou [1924] himself wrote, "All of the relative maxima are, as it were, the tops of hills higher than the surrounding country, but only one of them is the highest hill-top of all."

2.3 An alternative to Pigouvian policy

The main contribution of this paper is to provide an alternative, "super-Pigouvian" policy that retains the spirit of Pigouvian policy while avoiding coordination failure. Within this simple model, I now introduce super-Pigouvian policy in two steps.

First, consider the sectoral choice of the household who acts second. Our assumption that industrialization is efficient implies that it is efficient for this household to industrialize if the first household does. For simplicity, assume that it is otherwise efficient for the second household to choose the traditional sector, i.e. $A > 1$.

How can a planner set taxes to guarantee that the second household chooses the efficient sector? One simple way is to "sell the household the economy" by ensuring that the difference in post-tax wages across the two sectors is equal to the difference in aggregate output that results from her choice. This way, she is a full residual claimant to her action's effects and chooses the sector that maximizes output—i.e. whichever the first household has chosen.

The post-tax wages shown for the second household in the right panel of Figure 2 illustrate this logic. On one hand, if the first household chooses the traditional sector, then aggregate output is either $2A$ if the second household also chooses the traditional sector or $A+1$ if the second household chooses the industrial sector. So if the second household's post-tax wage in the traditional sector is A, then its post-tax wage in the industrial sector should be 1.[4] On the other hand, if the first household chooses the industrial sector, then aggregate output is either $1+A$ if the second household chooses the traditional sector or $2^{1+\alpha}$ if it chooses the industrial sector. So if the second household's post-tax wage in the traditional sector is A, then its post-tax wage in the industrial sector should be $2^{1+\alpha}-1$.

[3]Part of the problem is also that—due to the neoclassical assumption that households are price takers—each is unable to "coordinate with itself" and move to the industrial sector motivated by how doing so will affect its own productivity. However, the self-coordination issue disappears in the many household limit.

[4]The *level* of the post-tax wage does not affect behavior because households have inelastic total labor supply. What matters here is the *difference* in post-tax wages across sectors.

Second, consider the sectoral choice of the household who acts first. How can the planner set taxes to guarantee it chooses the efficient sector? One method is to follow the same approach as for the second household, setting taxes so as to ensure that the first household's difference in post-tax wages across sectors is equal to the difference in aggregate output that results from her choice of sector. This is not quite as simple as in the case of the second household, since the level of output that results from any sectoral choice of the first household also depends on how the second household responds. In order for to incentivize the first household to industrialize, taxes must reflect the knowledge that the second household will follow suit—as the taxes we specified for it imply.[5]

To see how this works, consider the first household's post-tax wages shown in the right panel of Figure 2. If the first household chooses the traditional sector, then—as the second will also choose the traditional sector in this case—total output is $2A$. If the first household instead chooses the industrial sector then—as the second household will also choose the industrial sector in this case—total output is $2^{1+\alpha}$. So if the second household's post-tax wage is A in the traditional sector, its post-tax wage in the industrial sector should be $2^{1+\alpha} - A$. This ensures that the first household chooses the industrial sector.

To summarize, we have decribed a policy in which each household is paid a post-tax wage that reflects its total contribution to output, taking as given past behavior, and taking into account effects on future behavior. Note that this can result in total post-tax wages that exceed aggregate output, since *both households* are full residual claimants of their effects on output. However, this does not present a challenge to the policy provided that the social planner can set lump-sum taxes to balance its budget after observing household behavior and—in the standard general equilibrium fashion— households take these transfers as given.

Note that this alternative policy retains elements of Pigouvian policy, as both regulate prices rather than quantities, and both compensate households for what is—in one sense or another—the contribution of their behavior to output. And, as we will later see, both policies offer incentives closely related to households' willingness to pay for one another's actions. Motivated by these similarities, I call the alternative policy "super-Pigouvian."

Super-Pigouvian policy differs from standard Pigouvian policy in two ways. First, it compensates households for the entire change in output induced by their actions—inclusive of indirect effects through changes in later actors' behavior. By contrast, Pigouvian policy only compensates households for the effect of changing their own actions on the margin of the equilibrium path. Second, whereas Pigouvian policy need not be specified away from the

[5]If, for instance, a policy reflected the belief that the second household would choose the traditional sector no matter what, then—in order to compensate the first household for its total effect on output—it would pay A in the traditional sector and 1 in the industrial sector. So the first household would choose the traditional sector.

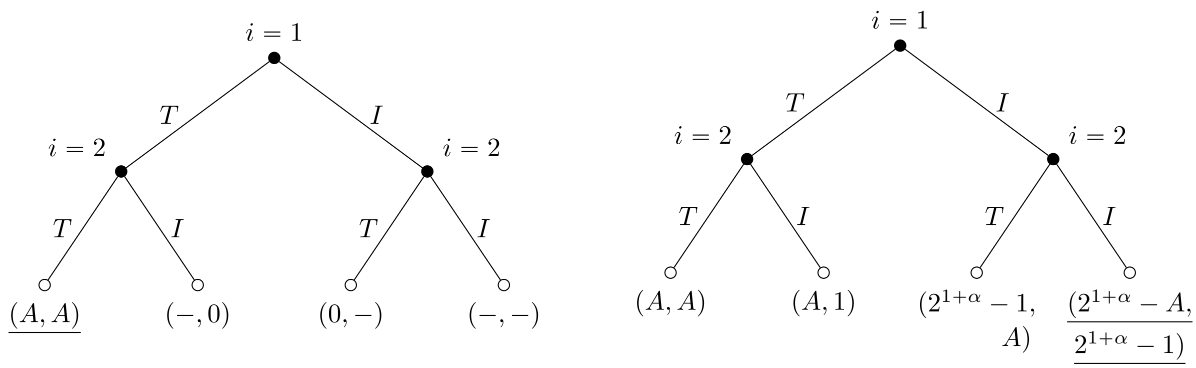

Figure 2: Left panel: Post-tax wages for each household in the Pigouvian equilibrium where both households work in the traditional sector. "-" represents wages that the policy need not specify, as they are more than one deviation away from the equilibrium path. Equilibrium outcome is underlined. Right panel: Post-tax wages for each household under super-Pigouvian policy. (I, I) is the unique equilibrium. "-" represents wages that the policy need not specify, as they are off the equilibrium path for either action taken by the first-actor. Equilibrium outcome is underlined.

equilibrium path, off-path super-Pigouvian policy disciplines what is meant by the "entire change in output" used in the definition of on-path Pigouvian policy.

The remainder of the paper formalizes, generalizes, and characterizes the idea of super-Pigouvian policy introduced informally above. Most importantly, I show that its ability to implement efficient outcomes is a general feature, rather than a quirk of this particular setting. Aside from this and other welfare properties, I also consider the information required for a planner to use Pigouvian or super-Pigouvian policy, showing that each can be implemented using knowledge of only households' willingnesses to pay for one another's choice of action. Finally, I turn to a quantitative application that gauges the potential welfare gains from super-Pigouvian policy in practice.

3 Model

I now lay out a general model in which to formalize the ideas introduced in Section 2.

I take as my starting point a standard model of competitive equilibrium with production externalities and taxation. However, this model is not rich enough to accomodate super-Pigouvian policy because it has no notion of how households behave off of the equilibrium path. To capture this idea while remaining as close as possible to competitive equilibrium, I study recursive competitive equilibria—a standard concept in macroeconomics, e.g. Prescott and Mehra [2005]. In a recursive competitive equilibrium, the path of the economy is well-defined starting from any initial "state" of the economy, but—unlike in, for example, the game-theoretic notion of subgame perfect equilibrium—each such path is a competitive equilibrium in which households take prices as given. I differ slightly from standard applications of recursive competitive equilibrium by taking the "state" to be the entire history of factor supply decisions to date, rather than a low-dimensional or Markovian variable.

3.1 Environment

The economy contains finitely many households $i \in \mathcal{I}$, finitely many factors $n \in \mathcal{N}$, one consumption good, and a finite or infinite number of discrete time periods $t \in \mathcal{T} = 0, 1, ...$

Preferences: At each time t, each household i supplies factors ℓ_t^i and consumes c_t^i units of the final good. Each household i has GHH preferences $u(c_t^i - v(\ell_t^i))$ over consumption and factor supply within each period t and discounts the future at a rate $\beta \in (0, 1)$. I assume u is increasing and v is differentiable.

While households can adjust their consumption freely, they are constrained in their ability to adjust their factor supply. Each household i has initial factor supply $\overline{\ell}_0^i$. At the end of each period t, exactly one household i_t may update her factor supply to a new level ℓ_{t+1}^i contained in a feasible set $\mathcal{A} \subset \mathbb{R}^{\mathcal{N}}$. In period $t+1$, i_t supplies ℓ_{t+1}^i and all other households j continue to supply $\ell_{t+1}^j = \ell_t^j$.

Technology: At each time t, a representative firm produces output Y_t using factor inputs L_t, according to production function $Y_t = F(L_t, L_t)$. The first argument of F represents inputs chosen by the firm, whereas the second represents a production externality that the firm takes as given. I assume that F is differentiable and that it is constant-returns-to-scale in its first argument.

Prices, taxes, and transfers: I normalize the final goods price to one in all periods. At each time t, firms face a vector of factor prices w_t. Each household i faces factor prices that differ from w_t by a vector of marginal taxes $\tau_t^i(\ell_t^i)$, which in general may depend on not only her identity but also her level of factor supply. I denote the corresponding post-tax factor prices by $\omega_t^i(\ell_t^i)$. Between periods, households have access to an exogenous interest rate $R > 1$, as in a small open economy. Finally, households receive lump-sum transfers T_t^i.

3.2 Equilibrium

A recursive competitive equilibrium is a set of competitive equilibria, each corresponding to the economy's path following a different history of factor supply. As in subgame perfect

Nash equilibrium, I require that the competitive equilibrium following a time-$(t+1)$ history coincides with the competitive equilibrium starting from the time-t sub-history provided that households take equilibrium actions at time t.

More formally, A history $h^t = (h_{t-1}, ..., h_0)$ is a sequence of factor supply updates, with each h_s representing the labor supply decision that i_s commits at time s to begin at time $s+1$. I denote by $h^0 = ()$ the empty history before any households have acted.

Definition 1. *A* ***recursive competitive equilibrium*** *is an on-path consistent profile of function* $\{\boldsymbol{c}^i, \boldsymbol{\ell}^i, \boldsymbol{Y}, \boldsymbol{L}, \boldsymbol{w}, \boldsymbol{\omega}^i, \tau^i, \boldsymbol{T}^i\}^{i\in\mathcal{I}}$ *that, when applied to any history* h^t*, returns a competitive equilibrium* $\{\boldsymbol{c}^i_s(h^t), \boldsymbol{\ell}^i_s(h^t), \boldsymbol{Y}_s(h^t), \boldsymbol{L}_s(h^t), \boldsymbol{w}_s(h^t), \boldsymbol{\omega}^i_s(h^t), \tau^i_s(h^t), \boldsymbol{T}^i_s(h^t)\}^{i\in\mathcal{I}}_{s\geqslant t}$ *starting from that history. Formally, conditions* (2)–(7) *hold.*

I now describe the equilibrium conditions for a competitive equilibrium following any history, as well as one additional condition that guarantees a recursive competitive equilibrium is self-consistent along the equilibrium path following any history.

First, the problem of household i following any history h^t is

$$\begin{aligned}\{\boldsymbol{c}^i_s(h^t), \boldsymbol{\ell}^i_s(h^t)\} \quad &\in \quad \underset{\text{feasible } \{c_s, \ell_s\}}{\arg\max} \sum_{s\geqslant t} \beta^{-(s-t)} u\Big(c_s - v(\ell_s)\Big) \\ &\qquad \text{s.t. } \; d^i(h^t) + \sum_{s\geqslant t} R^{-(s-t)} \left(c_s - \boldsymbol{\omega}^i_s(\ell_s; h^t)\cdot \ell_s - \boldsymbol{T}^i_s(h^t)\right) \leqslant 0,\end{aligned} \tag{2}$$

where $d^i(h^t)$ is household i's debt accumulated along the history h^t,[6] and where by "feasible" I mean that (a) $\ell_s \in \mathcal{A}$, (b) $\ell_{s+1} \neq \ell_s$ only if $i = i_s$, and (c) $\ell_t = h_{t'}$ for last time t' before t at which $i = i_{t'}$, or $\ell_t = \overline{\ell}^i_0$ if no such t' exists. Above and throughout, I omit indices when clear from context, e.g. above $\{c^i_s, \ell^i_s\}$ represents $\{c^i_s, \ell^i_s\}_{s\geqslant t}$.

Second, following any history h^t, the representative firm's problem at any time $s \geqslant t$ is

$$\boldsymbol{Y}_s(h^t), \boldsymbol{L}_s(h^t) \quad \in \quad \underset{Y,L}{\arg\max} \; Y - \boldsymbol{w}_s(h^t)\cdot L \qquad \text{s.t.} \qquad Y = F(L, \boldsymbol{L}_s(h^t)). \tag{3}$$

Third, following any history h^t, the schedule of factor prices faced by each household i at any time $s \geqslant t$ differ from the prices facing the firm by the schedule of marginal taxes:

$$\boldsymbol{\omega}^i_s(\cdot; h^t) \quad = \quad \boldsymbol{w}_s(h^t) - \boldsymbol{\tau}^i_s(\cdot; h^t). \tag{4}$$

[6] Fixing $h^t = (h_{t-1}, ..., h_0)$, let—for each $i \in \mathcal{I}$—$(\ell^i_s)_{s<t}$ be the path of labor supply defined by $\ell^i_0 = \overline{\ell}^i_0$, $\ell^i_{s+1} = \ell^i_s$ if $i \neq i_s$, and $\ell^i_{s+1} = h_s$ if $i = i_s$. Household i's private debt $d^i(h^t)$ is given by

$$d^i(h^t) = -\sum_{s<t} R^s \left(\boldsymbol{c}^i_s(h^s) - \boldsymbol{\omega}^i_s(\ell^i_s; h^s)\cdot \ell^i_s - \boldsymbol{T}^i_s(h^s)\right), \qquad \text{where} \qquad h^s = (h_r)_{0\leqslant r<s}.$$

Fourth, following any history h^t, factor markets clear at each time $s \geqslant t$:

$$\boldsymbol{L}_s(h^t) \quad = \quad \sum_{i \in \mathcal{I}} \boldsymbol{\ell}_s^i(h^t). \tag{5}$$

Fifth, following any history h^t, goods markets clear in net present value, taking as given aggregate debt accumulated up to h^t:[7] [8]

$$\sum_{s \geqslant t} R^{-(s-t)} \boldsymbol{Y}_s(h^t) \quad = \quad \sum_{s \geqslant t} R^{-(s-t)} \sum_{i \in \mathcal{I}} \boldsymbol{c}_s^i(h^t) \quad + \quad D(h^t). \tag{6}$$

Finally, in order for them to constitute a recursive competitive equilibrium, I require that the functions pertaining to each equilibrium quantity are consistent along the equilibrium path of factor supply following any history. That is, they make the same prescriptions at any history h^t as at the histories that follow from h^t on the equilibrium path—i.e. according to the function for factor supply. Formally, for any history $h^t = (h_{t-1}, ..., h_0)$,

$$\boldsymbol{x}_{t+1}(h^t) = \boldsymbol{x}_{t+1}\Big(\big(\boldsymbol{\ell}_t^{i_{t-1}}(h^t), h_{t-1}, ..., h_0\big)\Big), \tag{7}$$

for $\boldsymbol{x} = \boldsymbol{c}^i, \boldsymbol{\ell}^i, \boldsymbol{\omega}^i, \boldsymbol{\tau}^i, \boldsymbol{T}^i$ for all $i \in \mathcal{I}$ and for $\boldsymbol{x} = \boldsymbol{w}, \boldsymbol{Y}, \boldsymbol{L}$.

Moving forward, I refer to "recursive competitive equilibrium" as simply "equilibrium."

3.3 Discussion

Two elements of the model warrant further discussion.

First, I have required that at most one household may act at a time. This assumption is innocuous in settings—such as worker migration—where, even if some actors move simultaneously, each is small enough so that they do not affect the optimality of any other's decision. In settings—such as the formation of large firms—where this may not be the case, my focus on sequential actions rules out multiplicity due to self-fulfilling expectations about how others will act in the present. Still, it is possible to support multiplicity due to self-fulfilling

[7] Fixing $h^t = (h_{t-1}, ..., h_0)$, let—for each $i \in \mathcal{I}$—$(\ell_s^i)_{s<t}$ be the path of labor supply defined by $\ell_0^i = \overline{\ell}_0^i$, $\ell_{s+1}^i = \ell_s^i$ if $i \neq i_s$, and $\ell_{s+1}^i = h_s$ if $i = i_s$. Aggregate debt $D(h^t)$ is given by

$$D(h^t) = -\sum_{s<t} R^s \left(F\left(\sum_{i \in \mathcal{I}} \ell_t^i \sum_{i \in \mathcal{I}} \ell_t^i\right) - \sum_{i \in \mathcal{I}} \boldsymbol{c}_s^i(h^s) \right), \qquad \text{where} \qquad h^s = (h_r)_{0 \leqslant r < s}.$$

[8] The fact that goods markets need clear only in net present value—rather than in every period—reflects the presence of the savings technology (the source of the exogenous interest rate R) in the background. Together with households' budget constraints, goods market clearing implies the government satisfies a lifetime budget constraint, which I therefore omit.

expectations about how others will act in the future.

Second, I have assumed that households' consumption and factor supply strategies may condition on the entire history of factor supply, but may not condition on the history of consumption. While not critical, this assumption simplifies my analysis by allowing the planner to leave consumption un-taxed. I view it as sensible in the context of industrial policy, where, for example, workers forming beliefs about future employment in a sector are unlikely to need information on the path of consumption given they have information on the path of aggregate income and sectoral labor supply.

4 The limits of Pigouvian policy

I begin by analyzing standard Pigouvian policy. This highlights its strengths and brings its limitations into focus. In the next section, I will introduce a new policy that overcomes these limitations.

4.1 Defining Pigouvian policy

I begin by defining Pigouvian policy.

Definition 2. *An equilibrium is* ***Pigouvian*** *if for all times t, households i, and factor supplies ℓ,*

$$\boldsymbol{\tau}_t^i(\ell; h^0) \quad = \quad -F_{\overline{L}}\left(\boldsymbol{L}_t(h^0), \boldsymbol{L}_t(h^0)\right). \tag{8}$$

This definition is standard: Factor supply subsidies should be set equal to a household's marginal contribution to output through externalities, so that a household's post-tax wages equals its total, marginal contrbution to output. However, three details are worth higlighting. First, I have embedded Pigouvian policy in an environment that—unlike the standard neo-classical framework—explicitly models off-path behavior. Pigouvian policy does not specify how taxes are set off of the equilibrium path. Second, note that Pigouvian taxes only account for the externalities that a household generates directly, through its own behavior, and not any welfare gains it generates indirectly by influencing the behavior of others. Finally, note that that although marginal taxes in principle may depend on household identity and factor supply behavior, Pigouvian taxes do not.

4.2 Welfare properties of Pigouvian policy

Despite its inability to prevent coordination failure, Pigouvian policy has several positive welfare properties. As these properties are relatively well understood, I state them as a single result. I then discuss how they can guide our thinking in the design of a policy that shares these properties while also addressing coordination failure.

These properties all relate to a weak notion of Pareto efficiency that I now define.

Definition 3. *An allocation $\{c_t^i, \ell_t^i\}_{t\in\mathcal{T}}^{i\in\mathcal{I}}$ is **first-order Pareto efficient** if for all $\{\Delta c_t^i, \Delta\ell_t^i\}$ such that $\{c_t^i + \epsilon\Delta c_t^i, \ell_t^i + \epsilon\Delta\ell_t^i\}$ is feasible for small enough $\epsilon > 0$,*[9]

$$\left.\frac{d}{d\epsilon}\right|_{\epsilon=0} \sum_{j\in\mathcal{I}} e^j\Big(\left\{c_t^j + \epsilon\Delta c_t^j, \ell_t^j + \epsilon\Delta\ell_t^j\right\}_{t\in\mathcal{T}}, \left\{c_t^j, \ell_t^j\right\}_{t\in\mathcal{T}}\Big) \quad \leqslant \quad 0 \tag{9}$$

where $e^j(\{c_t'^j, \ell_t'^j\}_{t\in\mathcal{T}}, \{c_t^j, \ell_t^j\}_{t\in\mathcal{T}})$ is the additional expenditure i would require in the second profile in order to acheive the same utility as it does under the first.[10]

Proposition 1 (Pigouvian welfare properties)**.**

1. *Every Pigouvian equilibrium is first-order Pareto efficient.*
2. *Suppose factor supply disutility v and the action set $\mathcal{A}$ are both convex,*[11]*, the production function F is concave in its first argument, and consumption utility u is such that the household problem has well-defined value given the path of factor supply and total consumption expenditure.*[12] *Then any first-order Pareto efficient allocation can be implemented as a Pigouvian equilibrium.*
3. *Suppose an equilibrium implements a first-order Pareto efficient allocation. If taxes are differentiable in factor supply and equilibrium factor supplies are interior to the feasibility sets $\mathcal{A}$, then—local to the equilibrium path—taxes are effectively Pigouvian in the sense that they provide households with the same factor supply incentives.*[13]

[9] By feasible, I mean that (a) labor supply is within $\mathcal{A}$ and satisfies its timing constraints and (b) goods market clearing holds in net present value when output is given by F evaluated at the labor supply allocation.

[10] Because preferences are GHH, we have

$$e^j(\{c_t'^j, \ell_t'^j\}_{t\in\mathcal{T}}, \{c_t^j, \ell_t^j\}_{t\in\mathcal{T}}) = \left[\sum_{t\in\mathcal{T}} R^{-t}\left(\left(c_t'^j - v(\ell_t'^j)\right) - \left(c_t^j - v(\ell_t^j)\right)\right)\right]. \tag{10}$$

[11] On the surface, the combination of these assumptions may appear somewhat restrictive since they imply that, for example, a household who is indifferent between working in two different industries or locations must weakly prefer to split its time between them. However, $\mathcal{A}$'s convexity can be dispensed with in limits in which there are many copies of every agent, a convexification "trick" that dates back to Aumann [1965]. For this reason I do not dwell on it.

[12] I.e. the following is defined

$$\max_{\{c_t\}_{t\geqslant 0}} \sum_{t\geqslant 0} \beta^t u(c_t - v(\ell_t)) \quad \text{s.t.} \quad \sum_{t\geqslant 0} R^{-t} c_t \leqslant E \tag{11}$$

[13] Taxes can only violate (8) insofar as either (a) they shift payments across time differently or (b) they use the quantity-dependence of taxes to ensure that households' post-tax factor income varies with factor supply in the same as as under Pigouvian taxes. The proof formalizes these conditions.

In the second result above, and throughout the paper, I use "implement" in the Ramsey sense—i.e. to mean that there exists an equilibrium whose equilibrium path is the allocation in question—rather than the unique implementation sense from mechanism design.

The first result provides a fairly weak welfare guarantee for Pigouvian equilibria: They cannot be locally improved to first-order. A notable corollary is that, if the economy is convex in an appropriate sense—namely $\mathcal{A}$ convex and $F(L, L)$ convex in L—then Pigouvian equilibria are globally Pareto efficient, a first welfare theorem.[14]

So, where does Pigouvian taxation go wrong in non-convex economies? The second result provides an answer: Pigouvian taxation can implement any first-order Pareto efficient allocation—i.e. every "local optimum." In non-convex economies, there may be many such allocations, only one of which is globally efficient. In other words, the shortcoming of Pigouvian policy is not its inability to support efficient outcomes, but rather its ability to rule out some inefficient outcomes—as we saw in the example of Section 2 (where the economy was non-convex). Pigouvian taxation runs the risk that households will land on a Pareto-dominated local maximum, i.e. will fall into a *coordination failure*.

Given that Pigouvian policy implements not only global but also local optima, there is room for an alternative policy to do better. The third result narrows the scope for such improvements by clarifying that, *along the path of any policy that implements efficient allocations*, policy must (at least locally) be set according to the Pigouvian rule.

This leaves room for two places in which an alternative policy might differ from Pigou. First, it could differ off of the equilibrium path. However, this could not cause households—each of whom expect that the economy will remain on the equilibrium path whatever actions they take—to take different actions in the inefficient equilibria we would like to change. The second possibility is more subtle. On one hand, if we succeed in constructing an alternative policy that satisfies guarantees efficiency, then—as every equilibrium under that policy is efficient—the third result in Proposition 1 implies those equilibria must be Pigouvian. So the new policy must be Pigouvian in every strategy profile that is an equilibrium under the policy. On the other hand, Proposition 1 says nothing how taxes are set in hypothetical strategy profiles that are not equilibria. While, by definition, such profiles never arise in equilibrium, this may be precisely because of the incentives our alternative policy rule provides when applied to them. Given that the problem with Pigouvian policy is over- rather than under-inclusiveness, this is a promising approach.

Indeed, this is the avenue taken by the "super-Pigouvian" policy introduced next. Super-Pigouvian policy manages to set taxes so as to (a) be Pigouvian in efficient allocations and (b) prevent all other allocations from being consistent with household optimization. This

[14] A version of this "Pigouvian first welfare theorem for convex economies" appears in Starrett [1972].

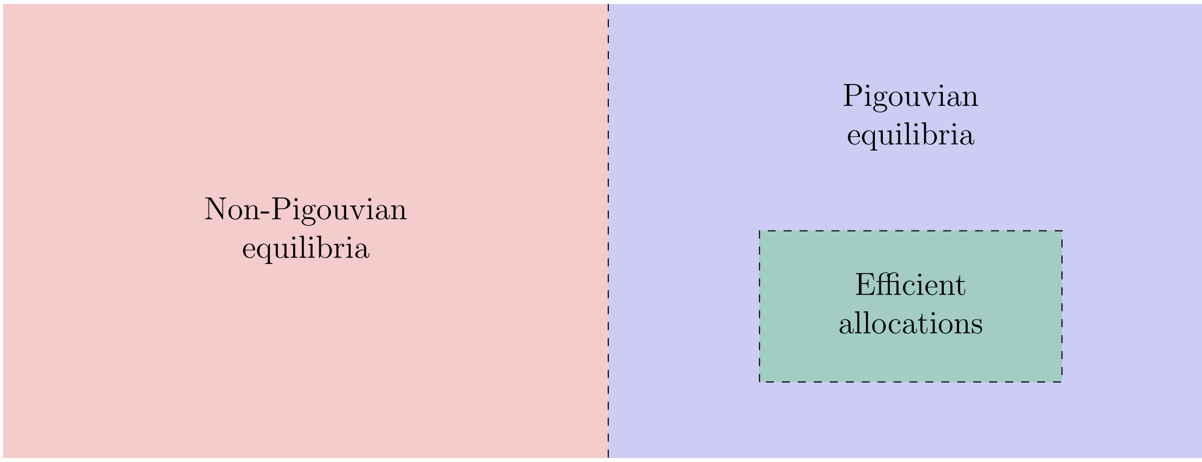

Figure 3: Summary of Pareto efficiency properties of Pigouvian equilibrium. Both efficient and inefficient allocations can be implemented as Pigouvian equilibria. However, all efficient allocations can be implemented as Pigouvian equilibria, and they (essentially) cannot be implemented in non-Pigouvian equilibria.

resolves the coordination failure problem of Pigouvian policy while, at the same time, policy remains Pigouvian in any strategy profile that actually arises as an equilibrium.

5 Super-Pigouvian policy

In the previous section, I showed that Pigouvian policy can implement any Pareto efficient allocation—and moreover that any policy implementing a Pareto efficient allocation must effectively be Pigouvian along the equilibrium path—and yet Pigouvian policy can also implement Pareto inefficient allocations.

In this section, I propose an alternative policy that, unlike Pigouvian policy, always guarantees an efficient outcome—i.e. satisfies a first welfare theorem, even in non-convex economies. For lack of a better name, I call this *super-Pigouvian policy*.

5.1 Defining super-Pigouvian policy

In order to highlight the similarities between Pigouvian and super-Pigouvian policy, I begin by rearranging the definition of Pigouvian policy. To start, since the Pigouvian post-tax wage is equal to a household's entire marginal contribution over time, adding these contributions starting from any time $t > 0$ until the next time t' at which i_{t-1} can adjust implies that for any $\ell \in \mathcal{A}$,

$$\sum_{s=t}^{t'} R^{-(s-t)} \boldsymbol{\omega}_s^{i_{t-1}}(\ell, h^0) \cdot \ell \quad = \quad \sum_{s=t}^{t'} R^{-(s-t)} \frac{d}{dL}\bigg|_{L=L_s(h^0)} F(L,L) \cdot \ell. \tag{12}$$

This implies that, for any change $\Delta\ell$ in i_{t-1}'s factor supply between t and t' (away from $\boldsymbol{\ell}_t^{i_{t-1}}(h^t)$) the change corresponding in i_{t-1}'s NPV income is

$$\approx \sum_{s=t}^{t'} R^{-(s-t)} F\left(L_s(h^0) + \Delta\ell, L_s(h^0) + \Delta\ell\right) \; - \; \sum_{s=t}^{t'} R^{-(s-t)} F\left(L_s(h^0), L_s(h^0)\right). \tag{13}$$

Finally, note that if we add to both sides the effect that this $\Delta\ell$ has on i_{t-1}'s factor supply disutility between t and t', then the left-hand side equals her full perceived change in consumption-equivalent utility from deviating by $\Delta\ell$. Let $\Delta U^{i_{t-1}}(\Delta\ell, h^t)$ denote this change starting from an arbitrary history h^t.[15] Similarly, if we add the change in i_{t-1}'s factor supply disutility to the RHS, then it represents the full effect of the $\Delta\ell$ deviation on the NPV of aggregate output less factor supply disutility over the entire future following t. This is because $\Delta\ell$ (a) only affects output until period t' and (b) does not affect the labor disutility of any households other than i_{t-1}. Let $W(\{\ell_s^j\})$ denote the NPV of aggregate output less factor supply disutility on any path of factor supply $\{\ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}$.[16] Putting these observations together, we obtain the following description of Pigouvian policy:

Lemma 1. *Suppose $\frac{d^2}{dL^2}F(L,L)$ exists and is bounded. If policy is Pigouvian, then for all histories $h^t = (\boldsymbol{\ell}_t^{i_{t-1}}(h^0), ..., \boldsymbol{\ell}_1^{i_0}(h^0))$ on the equilibrium path—we have*

$$\Delta U^{i_{t-1}}(\Delta\ell, h^t) \; = \; W\left(\left\{\boldsymbol{\ell}_s^j(h^t) + \mathbb{1}_{s\leqslant t'}^{j=i_{t-1}}\Delta\ell\right\}\right) \; - \; W\left(\left\{\boldsymbol{\ell}_s^j(h^t)\right\}\right) \; + \; o(\Delta\ell). \tag{14}$$

In words, Pigouvian policy ensures that, in adjusting their factor supply at a point along

[15] Formally, for any history $h^{t>0} = (h_{t-1}, ..., h_0)$, and any $\Delta\ell$ such that $h_{t-1} + \Delta\ell \in \mathcal{A}$, define

$$\begin{aligned}\Delta U^{i_{t-1}}(\Delta\ell, h^t) \equiv & \sum_{s=t}^{t'} R^{-(s-t)} \left[\boldsymbol{\omega}_s^{i_{t-1}}(\boldsymbol{\ell}_s^{i_{t-1}}(h^t) + \Delta\ell, h^t) \cdot (\boldsymbol{\ell}_s^{i_{t-1}}(h^t) + \Delta\ell) - v(\boldsymbol{\ell}_s^{i_{t-1}}(h^t) + \Delta\ell)\right] \\ & - \sum_{s=t}^{t'} R^{-(s-t)} \left[\boldsymbol{\omega}_s^{i_{t-1}}(\boldsymbol{\ell}_s^{i_{t-1}}(h^t), h^t) \cdot \boldsymbol{\ell}_s^{i_{t-1}}(h^t) - v(\boldsymbol{\ell}_s^{i_{t-1}}(h^t))\right],\end{aligned}$$

where t' is the next time after $t-1$ at which i_{t-1} can adjust her factor supply (or $|\mathcal{T}|$ if no such t' exists).

[16] Formally,

$$W\left(\left\{\ell_s^j\right\}_{s\geqslant t}^{j\in\mathcal{I}}\right) \equiv \sum_{s\geqslant t} R^{-(s-t)} \left[F\left(\sum_{j\in\mathcal{I}} \ell_s^j, \sum_{j\in\mathcal{I}} \ell_s^j\right) - \sum_{j\in\mathcal{I}} v(\ell_s^j)\right].$$

the equilibrium path, a household i's perceived private incentive equals her effect on the NPV of aggregate output less labor disutility —provided that all other households (including i in the future) behave as if i had played her equilibrium action.

The definition of super-Pigouvian policy below differs in two key ways. First, in measuring the social benefit of any deviation in factor supply, it accounts for how households who act after any deviation *change their behavior* in response, according to the competitive equilibrium that begins from the resulting history. Second, it considers the benefits of deviations away from *all possible histories*, not just those on the equilibrium path.

Definition 4. *An equilibrium is* ***super-Pigouvian*** *if, for all histories $h^{t>0} = (h_{t-1}, ..., h_0)$ and for all $\Delta\ell$ such that $h_{t-1} + \Delta\ell \in \mathcal{A}$,*[17]

$$\Delta U^{i_{t-1}}(\Delta\ell, h^t) \;=\; W\left(\left\{\boldsymbol{\ell}^j_s\Big((h_{t-1} + \Delta\ell, h_{t-2}, ..., h_0)\Big)\right\}\right) \;-\; W\left(\left\{\boldsymbol{\ell}^j_s(h^t)\right\}\right). \qquad (15)$$

The first essential difference between Pigouvian and super-Pigouvian policy is that the latter incentivizes households to change welfare not only directly, through the effects of their own actions, but also indirectly, through their actions' influence on other households behavior. As in the motivating example of Section 2, this incentivizes households to behave as if they understood their ability to change the behavior of others with whom they seek to coordinate. This helps prevent self-fulfilling expectations and coordination failure.

The second essential difference is that super-Pigouvian policy applies both on and off of the equilibrium path, whereas Pigouvian policy need only apply on the equilibrium path. The role of this distinction is rather subtle, as—since households believe that (however they behave) prices and transfers will follow the equilibrium path—taxes set off of the equilibrium path do not affect behavior on the equilibrium path. To the extent that we are only concerned with efficiency along the equilibrium path, therefore, it would appear that the off-path definition of super-Pigouvian policy is irrelevant. However, note that *the definition of super-Pigouvian policy on the equilibrium path depends on outcomes that occur off of the equilibrium path*, through off-path welfare. As a result, the precise way in which off-path welfare is constructed can affect the design of super-Pigouvian policy on path and therefore what allocation it implements. As I will later show, the assumption that super-Pigouvian policy holds off of the equilibrium path is precisely the right way to discipline off-path welfare in order to guarantee that the policy succeeds on path.[18]

[17] Although this definition only *implicitly* defines marginal taxes, one may rearrange it in order to obtain an explicit definition for marginal taxes using the definition of $\Delta U^{i_{t-1}}$. Intuitively, $\Delta U^{i_{t-1}}$ contains pre-tax wage payments, tax payments, and labor disutility; one may simply bring market wages and labor disutility to the other side of the equation. Note that taxes are still only defined up to the re-shuffling of payments over time between t and the next period when i_{t-1} may adjust her behavior.

[18] Pigouvian and super-Pigouvian policy also differ in a fourth, less essential, way that corresponds to

Note that whereas Pigouvian policy only depends on production externalities, the super-Pigouvian policy depends on all of an acition's impacts on production, as well as its effects on factor supply disutility. The reason these terms do not appear in the Pigouvian case is that a household already internalizes its own direct effect on production through its wage as well as its own direct effect on its factor supply disutility. This is no longer the case once one accounts for a household's impacts on *other* households.

5.2 Avoiding coordination failure with super-Pigouvian policy

In Section 4, I showed that Pigouvian policy can implement any efficient allocation, but also runs the risk of implementing inefficient allocations if there is coordination failure. Above, I defined a new policy—super-Pigouvian policy—that differs from Pigouvian policy in that (a) it compensates actors for both their direct effects on welfare and their indirect effects, through impacts on other households' behavior and (b) it requires that private and social incentives align both on and off the equilibrium path. Are these two modifications enough to rule out coordination failure while still supporting efficient equilibria?

Our next two results show that this is the case. First, super-Pigouvian policy satisfies a first welfare theorem.

Theorem 1 (Super-Pigouvian first welfare theorem)**.** *If an equilibrium is super-Pigouvian, then it is Pareto efficient starting from any history.*[19]

Both of super-Pigouvian policy's distinguishing features—the consideration of indirect effects on welfare and the requirement that policy be set off path—play key roles in this result. The first feature ensures that policy effectively "sells the future of the economy" to the household i who acts at t, by compensating her for all other households' willingness to pay for the future her action creates. So i acts in the public interest, taking as given her equilibrium belief about how the next actor will behave along the path she starts them on. The second feature ensures that this next actor, too, acts in the public interest, as do all who follow her. Since, at every history, these households act as the planner would, the principle of optimality from dynamic programming ensures the equilibrium path is efficient.

Apart from the design of super-Pigouvian policy, the other key building block of Theorem 1 is the assumption that only one household acts at each time. This rules out the possibility

the difference between the $o(\Delta\ell)$ term in (14). Namely, Pigouvian taxation compensates households for their marginal impact on production at existing quantities rather than their total impact. However, this difference vanishes in the large household limit, where any individual household's action is small relative to the curvature of the production function.

[19]This result relies on an additional technical assumption: The function $F(\sum_{i\in\mathcal{I}}\ell^i, \sum_{i\in\mathcal{I}}\ell^i) - \sum_{i\in\mathcal{I}}v(\ell^i)$ is bounded over across all $\{\ell^i\}^{i\in\mathcal{I}} \in \mathcal{A}^{\mathcal{I}}$.

of coordination failure between two households acting simultaneously. With one household acting at a time and perfect alignment of household incentives with social welfare, it is as if a social planner simply hires one household to "run the economy" in each period.

A potential concern with Theorem 1 is that it could hold vacuously by simply ensuring that no super-Pigouvian equilibria exist. My next result ensures that this is not the case.

Theorem 2 (Super-Pigouvian second welfare theorem)**.** *Suppose that, given any history h^t, some functions $\{\boldsymbol{c}^i, \boldsymbol{\ell}^i\}$ return Pareto efficient paths of consumption and factor supply following that history and satisfy the on-path consistency condition* (7)*. Then $\{\boldsymbol{c}^i, \boldsymbol{\ell}^i\}$ can be implemented as a super-Pigouvian equilibrium.*

5.3 Decentralization and super-Pigouvian policy

So far, I have shown that super-Pigouvian policy satisfies strong welfare properties, ruling out coordination failure and ensuring Pareto efficiency. However, skeptics will rightly point out that it is not the only such policy. For example, a planner could implement the efficient allocation by simply regulating quantities—i.e. prohibitively taxing all inefficient actions.

If its welfare properties alone do not uniquely characterize super-Pigouvian policy, what does? My next result answers this question by showing that any policy rule that *decentralizes* efficient decisions in an appropriate sense must be super-Pigouvian policy.

To state this result, I first introduce the notion of a generic policy rule. A policy rule is simply a method for setting history-contingent marginal taxes in response to (a) economic fundamentals and (b) households' equilibrium factor supply strategies.

Definition 5. *A **policy rule** is a function that, given any fundamentals θ and any on-path consistent factor supply strategies $\{\boldsymbol{\ell}^j\}$, returns some on-path consistent marginal tax functions $\{\boldsymbol{\tau}^j(\cdot;\cdot;\theta,\boldsymbol{\ell})\}$.*[20]

Super-Pigouvian taxation describes a particular policy rule characterized by Equation 15 holds.[21] In words, super-Pigouvian policy incentivizes households according to the full welfare impact of their behavior at fundamentals θ, accounting for reactions of other households according to their strategies $\boldsymbol{\ell}$. This policy rule has two stark properties: First, it guarantees that households make efficient choices, even if the behavior that follows them is inefficient. Since they are incentivized according to their actions' total welfare impact, taking as given the strategies using which the economy will respond, their actions are constrained Pareto efficient given those responses. Second, it compensates households for the social benefits of

[20]By on-path consistent, I mean that (7) holds at every history.

[21]Note that super-Pigouvian policy is only defined up to re-shuffling of payments to each household i between periods during which i cannot adjust her behavior, as such changes have no effect on incentives.

their actions, but leaves them to appropriately internalize their private costs (factor supply disutility) in a decentralized fashion. Concretely, taxes set according to the super-Pigouvian rule are, at any history, invariant to changes in the actor's current factor supply disutility.

Remark. *The model of Section 3 assumes that all households have the same factor supply disutility function, v. However, this makes it impossible to change the factor supply disutility of the current actor—as mentioned above—without changing disutilities of future actors, and therefore the efficient allocation. I therefore assume for the remainder of the section that households have heterogeneous and time-varying factor supply disutilities v_t^i. I discuss this extension formally in Section 5.5 and show that it does not affect my other results.*

My next result shows that these two properties are in fact unique to super-Pigouvian taxation: Any policy rule that implements efficient outcomes without depending on the current actor's factor-supply disutility must be super-Pigouvian.

Proposition 2. *Fix a policy rule. Suppose that for all fundamentals θ, factor supply strategies $\boldsymbol{\ell}$, and feasible histories h^{t-1}:*

- *Any solution to i_{t-1}'s problem of adjusting its factor supply for period t subject to the policy rule's taxes is Pareto efficient, given that future factor supply follows $\boldsymbol{\ell}$.*[22]
- *For any alternative fundamentals θ' that differ from θ only in terms of i_{t-1}'s factor supply disutility at t, taxes $\boldsymbol{\tau}_t^{i_{t-1}}(\cdot; h^{t-1}; \theta', \boldsymbol{\ell})$ are the same as at θ.*

Then taxes are super-Pigouvian at θ, $\boldsymbol{\ell}$, and h^{t-1}, i.e. (15) *holds at $h^t \equiv \boldsymbol{\ell}_t^{i_{t-1}}(h^{t-1}) \frown h^{t-1}$.*

The main idea behind the Proposition is that any non-super-Pigouvian policy rules offer households net incentives that are not aligned with social welfare. For some factor supply disutilities, the discrepancy may be innocuous, not interfering with the actor's optimal decision. However, there always exists a factor supply disutility for which—if policy cannot update in response to the specific disutility profile—it must result in an inefficient choice.

5.4 Informational requirements for super-Pigouvian policy

When comparing standard Pigouvian policy to full central planning, advocates point to an additional advantage Pigouvian policy: It can be implemented with limited information on the part of a social planner. In particular, a planner can evaluate whether an equilibrium is Pigouvian using only knowledge about the extent of externalities, by checking whether

[22] More formally, any profile constructed by combining this path for factor supply with a path of consumption that satisfies household optimization and market clearing is Pareto efficient among feasible paths that start from h^{t-1} and whose factor supply obeys $\boldsymbol{\ell}$ after time t.

subsidies indeed align with factors' marginal output contributions through the externality channel, and updating them if not. This does not require her to know the efficient allocation.

It is less obvious whether super-Pigouvian policy should satisfy a similar "test," since its definition depends on not only more primitives but also households' endogenous behavior. My next result shows a social planner may nevertheless verify whether taxes are super-Pigouvian using a simple test that only depends on households' willingnesses to pay for one another's factor supply adjustments.

Proposition 3. *An equilibrium is super-Pigouvian if and only if for all histories $h^{t>0} = (h_{t-1}, ..., h_0)$ and all alternative histories $h'^t = (h_{t-1} + \Delta\ell, h_{t-2}, ..., h_0)$ with $h_{t-1} + \Delta\ell \in \mathcal{A}$,*

$$\sum_{j \neq i_{t-1}} e^j \left(\left\{ \boldsymbol{c}_s^j(h'^t), \boldsymbol{\ell}_s^j(h'^t) \right\}_{s \geqslant t}, \left\{ \boldsymbol{c}_s^j\left(h^t\right), \boldsymbol{\ell}_s^j\left(h^t\right) \right\}_{s \geqslant t} \right) \quad + \quad \delta_{SP}(\Delta\ell, h^t) \quad = \quad 0, \qquad (16)$$

where $\delta_{SP}(\Delta\ell, h^t)$ is a "correction" term that goes to zero in any limit in which individual households' impacts on prices and transfers vanish.

In words, taxes are super-Pigouvian if and only if the aggregate willingness to pay of all households other than i_{t-1} for a change in her behavior is zero. This is the case when factor supply subsidies to i_{t-1}—which these households finance by accepting smaller lump-sum transfers—exactly offset the other benefits of i_{t-1}'s actions to other households, making them indifferent.

This condition is akin to the willingness-to-pay condition often used as the definition of Pigouvian taxation for externalities that manifest in households' utility functions (rather than production): Taxes on an action are Pigouvian if and only if they equal aggregate willingness to pay for that actions' effect on utility through the externalities it generates, holding all other behavior constant. Proposition 3 provides a similar condition, but allows for willingnesses to pay to account for indirect effects—through impacts on other households' behavior—and for all channels, rather than just the externality channel.

While super-Pigouvian policy can be implemented with relatively limited information on the part of a social planner, doing so requires that households have somewhat more information than neoclassical models assume. Typically, households are assumed to have rational expectations about the evolution of prices and transfers along the equilibrium path; this is what allows them to make optimal forward-looking decisions. However, for a planner to elicit from households their willingnesses to pay for one another's behavior, households must also have rational expectations about how prices and transfers will evolve *off of the equilibrium path.* More precisely, it must be the case that, should any history h^t be reached, all households will have rational expectations of prices and transfers both along the equilibrium path

	Laissez faire w/o externalities	Pigou	Full central planning	Super-Pigou
Welfare guarantee	Efficient	First-order efficient	Efficient	Efficient
Planner info requirement	None	Externalities or household WTPs	Know efficient allocation	Household WTPs
Household info requirement	Rational expectations	Rational expectations	None	Rat'l expectations on and off path

Table 1: Welfare guarantees of various policies and the information they require from the social planner and households. The first-column is not discussed in the main text; it refers to the first welfare theorem without market failures.

following h^t and along the equilibrium path that follows any one-step deviation away from h^t. This is a more demanding requirement.

Note, however, that this does not imply that any individual household is *so* knowledgable that, on its own, it knows the efficient path of the economy. Households only need know only the equilibrium path of prices and transfers starting from one step deviation away from any history they may reach—not the path of aggregate welfare. In this sense, the economy still aggregates information, achieving an allocation that no household could on its own.

Above, I have emphasized the possiblity of implementing super-Pigouvian policy without the social planner having any information other than household willingnesses to pay. The downside of this approach is that it imposes a larger informational burden on households than is standard in welfare economics, requiring them to have rational expectations both on and off of the equilibrium path. In practice, government institutions could share this burden with households.

As a point of comparison, consider how governments implement Pigouvian carbon pricing: First, they fund research and employ scientists who model the causal relationship between tons of CO2 that factories emit and actions in people's lives, such as higher temperatures and flooded homes. This step is purely descriptive. Second, policymakers assign a monetary cost to these outcomes. This step is normative, and—in order to implement efficient outcomes—must assign costs that align with households' willingness to pay to avoid adverse outcomes.

In a similar fashion, super-Pigouvian policy could be implemented using a mix of technocratic analysis and household preference reports. Social scientists could first conduct empirical analyses to gauge how any one household's decisions—say, to transition to a new sector—affect the sectoral choices of other households. They could similarly gauge how sectoral growth affects wages. Second, households could report their willingness to pay for higher wages in various sectors.[23] Such an approach could reduce the information needed by households without requiring either households or the government to know the efficient allocation. Of course, the viability of such an arrangement hinges on the assumption that social scientists can estimate a wide range of behavioral elasticities. While this may be challenging, the problem is analogous to estimating the technologies elasticities on which standard Pigouvian policy relies—also a difficult task in the case of, say, climate change.

5.5 Extensions

With an eye toward the quantitative application presented in the next section, I now explain how my characterization of super-Pigouvian policy extends to environments with various forms of heterogeneity, uncertainty, and imperfect competition. The same extensions also apply to my characterization of Pigouvian policy.

Heterogeneity and time variation: The model of Section 3 assumes that all households have the same intra- and inter-temporal preferences over consumption and labor supply. Additionally, it assumes that the production function and the interest rate are constant over time. All of my results continue to hold in the more general model where:

- Each household i has a discount factor $\beta^i \in (0,1)$, time-varying factor supply disutility functions $\{v_t^i\}$, time-varying consumption utility functions $\{u_t^i\}$, and time-varying action sets $\{\mathcal{A}_t^i\}$. Preferences are still GHH and satisfy the same technical properties, and action sets are consistent with the timing of behavior.[24]
- The production function F_t varies over time. The production function still satisfies the same technical properties within each period.
- The interest rate $R_t > 0$ varies over time. The quantity $\sum_{t=0}^{\infty}(R_t)^{-t}$ is still bounded.

Uncertainty: The model of Section 3 is completely deterministic. All of my results continue to hold in the more general model where:

- Household preferences, the production function, and the variable i_t that determines the timing of household actions are all stochastic processes, measurable with respect to a common filtration.

[23] To first order, such reports are unnecessary, as a household's WTP for a change its wage is simply the wage change itself. However, extrapolating WTPs for larger wage changes—which might be particularly relevant if small changes in one household's action triggered large shifts in aggregate outcomes—requires knowledge of preferences.

[24] That is, if i does not act between t and t', then $\mathcal{A}_t^i = \mathcal{A}_{t'}^i$.

- All households have accurate beliefs about these processes, and all uncertainty is public—there is no imperfect/asymmetric information.
- Households have access to perfect insurance at exogenous and risk-neutral prices.

In this model, histories must include all realizations of uncertainty to date, so that equilibrium variables—including taxes—are interpreted as filtration-measurable functions of the underlying state of the world.

Imperfect competition: It is well known that models of imperfect competition can be represented in reduced form as models of perfect competition with external economies of scale—i.e. externalities—when firms are monopolistically competitive within each sector and face CES demand [Kucheryavyy et al., 2016]. Below, I provide a new "recipe" for reinterpretting imperfect competition as externalities under far more general assumptions. Specifically, I require that firms make their taxable decisions—such as entry or technology adoption—competitively, but allow for arbitrary forms of competition given these decisions.

- Households' factor supply decisions are interpretted generally, so as to include decisions such as operating a firm of a particular type or adopting a new technology at such firms. Similarly, households' factor incomes are interpretted generally, so as to include income sources such as profits. Households take as given the post-tax income schedule that corresponds to these decisions.[25]
- To the extent that any household has additional sources of income stemming from her additional factor supply decisions—such as labor income and dividend income—these can be taxed jointly and with quantity-dependence.
- There exists a (differentiable) per-unit-supplied income function $\hat{w}(\cdot)$ such that at any factor supply profile $\{\ell_t^j\}^{j\in\mathcal{I}}$, each household i has pre-tax income $\hat{w}(\sum_{j\in\mathcal{I}} \ell_t^j) \cdot \ell_t^i$.
- At every factor supply profile, aggregate output equals aggregate pre-tax income.

Under these assumptions, the production side of the economy can simply be reinterpreted as deriving from a perfectly competitive model with externalities. In particular, the production function F of this perfectly competitive economy may be defined as

$$F(L, \overline{L}) \quad \equiv \quad \hat{w}(\overline{L}) \cdot L. \tag{17}$$

[25]For example, this formulation does not allow a household to open fewer firms out of concern that her marginal firm will decrease the profits of her inframarginal firms. However, it does allow these firms to themselves make imperfectly competitive choices—such as marking up prices or marking down wages.

6 Quantitative application

So far, I have compared the design and welfare properties of Pigouvian and super-Pigouvian policies in a general, theoretical setting. This leaves open an important practical question: Could super-Pigouvian industrial policies produce quantitatively meaningful welfare gains?

I now address this question by assessing the impacts of super-Pigouvian policy in a calibrated dynamic model of structural transformation in a small open economy. This model retains the basic spirit of the motivating example in Section 2 but captures several key institutional features that make it possible to calibrate empirically. In the model, industrial firms decide whether to pay a fixed cost to use a more productive technology. Importantly, fixed costs are paid in goods, labor supply to the industrial sector is somewhat elastic, and firms use intermediate inputs in production. At realistic parameter values, these features make investment decisions complementary and can generate multiple equibria. The model fits into the extensions described in Section 5.5, so my earlier theoretical results apply.

I calibrate this model using South Korea's targetted industrial policy to the heavy and chemical industries—which I interpret as the "industrial sector"—in the 1970s. In particular, I interpret the rapid expansion of these industries as a transition between steady states and use this assumption to back out the parameters of my model from statistics reported in Lane [2022]. Under this calibration, the model admits multiple, Pareto-ranked Pigouvian equilibria. Finally, I simulate the calibrated model in order to assess the gains from super-Pigouvian policy.

6.1 Model

Motivated by my application to South Korea in the 1970s, I develop a model of a structural transformation featuring an industrial sector with strong complementarities in investment. I interpret the industrial sector as a reduced form representation of the heavy and chemical industries (HCI), which were directly targetted during by "HCI drive" policies between 1973 and 1979. The static version of the model is similar to Buera et al. [2021] except without firm entry, but—as staggered actions over multiple time periods are essential to super-Pigouvian policy—I embed it in a dynamic setting where firms' technology adoption decisions are inertial. Similarly to Frankel and Pauzner [2000], I assume that productivity in the industrial sector follows a stochastic process, with the goal of "smoothing out" the responsibilty of individual households for dramatic transitions between steady states. However, I depart from Frankel and Pauzner [2000] by assuming that productivity shocks are *mean-reverting*, so as to ensure that the model can generate equilibrium multiplicity and, therefore, coordination failure.

6.1.1 Environment

A representative worker and a unit mass of entrepreneurs interact in a traditional and an industrial sector over a discrete, infinite horizon of periods $t \in \mathcal{T}$ spaced an amount Δ. At any time t, N_t entrepreneurs produce using a "new" technology and $1 - N_t$ use an "old" technology.

Households: The population contains two types of households: workers and entrepreneurs. While agents of both types are "households" in the sense of the general model in Section 3, the entrepreneurs are the main economic actors of interest and those whose behavior super-Pigouvian policy can influence in an interesting way.

First, a representative worker inelastically supplies $M\Delta$ total units of labor in each period. Each infinitesimal unit m of this labor has productivity 1 in the industrial sector and productivity $a(m)$ in the traditional sector, distributed according to the CDF

$$\mathbb{P}[a(m) \leqslant a] \quad = \quad \max\left[\left(\frac{a}{\bar{a}}\right)^{\varepsilon}, 1\right] \tag{18}$$

for some $\varepsilon, \bar{a} > 0$. I assume that labor is allocated within the market—i.e. its sectoral allocation cannot be taxed—and that it can move across sectors freely. Under a parametric assumption ensuring $\bar{a}$ is large enough, this implies supply curves for industrial labor L_t^I and efficiency units of traditional labor L_t^T:

$$L_t^I = M\left(\frac{w_t^I}{\bar{a} w_t^T}\right)^{\varepsilon}, \qquad L_t^T = \frac{\varepsilon}{\varepsilon + 1} M \left(\bar{a} - \frac{(w_t^I / w_t^T)^{1+\varepsilon}}{\bar{a}^{\varepsilon}}\right) \tag{19}$$

where w_t^I and w_t^T are the industrial and traditional sector wages per effective unit of labor.

Second, a finite set of entrepreneurs $i \in \mathcal{I}$ each own and operate a mass $1/|\mathcal{I}|$ of industrial firms. Entrepreneurs choose whether to operate each firm using the new or old technology, but face an adjustment friction. Concretely, each $i \in \mathcal{I}$ may adjust her technology decisions at the end of any period t if and only if $i_t = i$, where i_t is a scalar random variable. i_t is distributed independently across time and symmetrically across mobile households, taking each value i with probablity $\delta\Delta < 1/|\mathcal{I}|$.[26] The frictional nature of entrepreneurs' technology choice makes it akin to an investment—an interpretation I leverage in the model's calibration.

Intertemporally, households of all types save at an exogenous interest rate and buy in-

[26] i_t can also take the value 0, in which case no household may adjust.

surance in perfect and risk-neutral markets in order to maximize expected lifetime utility

$$\mathbb{E}\left[\sum_{t\in\mathcal{T}} \beta^{-t} u(c_t)\Delta\right], \tag{20}$$

where u is increasing and differentiable. Households' consumption strategy—as well as their technology decisions—may depend not only on the history of those decisions, but also on the history of the random process for mobility, i_t, and the random process for productivity, z_t, introduced below.

Production: In the traditional sector, a representative, competitive firm produces the final good one-for-one with effective units of labor and pays a competitive wage w_t^T per effective unit of labor.

At any time t, the industrial sector contains two types of firms: intermediate variety producers operating with either the new or old technology and a representative, competitive final goods producer that aggregates these varieties.

At each time t, each intermediate variety n pays a (flow) fixed cost $f(n)$ in units of the final good and produces output $q_t(n)$ according to a Cobb-Douglas production function

$$q_t(n) \quad = \quad z(n)\ l_t(n)^{\alpha}\ x_t(n)^{1-\alpha}, \tag{21}$$

where $\alpha \in (0,1)$ is the labor share, $l_t(n)$ is labor, and $x_t(n)$ is intermediate inputs. Which technology n uses determines its fixed cost, productivity shifter, and labor share. The new technology has higher fixed costs $f^N > f^O$ but also higher productivity $z^N > z^O$.

Each intermediate firm buys labor and intermediate inputs competitively at wages w_t^I and final goods price 1, but sets its selling price $p_t(n)$ to maximize profits given the final goods producer's demand function for inputs. The firm therefore earns profits

$$\pi_t(n) \quad = \quad p_t(n)q_t(n)\ -\ x_t(n)\ -\ w_t^I l_t(n)\ -\ f(n). \tag{22}$$

I assume that variety producers must operate at their fixed technology, even if they earn negative profits.

At each time t, a representative, competitive final goods producer produces gross output Q_t according to a CES production function

$$Q_t \quad = \quad z_t\ \left[\int_0^1 X_t(n)^{\frac{\eta-1}{\eta}} dn\right]^{\frac{\eta}{\eta-1}}, \tag{23}$$

where $\eta > 1$ is the elasticity of substitution across varieties, $X_t(n)$ is the quantity of each variety n used as an input, and z_t is a productivity shock. I assume $\log z_t$ follows an AR(1) process

$$\log z_t \quad = \quad \log z_{t-1} \quad - \quad \theta \Delta \log z_{t-1} \quad + \quad \sigma \sqrt{\Delta} \epsilon_t \tag{24}$$

where $\theta, \sigma > 0$ and $\epsilon_t \sim N(0,1)$ i.i.d. I normalize the price of the final good to one.

Policy: A government taxes and subsidizes firms' technology adoption decisions, but not the allocation of workers to sectors, which occurs within the market. Taxes on aggregate labor supply and firm entry are irrelevant as both are exogenously given and the government has access to lump-sum transfers. Taxes and transfers may condition not only on the history of technology adoption decisions, but also on the history of the random processes i_t and z_t and household identity. The government saves and trades in insurance so as to balance its budget intertemporally.

Market clearing: In each period, labor markets clear in both the traditional and industrial sectors, and input variety markets clear. Starting from any history, the final good market clears in net present value given debts accumulated along the history.

6.2 Investment complementarities and multiplicity

A key question in this model is whether Pigouvian policy can support multiple equilibria. If so, there is scope for coordination failure and therefore welfare gains via super-Pigouvian policy. Before explicitly computing the model's equiilbria, I discuss key economic mechanisms that determine whether there is multiplicity.

As a first step, I express aggregate net output as the sum of labor income and profit income net of fixed costs.

$$Y_t \quad = \quad \underbrace{\text{const} \ + \ \frac{w_t^I L_t^I}{1+\varepsilon}}_{\text{Aggregate labor income}} \quad + \quad \underbrace{\frac{w_t^I L_t^I}{\alpha(\eta-1)}}_{\text{Profits of industrial firms}} \quad - \ \underbrace{\left(f^N N_t \ + \ f^O(1-N_t)\right)}_{\text{Fixed costs}}. \tag{25}$$

The fact that labor income varies less than one-for-one with industrial sector income reflects that each unit of labor allocated to the industrial sector has an opportunity cost in the traditional sector. Using the industrial labor supply curve to express industrial labor supply in terms of wages and solving for wages using the optimization condition of final producers, we have that for some constant $\kappa > 0$,

$$Y_t = \text{const} + \kappa \cdot (z_t)^{\frac{1+\varepsilon}{\alpha}} \left[N_t(z^N)^{\eta-1} + (1-N_t)(z^O)^{\eta-1}\right]^{\frac{1}{\eta-1}\frac{1+\epsilon}{\alpha}} - \left(f^N N_t + f^O(1-N_t)\right). \tag{26}$$

Finally, recall that under Pigouvian investment subsidies, the difference in profit between firms operating new and old technologies is equal to the marginal impact of technology adoption on aggregate output:

$$\pi_t^N - \pi_t^O = \frac{dY_t}{dN_t}. \tag{27}$$

We can now address the question of whether firms' technology investments are complementary under Pigouvian policy. In other words, does investing generate more output when more other firms invest? Equation 27 makes clear that this is the case if and only if Y_t is convex in N_t, i.e.—by Equation 26—if and only if $1 + \varepsilon > \alpha(\eta - 1)$.

The parameters ε, α, and η all play intuitive roles in determining whether investments are complementary. First, the higher is elasticity of industrial labor supply, ε, the more an increase in industrial wages draws in workers from the traditional sector—so that one firm's productivity growth does not limit another's ability to hire workers. Second, when the labor share, α, is low, improvements in productivity increase gross output above and beyond their effects on value added, which raises the value of further productivity improvements (by Hulten's theorem). Finally, the lower is the elasticity of substitution across varieties, η, the less does any one firm's productivity improvement detract from the importance of improvements in other varieties.

The calibration in the next section finds that this inequality holds, i.e. technology investments are complementary. Still, the mere presence of such complementarities is not sufficient to guarantee there are multiple Pigouvian equilibria. In a *static* setting, complementarities in investment imply multiplicity so long as productivity is neither too high nor too low. However, this logic breaks down because the model is *dynamic* and the productivity process z_t can reach extreme values at which choosing either technology is optimal however other households will behave. Under such conditions, results from the global games literature imply there is still a unique equilibrium in the continuous time limit where $\Delta \to 0$ if productivity follows a random walk [Frankel and Pauzner, 2000]. Nonetheless, these results do not apply if productivity is mean reverting—as will be the case in my calibration—so it will indeed be the case that investment complementarities support equilibrium multiplicity.[27]

[27] Frankel et al. [2005] show that mean-reversion is consistent with equilibrium uniqueness if it only lasts for a certain number of periods. I assume the productivity process has a constant level of mean-reversion over an infinite horizon.

6.3 Calibration

6.3.1 Calibration strategy

I calibrate this model using estimates from Lane [2022] of the effects of South Korea's heavy and chemical industrial (HCI) drive from 1973 to 1979. This policy provided large, targetted subsidies to these industries in the form of subsidized lending, as well as exemptions from trade restrictions and tariffs on imports. I do not attempt to compare South Korean industrial policies to the Pigouvian or super-Pigouvian benchmarks, but rather use them as a source of variation that identifies key model parameters.

Two main assumptions underlie this approach. The first is that differential growth in the heavy and chemical industries was caused by investments in new technology, rather than differential exogenous TFP growth. This view is supported by the large increases in investment in these industries as well as extensive technology adoption contracts between South Korean and foreign firms [Lane, 2022, Choi and Shin, 2022]. Moreover, there is no particular reason to suspect that the heavy and chemical sectors were on the cusp of sudden technological improvements; the Park government championed their expansion largely for political, rather than economic, reasons.

Second, I assume that heavy and chemical industry productivity after the policy drive reflects technological upgrading among all, rather than some, major players in these industries. Here, I draw on the fact that investment in the heavy and chemical industries remained steady (relative to other manufacturing) after the end of the policy drive in 1979. Given my estimates of strong complementarities in technology adoption, this stability suggests there was little scope for further improvements.

Remark. *Whereas the model features endogenous TFP growth through investments in technology, it abstracts away from exogenous TFP growth. I have already justified my decision to ignore exogenous, sector-specific TFP growth, but I also ignore exogenous, aggregate TFP growth, which may have been significant during this period. I abstract from aggregate exogenous growth for the sake of tractability; ignoring aggregate productivity changes allows me to work in a stationary environment.*[28] *For this reason, the most literal interpretation of my quantitative exercise is that I study what the effects of industrial policy would have been in an economy similar to 1970s South Korea but which experienced no exogenous aggregate productivity growth.*

[28]The model can accomodate exogenous productivity growth without sacrificing stationarity if the cost of investment goods increases at the same growth rate. In this case one may simply interpret the model's interest rate as the true interest rate less the growth rate. However, high-tech investment goods were largely imported during this period, which—given South Korea's rapid development—suggests that their real cost fell over time [Choi and Shin, 2022].

6.3.2 Calibration details

I now explain how I calibrate each parameter of the model. In cases where the identification of a parameter from moments in the data requires significant rearrangement of the model's equilibrium conditions, I simply describe which moments are used and relegate the algebra to Appendix B.3. Except where I state otherwise, pre- and post-HCI drive statistics are from summary statistics tables (4-digit panel) provided to the author by Nathan Lane, and changes over time are from the 4-digit panel, doubly-robust DD estimates of Lane [2022], with controls.

I set the labor share α equal to the labor share of total input costs within heavy and chemical industries in the 1970 South Korean input-output table.[29]. I set the elasticity of substitution η to 5, a standard parameterization in industrial policy calibrations, based on estimates in the trade and industrial organization literatures [Broda and Weinstein, 2006, Hendel and Nevo, 2006].

I set productivity under the old technology, z^O, to match gross output per worker in 1972. I set productivity under the new technology, z^N, to match Lane [2022]'s estimates of the increase in gross output per worker during the HCI drive, relative to non-HCI manufacturing industries. I set the fixed costs of production under the old and new technologies, f^O and f^N, equal to HCI investment in 1972 and 1980, respectively.

I set the elasticity of labor supply into the industrial sector, ε, to match HCI employment growth given HCI wage growth implied by my estimates of z^O and z^N. I set the constant in the labor supply equation, $M\bar{a}^{-\varepsilon}$, to match HCI employment in 1972.[30]

I set the interest rate $R = 1.1$, an approximation of bank interest rates in South Korea over this period [Kim, 1991]. I set the frequency of adjustment, δ, to 0.11, a typical rate of capital depreciation in heavy manufacturing [Bureau of Economic Analysis]. Here, I draw on a natural interpretation of firms' adjustment friction: they each invest in capital, which functions properly until randomly breaking down at a memoryless rate δ.

Finally, I consider a range of parameter values for the levels of noise and mean-reversion in the productivity process, i.e. σ and θ. As discussed above, these parameters—especially θ—are key determinants of whether the economy can support multiple Pigouvian equilibria. My main estimates focus on high levels of mean reversion and low levels of noise in order

[29]Coarsely following Lane [2022], I classify the following industries as heavy and chemical: Metal ore, inorganic basic chemical products organic basic chemicals, chemical fertilizer, other chemical products, petroleum products, pig iron and crude steel, primary steel products, non-ferrous metal base and primary products, metal products general machinery, electric machines, transport machines, precision machinery and optical instruments.

[30]The separation of $M/\bar{a}^{\varepsilon}$ into M and $\bar{a}$ is undetermined but unnecessary for my purposes, except for that I maintain the assuption that $\bar{a}$ is sufficiently large; see Appendices B.1 and B.2.

to illustrate how multiplicity can arise under Pigouvian policy. Appendix B.4 explores the robustness of my results to these parameter choices.

Table 2 summarizes the calibration.

Parameter	Value	Strategy	Data source
α	0.32	Share of HCI inputs in HCI gross output	Bank of Korea (1970)
η	5	Follow trade, IO literatures	Broda and Weinstein [2006], Hendel and Nevo [2006]
z^O	4.29	Match HCI output per worker in 1972	Lane [2022]
z^N	4.85	Match HCI output / worker growth during HCI drive	Lane [2022]
f^O	$647,700$	Match aggregate HCI investment in 1972	Lane [2022]
f^N	$3,621,700$	Match aggregate HCI investment in 1980	Lane [2022]
ε	1.23	Match HCI emp. growth during HCI drive	Lane [2022]
$M\bar{a}^{-\varepsilon}$	25,738	Match HCI employment before HCI drive	Lane [2022]
R	1.1	Match interest rates during HCI drive	Kim [1991]
δ	0.11	Rate of heavy manufacuring capital depreciation	Bureau of Economic Analysis
σ	0.01	Illustrate multiplicity	-
θ	2	Illustrate multiplicity	-

Table 2: Calibration summary

Under this calibration, I simulate the model in the continuous time, many agent limit as $\Delta \to 0$ and $|\mathcal{I}| \to \infty$. This allows me to use efficient methods based on Moll [2017]. As there are many possible equilibria, I focus on the most-industrialized and least-industrialized equilibria—both under Pigouvian policy and laissez faire—as well as the social optimum—or equivalently the super-Pigouvian equilibrium. Here I leverage the fact that, since my calibration satisfies $\alpha(\eta - 1) < 1 + \varepsilon$, profits are increasing in the number of firms; a monotonicity argument in Frankel and Pauzner [2000] then implies that these extremal equilibria exist and take the form of cutoffs in the space of z_t and N_t for whether or not firms invest in the new technology.

6.4 Results

Using the calibration and computational approach discussed above, I compute the least-industrialized and most-industrialized Pigouvian equilibria, as well as the efficient strategy—which recall is a super-Pigouvian equiibrium. I discuss these findings positively and normatively below.

6.4.1 Pigouvian multiplicity

I focus on the class of equilibria in which the investment decision depends on a cutoff threshold in the space of productivity and the current number of high-tech firms. Because my model fits into the framework of Frankel and Pauzner [2000] their results imply that the equilibria with the most and least industrialization always take this simple, cutoffs form.[31] Concretely, in such an equilibrium there exists a function $z(N)$ such that firms at t invest if and only if $z_t > z(N_t)$. An equilibrium with low cutoffs reflects that firms are willing to invest in high-tech even when the exogenous component of productivity is low, because they expect many other firms to invest soon. Conversely, an equilibrium with high cutoffs reflects that firms are pessimistic about whether others will invest in high-tech and so are unwilling to do so themselves unless the exogenous component of productivity is very high.

Can these expectations about other firms be so important so as to be self-fulfilling—creating a coordination failure? I find that this is indeed the case. Figure 4 demonstrates this by plotting the equilibrium cutoffs that define the most- and least-industrializing equilibria, i.e. those in which the productivity cutoffs for investment in the new technology are the lowest and highest, respectively.[32] These cutoffs differ at low shares of new technology usage, including at the point—where the dotted line representing mean log productivity intersects the X axis—that represents the initial condition of the HCI drive. Since the stochastic process for log productivity is fairly stable, with its stationary distribution having a standard deviation of $\sigma/\sqrt{2\theta} = 0.005$, it is easy to interpret the path of the economy under these two equilibria: Following the lowest-HCI Pigouvian equilibrium, the economy (with high likelihood) stays at its initial level of development, with all firms adopting the old technology. Following the highest-HCI Pigouvian equilibrium, all firms invest in the new technology at their first opportunity and (with high likelihood) the economy industrializes.

The figure also plots the cutoffs that describe the first-best equilibrium. As there is a unique first-best strategy profile, Theorem 1 implies that these are the equilibrium cutoffs in

[31] This fact does not depend on whether z_t is a mean-reverting (which I assume and they do not).

[32] Figure 4 can be interpretted as a stochastic phase diagram. In the Y dimension, the share of firms who have invested drifts either up or down depending on whether productivity is above or below the cutoff. In the X dimension, productivity noisily drifts toward its mean. Appendix Figure A3 illustrates this interpretation.

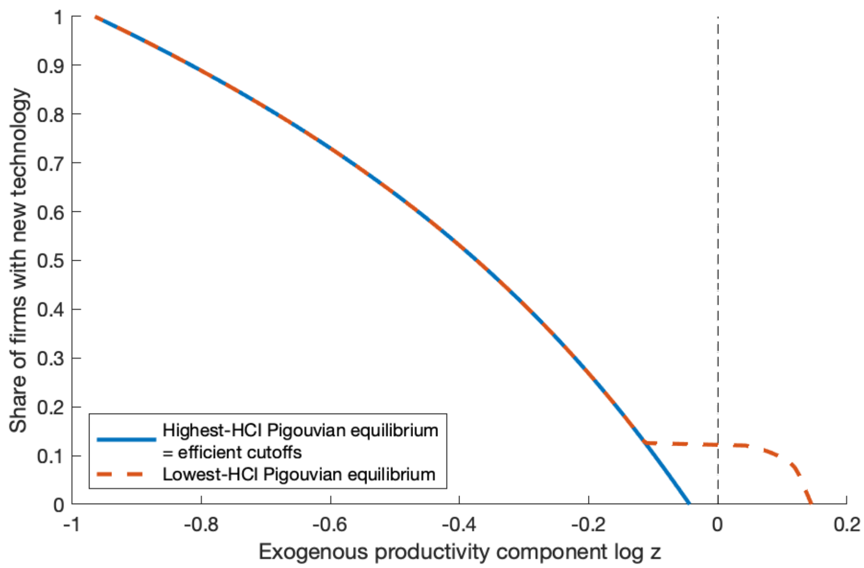

Figure 4: Productivity cutoffs for new technology adoption characterizing the highest- and lowest-industrializing equilibria of the calibrated model under Pigouvian policy.

any super-Pigouvian equilibrium. Interestingly, the efficient cutoffs coincide with the most-industrializing Pigouvian equilibrium. In other words—although this need not be the case in general—it is under this calibration efficient for households to have the higest level of industrial optimism that Pigouvian policy can sustain. Appendix Figure A2 shows a version of the figure that also contains the (unique) laissez-faire equilibrium.

6.4.2 Welfare gains from super-Pigouvian policy

I next consider welfare in each of the two equilibria depicted in Figure 4. Since labor supply is fixed and the planner has access to lump-sum transfers, welfare is simply the net present value of aggregate output. In principle, one could compare output across the two equilibria starting from any (z_t, N_t) initial condition. However, the calibrated productivity process is sufficiently concentrated around its mean that it suffices to simply perform this exercise while setting $\log z_t = 0$.

Figure 5 shows welfare gains—compared to laissez faire—from Pigouvian policy in the lowest- and highest-industrialization equilibria. Recall that the latter coincides with the welfare gains from super-Pigouvian policy. I consider these gains as a share of the annuitized value of HCI value added before the HCI drive. A value of 1% is therefore the welfare equivalent of a permanent 1% increase in value added in the heavy and chemical industries.

The figure shows that, depending on the initial share of firms using the new technology, the efficient allocation improves upon laissez faire by between between two and twelve percent

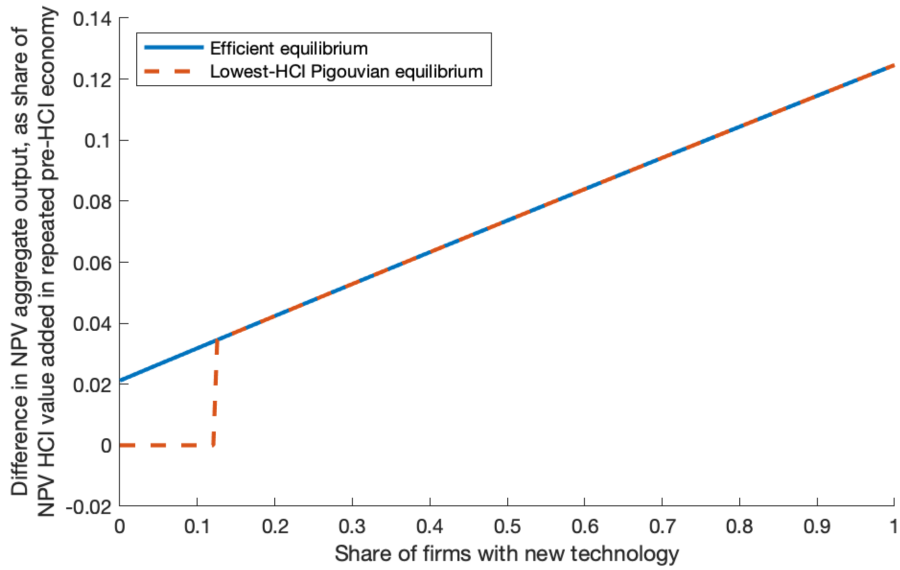

Figure 5: Welfare gains from the least-industrializing Pigouvian equilibrium and the efficient, most-industrializing Pigouvian equilibrium, at log productivity 0 and various initial shares of new technology adoption.

of HCI value added. These gains are smaller when fewer firms initially adopt the new technology, as in this case the (usually) efficient path—industrialization—is only slightly more efficient than the typical laissez-faire path—de-industrialization due to transition costs.

Most interesting is how welfare under the lowest-industrialization Pigouvian equilibrium compares to the laissez-faire and efficient welfare levels. Namely, whether this equilibrium delivers the laissez-faire welfare level or the efficient welfare level depends on the initial share of firms using the new technology. When there are sufficiently many, even the low-industrialization Pigouvian path leads to the efficient path, industrialization (with high likelihood); when there are too few, it leads to, it leads to the inefficient path, deindustrialization (with high likelihood).

The initial conditions of my calibration correspond to the left-most points in the figure. Here, the worst Pigouvian equilibrium is worse for welfare than the efficient equilibrium, implying there are gains from super-Pigouvian policy. However, these gains are modest, equivalent to about two percent value added in HCI industries. The modest size of these welfare gains do not reflect a failure of policy but rather the large cost of investments in new technology, which—in my calibration—almost completely offset the productivity gains they generate.

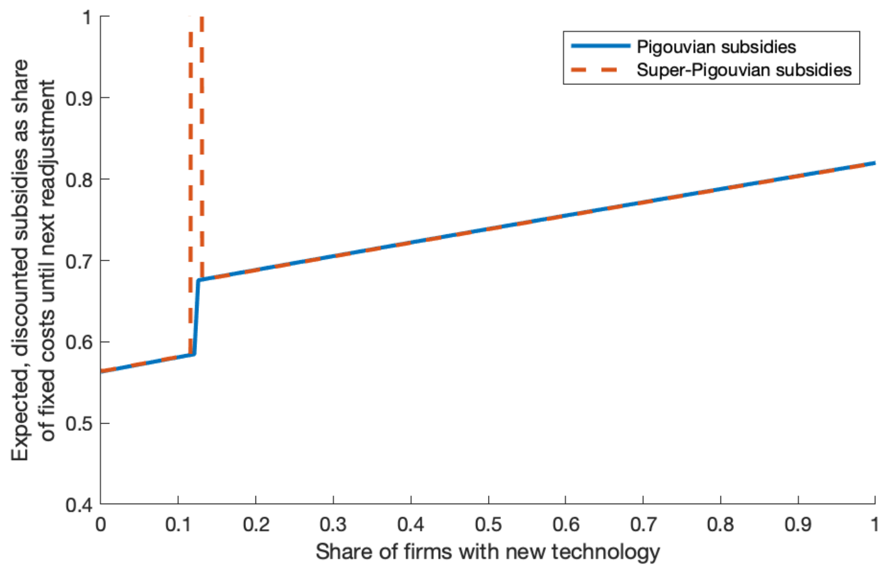

Figure 6: Subsidies to new-technology firms under Pigouvian and super-Pigouvian policy, at log productivity 0 and various initial shares of new technology adoption, in the lowest-HCI Pigouvian equilibrium. Y axis is expected discounted subsidies until next readjustment, as a share of expected discounted fixed costs until next readjustment.

6.4.3 Super-Pigouvian policy in inefficient Pigouvian equilibria

As discussed in Sections 4 and 5, super-Pigouvian policy is simply Pigouvian when evaluated in any efficient equilibrium. However, one benefit of the calibrated model is that I can use it to study the structure of super-Pigouvian policy in inefficient allocations—such as the lowest-HCI Pigouvian equilibrium discussed above. This both helps to clarify the mechanism behind super-Pigouvian policy and to shed light on whether a social planner can credibly commit to offer super-Pigouvian incentives.

Figure 6 compares Pigouvian and super-Pigouvian subsidies to HCI firms in the inefficient, lowest-HCI Pigouvian equilibrium discussed in the preceding sections. As in earlier sections, I study these subsidies at a range of values for the share of firms using new technology but assume log productivity is at its mean, 0. Intuitively, Pigouvian subsidies increase in the number of new technology users—i.e. the number of existing firms on whom an additional firm can have an externality—and increase sharply when around 12% of firms use the new technology, since this is the threshold above which the future actors will choose to do so (see Figure 4).

Super-Pigouvian policy coincides with Pigouvian policy almost exactly except for in a narrow region around the 12% threshold. Intuitively, this is because this cutoff is inefficiently high, and super-Pigouvian policy can prevent it from being consistent with individual optimization if it offers households a large incentive to enter at a slightly lower threshold.

Although strong incentives for coordination in a critical region are a general feature of super-Pigouvian policy, the strength of these incentives is particularly dramatic in this example. The magnitude of the super-Pigouvian subsidy dwarfs the Pigouvian subsidy around the critical threshold; if not truncated, the graph would show a maximum around 5 (see Appendix Figure A4). This extremeness presents two challenges for implementing super-Pigouvian policy: First, the government may not be able to commit to make sure large payments. Second, the government may not be able to acquire precise enough information to target the critical region, given it is very narrow. To the first point, it is important that, in fact, firms would be willing to use the new technology with much smaller subsidies: For example, it of course suffices to provide subsidy payment equal to the expected discounted cost of increased fixed costs from the new technology. This corresponds to a value of 1 in the Y axis of Figure 6. In this sense, the super-Pigouvian planner has no need to provide such generous subsidies. It wishes to, because the welfare contributions of the pivotal household are so large, but household behavior would be the same with subsidies that are capped at a much lower level. To the second point, it is worth noting that a "smoothed out" policy that spreads the super-Pigouvian subsidy over a larger region would still prevent the low-industrialization Pigouvian allocation from being an equilibrium, even if it created a small amount of misallocation on the margin. Finally, although the super-Pigouvian policy is somewhat extreme, the presence of noise in the productivity process prevents super-Pigouvian policy from recommending that the planner pay a finite fraction of GDP to the single, infinitesimal firm that pushes the economy over the cutoff threshold. Rather, the planner spreads payments across firms sufficiently close to this threshold. This point is illustrated more clearly at higher productivities, where the lowest-HCI Pigouvian equilibrium cutoffs are steeper in productivity $\times$ adoption-share space; see Appendix Figure A5.

7 Conclusion

Coordination failures are one of the primary justifications for industrial policy. Yet, they have no clear policy solution: Pigouvian taxes that address market failures on the margin risk arriving at a local, rather, than global optimum—whereas its known alternatives require a centralized approach based on regulating quantities rather than prices.

This paper has proposed an alternative, "super-Pigouvian" policy that can resolve coordination failures without requiring central planning. The idea behind this policy is to elicit households' willingness to pay for each others' actions—as in Pigouvian policy—but taking into account those actions' indirect effects on others' later behavior. When households know that policy has this structure both on and off the equilibrium path, they are guaranteed to take the same actions the social planner would. This guarantees Pareto efficiency.

While this paper has studied the application of the super-Pigouvian approach to industrial policy, there is room to apply similar ideas to other settings. At a high level, this paper's method for deriving welfare theorems—one that leans on arguments from dynamic programming rather than convex optimization—may be useful in applications where convexity is unlikely to hold. For instance, the literature on economic geography actively debates whether or not agglomeration forces lead to multiple equilibria, but has not yet proposed new policies that can ensure multiplicity is resolved efficiently [Bleakley and Lin, 2012, Allen and Donaldson, 2020]. This would be a natural setting for future work.

References

Daron Acemoglu, Philippe Aghion, Leonardo Bursztyn, and David Hemous. The environment and directed technical change. *American economic review*, 102(1):131–66, 2012.

Treb Allen and Dave Donaldson. Persistence and path dependence in the spatial economy. Technical report, National Bureau of Economic Research, 2020.

G-M Angeletos and Chen Lian. Incomplete information in macroeconomics: Accommodating frictions in coordination. In *Handbook of macroeconomics*, volume 2, pages 1065–1240. Elsevier, 2016.

Andrew Atkeson, Varadarajan V Chari, and Patrick J Kehoe. Sophisticated monetary policies. *The Quarterly Journal of Economics*, 125(1):47–89, 2010.

Robert J Aumann. Integrals of set-valued functions. *Journal of mathematical analysis and applications*, 12(1):1–12, 1965.

Dominick G Bartelme, Arnaud Costinot, Dave Donaldson, and Andres Rodriguez-Clare. The textbook case for industrial policy: Theory meets data. Technical report, National Bureau of Economic Research, 2019.

William J Baumol. On taxation and the control of externalities. *The American Economic Review*, 62(3):307–322, 1972.

Dirk Bergemann and Juuso Välimäki. The dynamic pivot mechanism. *Econometrica*, 78(2): 771–789, 2010.

Hoyt Bleakley and Jeffrey Lin. Portage and path dependence. *The quarterly journal of economics*, 127(2):587–644, 2012.

Christian Broda and David E Weinstein. Globalization and the gains from variety. *The Quarterly journal of economics*, 121(2):541–585, 2006.

Francisco J Buera, Hugo Hopenhayn, Yongseok Shin, and Nicholas Trachter. Big push in distorted economies. Technical report, National Bureau of Economic Research, 2021.

Jaedo Choi and Younghun Shin. Technology adoption and late industrialization. Technical report, 2022.

Antonio Ciccone. Input chains and industrialization. *The Review of Economic Studies*, 69 (3):565–587, 2002.

Edward H Clarke. Multipart pricing of public goods. *Public choice*, pages 17–33, 1971.

Donald R Davis and David E Weinstein. Bones, bombs, and break points: the geography of economic activity. *American economic review*, 92(5):1269–1289, 2002.

Pablo D Fajgelbaum and Cecile Gaubert. Optimal spatial policies, geography, and sorting. *The Quarterly Journal of Economics*, 135(2):959–1036, 2020.

David Frankel and Ady Pauzner. Resolving indeterminacy in dynamic settings: the role of shocks. *The Quarterly Journal of Economics*, 115(1):285–304, 2000.

David M Frankel, Krzysztof Burdzy, et al. Shocks and business cycles. *The BE Journal of Theoretical Economics*, 5(1):1–88, 2005.

Shota Fujishima. Evolutionary implementation of optimal city size distributions. *Regional Science and Urban Economics*, 43(2):404–410, 2013.

Holger Görg and Eric Strobl. The effect of r&d subsidies on private r&d. *Economica*, 74 (294):215–234, 2007.

Theodore Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.

Bernardo Guimaraes and Caio Machado. Dynamic coordination and the optimal stimulus policies. *The Economic Journal*, 128(615):2785–2811, 2018.

Bernardo Guimaraes and Ana Elisa Pereira. Qwerty is efficient. *Journal of Economic Theory*, 163:819–825, 2016.

Ann Harrison and Andrés Rodríguez-Clare. Trade, foreign investment, and industrial policy for developing countries. *Handbook of development economics*, 5:4039–4214, 2010.

Ricardo Hausmann and Dani Rodrik. Economic development as self-discovery. *Journal of development Economics*, 72(2):603–633, 2003.

Igal Hendel and Aviv Nevo. Measuring the implications of sales and consumer inventory behavior. *Econometrica*, 74(6):1637–1673, 2006.

Leonid Hurwicz. Optimality and informational efficiency in resource allocation processes. *Mathematical methods in the social sciences*, 1960.

Kwan S Kim. *The Korean miracle (1962-1980) revisited: myths and realities in strategy and development*. Helen Kellogg Institute for International Studies, University of Notre Dame ..., 1991.

Minho Kim, Munseob Lee, and Yongseok Shin. The plant-level view of an industrial policy: The korean heavy industry drive of 1973. Technical report, National Bureau of Economic Research, 2021.

Paul Krugman. History versus expectations. *The Quarterly Journal of Economics*, 106(2): 651–667, 1991.

Konstantin Kucheryavyy, Gary Lyn, and Andrés Rodríguez-Clare. Grounded by gravity: A well-behaved trade model with industry-level economies of scale. Technical report, National Bureau of Economic Research, 2016.

Roger Lagunoff and Akihiko Matsui. Asynchronous choice in repeated coordination games. *Econometrica: Journal of the Econometric Society*, pages 1467–1477, 1997.

Nathan Lane. Manufacturing revolutions: Industrial policy and industrialization in south korea. *Available at SSRN 3890311*, 2022.

Ernest Liu. Industrial policies in production networks. *The Quarterly Journal of Economics*, 134(4):1883–1948, 2019.

Rebecca Mandt, Kushal Seetharam, and Chung Hon Michael Cheng. Federal r&d funding: the bedrock of national innovation. *MIT Science Policy Review*, 1:44–54, 2020.

Eric Maskin. Nash equilibrium and welfare optimality. *The Review of Economic Studies*, 66 (1):23–38, 1999.

Kiminori Matsuyama. Increasing returns, industrialization, and indeterminacy of equilibrium. *The Quarterly Journal of Economics*, 106(2):617–650, 1991.

Ben Moll. Heterogeneous agent models in continuous time part i, 2017. URL `https://benjaminmoll.com/wp-content/uploads/2019/07/HACT.pdf`.

John Moore and Rafael Repullo. Subgame perfect implementation. *Econometrica: Journal of the Econometric Society*, pages 1191–1220, 1988.

Kevin M Murphy, Andrei Shleifer, and Robert W Vishny. Industrialization and the big push. *Journal of political economy*, 97(5):1003–1026, 1989.

William D Nordhaus. After kyoto: alternative mechanisms to control global warming. *American Economic Review*, 96(2):31–34, 2006.

Arthur Cecil Pigou. *The Economics of Welfare*. Macmillan, 1924.

Edward C Prescott and Rajnish Mehra. Recursive competitive equilibrium: The case of homogeneous households. In *Theory Of Valuation*, pages 357–371. World Scientific, 2005.

Francisco Rodríguez. Comment on "coordination failures, clusters, and microeconomic interventions" by Andres Rodriguez-Clare. *Economia*, 6(1):1–42, 2005.

Andres Rodriguez-Clare. The division of labor and economic development. *Journal of Development Economics*, 49(1):3–32, 1996.

Dani Rodrik. Getting interventions right: how south korea and taiwan grew rich. *Economic policy*, 10(20):53–107, 1995.

Dani Rodrik. Coordination failures and government policy: A model with applications to east asia and eastern europe. *Journal of international economics*, 40(1-2):1–22, 1996.

Dani Rodrik. Industrial policy for the twenty-first century. *Available at SSRN 666808*, 2004.

Paul N Rosenstein-Rodan. Problems of industrialisation of eastern and south-eastern europe. *The economic journal*, 53(210/211):202–211, 1943.

William H Sandholm. Evolutionary implementation and congestion pricing. *The Review of Economic Studies*, 69(3):667–689, 2002.

William H Sandholm. Negative externalities and evolutionary implementation. *The Review of Economic Studies*, 72(3):885–915, 2005.

William H Sandholm. Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory*, 132(1):367–382, 2007.

Edouard Schaal and Mathieu Taschereau-Dumouchel. Coordinating business cycles. *Available at SSRN 2567344*, 2015.

David A Starrett. Fundamental nonconvexities in the theory of externalities. *Journal of Economic Theory*, 4(2):180–199, 1972.

Nancy L Stokey, Robert E Lucas, and Edward C Prescott. *Recursive methods in economic dynamics*. Harvard University Press, 1989.

Bureau of Economic Analysis. Bea depreciation rates. URL https://apps.bea.gov/national/pdf/BEA_depreciation_rates.pdf.

Earl A Thompson et al. Underinvestment traps and potential cooperation. Technical report, UCLA Department of Economics, 1982.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.

Appendix

A Proofs

A.1 Lemmas characterizing Pareto efficiency, household optimality

Lemma 2. *Fix an equilibrium and a history h^t. The equilibrium profile is Pareto efficient starting from h^t is and only if its path of factor supply $\{\ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}$ maximizes $W(\{\ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}})$ among feasible allocations starting from h^t.*

Proof. Formally, a profile $\{c_s^j, \ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}$ is Pareto efficient starting from h^t if and only it satisfies the following conditions:

- For all $j \in \mathcal{I}$, we have (a) $\ell_s^j \in \mathcal{A}$ for all $s \geqslant t$, (b) $\ell_t^j = h_{t^j}^j$ for the last time t^j before t at which $i_{t^j} = j$, or $\ell_t^j = \overline{\ell}_0^j$ if no such time exists, and (c) for all $s \geqslant t$, $s < |\mathcal{T}|$, we have $\ell_s^j = \ell_{s+1}^j$ unless $j = i_s$.
- $\sum_{s\geqslant t} R^{-(s-t)}\left[\sum_{j\in\mathcal{I}} c_s^j - F\left(\sum_{j\in\mathcal{I}}\ell_s^j, \sum_{j\in\mathcal{I}}\ell_s^j\right)\right] + D(h^t) \geqslant 0.$
- Among all other profiles satisfying the first two bullets, none weakly improve welfare for all households and strictly improve welfare for at least one.

Since utility is non-satiated, these Pareto efficiency conditions are equivalent to

$$\begin{aligned}
\{c_s^j, \ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}} \quad &\in \quad \underset{\substack{\text{feasible}\\ \{\hat{c}_s^j,\hat{\ell}_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}}}{\arg\min} \quad \sum_{s\geqslant t} R^{-(s-t)}\left[\sum_{j\in\mathcal{I}}\hat{c}_s^j - F\left(\sum_{j\in\mathcal{I}}\hat{\ell}_s^j, \sum_{j\in\mathcal{I}}\hat{\ell}_s^j\right)\right] \\
&\qquad\qquad \text{s.t.} \quad \sum_{s\geqslant t}\beta^{s-t}u(\hat{c}_s^j - v(\hat{\ell}_s^j)) \quad \geqslant \quad \sum_{s\geqslant t}\beta^{s-t}u(c_s^j - v(\ell_s^j)) \\
&\qquad\qquad\qquad \Longleftrightarrow \\
\{c_s^j - v(\ell_s^j), \ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}} \quad &\in \quad \underset{\substack{\text{feasible}\\ \{\hat{C}_s^j,\hat{\ell}_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}}}{\arg\min} \quad \sum_{s\geqslant t} R^{-(s-t)}\left[\sum_{j\in\mathcal{I}}\left(\hat{C}_s^j + v(\hat{\ell}_s^j)\right) - F\left(\sum_{j\in\mathcal{I}}\hat{\ell}_s^j, \sum_{j\in\mathcal{I}}\hat{\ell}_s^j\right)\right] \\
&\qquad\qquad \text{s.t.} \quad \sum_{s\geqslant t}\beta^{s-t}u(\hat{C}_t^j) \quad \geqslant \quad \sum_{s\geqslant t}\beta^{s-t}u(c_s^j - v(\ell_s^j)) \\
&\qquad\qquad\qquad \Longleftrightarrow \\
\{\ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}} \quad &\in \quad \underset{\substack{\text{feasible}\\ \{\hat{\ell}_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}}}{\arg\max} \quad \sum_{s\geqslant t} R^{-(s-t)}\left[F\left(\sum_{j\in\mathcal{I}}\hat{\ell}_s^j, \sum_{j\in\mathcal{I}}\hat{\ell}_s^j\right) - \sum_{j\in\mathcal{I}} v(\hat{\ell}_s^j)\right] \qquad \text{(A1)} \\
\{c_s^j - v(\ell_s^j)\}_{s\geqslant t} \quad &\in \quad \underset{\{\hat{C}_s^j\}_{s\geqslant t}}{\arg\min} \quad \sum_{s\geqslant t} R^{-(s-t)}\sum_{j\in\mathcal{I}}\hat{C}_s^j \\
&\qquad\qquad \text{s.t.} \quad \sum_{s\geqslant t}\beta^{s-t}u(\hat{C}_t^j) \quad \geqslant \quad \sum_{s\geqslant t}\beta^{s-t}u(c_s^j - v(\ell_s^j)) \\
&\qquad\qquad\qquad \Longleftrightarrow
\end{aligned}$$

$$
\begin{aligned}
\{\ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}} \quad &\in \quad \underset{\substack{\text{feasible}\\ \{\hat{\ell}_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}}}{\arg\max} \quad \sum_{s\geqslant t} R^{-(s-t)}\left[F\left(\sum_{j\in\mathcal{I}}\hat{\ell}_s^j, \sum_{j\in\mathcal{I}}\hat{\ell}_s^j\right) - \sum_{j\in\mathcal{I}} v(\hat{\ell}_s^j)\right] \\
\{c_s^j\}_{s\geqslant t} \quad &\in \quad \underset{\{\hat{c}_s^j\}_{s\geqslant t}}{\arg\min} \quad \sum_{s\geqslant t} R^{-(s-t)}\sum_{j\in\mathcal{I}}\hat{c}_s^j \\
&\qquad\qquad \text{s.t.} \quad \sum_{s\geqslant t}\beta^{s-t}u(\hat{c}_s^j - v(\ell_s^j)) \quad \geqslant \quad \sum_{s\geqslant t}\beta^{s-t}u(c_s^j - v(\ell_s^j)) \\
&\qquad\qquad\qquad \Longleftrightarrow \\
\{\ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}} \quad &\in \quad \underset{\substack{\text{feasible}\\ \{\hat{\ell}_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}}}{\arg\max} \quad \sum_{s\geqslant t} R^{-(s-t)}\left[F\left(\sum_{j\in\mathcal{I}}\hat{\ell}_s^j, \sum_{j\in\mathcal{I}}\hat{\ell}_s^j\right) - \sum_{j\in\mathcal{I}} v(\hat{\ell}_s^j)\right],
\end{aligned}
$$

where "feasible" refers to the first bullet in the conditions for Pareto efficiency above. The second part of the third and fourth sets of conditions hold for all $j \in \mathcal{I}$. The final equivalence follows from household optimization with u strictly increasing, using that $\{c_s^j, \ell_s^j\}_{s\geqslant t}^{j\in\mathcal{I}}$ the path of consumption and factor supply in an equilibrium. $\square$

Lemma 3. *Fix a history h^t and a household i. The profile $\{c_s^i, \ell_s^i\}_{s\geqslant t}^{i\in\mathcal{I}}$ satisfies household optimization for i starting from h^t if and only if*

$$
\begin{aligned}
\{c_s^i\}_{s\geqslant t} \quad &\in \quad \underset{\{\hat{c}_s^i\}_{s\geqslant t}}{\arg\max}\sum_{s\geqslant t}\beta^{-(s-t)}u(\hat{c}_s^i - v(\ell_s^i)) \\
&\qquad \textit{s.t.} \quad d^i(h^t) + \sum_{s\geqslant t} R^{-(s-t)}\left(\hat{c}_s^i - \boldsymbol{\omega}_s^i(\ell_s^i; h^t)\cdot\ell_s^i - \boldsymbol{T}_s^i(h^t)\right) \leqslant 0 \\
\{\ell_s^i\}_{s\geqslant t} \quad &\in \quad \underset{\textit{feasible}\ \{\hat{\ell}_s^i\}_{s\geqslant t}}{\arg\max}\sum_{s\geqslant t} R^{-(s-t)}\left(\boldsymbol{\omega}_s^i(\hat{\ell}_s^i; h^t)\cdot\hat{\ell}_s^i - v(\hat{\ell}_s^i)\right).
\end{aligned}
\tag{A2}
$$

Proof. We use the GHH structure of household preferences to rearrange the condition for optimality:

$$
\begin{aligned}
\{c_s^i, \ell_s^i\}_{s\geqslant t} \quad &\in \quad \underset{\text{feasible }\{\hat{c}_s^i,\hat{\ell}_s^i\}_{s\geqslant t}}{\arg\max} \sum_{s\geqslant t} \beta^{-(s-t)} u\Big(\hat{c}_s^i - v(\hat{\ell}_s^i)\Big) \\
&\qquad \text{s.t.} \quad d^i(h^t) + \sum_{s\geqslant t} R^{-(s-t)} \left(\hat{c}_s^i - \boldsymbol{\omega}_s^i(\hat{\ell}_s^i; h^t)\cdot \hat{\ell}_s^i - \boldsymbol{T}_s^i(h^t)\right) \leqslant 0 \\
&\iff \\
\{c_s^i - v(\ell_s^i), \ell_s^i\}_{s\geqslant t} \quad &\in \quad \underset{\text{feasible }\{\hat{C}_s^i,\hat{\ell}_s^i\}_{s\geqslant t}}{\arg\max} \sum_{s\geqslant t} \beta^{-(s-t)} u(\hat{C}_s^i) \\
&\qquad \text{s.t.} \quad d^i(h^t) + \sum_{s\geqslant t} R^{-(s-t)} \left(\hat{C}_s^i + v(\hat{\ell}_s^i) - \boldsymbol{\omega}_s^i(\hat{\ell}_s^i; h^t)\cdot \hat{\ell}_s^i - \boldsymbol{T}_s^i(h^t)\right) \leqslant 0 \\
&\iff \\
\{c_s^i - v(\ell_s^i)\}_{s\geqslant t} \quad &\in \quad \underset{\{\hat{C}_s^i\}_{s\geqslant t}}{\arg\max} \sum_{s\geqslant t} \beta^{-(s-t)} u(\hat{C}_s^i) \\
&\qquad \text{s.t.} \quad d^i(h^t) + \sum_{s\geqslant t} R^{-(s-t)} \left(\hat{C}_s^i + v(\ell_s^i) - \boldsymbol{\omega}_s^i(\ell_s^i; h^t)\cdot \ell_s^i - \boldsymbol{T}_s^i(h^t)\right) \leqslant 0 \\
\{\ell_s^i\}_{s\geqslant t} \quad &\in \quad \underset{\text{feasible }\{\hat{\ell}_s^i\}_{s\geqslant t}}{\arg\min} \sum_{s\geqslant t} R^{-(s-t)} \left(v(\hat{\ell}_s^i) - \boldsymbol{\omega}_s^i(\hat{\ell}_s^i; h^t)\cdot \hat{\ell}_s^i\right).
\end{aligned}
\tag{A3}
$$

where the final equivalent statement uses that u is non-satiated. $\square$

A.2 Proof of Proposition 1

I prove each item in turn.

First welfare theorem

Suppose not, i.e. letting $\{c_t^i, \ell_t^i\}$ denote the equilibrium path, there exists some $\{\Delta c_t^i, \Delta\ell_t^i\}$ such that (a) for small enough $\epsilon > 0$, $\{c_t^i + \epsilon\Delta c_t^i, \ell_t^i + \epsilon\Delta\ell_t^i\}$ is feasible, i.e. $\ell_t^i + \epsilon\Delta\ell_t^i \in \mathcal{A}$, $\ell_t^i + \epsilon\Delta\ell_t^i = \ell_{t-1}^i + \epsilon\Delta\ell_{t-1}^i$ unless $i = i_{t-1}$, and

$$
\sum_{t\geqslant 0} R^{-t} \sum_{i\in\mathcal{I}} \left(c_t^i + \epsilon\Delta c_t^i\right) \quad \leqslant \quad \sum_{t\geqslant 0} R^{-t} F\left(\sum_{i\in\mathcal{I}} \ell_t^i + \epsilon\Delta\ell_t^i, \sum_{i\in\mathcal{I}} \ell_t^i + \epsilon\Delta\ell_t^i\right), \tag{A4}
$$

and (b) $\{c_t^i + \epsilon\Delta c_t^i, \ell_t^i + \epsilon\Delta\ell_t^i\}$ generates a first-order Pareto improvement

$$
\sum_{t\geqslant 0} R^{-t} \sum_{i\in\mathcal{I}} \left[\Delta c_t^i - v_\ell(\ell_t^i)\cdot\Delta\ell_t^i\right] > 0. \tag{A5}
$$

To proceed, note that since (A4) holds for all small enough $\epsilon > 0$ and F is differentiable,

$$
\sum_{t\geqslant 0} R^{-t} \sum_{i\in\mathcal{I}} \Delta c_t^i \quad \leqslant \quad \frac{d}{d\epsilon}\bigg|_{\epsilon=0} \sum_{t\geqslant 0} R^{-t} F\left(\sum_{i\in\mathcal{I}} \ell_t^i + \epsilon\Delta\ell_t^i, \sum_{i\in\mathcal{I}} \ell_t^i + \epsilon\Delta\ell_t^i\right). \tag{A6}
$$

Combining this with (A5) implies

$$\frac{d}{d\epsilon}\bigg|_{\epsilon=0} \sum_{t\geqslant 0} R^{-t} \sum_{i\in\mathcal{I}} \Bigg[\underbrace{\frac{d}{dL_t^i}\bigg|_{L_t^i=\sum_{i\in\mathcal{I}} \ell_t^i} F(L_t^i, L_t^i)}_{=\boldsymbol{\omega}_t^i(\ell_t^i, h^0)} - v_\ell(\ell_t^i) \Bigg] \cdot \Delta\ell_t^i > 0. \tag{A7}$$

where the relationship to $\boldsymbol{\omega}_t^i(\ell_t^i, h^0)$ follows from the definition of Pigouvian taxation

Next, for all i and t, let $t^{\text{next}}(i,t)$ denote the next time after t at which i can adjust its behavior, or $|\mathcal{T}|$ if no such time exists. In this notation, and using the fact that feasiblity implies $\Delta\ell_t^i = 0$ for $t < t^{\text{next}}(i,0)$, we can rewrite (A7) as

$$\sum_{t\geqslant 0} R^{-t} \sum_{s=t+1}^{t^{\text{next}}(i_t,t)} R^{-(s-t)} \left[\boldsymbol{\omega}_s^{i_t}(\ell_s^{i_t}, h^0) - v_\ell(\ell_s^{i_t})\right] \cdot \Delta\ell_s^{i_t} > 0. \tag{A8}$$

It follows that there exists at least one $t \geqslant 0$ such that

$$\begin{aligned} &\sum_{s=t+1}^{t^{\text{next}}(i_t,t)} R^{-(s-t)} \left[\boldsymbol{\omega}_s^{i_t}(\ell_s^{i_t}, h^0) - v_\ell(\ell_s^{i_t})\right] \cdot \Delta\ell_s^{i_t} > 0 \\ \implies \quad &\sum_{s=t+1}^{t^{\text{next}}(i_t,t)} R^{-(s-t)} \Bigg[\left(\boldsymbol{\omega}_s^{i_t}(\ell_s^{i_t} + \epsilon\Delta\ell_s^{i_t}, h^0) \cdot (\ell_s^{i_t} + \epsilon\Delta\ell_s^{i_t}) - v(\ell_s^{i_t} + \epsilon\Delta\ell_s^{i_t})\right) \\ &\qquad - \left(\boldsymbol{\omega}_s^{i_t}(\ell_s^{i_t}, h^0) \cdot \ell_s^{i_t} - v(\ell_s^{i_t})\right) \Bigg] > 0 \quad \text{for small enough } \epsilon > 0, \end{aligned} \tag{A9}$$

where the second line follows from the fact that, under Pigouvian policy, taxes are not quantity dependent. However, since, for small enough $\epsilon > 0$, $\ell_s^{i_t} + \epsilon\Delta\ell_s^{i_t}$ is feasible by assumption, (A9) implies that $\{\ell_s^{i_t}\}_{s\geqslant 0}$ is suboptimal for i_t, a contradiction.

Second welfare theorem

Fix a first-order Pareto efficient allocation $\{c_t^i, \ell_t^i\}$. We now construct a Pigouvian equilibrium as follows. First, for each $i \in \mathcal{I}$, fix any $\hat{\ell}^i \in \mathcal{A}$. Also denote by $\widetilde{H}$ the set of histories that follow $\{\ell_t^i\}$, i.e. $h^t = (\ell_t^{i_{t-1}}, ..., \ell_1^{i_0})$. For histories $h^t = (h_{t-1}, ..., h_0)$ not in $\widetilde{H}$, let $\tilde{t}(h^t)$ denote the first time t at which $h_{t-1} \neq \ell_t^{i_{t-1}}$. And for an arbitrary history $h^t = (h_{t-1}, ..., h_0)$, let $\ell^{\text{last}}(i, h^t)$ denote the last action to which i has commited along the history h^t, i.e. h_s if there exists $s \leqslant t$ with $i_t = s$ and $\overline{\ell}_0^i$ otherwise. Then set

$$
\begin{aligned}
\boldsymbol{\ell}_s^i(h^t) &= \begin{cases} \ell_s^i, & \text{if} \quad h^t \in \mathcal{H} \\ \ell^{\text{last}}(i, h^t), & \text{if} \quad h^t \notin \mathcal{H} \text{ and } t^{\text{next}}(i,t) > s \\ \hat{\ell}^i & \text{if} \quad h^t \notin \mathcal{H} \text{ and } t^{\text{next}}(i,t) \leqslant s \end{cases} \\
\boldsymbol{L}_s(h^t) &= \sum_{i \in \mathcal{I}} \boldsymbol{\ell}_s^i(h^t) \\
\boldsymbol{w}_s(h^t) &= F_L(\boldsymbol{L}_s(h^t), \boldsymbol{L}_s(h^t)) \\
\boldsymbol{\tau}_s^i(\tilde{\ell}; h^t) &= \begin{cases} -F_{\overline{L}}(\boldsymbol{L}_s(h^t), \boldsymbol{L}_s(h^t)), & \text{if} \quad h^t \in \widetilde{H} \\ \boldsymbol{w}_s(h^t) - (v(\tilde{\ell}) - v(0)) \frac{\tilde{\ell}}{||\mathbb{1}_{\tilde{\ell}=0} + \tilde{\ell}||^2} - \mathbb{1}_{\tilde{\ell} = \boldsymbol{\ell}_s^i(h^t)} & \text{otherwise} \end{cases} \\
\boldsymbol{\omega}_s^i(\tilde{\ell}; h^t) &= \boldsymbol{w}_s(h^t) - \boldsymbol{\tau}_s^i(\tilde{\ell}; h^t) \\
\boldsymbol{T}_s^i(h^t) &= \begin{cases} \sum_{s' \geqslant 0} R^{-s'} \left[c_{s'}^i - \boldsymbol{\omega}_{s'}^i(\ell_{s'}^i; h^0) \cdot \ell_{s'}^i \right], & \text{if} \quad s = t = 0, \\ 0, & \text{if} \quad h^t \in \widetilde{H} \text{ and } s > 0, \\ \frac{R-1}{R|\mathcal{I}|} \left(\sum_{s' \geqslant t} R^{-(s'-t)} \sum_{j \in \mathcal{I}} \boldsymbol{\tau}_{s'}^j(\boldsymbol{\ell}_{s'}^j(h^t); h^t) \cdot \boldsymbol{\ell}_{s'}^j(h^t) - D(h^t) \right) & \text{otherwise.} \end{cases} \\
\{\boldsymbol{c}_s^j(h^t)\}_{s \geqslant t} &= \begin{cases} c_s^j & \text{if} \quad h^t \in \widetilde{H} \\ \arg\max_{\{c_s\}} \sum_{s \geqslant t} \beta^{s-t} u(c_s - v(\boldsymbol{\ell}_s^i(h^t))) & \text{if} \quad h^t \notin \widetilde{H} \\ \text{s.t.} \quad \sum_{s \geqslant t} R^{-(s-t)} \left[c_s - \boldsymbol{\omega}_s^i(\boldsymbol{\ell}_s^i(h^t); h^t) \cdot \boldsymbol{\ell}_s^i(h^t) - \boldsymbol{T}_s^i(h^t) \right], \end{cases}
\end{aligned}
\tag{A10}
$$

where, in the $h^t \notin \widetilde{H}$ case, we select $\boldsymbol{c}_s^j(h^t)$ to be on-path consistent in the sense of (7) and where note that the arg max used to define consumption is non-empty by assumption.

It remains to show that the profile above constitutes a Pigouvian equilibrium. The "Pigouvian" qualifier is immediate from the definition of taxes on the equilibrium path. All equilibrium conditions (including on-path consistency) are immediate from the construction except goods market clearing and household optimization, which I now show.

First, consider goods market clearing. Starting from any history $h^t \in \widetilde{H}$, goods market clearing follows from the assumption that $\{c_t^i, \ell_t^i\}$ is Pareto efficient. Now instead consider a history $h^t \notin \widetilde{H}$. Adding household budget constraints and using the definitions of lump-sum transfers and aggregate debt implies:

$$
\begin{aligned}
0 &= \sum_{s\geqslant t} R^{-(s-t)}\left[\sum_{i\in\mathcal{I}} \boldsymbol{c}_s^i(h^t) - \sum_{i\in\mathcal{I}} \boldsymbol{w}_s(h^t)\cdot\boldsymbol{\ell}_s^i(h^t) + \sum_{i\in\mathcal{I}} \boldsymbol{\tau}_s^i(\boldsymbol{\ell}_s^i(h^t);h^t)\cdot\boldsymbol{\ell}_s^i(h^t) - \sum_{i\in\mathcal{I}} \boldsymbol{T}_s^i(h^t)\right] \\
&= \sum_{s\geqslant t} R^{-(s-t)}\left[\sum_{i\in\mathcal{I}} \boldsymbol{c}_s^i(h^t) - \boldsymbol{Y}(h^t)\right] + \sum_{s\geqslant t} R^{-(s-t)} \sum_{i\in\mathcal{I}} \boldsymbol{\tau}_s^i(\boldsymbol{\ell}_s^i(h^t);h^t)\cdot\boldsymbol{\ell}_s^i(h^t) \\
&- \sum_{s\geqslant t} R^{-(s-t)} \frac{R-1}{R}\left(\sum_{s'\geqslant t} R^{-(s'-t)} \sum_{j\in\mathcal{I}} \boldsymbol{\tau}_{s'}^j(\boldsymbol{\ell}_{s'}^j(h^t);h^t)\cdot\boldsymbol{\ell}_{s'}^j(h^t) - D(h^t)\right) \\
&= \sum_{s\geqslant t} R^{-(s-t)}\left[\sum_{i\in\mathcal{I}} \boldsymbol{c}_s^i(h^t) - \boldsymbol{Y}(h^t)\right] + \cancel{\sum_{s\geqslant t} R^{-(s-t)} \sum_{i\in\mathcal{I}} \boldsymbol{\tau}_s^i(\boldsymbol{\ell}_s^i(h^t);h^t)\cdot\boldsymbol{\ell}_s^i(h^t)} \\
&- \cancel{\sum_{s'\geqslant t} R^{-(s'-t)} \sum_{j\in\mathcal{I}} \boldsymbol{\tau}_{s'}^j(\boldsymbol{\ell}_{s'}^j(h^t);h^t)\cdot\boldsymbol{\ell}_{s'}^j(h^t)} + D(h^t) \\
&= R^t \sum_{s\geqslant 0} R^{-s}\left[\sum_{i\in\mathcal{I}} \boldsymbol{c}_s^i(h^t) - \boldsymbol{Y}(h^t)\right],
\end{aligned}
\tag{A11}
$$

where the second line follows from that production is CRS.

Finally, consider household optimality. For any $h^t \in \widetilde{H}$, the Pareto efficiency of $\{c_t^i, \ell_t^i\}$ and the set of equivalent conditions in the proof of Lemma 2 that for all $i \in \mathcal{I}$,

$$
\begin{aligned}
\{\ell_s^i\}_{s\geqslant t} \quad &\in \quad \underset{\substack{\text{feasible} \\ \{\hat{\ell}_s^i\}_{s\geqslant t}}}{\arg\max} \quad \sum_{s\geqslant t} R^{-(s-t)}\left[F\left(\hat{\ell}_s^i + \sum_{j\neq i}\ell_s^j, \hat{\ell}_s^i + \sum_{j\neq i}\ell_s^j\right) - v(\hat{\ell}_s^i) - \sum_{j\neq i} v(\ell_s^j)\right] \\
\{c_s^i\}_{s\geqslant t} \quad &\in \quad \underset{\{\hat{c}_s^j\}_{s\geqslant t}}{\arg\min} \quad \sum_{s\geqslant t} R^{-(s-t)} \sum_{j\in\mathcal{I}} \hat{c}_s^j \\
&\qquad\quad \text{s.t.} \quad \sum_{s\geqslant t} \beta^{s-t} u(\hat{c}_s^i - v(\ell_s^i)) \quad \geqslant \quad \sum_{s\geqslant t} \beta^{s-t} u(c_s^i - v(\ell_s^i)).
\end{aligned}
\tag{A12}
$$

By Lemma 3 and the fact that u is increasing, these conditions imply household optimality for i so long as factor supply maximizes the NPV of income less factor supply disutility. Toward a contradition, suppose that for some feasible deviation to i's factor supply, $\{\Delta\ell_s^i\}$, we have

$$
\begin{aligned}
&\sum_{s\geqslant t} R^{-(s-t)}\Big[\overbrace{\boldsymbol{\omega}_s^i(\ell_s^i;h^t)}^{\frac{d}{dL}\big|_{L=\sum_{k\in\mathcal{I}}\ell_s^k}F(L,L)}\cdot \ell_s^i - v(\ell_s^i)\Big] \\
<&\sum_{s\geqslant t} R^{-(s-t)}\Big[\underbrace{\boldsymbol{\omega}_s^i(\ell_s^i+\Delta\ell_s^i;h^t)}_{\frac{d}{dL}\big|_{L=\sum_{k\in\mathcal{I}}\ell_s^k}F(L,L)}\cdot(\ell_s^i+\Delta\ell_s^i) - v(\ell_s^i+\Delta\ell_s^i)\Big] \qquad \text{(A13)}\\
\implies\quad &\frac{d}{d\epsilon}\sum_{s\geqslant t} R^{-(s-t)}\Big[F(L+\epsilon\Delta\ell_s^i, L+\epsilon\Delta\ell_s^i) - v(\ell_s^i+\epsilon\Delta\ell_s^i)\Big] > 0
\end{aligned}
$$

where the second line follows from that $\mathcal{A}$ is convex and v is convex. Note that the final line implies one can construct a first-order Pareto improvement, a contradiction.

It remains to verify household optimality in histories $h^t \notin \widetilde{H}$. This follows immediately from the construction of marginal taxes and consumption, using Lemma 3.

Pigouvian necessity

Suppose that the equilibrium path $\{c_t^i, \ell_t^i\}$ is first-order Pareto efficient allocation. The absence of a first-order Pareto improvement implies that for all $t > 0$, if $\ell_t^{i_t}$ is in the interior of $\mathcal{A}$, then

$$
\sum_{s\geqslant t}^{t'-1} R^{-(s-t)}\left[\frac{d}{dL}\Big|_{L=\sum_{j\in\mathcal{I}}\ell_s^j} F(L,L) - v_\ell(\ell_t^{i_t})\right] = 0, \qquad \text{(A14)}
$$

where t' is the next time after t at which i_{t-1} can adjust its factor supply (and $|\mathcal{T}|$ if no such time exists).

At the same time, by Lemma 3, household optimality for i_{t-1}—given that $\ell_t^{i_t}$ is in the interior of $\mathcal{A}$—implies that if the prices i_{t-1} faces between t and t' are differentiable in quantity, then

$$
\sum_{s\geqslant t}^{t'-1} R^{-(s-t)}\left[\frac{d}{d\ell}\Big|_{\ell=\ell_t^{i_t}}\left[\boldsymbol{\omega}_s^{i_{t-1}}(\ell;h^0)\cdot\ell\right] - v_\ell(\ell_t^{i_t})\right] = 0. \qquad \text{(A15)}
$$

We conclude that

$$
\begin{aligned}
\sum_{s\geqslant t}^{t'-1} R^{-(s-t)}\frac{d}{d\ell}\Big|_{\ell=\ell_t^{i_t}}\left[\boldsymbol{\omega}_s^{i_{t-1}}(\ell;h^0)\cdot\ell\right] &= \sum_{s\geqslant t}^{t'-1} R^{-(s-t)}\frac{d}{dL}\Big|_{L=\sum_{j\in\mathcal{I}}\ell_s^j}F(L,L). \\
\sum_{s\geqslant t}^{t'-1} R^{-(s-t)}\boldsymbol{\tau}_s^{i_{t-1}}(\ell_t^{i_t};h^0) &= -\sum_{s\geqslant t}^{t'-1} R^{-(s-t)}F_{\overline{L}}\left(\sum_{j\in\mathcal{I}}\ell_s^j, \sum_{j\in\mathcal{I}}\ell_s^j\right) \qquad \text{(A16)}\\
&\quad+\sum_{s\geqslant t}^{t'-1} R^{-(s-t)}\frac{d}{d\ell}\Big|_{\ell=\ell_t^{i_t}}\left[\boldsymbol{\omega}_s^{i_{t-1}}(\ell;h^0)\right]\cdot\ell_t^{i_t}.
\end{aligned}
$$

A.3 Proof of Theorem 1

Before we begin, a few pieces of notation:

- For any history of factor supply adjustments $h^t = (h_{t-1}, ..., h_0)$, let $\{\widehat{\ell}^j_s(h^t)\}^{j\in\mathcal{I}}_{s\leqslant t}$ be the corresponding history of factor supply, i.e. for all $j \in \mathcal{I}$, $\widehat{\ell}^j_s(h^t) = \overline{\ell}^j_0$, and for $s = 1, ..., t$, $\widehat{\ell}^j_s(h^t) = \mathbb{1}_{j=i_{s-1}} h_{s-1} + \mathbb{1}_{j\neq i_{s-1}} \widehat{\ell}^j_{s-1}(h^t)$.
- Similarly, for any feasible history of factor supply $\{\ell^j_s\}^{j\in\mathcal{I}}_{s\leqslant t}$, let $\widehat{h}^t(\{\ell^j_s\}^{j\in\mathcal{I}}_{s\leqslant t}) = (\ell^{i_{t-1}}_t, ..., \ell^{i_0}_1)$ be the corresponding history of factor supply adjustments.
- For any history h^t, let $\Pi(h^t)$ denote the set of all feasible factor supply paths that pass through h^t, i.e.

$$\Pi(h^t) \equiv \left\{ \{\ell^j_s\}^{j\in\mathcal{I}}_{s\in\mathcal{T}} \;\middle|\; \{\ell^j_s\}^{j\in\mathcal{I}}_{s\leqslant t} = \{\widehat{\ell}^j_s(h^t)\}^{j\in\mathcal{I}}_{s\leqslant t}, \text{ and } t' > t,\ \{\ell^j_{t'}\}^{j\in\mathcal{I}} \in \Gamma\left(\{\ell^j_s\}^{j\in\mathcal{I}}_{s<t'}\right) \right\},$$
$$\text{where} \quad \Gamma\left(\{\ell^j_s\}^{j\in\mathcal{I}}_{s<t'}\right) \equiv \left\{ \{\hat{\ell}^j\}^{j\in\mathcal{I}} \;\middle|\; \forall j\in\mathcal{I},\ \hat{\ell}^j = \ell^j_{t'-1} \text{ unless } j = i_{t'-1} \right\}. \tag{A17}$$

Fix a history h^t. In order to establish that the equilibrium path following h^t is Pareto efficient, it suffices by Lemma 2 to show that

$$\begin{aligned} \{\boldsymbol{\ell}^j_s(h^t)\}^{j\in\mathcal{I}}_{s\geqslant t} &\in \underset{\{\hat{\ell}^j_s\}^{j\in\mathcal{I}}_{s\geqslant t}\in\Pi(h^t)}{\arg\max}\; W\left(\{\ell^j_s\}^{j\in\mathcal{I}}_{s\geqslant t}\right) \\ &= \underset{\{\hat{\ell}^j_s\}^{j\in\mathcal{I}}_{s\geqslant t}\in\Pi(h^t)}{\arg\max}\; \sum_{s\geqslant t} R^{-(s-t)} \left[F\left(\sum_{j\in\mathcal{I}} \ell^j_s, \sum_{j\in\mathcal{I}} \ell^j_s\right) - \sum_{j\in\mathcal{I}} v(\ell^j_s) \right]. \end{aligned} \tag{A18}$$

By Theorems 4.3 and 4.5 of Stokey et al. [1989] the boundedness assumption in Footnote , it suffices to show there exists a function $\widetilde{V}$ that, for all histories $h^{t'}$ that can feasibly follow h^t, satisfies

$$\begin{aligned} &\widetilde{V}\left(\{\widehat{\ell}^j_s(h^{t'})\}^{j\in\mathcal{I}}_{s\leqslant t'}\right), \boldsymbol{\ell}_{t'+1}(h^{t'}) \\ &\quad \in \max, \underset{\ell\in\Gamma(\{\widehat{\ell}^j_s(h^{t'})\}^{j\in\mathcal{I}}_{s\leqslant t'})}{\arg\max}\; F\left(\sum_{j\in\mathcal{I}} \widehat{\ell}^j_{t'}(h^{t'}), \sum_{j\in\mathcal{I}} \widehat{\ell}^j_{t'}(h^{t'})\right) - \sum_{j\in\mathcal{I}} v\left(\widehat{\ell}^j_{t'}(h^{t'})\right) \\ &\qquad\qquad + R^{-1}\widetilde{V}\left(\{\widehat{\ell}^j_s(h^{t'})\}^{j\in\mathcal{I}}_{s\leqslant t'} \frown \ell\right), \end{aligned} \tag{A19}$$

or equivalently,

$$\begin{aligned} &\widetilde{V}\left(\{\widehat{\ell}^j_s(h^{t'})\}^{j\in\mathcal{I}}_{s\leqslant t'}\right),\ \boldsymbol{\ell}^{i_{t'}}_{t'+1}(h^{t'}) \\ &\quad = \max, \underset{\ell^{i}{}_{t'}\in\mathcal{A}}{\arg\max}\; F\left(\sum_{j\in\mathcal{I}} \widehat{\ell}^j_{t'}(h^{t'}), \sum_{j\in\mathcal{I}} \widehat{\ell}^j_{t'}(h^{t'})\right) - \sum_{j\in\mathcal{I}} v\left(\widehat{\ell}^j_{t'}(h^{t'})\right) \\ &\qquad\qquad + R^{-1}\widetilde{V}\left(\{\widehat{\ell}^j_s(h^{t'})\}^{j\in\mathcal{I}}_{s\leqslant t'} \frown \{\widehat{\ell}^j_{t'}(h^{t'})\}^{j\neq i_{t'}} \frown \ell^{i_{t'}}\right). \end{aligned} \tag{A20}$$

To this end, for all feasible histories of factor supply $\{\ell^j_s\}^{j\in\mathcal{I}}_{s\leqslant t'}$, define $\widetilde{V}(\{\ell^j_s\}^{j\in\mathcal{I}}_{s\leqslant t'})$ as the NPV of future aggregate output less factor supply disutility on the equilibrium path:

$$\widetilde{V}\left(\{\ell_s^j\}_{s\leqslant t'}^{j\in\mathcal{I}}\right) \;\equiv\; W\left(\left\{\boldsymbol{\ell}_{t''}^i\left(\widehat{h}\left(\{\ell_s^j\}_{s\leqslant t'}^{j\in\mathcal{I}}\right)\right)\right\}_{t''\geqslant t'}^{i\in\mathcal{I}}\right). \tag{A21}$$

By (a) household optimality, (b) the definition of super-Pigouvian policy, and (c) the definition of $W(\cdot)$, we have that for any history $h^{t'} = (h_{t'-1}, ..., h_0)$ that can feasibly follow h^t,

$$\begin{aligned}
\boldsymbol{\ell}_{t'+1}^{i_{t'}}(h^{t'}) \quad &\in \quad \underset{\ell^{i}{}_{t'}\in\mathcal{A}}{\arg\max}\; \widehat{U}^{i_{t'}}(\ell^{i_{t'}} - \boldsymbol{\ell}_{t'+1}^{i_{t'}}(h^{t'}), (\boldsymbol{\ell}_{t'}^{i_{t'}}(h^{t'}), h_{t'-1}, ..., h_0)) \\
&= \quad \underset{\ell^{i}{}_{t'}\in\mathcal{A}}{\arg\max}\; W\left(\left\{\boldsymbol{\ell}_s^j\left((\ell^{i_{t'}}, h_{t'-1}, h_{t'-2}, ..., h_0)\right)\right\}_{s\geqslant t'+1}^{j\in\mathcal{I}}\right) \\
\boldsymbol{\ell}_{t'+1}^{i_{t'}}(h^{t'}) \quad &\in \quad \underset{\ell^{i}{}_{t'}\in\mathcal{A}}{\arg\max}\; F\left(\sum_{j\in\mathcal{I}}\widehat{\ell}_{t'}^j(h^{t'}), \sum_{j\in\mathcal{I}}\widehat{\ell}_{t'}^j(h^{t'})\right) - \sum_{j\in\mathcal{I}} v\left(\widehat{\ell}_{t'}^j(h^{t'})\right) \\
&\qquad + R^{-1}\underbrace{W\left(\left\{\boldsymbol{\ell}_s^j\left((\ell^{i_{t'}}, h_{t'-1}, h_{t'-2}, ..., h_0)\right)\right\}_{s\geqslant t'+1}^{j\in\mathcal{I}}\right)}_{=\widetilde{V}\left(\{\widehat{\ell}_s^j(h^{t'})\}_{s\leqslant t'}^{j\in\mathcal{I}} \frown \{\widehat{\ell}_{t'}^j(h^{t'})\}^{j\neq i_{t'}} \frown \ell^{i_{t'}}\right)}
\end{aligned} \tag{A22}$$

Moreover, the definition of $W(\cdot)$ implies that the right-hand side of (A22) evaluated at $\ell^{i_{t'}} = \boldsymbol{\ell}_{t'+1}^{i_{t'}}(h^{t'})$ is equal to $W(\{\boldsymbol{\ell}_{t''}^i(h^{t'})\}_{t''\geqslant t'}^{i\in\mathcal{I}})$. So (A20) holds, completing the proof.

A.4 Proof of Theorem 2

Suppose that the functions $\{\boldsymbol{c}^i, \boldsymbol{\ell}^i\}^{i\in\mathcal{I}}$ satisfy (7) and generate paths $\{\boldsymbol{c}_s^i(h^t), \boldsymbol{\ell}_s^i(h^t)\}_{s\geqslant t}^{i\in\mathcal{I}}$ starting from any history h^t that are Pareto efficient taking as given the labor allocation at t in that history and aggregate debt $D(h^t)$ accumulated along that history.

Starting from any history h^t, (a) the fact that u is non-satiated and (b) Lemma 2 imply that

$$\begin{aligned}
\{\boldsymbol{\ell}_s^i(h^t)\}_{s\geqslant t}^{i\in\mathcal{I}} \quad &\in \quad \underset{\substack{\text{feasible } \{\ell_s^i\}_{s\geqslant t}^{i\in\mathcal{I}} \\ \text{given } h^t}}{\arg\max}\; W\left(\{\ell_s^i\}_{s\geqslant t}^{i\in\mathcal{I}}\right) \\
\forall i\in\mathcal{I}, \quad \{\boldsymbol{c}_s^i(h^t)\}_{s\geqslant t}^{i\in\mathcal{I}} \quad &\in \quad \underset{\{c_s^i\}_{s\geqslant t}^{i\in\mathcal{I}}}{\arg\max}\; \sum_{s\geqslant t}\beta^{s-t}u(c_s^i - v(\boldsymbol{\ell}_s^i(h^t))) \\
&\qquad \text{s.t.} \quad \sum_{s\geqslant t}R^{-(s-t)}c_s^i \leqslant \sum_{s\geqslant t}R^{-(s-t)}\boldsymbol{c}_s^i(h^t)
\end{aligned} \tag{A23}$$

Note that the above condition for factor supply in particular implies that for all histories $h^t = (h_{t-1}, ..., h_0)$,

$$\ell_t^{i_{t-1}}(h^t) \quad \in \quad \underset{\ell\in\mathcal{A}}{\arg\max}\; W\left((\boldsymbol{\ell}_{s\geqslant t}^j(\ell, h_{t-2}, ..., h_0))_{s\geqslant t}^{j\in\mathcal{I}}\right) \tag{A24}$$

We now construct the equilibrium as follows: Let consumption and factor supply be given by $\{\boldsymbol{c}^i, \boldsymbol{\ell}^i\}^{i\in\mathcal{I}}$. For any history h^t and any $s \geqslant t$,

$$
\begin{aligned}
\boldsymbol{L}_s(h^t) &= \sum_{i\in\mathcal{I}} \boldsymbol{\ell}^i_s(h^t) \\
\boldsymbol{Y}_s(h^t) &= F\left(\boldsymbol{L}_s(h^t), \boldsymbol{L}_s(h^t)\right) \\
\boldsymbol{w}_s(h^t) &= F_L\left(\boldsymbol{L}_s(h^t), \boldsymbol{L}_s(h^t)\right) \\
\boldsymbol{\tau}^i_s(\ell; h^t)\cdot\ell &= \begin{cases} -W\left(\left\{\boldsymbol{\ell}^j_{t'}\left(\ell \frown \{\boldsymbol{\ell}^{i_{t''}}_{t''+1}(h^t)\}_{t''=t,\dots,s-1} \frown h^t\right)\right\}^{j\in\mathcal{I}}_{t'\geqslant s+1}\right) & \\ + \sum_{t'=s+1}^{t^{\text{next}}(i,s)} R^{-(t'-(s+1))}\left[\boldsymbol{w}_{t'}(h^t)\cdot\ell - v(\ell)\right] & \text{if } i = i_s, \\ 0 & \text{otherwise} \end{cases} \\
\boldsymbol{\omega}^i_s(\ell; h^t)\cdot\ell &= \boldsymbol{w}_s(h^t)\cdot\ell - \boldsymbol{\tau}^i_s(\ell; h^t)\cdot\ell \\
\boldsymbol{T}^i_s(h^t) &= \begin{cases} \sum_{t'\geqslant t} R^{-(t'-t)}\left[\boldsymbol{c}^i_{t'}(h^t) - \boldsymbol{\omega}^i_{t'}(\boldsymbol{\ell}^i_{t'}(h^t); h^t)\cdot\boldsymbol{\ell}^i_{t'}(h^t)\right] + d^i(h^t) & \text{if } s = t \\ 0 & \text{otherwise,} \end{cases}
\end{aligned}
\tag{A25}
$$

where $t^{\text{next}}(i,s)$ is the next period after s with $i_s = i_{t^{\text{next}}(i,s)}$ (and $\mathcal{T}$ if no such period exists), and where note that $\boldsymbol{T}^i_s(h^t)$ is defined recursively, since $d^i(h^t)$ depends on transfers at times before t. Note that $\boldsymbol{\tau}^i_s(\ell; h^t)$ is constructed so that for any history h^t and any time $t' \geqslant t$ with $i = i_{t'}$,

$$
\begin{aligned}
&\sum_{s=t'+1}^{t^{\text{next}}(i,t')} R^{-(s-(t'+1))}\left[\left(\boldsymbol{w}_s(h^t) - \boldsymbol{\tau}^i_s(\ell; h^t)\right)\cdot\ell - v(\ell)\right] \\
&\qquad = W\left(\{\boldsymbol{\ell}^j_s(\ell \frown \{\boldsymbol{\ell}^{i_{t''}}_{t''+1}(h^t)\}_{t''=t,\dots,t'-1} \frown h^t)\}^{j\in\mathcal{I}}_{s\geqslant t'+1}\right).
\end{aligned}
\tag{A26}
$$

As the equilibrium consumption and factor supply profile is Pareto efficient starting from any history by construction, it suffices to verify that this profile is a super-Pigouvian equilibrium, i.e. (7) and (2)–(6) hold. (7) holds for consumption and factor supply by assumption, and therefore holds for all other equilibrium functions because they are constructed from the consumption and factor supply functions.

Of the conditions (2)–(6), all are immediate from their construction and—in the case of goods market clearing—the fact that for a profile to be Pareto efficient from any history it must satisfy goods market clearing from any history, except for household optimality. By Lemma 3, it suffices to show, for every history h^t and household i that

$$
\begin{aligned}
\{\boldsymbol{c}^i_s(h^t)\}_{s\geqslant t} &\in \underset{\{\hat{c}^i_s\}_{s\geqslant t}}{\arg\min} \sum_{s\geqslant t} \beta^{s-t} u(\hat{c}^i_s - v(\boldsymbol{\ell}^i_s(h^t))) \\
&\quad \text{s.t.} \quad \sum_{s\geqslant t} R^{-(s-t)}\hat{c}^i_s \leqslant \sum_{s\geqslant t} R^{-(s-t)}\left[\omega^i_s(\boldsymbol{\ell}^i_s(h^t); h^t)\cdot\boldsymbol{\ell}^i_s(h^t) + \boldsymbol{T}^i_s(h^t)\right] - d^i(h^t) \\
\{\boldsymbol{\ell}^i_s(h^t)\}_{s\geqslant t} &\in \underset{\text{feasible } \{\hat{\ell}^i_s\}_{s\geqslant t}}{\arg\min} \sum_{s\geqslant t} R^{-(s-t)}\left(\boldsymbol{\omega}^i_s(\hat{\ell}^i_s; h^t)\cdot\hat{\ell}^i_s - v(\hat{\ell}^i_s)\right)
\end{aligned}
\tag{A27}
$$

To see that the condition for consumption is true, note that by the definition of transfers, the budget constraint can be replaced by $\sum_{s\geqslant t} R^{-(s-t)}\hat{c}^i_s \leqslant \sum_{s\geqslant t} R^{-(s-t)}\boldsymbol{c}^i_s(h^t)$. The fact that consumption satisfies this condition then follows immediately from the Pareto efficiency of the profile with which we began the proof. Next, consider the condition for labor supply. Note that it is equivalent to the condition that for all times $t' \geqslant t$ at which $i = i_{t'}$,

$$\boldsymbol{\ell}^i_{t'+1}(h^t) \;\in\; \arg\min_{\ell\in\mathcal{A}} \sum_{s=t'+1}^{t^{\text{next}}(i,t')} R^{-(s-(t'+1))}\left(\boldsymbol{\omega}^i_s(\ell;h^t)\cdot\ell - v(\ell)\right)$$

$$\Longleftrightarrow$$

$$\boldsymbol{\ell}^i_{t'+1}\left(\{\boldsymbol{\ell}^{i_{t''}}_{t''+1}(h^t)\}_{t''=t,\dots,t'-1} \frown h^t\right) \in \arg\min_{\ell\in\mathcal{A}} W\left(\{\boldsymbol{\ell}^j_s(\ell \frown \{\boldsymbol{\ell}^{i_{t''}}_{t''+1}(h^t)\}_{t''=t,\dots,t'-1} \frown h^t)\}^{j\in\mathcal{I}}_{s\geqslant t'+1}\right) \tag{A28}$$

where the equivalence follows from—on the LHS—the fact that the profile we began with satisfies (7) and—on the RHS—(A26). As desired, the last condition holds by (A24).

A.5 Proof of Proposition 2

Fix a policy rule, fundamentals θ, factor supply strategies $\boldsymbol{\ell}$, and a feasible history h^{t-1}. Let t' be the next time at which i_{t-1} can adjust its behavior (or $|\mathcal{T}|$ if no such time exists). For $\ell \in \mathcal{A}$, define

$$U^{i_{t-1}}(\ell;\theta) \equiv \sum_{s=t}^{t'-1} R^{-(s-t)}\left[\boldsymbol{\omega}^{i_{t-1}}_s(\ell;h^{t-1};\theta,\boldsymbol{\ell})\cdot\ell - v^{i_{t-1}}_t(\ell;\theta)\right], \tag{A29}$$

where I have suppressed the dependence on θ of variables (such as i_{t-1} and t') that we will hold fixed throughout, and where post-tax prices are equal to market prices given factor supply follows the equilibrium path, less marginal taxes under the policy rule:

$$\boldsymbol{\omega}^{i_{t-1}}_s(\ell;h^{t-1};\theta,\boldsymbol{\ell}) = F_L\left(\sum_{j\in\mathcal{I}}\boldsymbol{\ell}^j_t(h^{t-1}), \sum_{j\in\mathcal{I}}\boldsymbol{\ell}^j_t(h^{t-1})\right) - \boldsymbol{\tau}^{i_{t-1}}_s(\ell;h^{t-1};\theta,\boldsymbol{\ell}). \tag{A30}$$

Toward a contradiction, suppose that the policy rule is not super-Pigouvian, i.e. there exist $\ell,\ell' \in \mathcal{A}$ such that

$$U^{i_{t-1}}(\ell,\theta) - U^{i_{t-1}}(\ell',\theta) \quad > \quad W(\{\boldsymbol{\ell}^j_s(\ell\frown h^t)\},\theta) - W(\{\boldsymbol{\ell}^j_s(\ell'\frown h^t)\},\theta), \tag{A31}$$

where W is as defined in Footnote 16. Expanding definitions, this implies that

$$\begin{aligned}
G(\ell;\theta) - G(\ell';\theta) \;&>\; G^*(\ell;\theta) - G^*(\ell';\theta) \\
\text{where} \quad G(\tilde{\ell};\theta) &\equiv \sum_{s=t}^{t'-1} R^{-(s-t)} \boldsymbol{\omega}_s^{i_{t-1}}(\tilde{\ell}; h^{t-1}; \theta, \boldsymbol{\ell}) \cdot \tilde{\ell} \\
\text{and} \quad G^*(\tilde{\ell};\theta) &\equiv \sum_{s\geqslant t} R^{-(s-t)} \Bigg[F\Bigg(\sum_{j\in\mathcal{I}} \boldsymbol{\ell}_s^j(\tilde{\ell} \frown h^{t-1}), \sum_{j\in\mathcal{I}} \boldsymbol{\ell}_s^j(\tilde{\ell} \frown h^{t-1}) \Bigg) \\
&\qquad - \sum_{j\in\mathcal{I}} \mathbb{1}_{j\neq i_{t-1} \text{ or } s\geqslant t'} v_s^j(\boldsymbol{\ell}_s^j(\tilde{\ell} \frown h^{t-1})) \Bigg].
\end{aligned} \tag{A32}$$

Note that $G(\tilde{\ell};\theta)$ and $G^*(\tilde{\ell};\theta)$ are both invariant to changes in θ that only affect $v_s^{i_{t-1}}$ for $s \geqslant t$ and $s < t'$.

Now, construct a new fundamental θ' as follows: Keep all quantities the same as under θ except for $v_t^{i_{t-1}}$, defined as:

$$\begin{aligned}
v_t^{i_{t-1}}(\tilde{\ell};\theta') \;&\equiv\; v_t^{i_{t-1}}(\tilde{\ell};\theta) \;+\; U^{i_{t-1}}(\tilde{\ell},\theta) \;-\; \mathbb{1}_{\tilde{\ell}=\ell}\epsilon \;-\; \mathbb{1}_{\tilde{\ell}=\ell'}\epsilon/2 \\
\text{where} \quad \epsilon \;&\equiv\; [G(\ell;\theta) - G(\ell';\theta)] \;-\; [G^*(\ell;\theta) - G^*(\ell';\theta)].
\end{aligned} \tag{A33}$$

To obtain a contradiction and therefore complete the proof, I now claim that at fundamentals θ',

- ℓ solves i_{t-1}'s problem at h^{t-1} when she faces taxes given by the policy rule, and
- ℓ is Pareto inefficient given that future behavior follows $\boldsymbol{\ell}$.

The first point is immediate from the construction of $v_t^{i_{t-1}}(\cdot,\theta')$, the fact that $\epsilon > 0$, and the assumption that i_{t-1}'s post-tax wages at h_{t-1} are the same under θ and θ'. To see the second point, note that it suffices by the forward direction of the proof of Lemma 2 to show that $W(\{\boldsymbol{\ell}_s^j(\ell \frown h^{t-1})\},\theta') < W(\{\boldsymbol{\ell}_s^j(\ell' \frown h^{t-1})\},\theta')$. Indeed, the definitions of θ', G, and G^* and (A31) imply

$$\begin{aligned}
&W(\{\boldsymbol{\ell}_s^j(\ell' \frown h^{t-1})\},\theta') - W(\{\boldsymbol{\ell}_s^j(\ell \frown h^{t-1})\},\theta') \\
&= \big[W(\{\boldsymbol{\ell}_s^j(\ell' \frown h^{t-1})\},\theta) - U^{i_{t-1}}(\ell';\theta) + \epsilon/2\big] \\
&- \big[W(\{\boldsymbol{\ell}_s^j(\ell \frown h^{t-1})\},\theta) - U^{i_{t-1}}(\ell;\theta) + \epsilon\big] \\
&= [G^*(\ell';\theta) - G(\ell';\theta)] - [G^*(\ell;\theta) - G(\ell;\theta)] - \epsilon/2 = \epsilon/2 > 0.
\end{aligned} \tag{A34}$$

A.6 Proof of Proposition 3

Below, I state a full version of the proposition that defines the δ_{SP} term used in the main text. To do so, I introduce the following notation: Fix a history $h^{t>0} = (h_{t-1}, ..., h_0)$. First, denote by $\mathrm{O}^{i_{t-1}}(\Delta\ell, h^t)$ i_{t-1}'s utility loss—in NPV consumption terms, starting from t—from misoptimizing along the path where prices and transfers follow h^t but i_{t-1}'s initial labor supply is $h_{t-1} + \Delta\ell$ that results from i_{t-1} playing according to her strategy that is optimal if prices and transfers follow her deviation. Second, denote by $\mathrm{I}^{i_{t-1}}(\Delta\ell, h^t)$ the internality bias—again, in NPV consumption units starting from h^t—that i_{t-1} has toward her equilibrium behavior according to the deviated-by-$\Delta\ell$ path following h^t because she believes that prices and transfers evolve as they do on the equilibrium path following h^t rather than on the deviated-by-$\Delta\ell$ path.[33]

Importantly, note that both of these terms disappear in any limit where i_{t-1} has no market power, i.e. cannot influence prices or transfers (except through the quantity-dependence of marginal taxes). For $O^{i_{t-1}}$, this is because if i_{t-1}'s behavior has no effect on prices, then her optimal behavior following h^t and $(h_{t-1}+\Delta\ell, h_{t-2}, ..., h_0)$ only differs due to the change in initial condition. For $I^{i_{t-1}}$, this is because i_{t-1} has no internality if her incorrect belief that she does not affect prices or transfers is, in fact, true.

Proposition 3'. *An equilibrium is super-Pigouvian if and only if for all histories $h^{t>0} = (h_{t-1}, ..., h_0)$ and all alternative histories $h'^t = (h_{t-1}+\Delta\ell, ..., h_0)$ with $h_{t-1}+\Delta\ell \in \mathcal{A}$,*

$$\sum_{j \neq i_{t-1}} e^j \left(\left\{ \boldsymbol{c}_s^j(h'^t), \boldsymbol{\ell}_s^j(h'^t) \right\}_{s \geqslant t}, \left\{ \boldsymbol{c}_s^j\left(h^t\right), \boldsymbol{\ell}_s^j\left(h^t\right) \right\}_{s \geqslant t} \right) - O^i(\Delta\ell, h) - I^i(\Delta\ell, h) = 0. \quad \text{(A36)}$$

Fix a history $h^t = (h_{t-1}, ..., h_0)$ and for notational convenience, let $i \equiv i_{t-1}$. Fix some $\Delta\ell$ such that $h_{t-1}+\Delta\ell \in \mathcal{A}$, and let $h'^t = (h_{t-1}+\Delta\ell, h_{t-2}, ..., h_0)$. Finally, let t' denote the next time after $t-1$ at which $i = i_{t'}$, or $|\mathcal{T}|$ if no such time exists. We then have

$$\begin{aligned}
&\sum_{j \neq i} e^j \left(\left\{ \boldsymbol{c}_s^j(h'^t), \boldsymbol{\ell}_s^j(h'^t) \right\}_{s \geqslant t}, \left\{ \boldsymbol{c}_s^j\left(h^t\right), \boldsymbol{\ell}_s^j\left(h^t\right) \right\}_{s \geqslant t} \right) \\
= &\sum_{s \geqslant t} R^{-(s-t)} \sum_{j \neq i} \Big[\left(\boldsymbol{c}_s^j(h'^t) - v(\boldsymbol{\ell}_s^j(h'^t)) \right) \\
&\qquad - \left(\boldsymbol{c}_s^j(h^t) - v(\boldsymbol{\ell}_s^j(h^t)) \right) \Big] \\
= &\sum_{s \geqslant t} R^{-(s-t)} \sum_{j \neq i} \Big[\left(\boldsymbol{\omega}_s^j(\boldsymbol{\ell}_s^j(h'^t); h'^t) \cdot \boldsymbol{\ell}_s^j(h'^t) + \boldsymbol{T}_s^j(h'^t) - v(\boldsymbol{\ell}_s^j(h'^t)) \right) \\
&\qquad - \left(\boldsymbol{\omega}_s^j(\boldsymbol{\ell}_s^j(h^t); h^t) \cdot \boldsymbol{\ell}_s^j(h^t) + \boldsymbol{T}_s^j(h^t) - v(\boldsymbol{\ell}_s^j(h^t)) \right) \Big] \\
= &\sum_{s \geqslant t} R^{-(s-t)} \sum_{j \neq i} \Big[\boldsymbol{w}_s^j(h'^t) \cdot \boldsymbol{\ell}_s^j(h'^t) - v(\boldsymbol{\ell}_s^j(h'^t)) - \boldsymbol{w}_s^j(h^t) \cdot \boldsymbol{\ell}_s^j(h^t) + v(\boldsymbol{\ell}_s^j(h^t)) \\
&\qquad - \boldsymbol{\tau}_s^j(\boldsymbol{\ell}_s^j(h'^t); h'^t) \cdot \boldsymbol{\ell}_s^j(h'^t) + \boldsymbol{T}_s^j(h'^t) + \boldsymbol{\tau}_s^j(\boldsymbol{\ell}_s^j(h^t); h^t) \cdot \boldsymbol{\ell}_s^j(h^t) - \boldsymbol{T}_s^j(h^t) \Big] \\
= &\sum_{s \geqslant t} R^{-(s-t)} \Big[F(\boldsymbol{L}_s(h'^t), \boldsymbol{L}_s(h'^t)) - \boldsymbol{w}_s^i(h'^t) \cdot \boldsymbol{\ell}_s^i(h'^t) - F(\boldsymbol{L}_s(h^t), \boldsymbol{L}_s(h^t)) + \boldsymbol{w}_s^i(h^t) \cdot \boldsymbol{\ell}_s^i(h^t)
\end{aligned}$$

[33] Formally, fix any history $h^t = (h_{t-1}, ..., h_0)$ and any $\Delta\ell$ such that $h_{t-1}+\Delta\ell \in \mathcal{A}$. Let $i \equiv i_{t-1}$. Then

$$\begin{aligned}
O^i(\Delta\ell, h^t) &= \max_{\substack{\text{feasible } \{\ell_s^i\}_{s \geqslant t} \\ \text{given } \ell_s^i = h_{t-1}+\Delta\ell}} \sum_{s \geqslant t} R^{-(s-t)} \left[\left(\boldsymbol{\omega}_s^i(\hat{\ell}_s^i; h^t) \cdot \hat{\ell}_s^i - v(\hat{\ell}_s^i) \right) - \left(\boldsymbol{\omega}_s^i({\ell'}_s^i; h^t) \cdot {\ell'}_s^i - v({\ell'}_s^i) \right) \right], \\
I^i(\Delta\ell, h^t) &= \sum_{s \geqslant t} R^{-(s-t)} \left[\boldsymbol{\omega}_s^i({\ell'}_s^i; h^t) \cdot {\ell'}_s^i + \boldsymbol{T}_s^i(h^t) - \boldsymbol{\omega}_s^i({\ell'}_s^i; h'^t) \cdot {\ell'}_s^i - \boldsymbol{T}_s^i(h'^t) \right],
\end{aligned} \quad \text{(A35)}$$

where $\ell_s^i = \boldsymbol{\ell}_s^i(h^t)$, ${\ell'}_s^i = \boldsymbol{\ell}_s^i(h'^t)$, and $h'^t = (h_{t-1}+\Delta\ell, h_{t-2}, ..., h_0)$.

$$
\begin{aligned}
&\quad -\sum_{j\neq i}\Big(v(\boldsymbol{\ell}_s^j(h'^t)) - v(\boldsymbol{\ell}_s^j(h^t))\Big) \\
&\quad +\boldsymbol{\tau}_s^i(\boldsymbol{\ell}_s^j(h'^t);h'^t)\cdot\boldsymbol{\ell}_s^i(h'^t) - \boldsymbol{T}_s^i(h'^t) - \boldsymbol{\tau}_s^i(\boldsymbol{\ell}_s^i(h^t);h^t)\cdot\boldsymbol{\ell}_s^i(h^t) + \boldsymbol{T}_s^i(h^t)\Big] \\
&= \sum_{s\geqslant t} R^{-(s-t)}\Bigg[F(\boldsymbol{L}_s(h'^t),\boldsymbol{L}_s(h'^t)) - F(\boldsymbol{L}_s(h^t),\boldsymbol{L}_s(h^t)) - \sum_{j\in\mathcal{I}}\Big(v(\boldsymbol{\ell}_s^j(h'^t)) - v(\boldsymbol{\ell}_s^j(h^t))\Big) \\
&\quad -\boldsymbol{\omega}_s^i(\boldsymbol{\ell}_s^j(h'^t);h'^t)\cdot\boldsymbol{\ell}_s^i(h'^t) - \boldsymbol{T}_s^i(h'^t) + v(\boldsymbol{\ell}_s^j(h'^t)) \\
&\quad +\boldsymbol{\omega}_s^i(\boldsymbol{\ell}_s^i(h^t);h^t)\cdot\boldsymbol{\ell}_s^i(h^t) + \boldsymbol{T}_s^i(h^t) - v(\boldsymbol{\ell}_s^j(h^t))\Bigg] \qquad \text{(A37)} \\
&= I^i(\Delta\ell,h^t) + W\left(\left\{\boldsymbol{\ell}_s^j\left((h_{t-1}+\Delta\ell,h_{t-2},...,h_0)\right)\right\}_{s\geqslant t}^{j\in\mathcal{I}}\right) - W\left(\left\{\boldsymbol{\ell}_s^j\left(h^t\right)\right\}_{s\geqslant t}^{j\in\mathcal{I}}\right) \\
&+\sum_{s\geqslant t} R^{-(s-t)}\Big[-\boldsymbol{\omega}_s^i(\boldsymbol{\ell}_s^j(h'^t);h^t)\cdot\boldsymbol{\ell}_s^i(h'^t) - \boldsymbol{T}_s^i(h^t) + v(\boldsymbol{\ell}_s^j(h'^t)) \\
&\quad +\boldsymbol{\omega}_s^i(\boldsymbol{\ell}_s^i(h^t);h^t)\cdot\boldsymbol{\ell}_s^i(h^t) + \boldsymbol{T}_s^i(h^t) - v(\boldsymbol{\ell}_s^j(h^t))\Big] \\
&= I^i(\Delta\ell,h^t) + O^i(\Delta\ell,h^t) + W\left(\left\{\boldsymbol{\ell}_s^j\left((h_{t-1}+\Delta\ell,h_{t-2},...,h_0)\right)\right\}_{s\geqslant t}^{j\in\mathcal{I}}\right) - W\left(\left\{\boldsymbol{\ell}_s^j\left(h^t\right)\right\}_{s\geqslant t}^{j\in\mathcal{I}}\right) \\
&+\sum_{s\geqslant t} R^{-(s-t)}\Big[-\boldsymbol{\omega}_s^i((\boldsymbol{\ell}_s^j(h^t)+\mathbb{1}_{s\leqslant t'}\Delta\ell);h^t)\cdot(\boldsymbol{\ell}_s^j(h^t)+\mathbb{1}_{s\leqslant t'}\Delta\ell) - \boldsymbol{T}_s^i(h^t) + v(\boldsymbol{\ell}_s^j(h^t)+\mathbb{1}_{s\leqslant t'}\Delta\ell) \\
&\quad +\boldsymbol{\omega}_s^i(\boldsymbol{\ell}_s^i(h^t);h^t)\cdot\boldsymbol{\ell}_s^i(h^t) + \boldsymbol{T}_s^i(h^t) - v(\boldsymbol{\ell}_s^j(h^t))\Big] \\
&= I^i(\Delta\ell,h^t) + O^i(\Delta\ell,h^t) + W\left(\left\{\boldsymbol{\ell}_s^j\left((h_{t-1}+\Delta\ell,h_{t-2},...,h_0)\right)\right\}_{s\geqslant t}^{j\in\mathcal{I}}\right) - W\left(\left\{\boldsymbol{\ell}_s^j\left(h^t\right)\right\}_{s\geqslant t}^{j\in\mathcal{I}}\right) \\
&\quad -\Delta U^i(\Delta\ell,h^t),
\end{aligned}
$$

Above, the first equality is by the definition of e^j. The second equality is by the household budget constraint. The third and fifth inequalities are by the equilibrium condition relating pre- and post-tax wages. The fourth inequality is by (a) the fact that production is constant returns to scale and (b) the government budget constraint (which recall is implied by Walras' law) and the fact that h^t and h'^t are the same up until $t^+(h)$. The sixth equality is by the definitions of I^i and W^i. The seventh equality is by the definition of O^i and the observation that—conditional on playing $h_{t-1}+\Delta\ell$ at t—$\boldsymbol{\ell}_s^i(h^t)+\mathbb{1}_{s\leqslant t'}\Delta\ell$ is optimal behavior given prices following $\tilde{h}^t$, since by household optimality $\boldsymbol{\ell}_s^i(h^t)$ is optimal ignoring the constraint of playing ℓ_t^i at t. The seventh equality is by the definition of ΔU^i.

Moving $I^i(\Delta\ell,h^t)$ and $\widetilde{O}^i(\Delta\ell,h^t)$ to the left-hand side reveals that (A36) holds (i.e. the LHS is zero) for all histories h^t if and only if (15) holds (i.e. the RHS is zero) for all histories h^t.

B Analysis of quantitative model

B.1 Characterization of sectoral laboral labor supply

Each unit of labor m is assigned to the traditional sector if and only if $w_t^T a(m) > w_t^I$. So

$$L_t^I = M\mathbb{P}[a \leqslant w_t^I/w_t^T] = M\left(\frac{w_t^I}{\bar{a}w_t^T}\right)^{\varepsilon}, \tag{A38}$$

provided $w_t^I/w_t^T \leqslant \bar{a}$. Similarly,

$$L_t^T = M\int_{w_t^I/w_t^T}^{\bar{a}} a\frac{d}{da}\mathbb{P}[\tilde{a} \leqslant a]da = M\int_{w_t^I/w_t^T}^{\bar{a}} a\varepsilon\frac{a^{\varepsilon-1}}{\bar{a}^{\varepsilon}}da = M\frac{\varepsilon}{\varepsilon+1}\frac{1}{\bar{a}^{\varepsilon}}\left(\bar{a}^{1+\varepsilon} - \left(\frac{w_t^I}{w_t^T}\right)^{1+\varepsilon}\right). \tag{A39}$$

B.2 Characterization of equilibrium given investment

I now characterize the equilibrium of the quantitative model of Section 6 at any period t, given some fixed N_t and M.

First, wages are pinned down by the optimal price-setting behavior of intermediate producers and the normalization of the final good's price to one.

$$\begin{aligned} 1 &= N_t\left(\frac{\eta}{\eta-1}\frac{(w_t^I)^{\alpha}}{\alpha^{\alpha}(1-\alpha)^{1-\alpha}z^N z_t}\right)^{1-\eta} + (1-N_t)\left(\frac{\eta}{\eta-1}\frac{(w_t^I)^{\alpha}}{\alpha^{\alpha}(1-\alpha)^{1-\alpha}z^O z_t}\right)^{1-\eta} \\ \Longrightarrow \quad w_t^I &= \left[\frac{\eta-1}{\eta}\alpha^{\alpha}(1-\alpha)^{1-\alpha}z_t\right]^{\frac{1}{\alpha}}\left[N_t(z^N)^{\eta-1} + (1-N_t)(z^O)^{\eta-1}\right]^{\frac{1}{\alpha(\eta-1)}}. \end{aligned} \tag{A40}$$

Industrial labor supply is then given by the industrial labor supply curve:

$$L_t^I = \frac{M}{\bar{a}^{\varepsilon}}(w_t^I)^{\varepsilon}. \tag{A41}$$

Cost minimization by the final producer implies a relationship between each variety producer's prices and its gross output, which—given industrial wages and labor supply—

pins down the production of each firm as well as industrial input usage and gross output.

$$
\begin{aligned}
& w_t^I l_t(n) \propto p_t(n) q_t(n) \propto p_t(n)^{1-\eta} \propto z_t(n)^{\eta-1} \\
\Longrightarrow \quad & w_t^I l_t(n) = \frac{z_t(n)^{\eta-1}}{N_t (z^N)^{\eta-1} + (1-N_t)(z^O)^{\eta-1}} w_t^I L_t^I \\
\Longrightarrow \quad & x_t(n) = \frac{z_t(n)^{\eta-1}}{N_t (z^N)^{\eta-1} + (1-N_t)(z^O)^{\eta-1}} \frac{1-\alpha}{\alpha} w_t^I L_t^I, \\
& p_t(n) q_t(n) = \frac{z_t(n)^{\eta-1}}{N_t (z^N)^{\eta-1} + (1-N_t)(z^O)^{\eta-1}} \frac{\eta}{\eta-1} \frac{w_t^I L_t^I}{\alpha} \\
\Longrightarrow \quad & X_t = \frac{1-\alpha}{\alpha} w_t^I L_t^I, \qquad Q_t = \frac{\eta}{\eta-1} \frac{w_t^I L_t^I}{\alpha}.
\end{aligned}
\tag{A42}
$$

Aggregate net output is the sum of traditional sector output (or equivalently, effective labor supply in the traditional sector) with industrial gross output, less intermediate goods and investment expenditures.

$$
\begin{aligned}
Y_t &= L_t^T + Q_t - X_t - (f^N N_t + f^O(1-N_t)). \\
&= \text{const} - \frac{\varepsilon}{\varepsilon+1} w_t^I L_t^I + \frac{\eta}{\eta-1} \frac{w_t^I L_t^I}{\alpha} - \frac{1-\alpha}{\alpha} w_t^I L_t^I - (f^N N_t + f^O(1-N_t)). \\
&= \text{const} + \left(\frac{1}{\varepsilon+1} + \frac{1}{\alpha(\eta-1)} \right) w_t^I L_t^I - (f^N N_t + f^O(1-N_t))
\end{aligned}
\tag{A43}
$$

for some exogenous constant.

Given this expression for aggregate output, it is easy to compute the contribution of an additional firm's use of the new technology to output—or in other words a firm's total renumeration for investment under Pigouvian policy.

$$
\begin{aligned}
\frac{dY_t}{dN_t} &= \left(\frac{1}{\varepsilon+1} + \frac{1}{\alpha(\eta-1)} \right) \frac{d}{dN_t} \frac{M}{\bar{a}^\varepsilon} \left(w_t^I \right)^{1+\varepsilon} - (f^N - f^O) \\
&= \left(1 + \frac{1+\varepsilon}{\alpha(\eta-1)} \right) w_t^I L_t^I \frac{d \log w_t^I}{dN_t} - (f^N - f^O), \\
\text{where} \quad \frac{d \log w_t}{dN_t} &= \frac{1}{\alpha(\eta-1)} \frac{(z^N)^{\eta-1} - (z^O)^{\eta-1}}{N_t (z^N)^{\eta-1} + (1-N_t)(z^O)^{\eta-1}}.
\end{aligned}
\tag{A44}
$$

One other useful quantity is the difference in profits between new- and old- technology firms under laissez-faire investment policy:

$$
\begin{aligned}
\pi_t^{LF,N} - \pi_t^{LF,O} &= \left(\frac{p_t^N q_t^N}{\eta} - f^N \right) - \left(\frac{p_t^O q_t^O}{\eta} - f^O \right) \\
&= \frac{1}{\alpha(\eta-1)} \frac{(z^N)^{\eta-1} - (z^O)^{\eta-1}}{N_t (z^N)^{\eta-1} + (1-N_t)(z^O)^{\eta-1}} w_t^I L_t^I - \left(f^N - f^O \right) \\
&= w_t^I L_t^I \frac{d \log w_t^I}{dN_t} - \left(f^N - f^O \right).
\end{aligned}
\tag{A45}
$$

This implies that the Pigouvian investment subsidy is

$$\frac{dY_t}{dN_t} - \left(\pi_t^{LF,N} - \pi_t^{LF,O}\right) = \frac{1+\varepsilon}{\alpha(\eta-1)} w_t^I L_t^I \frac{d\log w_t^I}{dN_t} \tag{A46}$$

B.3 Expressions used for calibration

In this appendix I go through the mathematical steps required to back out various model parameters from moments of the data. I rely on expressions for various model relationships derived in Appendix B.2.

I set z^O to match HCI output per worker in 1972, given α and η and using (a) the observation that $Q_t = \frac{\eta}{\eta-1}\frac{w_t^I L_t^I}{\alpha}$ and (b) the expression for wages when all households use the old technology. Combining (a) and (b) implies

$$\left[\frac{\eta-1}{\eta}\alpha^\alpha(1-\alpha)^{1-\alpha}z^O\right]^{\frac{1}{\alpha}}\frac{L_t^I}{Q_t} = \frac{\eta-1}{\eta}\alpha \implies z^O = \left(\frac{\eta}{\eta-1}\Big/(1-\alpha)\right)^{1-\alpha}\left(\frac{Q_0}{L_0^I}\right)^\alpha. \tag{A47}$$

Applying the same steps, I set z^N to match Lane [2022]'s estimates of growth in HCI gross output and employment, relative to non-HCI manufacturing.

$$z^N = \left(\frac{\eta}{\eta-1}\Big/(1-\alpha)\right)^{1-\alpha}\left(\frac{Q_0}{L_0^I}e^{\Delta\log Q - \Delta\log L^I}\right)^\alpha. \tag{A48}$$

I set ε to match the relative growth of HCI employment and gross output, using the fact that $Q_t = \frac{\eta}{\eta-1}\frac{w_t^I L_t^I}{\alpha}$.

$$\begin{aligned}\Delta\log Q_t - \Delta\log L_t^I = \Delta\log w_t^I &= \frac{1}{\varepsilon}\Delta\log L_t^I \\ \implies \varepsilon &= \frac{\Delta\log L_t^I}{\Delta\log Q_t - \Delta\log L_t^I}.\end{aligned} \tag{A49}$$

I set $M\bar{a}^{-\varepsilon}$ to match HCI employment in 1972

$$L_0^I = M\bar{a}^{-\varepsilon}(w_0^I)^\varepsilon \implies M\bar{a}^{-\varepsilon} = \frac{L_0^I}{(w_0^I)^\varepsilon} = \frac{L_0^I}{\left(\frac{\eta-1}{\eta}\alpha\frac{Q_0}{L_0^I}\right)^\varepsilon}. \tag{A50}$$

B.4 Robustness to alternative productivity processes

The calibration used in the main text assumes that the productivity process has noise parameter $\sigma = 0.01$ and mean reversion parameter $\theta = 2$. This serves to illustrate the potential for equilibrium multiplicity under Pigouvian policy.

However, it is well known that with no mean reversion, models (including mine) that fit within the framework of Frankel and Pauzner [2000] have a unique dynamic, stochastic equilibrium. The essential feature of such models is that productivity can reach extreme enough values that either choice for firms (adopting the old or new technology) is a dominant strategy—i.e. it is individually rational, however other firms will behave. Intuitively, if firms know that others will adopt the new technology if productivity reaches a high enough level, they become more willing to adopt it themselves, which makes other firms more willing to adopt it, etc. When productivity follows a random walk, this "contagion" from the dominance regions eventually leaves only one rationalizable strategy. The fact that Pigouvian policy can always implement the first best then implies that this unique equilibrium must be Pareto efficient. In particular, there is no scope for coordination failure and therefore no room for welfare improvements via super-Pigouvian policy.

Away from the case where productivity follows a perfect random walk, it is an open research question in what models multiple equilibria exist [Frankel et al., 2005]. Arguing by analogy to the case of static global games, Angeletos and Lian [2016] suggest that multiplicity is the most likely when the productivity process is strongly mean reverting and has little noise. Intuitively, such cases are the closest to the static model with productivity at the mean—where we know that multiplicity is possible because investment is complementary.

To gauge the sensitivity of my conclusions to the parameterization of the productivity process, I simulate my model under a range of alternative assumptions about the mean reversion parameter θ and the noise parameter σ. Figure A1 shows at what parameter values I find unique equilibria and at what parameter values I find multiple Pigouvian equilibria. In particular, it shows that multiplicity requires sufficiently high mean reversion and low noise, as Angeletos and Lian [2016] suggest.

In all cases where I find multiplicity, the lowest- and highest-industrializing Pigouvian equilibria are qualitatively similar to those in the calibration discussed in the main text. In particular, starting from the pre-HCI-drive initial condition—i.e. productivity at its mean and no firms adopting the new technology—firms choose the new technology in the highest-industrializing equilibrium and choose the old technology in the lowest-industrializing equilibrium. This implies that, with high probability, the economy industrializes in the former case and does not industrialize in the latter. As a result, welfare gains from super-Pigouvian policy are quantitatively similar across the cases where I find multiplicity. When there is no Pigouvian multiplicity, there are no welfare gains from super-Pigouvian policy, as discussed above.

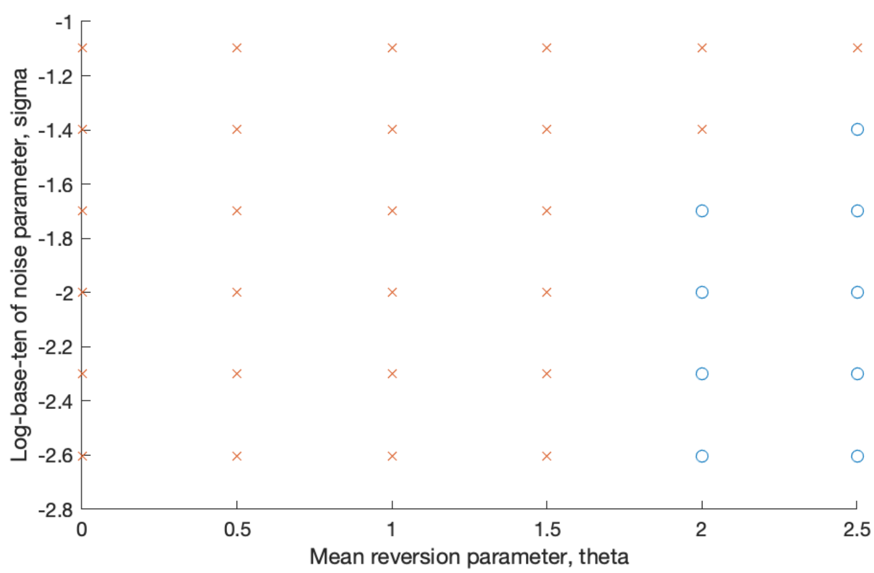

Figure A1: Presence of multiple Pigouvian equilibria at different levels of noise and mean-reversion in the productivity process. Red "X" = unique equilibrium. Blue "O" = multiple equilibria.

C Additional figures

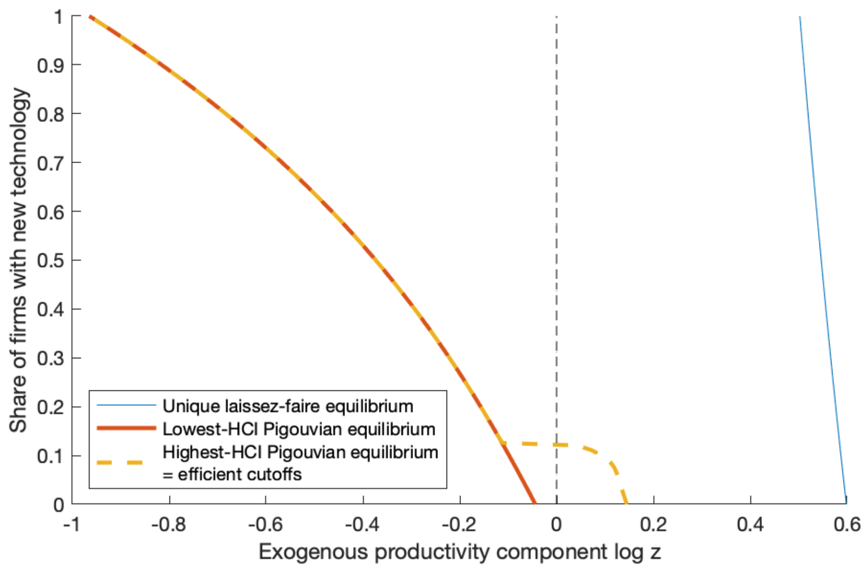

Figure A2: Productivity cutoffs for new technology adoption characterizing the highest- and lowest-industrializing equilibria of the calibrated model under Pigouvian policy, as well as the unique laissez-faire equilibrium.

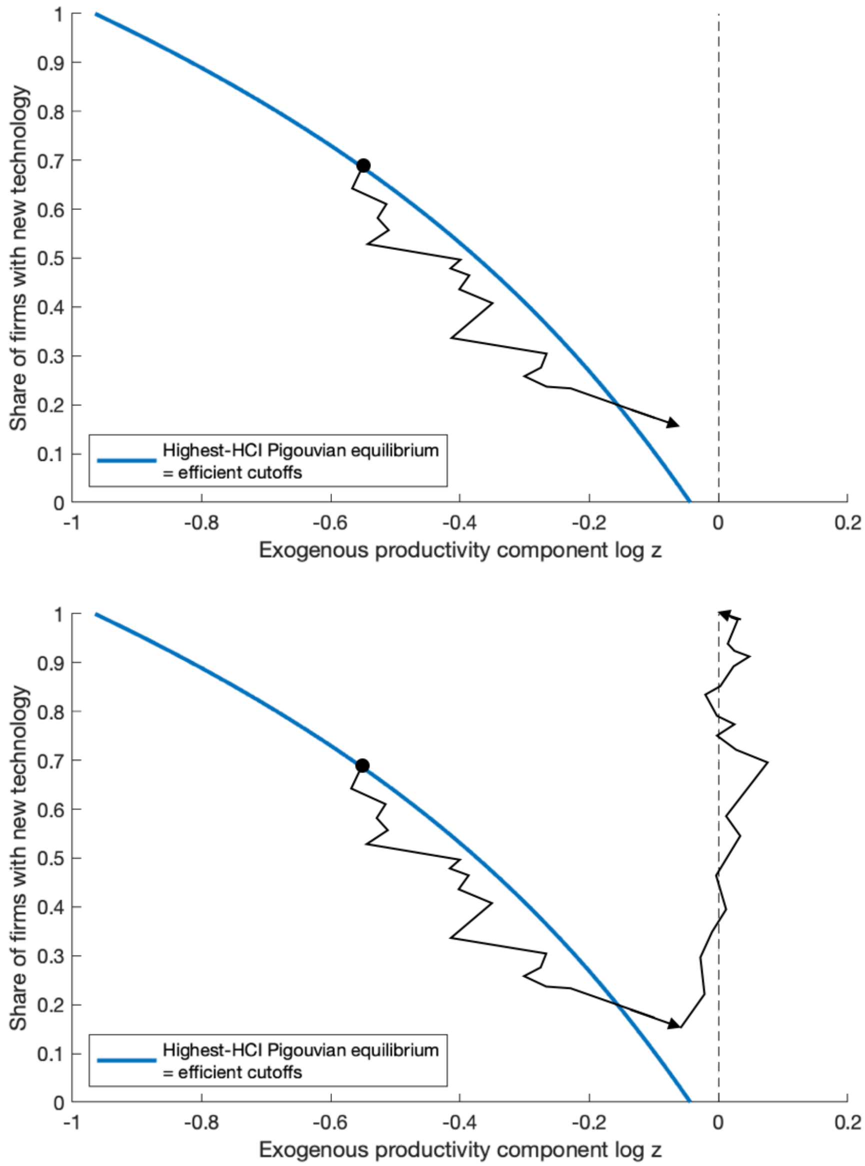

Figure A3: Interpretation of Figure 4 as a stochastic phase diagram, showing one possible path productivity and the share of firms using the new technology. Upper panel: Starting from an initial condition at which productivity is below the adoption, the share of firms using the new technology deterministically drifts downwards while productivity noisily drifts towards its mean. Eventually, the path may cross the cutoffs. Lower panel: After productivity crosses the cutoff level, the number of firms using the new technology deterministically drifts upwards while productivity continues to noisily drift toward its mean.

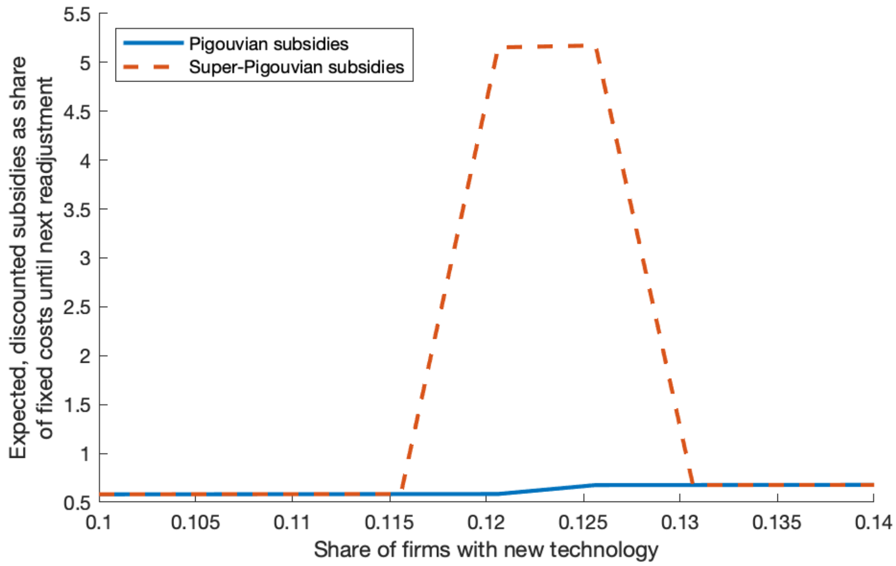

Figure A4: Subsidies to new-technology firms under Pigouvian and super-Pigouvian policy, at log productivity 0 and various initial shares of new technology adoption, in the lowest-HCI Pigouvian equilibrium. Y axis is expected discounted subsidies until next readjustment, as a share of expected discounted fixed costs until next readjustment. See Figure 6 for a more “zoomed-in” view.

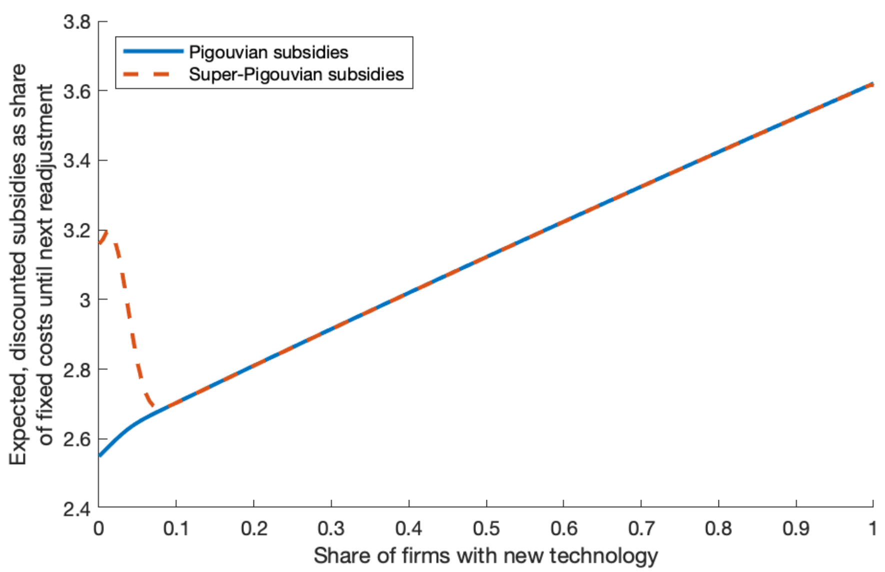

Figure A5: Subsidies to new-technology firms under Pigouvian and super-Pigouvian policy, at log productivity 0.69 and various initial shares of new technology adoption, in the lowest-HCI Pigouvian equilibrium. Y axis is expected discounted subsidies until next readjustment, as a share of expected discounted fixed costs until next readjustment.