

An Approximate General Theory of Second Best*

David R. Baqaee
UCLA

Ariel Burstein
UCLA

Joao Guerreiro
UCLA
FRB Minneapolis

May 31, 2026

Abstract

We analyze second-best optimal policy in heterogeneous-agent economies where pre-existing distortions are present but small. We study which policy maximizes aggregate-productivity, as defined in Baqaee and Burstein (2025a). We show that optimal policy can be approximated by minimizing the sum of deadweight-loss triangles. As in Woodford (2003), this reduces the policy problem to a linear-quadratic optimization problem, whose solution provides a first-order approximation to the optimal policy. We fully characterize these deadweight-loss triangles in terms of initial expenditures, wedges, microeconomic price elasticities, and the elasticity of wedges to policy instruments. We also consider cases where compensating transfers from winners to losers are themselves distortionary. We apply the framework to several analytical examples, including optimal insurance in a fiscal union and optimal monetary policy in a two-agent economy.

*We thank Elisa Rubbo for comments.

The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1 Introduction

The general theory of second best, Lipsey and Lancaster (1956), delivers one of the central insights of welfare economics: once an economy is already distorted, policies that appear desirable in isolation need not improve welfare, and removing one wedge need not move the economy closer to efficiency. But that insight is mostly cautionary. It tells us that seemingly good policies can be bad in the presence of pre-existing distortions, while offering relatively little guidance about how second-best policy should actually be conducted.

This paper provides some guidance. We study optimal policy in heterogeneous-agent economies using the aggregate-productivity criterion of Baqaee and Burstein (2025a). Under this criterion, a policy raises aggregate productivity if, after implementing it, one could scale down factor-augmenting productivity and still keep every household at least as well off as under the status quo. Our main result is that, as long as pre-existing distortions are small enough, the optimal policy can be approximated by choosing instruments to minimize a quadratic measure of misallocation: the sum of Harberger deadweight-loss triangles.

This characterization implies that optimal policy can be computed, to a first order, from initial expenditure shares, wedges, price elasticities, and the elasticity of wedges with respect to policy instruments. We provide a closed-form first-order approximation of optimal policy for economies with arbitrary heterogeneity in technologies and preferences in terms of these familiar and, in principle, observable statistics without solving the full model globally.

Our approach builds on the linear-quadratic methodology pioneered by Woodford and others in monetary economics. In that tradition, approximate optimal policy is obtained by maximizing a second-order approximation to the utility function, around the efficient allocation (i.e., the first best), subject to first-order approximations of the equilibrium conditions. We reformulate that approach and generalize it to heterogeneous-agent economies with arbitrary neoclassical preferences, production structures, and tax-like wedges. This allows this methodology to be used in a broad range of general equilibrium settings far from the traditional monetary settings in which it is normally applied. However, as with the representative-agent LQ methodology, our results are local: they apply when pre-existing distortions are small, or equivalently when the economy is close to a point at which first best is attainable. The approximation may remain useful farther from that point, but its accuracy is then an empirical or numerical matter rather than a theorem.

We begin with a representative-agent economy. There we show that, locally, maximizing utility is equivalent to minimizing deadweight-loss triangles. These triangles admit a simple characterization in terms of observed expenditure shares and compensated price elasticities. This gives the quadratic objective a clear economic interpretation: it is the second-order approximation to the efficiency cost of distortions. Unlike a second order approximation of utility, the objective function we define is invariant to monotone transformations of utility, and its magnitude is interpretable. The value of the objective function is approximately equal to the consumption-equivalent variation associated with different policies. Furthermore, we show that conditional on matching the same expenditure shares and price elasticities, non-homotheticities are irrelevant for both the objective function (to a second-order) and for the optimal policy (to a first-order).

This quadratic objective also gives a simple intuition for optimal second-best policy, with a tax-smoothing flavor in the spirit of Barro (1979). If the policymaker can affect the quantity of one good by changing the wedge on another, then it may be optimal to introduce or increase a distortion in the second market in order to reduce the deadweight loss in the first. The intuition is that welfare losses are convex in the size of deadweight-loss triangles: holding fixed the overall effect on quantities, it is better to spread distortions across markets than to concentrate them in a single one. Two smaller triangles can therefore be preferable to one very large triangle and one very small one. Of course, the force of this logic depends on the size of the wedges, the relevant quantity elasticities with respect to wedges, and the size of each market, as measured by expenditure shares. For example, if quantity in the first market responds only weakly to the wedge in the second, then the incentive to offset distortions in the first market by distorting the second is correspondingly limited. Likewise, if the second market is much larger than the first, then because each triangle is weighted by market size, the incentive to raise the wedge in the second market in order to shrink the triangle in the first is also reduced.

Having recast the representative-agent problem as one of minimizing deadweight-loss triangles, we then show that this formula also extends to heterogeneous-agent economies. If there are lump-sum transfers, so that winners from a policy change can compensate the losers, then we show that optimal policy is still approximated by minimizing deadweight-loss triangles, exactly as in the representative-agent case. If there are no lump-sum transfers, then optimal policy can be approximated by minimizing deadweight-loss triangles subject to an additional linear no-winners-or-losers constraint from the policy change. With a single agent, the heterogeneous-agent analysis collapses to the representative-agent one.

Throughout the paper, we provide pen-and-paper examples to illustrate how the for-

mulas work in practice and how the interaction between policy instruments and pre-existing wedges shapes the optimal policy. Besides some basic examples, we also work through two more detailed examples. In the first, a policymaker uses monetary policy to respond to a cost-push shock, where some households own equity in firms and others do not. The second example considers a fiscal union that can use state-dependent linear labor income taxes/subsidies to provide insurance across regions that are otherwise in financial autarky. In both cases, we derive an approximation to optimal policy that maximizes the aggregate surplus and allows for Pareto improvements, with or without lump-sum transfers.

While in this paper we work with simple fully specified models that can also be solved exactly using numerical methods, the usefulness of the approximation is not limited to such environments. Our approach has several practical advantages relative to solving the full nonlinear problem.

First, the first-order problem is considerably easier to solve in quantitative environments where the exact nonlinear problem is prohibitively difficult or infeasible to solve. For example, in a companion paper, Baqaee et al. (2026), we use the methodology developed here to solve for optimal monetary policy in a quantitative heterogeneous-agent New Keynesian economy with sticky prices, sticky wages, precautionary savings, and cyclical idiosyncratic income risk. More broadly, that paper illustrates how the results in this paper, which are primarily theoretical, can be used as an approximation theory for optimal policy in quantitatively realistic economies with many wedges, constrained instruments, and heterogeneous agents.¹

Second, computing the approximation does not require specifying and solving the entire model globally. Instead, it is enough to know expenditure shares, initial wedges, the Jacobian of wedges with respect to policy instruments, and the relevant price elasticities of demand. This substantially reduces the information required to characterize optimal policy.

Finally, even when a nonlinear solution is available, the first-order approximation can still provide useful intuition about the structure of the exact solution.

Outline of Paper. Section 2 sets up and loglinearizes a flexible general equilibrium environment. Section 3 studies optimal policy in this setting with a representative-agent. Section 4 extends the analysis to heterogeneous agents with lump-sum redistribution. Section 5 considers the case with heterogeneous agents but without lump-sum transfers.

¹In that paper, we also show formally that, in the context of a standard New Keynesian model, our approach collapses to the one in Woodford (2003) when there is a single agent.

We end with an application to an insurance via a fiscal union problem in Section 6. Examples throughout illustrate the general results and show how the framework can be applied in practice.

Related Literature. This paper is at the intersection of the theory of second best, the literature on misallocation, and the linear-quadratic approach to optimal policy. Starting with Lipsey and Lancaster (1956), the classical theory of second best emphasizes that when some distortions cannot be eliminated, moving one margin toward its first-best level need not raise welfare. Our paper complements that largely cautionary insight with a constructive local characterization: when pre-existing distortions are small, the optimal second-best response can be characterized to a first order by solving a linear-quadratic problem, or equivalently by minimizing a quadratic approximation to deadweight-loss subject to linearized equilibrium constraints. In that sense, the paper provides an approximate general theory of second best.

Our paper is also related to the misallocation literature. We provide a justification for using a second-order Harberger (1964) triangles approximation of misallocation as the objective function when choosing optimal policy. We do this by solving for the optimal policy that maximizes the change in aggregate productivity as defined by Baqaee and Burstein (2025a). We show that the optimal policy can be approximated by minimizing deadweight-loss triangles. When there is a representative agent, the deadweight-loss triangles objective is the same as the characterization of misallocation in Baqaee and Farhi (2020), connecting this paper to the large literature on misallocation in representative agent models, like Restuccia and Rogerson (2008), Hsieh and Klenow (2009), Bigio and La’O (2016), and Edmond et al. (2018) among others. However, our approach also extends to economies with heterogeneous agents and non-homothetic preferences.

Methodologically, the paper builds most directly on the welfare-based linear-quadratic approach developed in monetary economics by Rotemberg and Woodford (1997), Woodford and Benigno (2004) and Woodford (2003). That literature shows how to derive a policy problem with a quadratic welfare objective and linearized equilibrium conditions such that the solution to the approximate problem yields a correct local linear approximation to optimal policy. In that literature, when there is a representative agent, it is understood that the objective of the central bank is equivalent to minimizing misallocation (see, e.g., La’O and Tahbaz-Salehi, 2022; Rubbo, 2020). The closest formal reference for our main theorem is Benigno and Woodford (2012). Our contribution is to reformulate this logic in non-monetary contexts and generalize it from the representative-agent benchmark to heterogeneous-agent economies, with and without redistributive instru-

ments. At the same time, we also provide a new interpretation of the quadratic objective by showing that it coincides with the sum of Harberger deadweight-loss triangles.

Outside monetary economics, the paper is related to earlier uses of local linear-quadratic and perturbation methods in growth and public finance, especially the analysis of optimal tax-smoothing by Barro (1979), recursive linear-quadratic models studied by Hansen and Sargent (2013), and the application of the LQ method to optimal taxation in an RBC model by Benigno and Woodford (2006). It is also adjacent to the broader perturbation literature, including Schmitt-Grohé and Uribe (2004) and Kim and Kim (2003), which stresses that welfare calculations generally require second-order accuracy and that naive linearization can mis-rank policies. Relative to that literature, our emphasis is not only on approximation accuracy, but also on economic interpretation: to a first-order approximation in distortions, optimal policy minimizes the sum of deadweight-loss triangles subject to linearized feasibility and redistribution constraints.

Our approach differs from most of the literature on optimal policy in heterogeneous-agent economies, which evaluates policy using a social welfare function. Social welfare functions combine efficiency and redistributive objectives, whereas our criterion isolates only potential Pareto improvements. As a result, if the initial equilibrium is Pareto-efficient, the optimal policy is to not change anything. In contrast, the optimal policy according to a social welfare function may involve introducing distortions to redistribute. In Appendix A, we show that this distinction remains even if Social-Welfare/Pareto weights are chosen so as to try to eliminate redistributive motives as much as possible.² Thus, even with full freedom to choose Pareto weights, our approach is not equivalent to social-welfare maximization.

2 Preliminaries of the Economic Environment

In this section, we set up a fairly flexible general equilibrium economy and loglinearize its equilibrium with respect to changes in wedges. We assume quantities are differentiable throughout. The remainder of the paper characterizes optimal policy within this environment. Although the notation is general, the key objects that will matter for the policy results are expenditure shares, wedges, and elasticities of quantities with respect to wedges in the initial equilibrium.

²One way to do this is to choose Pareto weights so that the observed distribution of income maximizes the social welfare function holding prices fixed. This pins down the slope of social indifference curves at the initial allocation, but not their curvature. As a result, monotone transformations of utility generally change the second-best policy prescription because they change the curvature of social indifference curves.

2.1 Setup

Consider a general equilibrium with heterogeneous agents, arbitrary neoclassical production functions, and arbitrary distorting tax-like wedges. To analyze optimal policy, we later assume that these wedges are affected by policy instruments.

Households. Households are indexed by $h \in H$. Agent h has ordinal preferences \succeq_h over commodity vectors $c_h \in \mathbb{R}^N$, where N is the number of goods. Commodities are indexed by whatever attributes make them unique, including, e.g., time, location, and state of nature. A *consumption allocation* is a matrix $c \in \mathbb{R}^{H \times N}$ whose h th row, denoted by c_h , equals the consumption vector of agent h . Each household maximizes utility subject to a budget constraint

$$\max_{c_h} u_h(c_h) \text{ such that } \sum_i p_i c_{hi} \leq \sum_f Z w_f L_{hf} + \sum_i \omega_{hi} p_i y_i (1 - 1/\mu_i) + T_h, \quad (1)$$

where the budget constraint enforces that total expenditures not exceed total income. On the left-hand side of the budget constraint, p_i is the price of i and c_{hi} is the quantity of good i purchased by household h . On the right-hand side, households derive income from factor endowments and lump-sum transfers. Household h owns an endowment L_{hf} of factor f with wage w_f , earns a fraction ω_{hi} of wedge revenues generated by the wedge μ_i on the sales $p_i y_i$ of i , transfers are T_h , and Z is an aggregate factor-augmenting productivity shifter, which we normalize to be one in the initial equilibrium.

Remark (Labor-leisure choice). If there is labor-leisure choice, we model this by assuming that each household has a fixed endowment of time and must purchase leisure like any other good.

Producers. Producer i chooses its inputs to minimize costs

$$\min_{y_{ij}, l_{if}} \sum_j p_j y_{ij} + \sum_f w_f l_{if}, \text{ such that } y_i = F_i(\{y_{ij}\}, \{l_{if}\}), \quad (2)$$

where y_i is the quantity of output, F_i is a constant-returns production function, y_{ij} are intermediate inputs used by i produced by j , and l_{if} are primary factors used by i . The assumption that F_i has constant-returns is without loss of generality, since we can capture non-constant returns using producer-specific factors. The price of i is equal to a markup

or tax, $\mu_i > 0$, times i 's marginal cost of production

$$p_i = \mu_i mc_i. \quad (3)$$

That is, the price of i is inclusive of the wedge on i 's output.

Remark (Buyer-seller-specific productivity and wedges). Although we assume that wedges are on gross output only, and do not vary at the buyer-seller level, this is a notational convention rather than an assumption. We can recreate buyer-seller specific wedges by relabeling the producers. Specifically, we can treat firm or household i 's purchases of an input from j as a distinct good (made linearly using j 's output). A wedge on this good is then isomorphic to a buyer-seller specific wedge.

Resource constraints. The resource constraint for good $i \in N$ and factor $f \in F$ is

$$\sum_j y_{ji} + \sum_h c_{hi} \leq y_i, \quad \text{and} \quad \sum_i l_{if} \leq Z \sum_h L_{hf}, \quad (4)$$

where Z is the aggregate factor-augmenting productivity shifter. Finally, net transfers across households are equal to zero:

$$\sum_h T_h = 0. \quad (5)$$

We now define a general equilibrium with wedges as in Baqaee and Farhi (2020).

Definition 1 (Decentralized Equilibrium with Wedges). A *decentralized equilibrium with wedges* is the collection of prices and quantities such that: (1) the price of each good i equals its marginal cost times a wedge μ_i ; (2) each producer chooses quantities to minimize costs taking prices as given; (3) each household chooses consumption quantities to maximize utility taking prices, consumption taxes, and income as given; (4) net transfers across households are equal to zero; (5) all resource constraints are satisfied.

Remark (Non-Walrasian Economies and Endogeneous wedges). Although we define a Walrasian equilibrium, the presence of the wedges allow us to replicate non-Walrasian economies, as in Chari et al. (2007). In particular, the wedges μ_i could themselves be functions of other endogenous variables. This allow us to capture, for example, production inefficiencies (as in Hsieh and Klenow, 2009), nominal rigidities (as in Rubbo, 2020), variable markups (as in Baqaee et al., 2024), and financial market incompleteness (as in Baqaee and Burstein, 2025c).

2.2 Loglinearized Equilibrium

This section solves for how equilibrium quantities change as wedges change. For our purposes, it suffices to study the first-order changes in equilibrium variables around the efficient point, where wedges $\boldsymbol{\mu} = \mathbf{1}$.

Notation. For each producer $i \in N$ denote the cross-price elasticity of its demand (holding output constant) for input j with respect to the price of input k by

$$\theta_{jk}^i = \frac{\partial \log y_{ij}}{\partial \log p_k}.$$

For household $h \in H$, denote the income elasticity of demand for good i by

$$\epsilon_i^h = \frac{\partial \log c_{hi}}{\partial \log I_h}.$$

Finally, denote household h 's compensated cross-price elasticity of demand for good j with respect to the price of k by

$$\theta_{jk}^h = \frac{\partial \log c_{hj}}{\partial \log p_k} + \epsilon_j^h \frac{p_k c_{hk}}{I_h},$$

where the right-hand side uses Slutsky's equation and $\partial \log c_{hj} / \partial \log p_k$ is a Marshallian price elasticity (holding income and all other prices constant).

Using this notation, we can provide a loglinearization of every equilibrium price and quantity in response to changes in wedges.

Proposition 1 (Loglinearized Equilibrium). *The response of equilibrium variables to a perturbation in wedges $d \log \boldsymbol{\mu}$, around $\boldsymbol{\mu} = \mathbf{1}$, satisfies:*

$$d \log y_i = \sum_h \frac{c_{hi}}{y_i} d \log c_{hi} + \sum_{j \in N} \frac{y_{ji}}{y_i} d \log y_{ji}, \quad (6)$$

$$d \log y_{ij} = \sum_{k \in N} \theta_{jk}^i d \log p_k + \sum_{k \in F} \theta_{jk}^i d \log w_k + d \log y_i, \quad (7)$$

$$d \log l_{if} = \sum_{k \in N} \theta_{fk}^i d \log p_k + \sum_{k \in F} \theta_{fk}^i d \log w_k + d \log y_i, \quad (8)$$

$$d \log c_{hi} = \sum_{k \in N} \theta_{ik}^h d \log p_k + \epsilon_i^h \left[d \log I_h - \sum_{k \in N} \frac{p_k c_{hk}}{I_h} d \log p_k \right], \quad (9)$$

$$d \log p_i = d \log \mu_i + \sum_{j \in N} \frac{p_j y_{ij}}{p_i y_i} d \log p_j + \sum_{j \in F} \frac{w_j l_{ij}}{p_i y_i} d \log w_j, \quad (10)$$

$$0 = \sum_{i \in N} l_{if} d \log l_{if}, \quad (11)$$

$$dI_h = \sum_{f \in F} w_f L_{hf} d \log w_f + \sum_{i \in N} \omega_{hi} p_i y_i d \log \mu_i + dT_h, \quad (12)$$

$$\sum_{h \in H} dT_h = 0. \quad (13)$$

Given a feasible choice of transfers dT_h , and a choice of numeraire, these equations collectively pin down how every price and quantity responds in equilibrium to changes in wedges.

The loglinearization above holds the technology parameters, including the factor-augmenting productivity shifter Z , constant. In the true decentralized equilibrium of the model, these parameters may be changing, but for our purposes, we only need to loglinearize equilibrium quantities with respect to changes in the wedges.

There are a lot of equations, but they are all relatively straightforward. Equation (6) is the resource constraint for good i . Equations (7) to (9) are loglinearized demand curves for intermediates and factor services by firms and consumption goods by households. The elasticity of firms' input demand with respect to output is unity since, without loss of generality, we assume production has constant returns to scale. Equation (10) is the loglinearized supply curve for good i (which does not depend directly on quantity due to constant returns). We express the elasticity of i 's marginal cost to input j in terms of i 's expenditure share on j using Shephard's lemma. Equation (11) is the resource constraint for factor endowments. Equation (12) is household h 's budget constraint, and (13) is aggregate budget balance.

The most important fact about Proposition 1 is that computing changes in quantities requires knowledge of the initial expenditures, initial wedges, and price elasticities of demand. In the rest of the paper, we use this to characterize optimal policy in terms of these microeconomic primitives.

3 Optimal Policy for a Representative Agent

We begin with a single agent problem, where $|H| = 1$. Suppose wedges are some function $\mu(\tau, \sigma)$ of two groups of parameters τ and σ . The first group of parameters, τ , are instruments controlled by the policymaker. We assume these instruments, τ , only alter equilibrium wedges and do not alter preferences or technologies directly. The second

group of parameters, σ , are exogenous to the policymaker and potentially affect both the wedges and the preferences and technologies. Without loss of generality, we assume that σ is a scalar which indexes all relevant parameters affecting wedges that are outside of the policymaker's direct control.

The consumption allocation $c(\boldsymbol{\mu}; \sigma)$ depends on $\boldsymbol{\tau}$ and σ through both the wedges $\boldsymbol{\mu}(\boldsymbol{\tau}, \sigma)$ and through the direct effects of σ on technologies and preferences. Similarly, utility $u(c; \sigma)$ may depend directly on σ if σ affects preferences. The objective function of the policymaker can be written as:

$$W(\boldsymbol{\mu}(\boldsymbol{\tau}, \sigma); \sigma) = u(c(\boldsymbol{\mu}(\boldsymbol{\tau}, \sigma); \sigma); \sigma).$$

This objective function potentially depends on σ through two channels — through the wedges and through σ 's direct effect on preferences and technologies. On the other hand, the objective function depends on $\boldsymbol{\tau}$ only through its influence on wedges. The following example demonstrates.

Example 1 (Monopolistic competition). Consider an economy with consumption goods, indexed by i , and a leisure good. There is a representative household with preferences

$$u(\{c_i\}_i, \ell) = \left[\sum_i c_i^{1-\frac{1}{\eta}} + \ell^{1-\frac{1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where c_i is a composite consumption of good and ℓ is the quantity of leisure. There is a unit endowment of time, which can be used for labor or leisure. The budget constraint is

$$\sum_i p_i c_i = \mu_\ell w(1 - \ell) + T,$$

where $\mu_\ell = 1/\tau$ and τ is a labor income tax controlled by the policymaker and T is wedge revenues. Each good i is a CES aggregator over varieties of i :

$$y_i = \left(\int y_{i,\omega}^{\frac{\theta_i-1}{\theta_i}} d\omega \right)^{\frac{\theta_i}{\theta_i-1}},$$

where ω indexes differentiated varieties of i and θ_i is the elasticity of substitution across varieties. Each variety ω is produced linearly from labor

$$y_{i,\omega} = z_{i,\omega} l_{i,\omega},$$

where $z_{i,\omega}$ is a variety-specific productivity shifter. Each variety is sold at an optimal profit-maximizing markup $\mu_{i,\omega} = \theta_i/(\theta_i - 1)$ over marginal cost. Without loss of generality, write $\log \theta_i/(\theta_i - 1) = \log \theta/(\theta - 1) + \sigma m_i$. In this expression, when $\sigma = 0$, then $\theta_i = \theta$ for every i . As σ increase from 0, the elasticities of substitution for each i moves away from the common θ in the direction specified by m_i .

The quantity of labor used by differentiated varieties of good i , $l_i = \int l_{i,\omega} d\omega$, is

$$l_i = \frac{1}{z_i} \left[\frac{\mu_i/z_i}{\mu_\ell} \right]^{-\eta} \ell, \quad \text{where} \quad z_i = \left[\int z_{i,\omega}^{\theta_i-1} d\omega \right]^{\frac{1}{\theta_i-1}}.$$

Substituting this into the resource constraint for leisure, we can write

$$\ell = \frac{1}{1 + \sum_i \frac{1}{z_i} \left(\frac{\mu_i/z_i}{\mu_\ell} \right)^{-\eta}},$$

where the quantity of leisure is inversely related to the effective markup on leisure relative to goods. Substituting these into the utility function of the representative consumer yields the objective function of a benevolent policymaker, as a function of wedges, $\mu(\tau, \sigma)$ and σ :

$$W(\mu; \sigma) = \frac{1 + \sum_i \left[\frac{\mu_i/z_i}{\mu_\ell} \right]^{1-\eta}}{\left[1 + \sum_i \frac{1}{z_i} \left[\frac{\mu_i/z_i}{\mu_\ell} \right]^{-\eta} \right]^{\frac{\eta-1}{\eta}}}. \quad (14)$$

In this example, the objective of the policymaker W depends on σ through two channels: (1) an indirect effect through wedges: $\mu_i = \theta_i/(\theta_i - 1)$, which depends on σ , and (2) a direct effect because σ changes technology parameters: $z_i = \left[\int z_{i,\omega}^{\theta_i-1} d\omega \right]^{\frac{1}{\theta_i-1}}$.

We define optimal policy to be the policy that maximizes utility of the representative-agent.

Definition 2 (Optimal Policy). The optimal policy, for a given σ , solves

$$\max_{\tau} W(\mu(\tau, \sigma); \sigma).$$

We denote the *optimal policy* by $\tau^*(\sigma)$.

The following example demonstrates optimal policy in the previous example.

Example 2 (Optimal policy with monopolistic competition). Returning to Example 1, the optimal policy $\tau^*(\sigma)$ is the value of the labor income tax $\tau_\ell = 1/\mu_\ell$ that maximizes (14). The optimal policy sets μ_ℓ to be equal to a weighted-average of markups μ_i , where the weights depend on the share of labor employed in producing each i . That is,

$$1/\tau^*(\sigma) = \mu_\ell^* = \sum_i \frac{l_i}{\sum_{i'} l_{i'}} \mu_i = \sum_i \frac{l_i}{\sum_{i'} l_{i'}} \frac{\theta}{\theta - 1} \exp(\sigma m_i), \quad (15)$$

where $l_i / \sum_{i'} l_{i'} = \frac{\mu_i^{-\eta} z_i^{\eta-1}}{\sum_{i'} \mu_{i'}^{-\eta} z_{i'}^{\eta-1}}$. Hence, the policymaker subsidizes labor to offset the effects of monopoly power on labor supply. The optimal policy depends on σ both through the value of the wedges μ_i and the through the effect of σ on the technology parameters z_i (which controls the employment-share weights).

Because of the first welfare theorem, for any preferences and technologies, we know that setting all wedges equal to 1 is optimal. This gives the following simple benchmark.

Proposition 2 (Closing All Wedges is First-Best). *For every σ , the value of any policy, τ , is bounded above by setting all wedges equal to one:*

$$\max_{\tau} W(\boldsymbol{\mu}(\boldsymbol{\tau}, \sigma); \sigma) \leq W(\mathbf{1}; \sigma).$$

If $\tau^(\sigma)$ can implement $\boldsymbol{\mu}(\boldsymbol{\tau}^*(\sigma), \sigma) = \mathbf{1}$, then we say that $\boldsymbol{\tau}^*(\sigma)$ is a first-best policy.*

First-best policy is typically simple to characterize: set every wedge to equal one. However, in many problems of interest, first-best policy is infeasible. This forces the policymaker to confront trade-offs where moving a wedge away from its first-best value may help offset distortions caused by other wedges beyond the policymaker's control.

To make some progress on this problem, we study the optimal second-best policy when the pre-existing wedges are small. To formalize this idea, specify some status-quo allocation, which corresponds to the observed equilibrium, defined by some baseline policy choice $\boldsymbol{\tau}^0(\sigma)$ and parameter value σ .

Define the change in policy, relative to the status quo, to be $\Delta \boldsymbol{\tau}^* = \boldsymbol{\tau}^*(\sigma) - \boldsymbol{\tau}^0(\sigma)$. Suppose that there is some value of σ , denoted by $\sigma = 0$, where the policymaker is able to achieve the first-best outcome. If status quo policy coincides with first best $\boldsymbol{\tau}^*(0) = \boldsymbol{\tau}^0(0)$ at $\sigma = 0$, then, as long as everything is differentiable, we can say that for σ close to 0, the pre-existing distortions are small (i.e. order σ). Under this assumption, we can

approximate second-best policy as:

$$\underbrace{\tau^*(\sigma)}_{\text{second-best}} \approx \underbrace{\tau^0(\sigma)}_{\text{observed policy}} + \left. \frac{d\Delta\tau^*}{d\sigma} \right|_{\sigma=0} \Delta\sigma.$$

We now proceed to show how to derive the right-hand side as a function of observable expenditure shares and status quo price elasticities, wedges, and the Jacobian of wedges with respect to policy instruments. To do so, consider the following linear quadratic problem.

Definition 3 (Linear-Quadratic Problem). Consider the following quadratic problem where

$$\Delta\tau^{LQ} = \arg \max_{\Delta\tau} \Delta \log \mu^T \left. \frac{\partial^2 \log W}{\partial \log \mu^2} \right|_{(0, \tau^0(0))} \Delta \log \mu$$

where

$$\Delta \log \mu = \left. \frac{\partial \log \mu}{\partial \sigma} \right|_{(0, \tau^0(0))} \Delta\sigma + \left. \frac{\partial \log \mu}{\partial \tau} \right|_{(0, \tau^0(0))} \Delta\tau.$$

Denote the linear-quadratic (LQ) optimal policy, denoted by $\Delta\tau^{LQ}$.

The following, which builds on the linear-quadratic approach in Woodford (2003), shows that the optimal second-best policy can be approximated by solving the LQ problem instead.

Theorem 1 (LQ Policy is Approximately Optimal). *If $\tau^0(0)$ is first best, then the LQ optimal policy approximates the optimal policy to a first order away from that point in σ :*

$$\tau^*(\sigma) \approx \tau^0(\sigma) + \Delta\tau^{LQ}.$$

The error is order $(\Delta\sigma)^2$.

In words, if the assumptions hold, then the solution to the LQ problem is a first-order approximation to optimal second-best policy. The next proposition shows that solving the LQ problem is equivalent to choosing policy to minimize misallocation as measured by the sum of deadweight-loss triangles.

Proposition 3 (Approximate Optimal Policy Minimizes deadweight-loss). *Suppose that $\tau^0(0)$ is first-best. Then, the change in utility due to changes in wedges is, to a second-order approximation in σ and τ given by*

$$\Delta \log \mu^T \left. \frac{\partial^2 \log W}{\partial \log \mu^2} \right|_{(0, \tau^0(0))} \Delta \log \mu = -const \times \mathcal{H}(\Delta\sigma, \Delta\tau),$$

where the constant is the elasticity of utility with respect to total wealth, and

$$\begin{aligned}\mathcal{H}(\Delta\sigma, \Delta\tau) &= -\frac{1}{2} \sum_i \frac{p_i y_i}{\sum_{i' \in N} p_{i'} c_{i'}} \Delta \log y_i \Delta \log \mu_i, \\ \Delta \log y_i &= \left(\sum_j \frac{\partial \log y_i}{\partial \log \mu_j} \Delta \log \mu_j \right), \\ \Delta \log \mu &= \frac{\partial \log \mu}{\partial \sigma} \Big|_{(0, \tau^*(0))} \Delta \sigma + \frac{\partial \log \mu}{\partial \tau} \Big|_{(0, \tau^*(0))} \Delta \tau.\end{aligned}$$

The partial derivative $\partial \log \mu / \partial \sigma$ includes changes in wedges due to changes in the status quo policy $\tau^0(\sigma)$. The partial derivatives $\partial \log y_i / \partial \log \mu_j$ are given by Proposition 1.

Hence, utility losses from changes in wedges from first-best, $\Delta \log W$, are proportional to the sum of deadweight-loss triangles $\mathcal{H}(\Delta\sigma, \Delta\tau)$. Each deadweight-loss triangle is weighted by its sales relative to GDP, and its area is the product of each wedge $\Delta \log \mu_i$ (the height of the triangle) and the distortion in quantity $\Delta \log y$ (the base of the triangle) divided by 2. In this proposition, and throughout the rest of the propositions, the elasticities and expenditure shares in any loglinearized expression can be evaluated at the observed status-quo equilibrium (i.e. the data).

Since the marginal utility of total wealth is just some positive constant, Proposition 3 shows that if we replace the second-order approximation of the utility function in the LQ problem with $-\mathcal{H}(\Delta\sigma, \Delta\tau)$, then we still obtain an approximately optimal policy. This gives the following corollary.

Corollary 1 (Second-best Approximately Minimizes Deadweight Loss Triangles). *Suppose $\tau^0(0)$ is first best. Let*

$$\Delta\tau^{LQ} = \arg \max_{\Delta\tau} -\mathcal{H}(\Delta\sigma, \Delta\tau).$$

Then, to a first-order approximation in σ , we have

$$\tau^*(\sigma) \approx \tau^0(\sigma) + \Delta\tau^{LQ}.$$

Corollary 1 implies that the optimal second-best policy responds to pre-existing distortions by attempting to make the sum of deadweight loss triangles as small as possible. In particular, if the policymaker can affect the quantity of one good by changing a wedge on another, then it may be optimal to introduce or increase a distortion in the second market in order to move quantities, and hence reduce the deadweight loss triangle, in the first. The cross-market interactions are embedded in the general-equilibrium responses $\Delta \log y_i$, which can depend on all wedges simultaneously. This ties to the Lipsey and Lan-

caster (1956) insight that minimizing distortions in isolation (ignoring how that affects quantities in other markets through equilibrium linkages) can increase total deadweight loss.

As we will show in Section 4, the magnitude of $\mathcal{H}(\Delta\sigma, \Delta\tau)$ is interpretable — it measures the cost of misallocation in terms of foregone aggregate (factor-endowment augmenting) productivity. This is unlike $\Delta \log W$ whose magnitude is not generally interpretable (utility is only pinned down up to monotone increasing transformations). The arbitrariness of the magnitude $\Delta \log W$ is reflected in the presence of the constant term in Proposition 3, which converts misallocation as measured by $\mathcal{H}(\Delta\sigma, \Delta\tau)$ into utility losses given a specific cardinal utility function.

The following example illustrates how to apply Proposition 3 using a simple example.

Example 3 (Equivalence of LQ optimal and optimal policy). Return to Example 1 and consider $\sigma = 0$. At $\sigma = 0$, $\theta_i = \theta$ for every i and first-best policy is attainable by setting the labor income tax to offset the common markup $\tau^*(0) = (\theta - 1)/\theta$. This policy effectively taxes leisure at the same rate as every other consumption good, and hence, restores efficiency. The linear quadratic objective is to choose $\Delta\tau$ to maximize

$$-\mathcal{H}(\Delta\sigma, \Delta\tau) = -\frac{\eta}{2} \text{Var}_{(l, \ell)} [\Delta \log \mu],$$

where the variance of $\Delta \log \mu = (\Delta \log \mu_1, \dots, \Delta \log \mu_N, \Delta \log \mu_\ell)$ is computed using the share of time as the weights. This formula follows from Baqaee and Farhi (2020). Written more explicitly, the problem is

$$\max_{\Delta\tau} -\frac{\eta}{2} \left[\sum_i l_i (\Delta \log \mu_i)^2 + \ell (\Delta \log \mu_\ell)^2 - \left(\sum_{i'} l_{i'} \Delta \log \mu_{i'} + \ell \Delta \log \mu_\ell \right)^2 \right].$$

The wedges are given by

$$\Delta \log \mu_i = m_i \Delta\sigma$$

for goods, and

$$\Delta \log \mu_\ell = -\frac{\Delta\tau}{\tau^*(0)},$$

for leisure. Solving the LQ problem gives the optimal policy

$$\Delta \log \mu_\ell = \sum_i \frac{l_i}{\sum_{i'} l_{i'}} \Delta \log \mu_i,$$

which after substituting in the two constraints yields

$$\frac{\Delta\tau}{\tau^*(0)} = -\sum_i \frac{l_i}{\sum_{i'} l_{i'}} m_i \Delta\sigma.$$

Hence,


$$\frac{\Delta\tau^{LQ}}{\Delta\sigma} = -\tau^*(0) \sum_i \frac{l_i}{\sum_{i'} l_{i'}} m_i.$$

Therefore, we can write

$$\log \tau^*(\sigma) \approx \log \tau^*(0) + \frac{1}{\tau^*(0)} \frac{\Delta\tau^{LQ}}{\Delta\sigma} \sigma = -\sum_i \frac{l_i}{\sum_{i'} l_{i'}} \log \mu_i, \quad (16)$$

where we use $\tau^*(0) = (\theta - 1)/\theta$ and $\log \mu_i = \log \theta/(\theta - 1) + \sigma m_i$. Hence, the approximate optimal policy is to subsidize labor to offset the average employment-weighted log markup, matching the nonlinear solution in equation (15) to a first order in σ .

Figure 1: Exact policy from Example 2 and the first-order approximation from Example 3.



Notes: The calibration uses $\theta = 4$, $\eta = 2$, and $m = (0.45, 0.10, 0.20, 0.55)$. For each i , there are two-varieties with productivity $(1.50, 0.80)$, $(1.20, 0.95)$, $(1.35, 0.60)$, and $(1.10, 0.90)$. The grid for σ runs from 0 to 2. At $\sigma = 2$, $\tau_{\text{exact}} = 0.3755$ and $\tau_{\text{approx}} = 0.3915$, so the error is -0.0160 .

Figure 1 illustrates this with a numerical example. Even for large values of σ , where the economy is far from the efficient frontier (see Figure 1a), the first-order approximation provides a fairly good approximation of the optimal policy (see Figure 1b).

An interesting fact about the approximate optimal policy is that it does not depend on income elasticities of demand. In particular, conditional on matching expenditure shares and compensated price elasticities, non-homotheticities are irrelevant for optimal policy

up to a first-order approximation.

Proposition 4 (Approximate Irrelevance of Income Elasticities). *The value of the income elasticities ϵ_i^h are irrelevant, to a first-order approximation, for the optimal policy.*

This is a consequence of the envelope theorem. Changes in wedges around $\sigma = 0$ do not have first-order income effects, hence income elasticities are irrelevant to a first-order. We close by illustrating this irrelevance result with a simple example.

Example 4 (Irrelevance of non-homotheticities). Consider the same economy as in Example 1 but suppose that the utility function is non-homothetic CES, implicitly defined via

$$u(\{c_i\}_i, \ell) = \left[\sum_i c_i^{1-\frac{1}{\eta}} + u(\{c_i\}_i, \ell)^\xi \ell^{1-\frac{1}{\eta}} \right]^{\frac{\eta}{\eta-1}}.$$

If $\xi = 0$, we recover the preferences in Example 1. If $\xi > 0$, then preferences are non-homothetic and leisure is a luxury good relative to consumption — a richer agent spends a relatively higher share of their aggregate time endowment on leisure. We keep the rest of the model exactly as in Example 3.

For these preferences, the compensated price elasticities of demand depend on expenditure shares and the parameters η and θ_i — conditional on these, the value of ξ does not directly matter for the price elasticity of compensated demand. Hence, we can invoke Proposition 4 to deduce that the optimal policy is, to a first order, identical to equation (16) in Example 3 as long as expenditure shares, initial wedges, and substitution elasticities are calibrated to be the same.

4 Heterogeneous Agents with Lump-Sum Transfers

In this section, we extend the analysis in the previous section to allow for heterogeneous agents. When there are heterogeneous agents, it is no longer obvious what objective function the policymaker is maximizing. The most common approach in the literature is to assume that the policymaker maximizes some social welfare function. However, as is well-understood, social welfare functions mix efficiency considerations with redistributive ones. That is because social-welfare-based optimal policy with heterogeneous agents also take a stance on the optimal distribution of wealth across agents—a problem that is absent in the representative agent case. This can cause optimal policy to look quite different with heterogeneous agents, even if the underlying nature of the economic distortions is similar.

In this paper, instead of using a social welfare function, we assume that the policymaker has a pure efficiency mandate — setting aside any redistribution concerns. The efficiency mandate is motivated by a cost-benefit-type analysis that maximizes the resources left-over after the winners from a policy reform compensate the losers from that reform. Given this objective, we show that, just like the representative-agent case, optimal policy approximately minimizes the sum of deadweight-loss triangles. At the end of this section, we relate our approach to one using social welfare functions.

To define this aggregate objective, specify some status-quo allocation, which corresponds to the observed equilibrium, defined by some baseline policy choice $\tau^0(\sigma)$. The status-quo consumption allocation $c^0(\sigma)$ is the consumption allocation given primitives σ and status quo policy $\tau^0(\sigma)$. Define the change in aggregate productivity, following Baqaee and Burstein (2025a), as follows.

Definition 4 (Aggregate Productivity). The change in *aggregate productivity*, due to changing policy from $\tau^0(\sigma)$ to τ , is measured by the maximum reduction in factor-augmenting productivity Z such that, given τ and one-time lump-sum transfers, every agent can be kept at least indifferent to the status quo allocation. Formally,

$$A(\tau; \sigma) \equiv \max \left\{ \begin{array}{l} \text{there is a feasible } c \text{ given } \mu(\tau, \sigma), \text{ one-time lump-sum transfers, } \{T_h\} \\ Z \in \mathbb{R} : \text{and aggregate productivity shifter } 1/Z, \\ \text{such that } c_h \succeq_h c_h^0 \text{ for every } h \end{array} \right\}. \quad (17)$$

The scalar $A(\tau; \sigma)$ measures the distance of the status quo allocation $c^0(\sigma)$ from the consumption frontier implied by policy τ , measured in terms of wasted factor endowments. At the status quo policy, $\log A(\tau^0; \sigma) = 0$. The magnitude of $A(\tau; \sigma)$ is interpretable. For example, if $\log A(\tau; \sigma) = 0.1$, this implies that given policy τ and appropriate lump-sum transfers, it is possible to contract the productivity of all factor endowments by 10 log points and keep every household indifferent. In this sense, $A(\tau; \sigma)$ is a measure of the surplus left over, if we choose policy τ , after winners compensate the losers (in general equilibrium).³

We assume that the policymaker aims to set policy in a way that maximizes the increase in aggregate productivity. Since aggregate productivity is a measure of economic efficiency, we call the resulting optimal policy the *efficiency-mandate* optimal policy.

³This definition of aggregate productivity measures the value of policy τ relative to the status quo in total-factor-productivity equivalent terms, by scaling the productivity of every factor uniformly. Below we discuss an alternative definition based on biased factor-productivity shifters. We show in Appendix B that the approximate optimal policy is invariant to the choice of which factors are augmented by Z .

Definition 5 (Efficiency-Mandate Optimal Policy). With heterogeneous agents, the policymaker maximizes aggregate productivity:

$$\max_{\tau} A(\tau; \sigma).$$

We denote the *efficiency mandate* optimal policy by $\tau^*(\sigma)$.

If τ implements first-best policy, i.e. $\mu(\tau, \sigma) = \mathbf{1}$, then $A(\tau; \sigma)$ is the distance of the status-quo allocation from the Pareto efficient frontier as defined by Debreu (1951). That is, in this case, $A(\tau; \sigma)$ measures the economic waste associated with status-quo policy τ^0 relative to first-best policy in terms of factor endowments. Of course, such a first-best policy is not available in applications we are interested in. Nevertheless, it will be useful to define the amount of misallocation in the status quo (distance from Pareto efficient frontier). We do this using the Debreu (1951) approach.

Definition 6 (Distance to Frontier). The *distance to the Pareto efficient frontier* is measured by the reduction in TFP Z such that it is technologically feasible to keep every agent at least indifferent to the status quo allocation. Formally,

$$\mathcal{D}(\sigma) \equiv \max \left\{ Z \in \mathbb{R} : \begin{array}{l} \text{there is a technologically feasible } \mathbf{c} \text{ given productivity } 1/Z \\ \text{such that } \mathbf{c}_h \succeq_h \mathbf{c}_h^0 \text{ for every } h \end{array} \right\}.$$

Relative to $A(\tau; \sigma)$, which assumes the policymaker has chosen some feasible policy τ with access to lump-sum transfers, the definition of $\mathcal{D}(\sigma)$ assumes that the policymaker can choose any technologically feasible allocation.

The aggregate productivity gain $A(\tau; \sigma)$ associated with any policy τ is always bounded above by total misallocation as measured by $\mathcal{D}(\sigma)$. Furthermore, if a policy choice can eliminate all wedges, then the increase in aggregate productivity is equal to the amount of misallocation in the status quo.

Proposition 5 (Closing All Wedges is First-Best). *Aggregate productivity, given any wedges, preferences, technologies, and status quo, is always bounded above by the distance to the frontier*

$$A(\tau; \sigma) \leq \mathcal{D}(\sigma), \quad (\text{for every } \tau \text{ and } \sigma)$$

If there exists some policy, τ^{FB} , that can set all wedges equal to one, $\mu(\tau^{FB}, \sigma) = \mathbf{1}$, then we call this first-best policy. The increase in aggregate productivity caused by a first-best policy is equal to the distance to the Pareto frontier: $A(\tau^{FB}, \sigma) = \mathcal{D}(\sigma)$.

Intuitively, if a policy can set $\mu(\tau^{FB}, \sigma) = \mathbf{1}$, then by the second welfare theorem, the consumption possibility set associated with τ^{FB} is the Pareto-efficient frontier. Hence, if the policymaker can eliminate all wedges, $\mu(\tau, \sigma) = \mathbf{1}$, by picking the instrument τ appropriately, then this is optimal. Of course, τ^{FB} is typically not a feasible policy, but nevertheless, it serves as a useful benchmark.

Proposition 6. *Suppose that at $\sigma = 0$, status quo policy $\tau^0(0) = \tau^*(0)$ attains first-best. Then, to a second-order approximation, the change in aggregate productivity due to a change in policy $\Delta\tau = \tau - \tau^0(\sigma)$ is:*

$$\log A(\tau; \sigma) \approx \log \mathcal{D}(\Delta\sigma) - \mathcal{H}(\Delta\sigma, \Delta\tau),$$

where

$$\begin{aligned} \mathcal{H}(\Delta\sigma, \Delta\tau) &= -\frac{1}{2} \sum_i \frac{p_i y_i}{\sum_{i'} p_{i'} c_{i'}} \Delta \log y_i \Delta \log \mu_i, \\ \Delta \log y_i &= \sum_j \frac{\partial \log y_i}{\partial \log \mu_j} \Delta \log \mu_j, \\ \Delta \log \mu &= \left. \frac{\partial \log \mu}{\partial \sigma} \right|_{(0, \tau^*(0))} \Delta\sigma + \left. \frac{\partial \log \mu}{\partial \tau} \right|_{(0, \tau^*(0))} \Delta\tau. \end{aligned}$$

The partial derivative $\partial \log y_i / \partial \log \mu_j$ is given by Proposition 1 with transfers that satisfy:

$$\Delta T_h^{comp} = \sum_i p_i c_{hi} \Delta \log p_i - \sum_f w_f L_{hf} \Delta \log w_f - \sum_i \omega_{hi} p_i y_i \Delta \log \mu_i. \quad (18)$$

Note our convention that the term $\partial \log \mu / \partial \sigma \Delta\sigma$ incorporates changes in wedges due to changes in the status quo policy as σ changes (since the status quo policy $\tau^0(\sigma)$ changes with σ). Proposition 6 shows that $\log A(\tau; \sigma)$ is approximately the distance of the status quo allocation from the Pareto frontier $\mathcal{D}(\Delta\sigma)$ minus the deadweight-loss triangles $\mathcal{H}(\Delta\sigma, \Delta\tau)$ associated with the policy choice $\tau = \tau^0(0) + \Delta\tau$.

The changes in quantities for the deadweight-loss triangles are computed assuming compensating lump-sum transfers, given by ΔT_h^{comp} , that keep every household indifferent to the status quo as wedges change. When there is only one agent, $|H| = 1$, then (18) holds trivially and we recover the expression in Proposition 3.

Corollary 2 provides an explicit characterization of $\Delta \log y_i$ in terms of expenditure shares and price elasticities of demand in the status quo. This is done by substituting (18) into the system of equations in Proposition 1. This gives a smaller system, which we reproduce below.

Corollary 2 (Compensated Quantities). *The changes in compensated quantities, $\Delta \log y_i$, used to compute Harberger triangles solve the following system of equations:*

$$\begin{aligned}\Delta \log y_i &= \sum_h \frac{c_{hi}}{y_i} \Delta \log c_{hi} + \sum_{j \in N} \frac{y_{ji}}{y_i} \Delta \log y_{ji}, \\ \Delta \log y_{ij} &= \sum_{k \in N} \theta_{jk}^i \Delta \log p_k + \sum_{k \in F} \theta_{jk}^i \Delta \log w_k + \Delta \log y_i^{comp}, \\ \Delta \log l_{if} &= \sum_{k \in N} \theta_{fk}^i \Delta \log p_k + \sum_{k \in F} \theta_{fk}^i \Delta \log w_k + \Delta \log y_i^{comp}, \\ \Delta \log c_{hi} &= \sum_{k \in N} \theta_{ik}^h \Delta \log p_k, \\ \Delta \log p_i &= \Delta \log \mu_i + \sum_{j \in N} \frac{p_j y_{ij}}{p_i y_i} \Delta \log p_j + \sum_{j \in F} \frac{w_j l_{ij}}{p_i y_i} \Delta \log w_j, \\ 0 &= \sum_{i \in N} l_{if} d \log l_{if},\end{aligned}$$

for any choice of numeraire. Moreover, the system of equations above implies that, for every $h \in H$, we have $\sum_i p_i c_{hi} \Delta \log c_{hi} = 0$.

That is, (18) allows us to eliminate the terms involving income changes, dI_h , from the system. Just like the representative agent case, income elasticities are again irrelevant for changes in compensated quantities. Since the compensations ensure that every agent is kept on their initial indifference curve, only compensated price elasticities matter, and income elasticities disappear from the calculation. Intuitively, at $\sigma = 0$, the introduction of wedges do not reduce aggregate productivity to a first order due to the envelope theorem. Since aggregate productivity is unaffected to a first order, compensating transfers can ensure that every agent is kept indifferent in response to the introduction of the wedges. To a first order, keeping every agent indifferent implies that $\sum_i p_i c_{hi} \Delta \log c_{hi} = 0$.

An immediate consequence of Proposition 6 is the following.

Corollary 3 (Distance to the Frontier). *Suppose that at $\sigma = 0$, status quo policy $\tau^0(0) = \tau^*(0)$ attains first-best. Then, to a second-order approximation in σ , the distance to the Pareto frontier at the status quo is*

$$\log \mathcal{D}(\Delta\sigma) \approx \mathcal{H}(\Delta\sigma, 0).$$

Hence, the distance to the frontier also has a deadweight-loss triangle representation.⁴ This corollary follows from combining Proposition 6 with the fact that the change in aggregate productivity is zero if policy is equal to its status-quo value: $\log A(\tau^0(\sigma); \sigma) = 0$.

⁴See also Proposition 6 from Baqaee and Burstein (2025a).

Corollary 3 implies that the change in aggregate productivity from the change in policy, $\Delta\tau = \tau - \tau^0(\sigma)$, is approximately equal to the difference between deadweight-loss triangles associated with the new policy, $\mathcal{H}(\Delta\sigma, \Delta\tau)$, and the triangles associated with the status quo policy, $\mathcal{H}(\Delta\sigma, 0)$. That is,

$$\log A(\tau; \sigma) \approx \mathcal{H}(\Delta\sigma, 0) - \mathcal{H}(\Delta\sigma, \Delta\tau).$$

The first term does not depend on the policy choice $\Delta\tau$, and hence can be ignored when solving for the optimal policy.

The next result extends Theorem 1 to allow for heterogenous agents.

Theorem 2 (LQ Policy is Approximately Optimal). *Suppose that at $\sigma = 0$, the status quo policy $\tau^0(0) = \tau^*(0)$ attains first-best. Consider the LQ problem that minimizes deadweight-loss triangles:*

$$\Delta\tau^{LQ} \in \arg \max_{\Delta\tau} -\mathcal{H}(\Delta\sigma, \Delta\tau)$$

The LQ optimal policy associated with this problem approximates the optimal policy to a first order at $\sigma = 0$:

$$\tau^*(\sigma) \approx \tau^0(\sigma) + \Delta\tau^{LQ}$$

Since the compensated quantities in Corollary 2 do not depend on income elasticities, neither does the optimal policy, as shown by Corollary 4. This generalizes Proposition 4 to heterogeneous agents.

Corollary 4 (Approximate Irrelevance of Income Elasticities with Heterogeneous Agents). *The value of the income elasticities ϵ_i^h are irrelevant, up to a first-order approximation, for the optimal policy.*

Remark (Variations on the Definition of Aggregate Productivity). Following Debreu (1951) and Baqaee and Burstein (2025a), we define the productivity gain from policy τ via a uniform proportional contraction of all factor endowments, which treats all commodities symmetrically. Alternative definitions are also possible. For instance, one could scale only the productivity of some factors, rather than that of all factor endowments. In Appendix B, we show that this is equivalent, up to a second order, to dividing the objective function by the income share of the selected factors. Hence, such a change has no effect on the first-order properties of optimal policy, and so it does not alter our characterization of the approximately optimal policy.

The following example provides an application of Theorem 2.

Example 5 (Labor Taxes in the Presence of Monopoly). Consider a multi-agent version of Example 1. Each household, h , has preferences over consumption and leisure

$$u(\{c_{h,i}\}_i, \ell_h) = \left[\sum_{i \in N_h} c_{h,i}^{1-\frac{1}{\eta}} + \ell_h^{1-\frac{1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where $c_{h,i}$ is the composite consumption of goods i and ℓ_h is the quantity of leisure. We partition the set of consumption goods into N_h by household h . This means that each h consumes a mutually exclusive collection of goods. Household h has an h -specific efficiency units of labor, χ_h , which can be used for labor or leisure/home-production. The budget constraint is

$$\sum_i p_i c_{h,i} = \mu_\ell w (Z \chi_h - \ell_h) + D_h,$$

where Z is a factor-augmenting productivity shifter and $\mu_\ell = 1/\tau$ is a labor income tax controlled by the policymaker and D_h is the household's non-labor income from wedges and transfers.

As before, each composite good i is produced by a perfectly competitive retailer that combines varieties of ω using a CES technology:

$$y_i = \left(\int y_{i,\omega}^{\frac{\theta_i-1}{\theta_i}} d\omega \right)^{\frac{\theta_i}{\theta_i-1}},$$

Each variety ω is produced linearly from labor with productivity $z_{i,\omega}$ and sold at an optimal profit-maximizing markup $\mu_{i,\omega} = \theta_i/(\theta_i - 1)$ over marginal cost. Feasibility of transfers implies that total non-labor income must total equal wedge revenue:

$$\sum_h D_h = \sum_i \left(1 - \frac{1}{\mu_{i,\omega}} \right) \int p_{i,\omega} y_{i,\omega} d\omega + \left(1 - \frac{1}{\mu_\ell} \right) \mu_\ell w \sum_h (Z \chi_h - \ell_h).$$

As before, we parameterize $\log \frac{\theta_i}{\theta_i-1} = \log \frac{\theta}{\theta-1} + \sigma m_i$. The vector of wedges $\boldsymbol{\mu}$ consists of wedges on varieties of good i , whose value is controlled by σ , as well as wedge on labor income, whose value is directly controlled by the policymaker τ .

Aggregate productivity $A(\tau; \sigma)$ depends on the policy choice τ , as well as σ . In words, $A(\tau; \sigma)$ is the maximum contraction of TFP given policy τ such that it is feasible to compensate every household relative to the status quo τ^0 .

The policymaker's objective is to choose $\mu_\ell = 1/\tau$ in such a way that maximizes

$$\max_{\tau} A(\tau; \sigma).$$

The deadweight-loss triangle formula, applying Proposition 6, is

$$\mathcal{H}(\Delta\sigma, \Delta\tau) = \frac{1}{2} \sum_i \frac{p_i y_i}{\sum_{i'} p_i c_i} \Delta \log y_i \Delta \log \mu_i = -\frac{\eta}{2} \mathbb{E}_\chi \left[\text{Var}_{(l, \ell)} [\log \mu | h] \right],$$

which is the χ -weighted average of the single agent formula, where χ_h is the productivity of agent h .⁵ The expected variance formula written out explicitly is:

$$-\frac{\eta}{2} \sum_{h \in H} \frac{\chi_h}{\sum \chi_{h'}} \left[\sum_{i \in N_h} \frac{l_i}{\chi_h} (\Delta \log \mu_i)^2 + \frac{\ell_h}{\chi_h} (\Delta \log \mu_\ell)^2 - \left(\sum_{i' \in N_h} \frac{l_{i'}}{\chi_h} \Delta \log \mu_{i'} + \frac{\ell_h}{\chi_h} \Delta \log \mu_\ell \right)^2 \right].$$

By Theorem 2, the optimal policy $\Delta \log \tau_\ell^* = -\Delta \log \mu_\ell$ maximizes this objective. The solution is

$$\Delta \log \mu_\ell = \frac{\sum_{h \in H} \chi_h \frac{\ell_h}{\chi_h} \left(1 - \frac{\ell_h}{\chi_h}\right) \mathbb{E}_l [\Delta \log \mu_{i'} | h]}{\sum_{h \in H} \chi_h \frac{\ell_h}{\chi_h} \left(1 - \frac{\ell_h}{\chi_h}\right)},$$

where $\mathbb{E}_l [\Delta \log \mu_i | h] = \sum_{i \in N_h} \frac{l_i}{\chi_h - \ell_h} \Delta \log \mu_i$ is the average markup paid by agent h (excluding leisure).

Hence the optimal policy sets the labor subsidy according to some average of each household's average markup (excluding leisure). Each household's average markup is weighted according to its total wealth χ_h and how intermediate its labor-leisure choice is $\ell_h/\chi_h(1 - \ell_h/\chi_h)$. The total wealth, χ_h , matters because eliminating distortions on the labor-leisure choice of a high-productivity agent is very valuable from an efficiency perspective because that agent can compensate the rest using lump-sum transfers. The product $\ell_h/\chi_h(1 - \ell_h/\chi_h)$ matters because it controls the effective elasticity of labor-leisure choice with respect to the labor subsidy. If household h consumes no leisure $\ell_h = 0$ or only leisure $\ell_h = \chi_h$, then a labor subsidy accomplishes nothing and acts like a pure transfer or tax to that agent. In the case with only one agent, we recover the optimal policy from Example 3.

We end this section, with a brief discussion of how our approach differs from the standard social-welfare-function approach to optimal policy.

⁵See also Baqaee and Burstein (2025a) who provide another derivation of this formula.

Comparison to Maximizing a Social Welfare Function An alternative to our approach is to use a social welfare function. If one wishes to focus only on efficiency considerations, leaving aside redistributive concerns, then one approach is to choose Pareto weights so that the observed distribution is optimal at the initial allocation, following the approach in Negishi (1960). However, this does not generally eliminate redistributive considerations from second-best policy. The weights pin down the slope of social indifference curves at the initial allocation, but not their curvature; as a result, the implied second-best policy can depend on arbitrary monotone transformations of individual utility. Appendix A makes this point formally and contrasts the resulting SWF-optimal policy with the efficiency-mandate policy studied here. In that appendix, we also provide a version of the linear-quadratic approach that can be applied to study optimal policy according to a social welfare function. In contrast to our approach, the optimal policy in that case typically depends on non-empirical objects like Pareto weights and cardinal utilities.

5 Heterogeneous Agents without Lump-Sum Transfers

In this section, we extend the analysis in the previous section to do away with compensating lump-sum transfers. To do this, we simply drop the lump-sum transfers from the definition of $A(\tau; \sigma)$, i.e., $T_h = 0$ for all h .

Definition 7 (Aggregate Productivity without Lump-Sum Redistribution). The change in *aggregate productivity without lump-sum transfers*, due to changing policy to τ from τ^0 , is the maximum reduction in factor-augmenting productivity Z such that, given policy τ , every agent can be kept at least indifferent to the status quo allocation. Formally,

$$A^{\text{costly}}(\tau; \sigma) \equiv \max \left\{ Z \in \mathbb{R} : \begin{array}{l} \text{there is a feasible } c \text{ given } \mu(\tau, \sigma), \\ \text{and aggregate productivity shifter } 1/Z, \\ \text{such that } c_h \succeq_h c_h^0 \text{ for every } h \end{array} \right\}. \quad (19)$$

Once we rule out lump-sum transfers, it may become much harder or even impossible for the policymaker to compensate the losers. For example, if changing policy necessarily creates some losers, and the policymaker has no way to compensate those losers, then $A^{\text{costly}}(\tau; \sigma)$ is maximized at the status quo $\tau = \tau^0(\sigma)$. In other words, in this case, the status quo is constrained Pareto-efficient since, given the constraints on policy instruments, it is not possible to engineer a Pareto improvement.

We study optimal policy that maximizes $A^{\text{costly}}(\tau; \sigma)$.

Definition 8 (Optimal Policy with Heterogeneous Agents). Optimal policy without lump-sum transfers is $\tau^*(\sigma)$ that solves the following problem:

$$\max_{\tau} A^{\text{costly}}(\tau; \sigma).$$

Without lump-sum transfers, it is no longer guaranteed that setting all wedges equal to one is the upper-bound of what policy can accomplish. Intuitively, setting all wedges equal to one causes the resulting allocation to be Pareto efficient, but without the lump-sum transfers, it may not be possible to compensate the losers.

Nevertheless, Theorem 2 can still be generalized to this setting as the following proposition shows.⁶

Theorem 3 (Approximate Optimal Policy Minimizes deadweight-loss). *Suppose there exists some $\sigma = 0$, where status-quo policy $\tau^0(0) = \tau^*(0)$ attains first-best. Consider the LQ problem that minimizes deadweight-loss triangles*

$$\Delta\tau^{LQ} \in \arg \max_{\Delta\tau} -\mathcal{H}(\Delta\sigma, \Delta\tau),$$

subject to linear constraints

$$\sum_i \frac{p_i c_{hi}}{I_h} \sum_j \frac{\partial \log c_{hi}}{\partial \log \mu_j} \frac{d \log \mu_j}{d\tau} \cdot \Delta\tau = 0 \quad (h \in H), \quad (20)$$

where $\partial \log c_{hi} / \partial \log \mu_j$ is given by Proposition 1 in the decentralized equilibrium. If the matrix defined by (20) has rank $|H| - 1$, then the LQ optimal policy approximates the optimal policy to a first order at $\sigma = 0$:

$$\tau^*(\sigma) \approx \tau^0(\sigma) + \Delta\tau^{LQ}$$

Theorem 3 is identical to Theorem 2 except that the policymaker faces an additional set of constraints (20). These additional constraints ensure that, after a change in policy and a contraction total factor productivity, every household can still be kept exactly indifferent to the status quo. The rank condition effectively requires the policymaker has access to enough instruments to ensure that losers *can* be compensated for a policy change. If this

⁶To prove Theorem 3 we impose an extra technical assumption. Let $\mu(\tau; \text{technologies, preferences})$ be the function mapping policies τ , technologies, and preferences into wedges. For Theorem 3, we assume that if there is a single representative agent with homothetic preferences, then $\mu(\tau; \text{technologies, preferences})$ is invariant to changes in the factor-augmenting productivity shifter Z . This assumption is very mild, and while we require it for the version of the proof in the appendix, it is possible that a more general version of the proof could relax this assumption.

rank condition fails, then the optimal policy still exists, but it may not coincide with the LQ problem.

Unlike the general equilibrium elasticities in the Harberger triangles $\mathcal{H}(\Delta\sigma, \Delta\tau)$ which are computed exactly as in Theorem 2 according to Proposition 4, the general equilibrium elasticities in constraint (20), $\partial \log c_{hi} / \partial \log \mu_j$, are calculated using the decentralized equilibrium responses in Proposition 1 (without any compensating lump-sum transfers).

Given the fact that $A^{\text{costly}}(\tau; \sigma)$ is the same as $A(\tau; \sigma)$ with additional constraints, the following is immediate.

Corollary 5. *Optimal aggregate productivity without lump-sum transfers is always weakly less than optimal aggregate productivity with lump-sum transfers, and both are smaller than the gains from implementing first-best:*

$$\max_{\tau} A^{\text{costly}}(\tau; \sigma) \leq \max_{\tau} A(\tau; \sigma) \leq D(\sigma).$$

Furthermore, if Theorem 3 applies, then, just as before, the optimal LQ policy does not depend on income elasticities.

Corollary 6 (Approximate Irrelevance of Income Elasticities with Heterogeneous Agents). *The value of the income elasticities ϵ_i^h are irrelevant for the optimal policy chosen by Theorem 3 if (20) holds with equality.*

Once again, the restriction on policy in (20) implies that agents are kept on their initial indifference curves in response to policy changes to a first-order approximation. This means that the policy change does not trigger any income effects. Then, for the same reasons as before, the optimal policy does not depend on income elasticities.

The next example applies Theorems 2 and 3 to characterize optimal monetary policy in a two-agent economy. In this example, agents are differentially affected by monetary policy since only one agent owns the firms and monetary policy affects the split between labor and profit income. We analyze how optimal monetary policy responds to a cost-push shock with and without lump-sum transfers.

Example 6 (Optimal Monetary Policy with Two Agents). Consider a two-period economy, where time is indexed by $t \in \{1, 2\}$, populated by two types of households: S and B . Let χ_S and χ_B denote their population shares, with $\chi_S + \chi_B = 1$. Each household has one unit of time each period; leisure is ℓ_{it} and labor supply is $1 - \ell_{it}$.

Preferences. For $i \in \{S, B\}$, preferences are

$$u_i = (1 - \beta)[\log c_{i1} + \log \ell_{i1}] + \beta[\log c_{i2} + \log \ell_{i2}],$$

Prices and interest rate. Let p_t be the nominal price of consumption in period t , w_t the nominal wage, and R the gross nominal interest rate between periods 1 and 2.

Budget Constraints. We assume complete markets. Using the opportunity-cost representation for leisure, the intertemporal budget constraint for agent $i \in \{S, B\}$ is

$$[p_1 c_{i1} + w_1 \ell_{i1}] + \frac{1}{R} [p_2 c_{i2} + w_2 \ell_{i2}] = w_1 + \frac{w_2}{R} + (\pi_1 + \frac{\pi_2}{R}) \mathbf{1}\{i = S\},$$

where the left-hand side is total discounted value of consumption expenditures, including the opportunity cost of leisure, and the right-hand side is the total wealth, adding up the net-present value of the time endowment and firm profits if the household owns firms.

Technology and market clearing. In each period, a representative firm produces linearly from aggregate labor:

$$Y_t = \chi_S(1 - \ell_{St}) + \chi_B(1 - \ell_{Bt}).$$

Goods market clearing requires total consumption equal total production:

$$\chi_S c_{St} + \chi_B c_{Bt} = Y_t.$$

Nominal rigidity and monetary policy. The nominal price p_t in each period is exogenous (e.g. it is pre-determined by *ex-ante* expectations). The wage is flexible. Profits in each period are $\pi_t = (p_t - w_t)Y_t$. There is a cash-in-advance constraint each period,

$$p_t Y_t = \tau_t,$$

and the central bank controls the nominal money supply τ_t .

Point of approximation. Consider an equilibrium where prices are equal to marginal costs (no cost-push shocks). The first-best is attainable by the central bank at this point — if \bar{p} is the steady-state price, then set $\tau_t = \bar{p}/2$. We consider optimal monetary policy as we move away from this point. To do this, index a cost-push shock in period 1 by

$\log p_1 = \log \bar{p} + \sigma m_1$, where \bar{p} is the steady-state without shocks. We assume that status quo monetary policy is passive — that is $\tau_t^0(\sigma) = \bar{\tau}$ for both t . We also assume that, at $\sigma = 0$, where $\log p_t = \log \bar{p}$, status quo policy achieves first-best. That is, $\tau^0(0) = \bar{p}/2$, which ensures employment is efficient $Y_t = 1/2$. Given this choice of policy, at $\sigma = 0$, we have $c_{St} = c_{Bt} = \ell_{St} = \ell_{Bt} = Y_t = 1/2$, which is a point on the Pareto frontier.

Walrasian implementation via wedges. The monetary equilibrium allocation can be replicated in a Walrasian economy with two period-specific goods markups, μ_1 and μ_2 , defined as the wedge between the price of consumption and its marginal cost. Revenues from μ_1 and μ_2 are rebated lump-sum to the saver (the firm owner).

The goods markup in period t is

$$\mu_t = p_t/\tau_t - 1 = \bar{p} \exp(\sigma m_t)/\tau_t - 1.$$

This comes from combining the cash-in-advance constraint with the labor-leisure condition.⁷ Because prices are rigid, monetary policy (via τ_t) can implement any desired values of goods markups μ_1 and μ_2 .

Deadweight-loss triangles. Consider a second-order approximation of $A(\tau; \sigma)$ as τ and σ are perturbed from $\sigma = 0$ and $\tau^0(\sigma)$. A second-order approximation to deadweight-loss, using Corollary 2, is

$$-\mathcal{H}(\Delta\sigma, \Delta\tau) = -\frac{1-\beta}{2} [\Delta \log \mu_1]^2 - \frac{\beta}{2} [\Delta \log \mu_2]^2.$$

The first deadweight-loss triangle is caused by the goods markup in period 1 and the second is caused by the goods markup in period 2. The changes in wedges as we perturb σ and τ are given, to a first order, by

$$\begin{aligned} \Delta \log \mu_1 &= \frac{\bar{p}}{\bar{\tau}} [m_1 \Delta\sigma - \Delta \log \tau_1], \\ \Delta \log \mu_2 &= \frac{\bar{p}}{\bar{\tau}} [-\Delta \log \tau_2] \end{aligned}$$

⁷The labor-leisure condition implies that $p_t Y_t = w_t (\chi_S \ell_{St} + \chi_B \ell_{Bt})$. The cash-in-advance constraint implies that $Y_t = \frac{\tau_t}{p_t}$. The production function implies that $Y_t = 1 - (\chi_S \ell_{St} + \chi_B \ell_{Bt})$. Combine these three equations and solve for $\mu_t = p_t/w_t$ to get this result.

Here, m_1 is the exogenous cost-push shock. As σ increases from $\sigma = 0$, the cost-push shock raises the markup in the first period. The terms τ_t summarizes the policy-induced component of the goods-markup wedge.

Optimal policy with lump-sum compensation. With lump-sum compensations available, the optimal monetary policy minimizes the deadweight-loss triangles without any other constraints. The resulting policy is

$$\Delta \log \tau_1^{LQ} = m_1 \Delta \sigma, \quad \Delta \log \tau_2^{LQ} = 0.$$

That is, the central bank sets money supply τ_t to eliminate the cost-push shock μ_t period-by-period. Theorem 2 guarantees that to a first order at $\sigma = 0$, this linear quadratic optimum $\Delta \log \tau^{LQ}$ coincides with the true optimal policy that maximizes $A(\tau; \sigma)$. In this simple example, the central bank can implement first-best.

Optimal policy without lump-sum compensation. Without lump-sum compensations, per Theorem 3, we add an additional constraint to the problem:

$$\beta \Delta \log \tau_2 + (1 - \beta) \Delta \log \tau_1 = 0.$$

This constraint effectively ensures that when monetary policy moves in period 1 that must be offset in period 2 to ensure that the saver is compensated for the change in profit income. Under this constraint, the optimal policy is

$$\Delta \log \tau_1^{LQ} = \beta m_1 \Delta \sigma, \quad \Delta \log \tau_2^{LQ} = -(1 - \beta) m_1 \Delta \sigma.$$

Relative to the case with lump-sum compensations, the optimal policy now reacts less strongly to cost-push shock in period one. The extent of this attenuation depends on the length of period 1 to period 2, as captured by β . When β is close to one, period 1 is very short compared to period 2.

Offsetting the recession in period 1 also requires inducing a small recession in period 2, to compensate the saver for reduced profit income in period 1. Intuitively, as is standard in sticky price models, profits are countercyclical. By Theorem 3, this policy is a first-order approximation to the optimal policy, and by following it, the central bank can engineer a Pareto improvement (make both households better off) without any lump-sum compensations.

6 Application to Insurance Through a Fiscal Union

We end the paper by studying a non-monetary example. In this section, the policymaker can levy distortionary labor income taxes and subsidies across regions and must run balanced budgets in every period and state. These taxes and subsidies can be used to generate Pareto-improving insurance across regions without compensating lump-sum transfers. The example in this section assumes that regions are in financial autarky for analytical simplicity.

While the simple structure in this section preserves the insurance motives for fiscal unions, it does abstract from many of the issues that are important in practice, such as nominal rigidities, non-traded goods, and trade in financial assets. These ingredients are important elements of the literature on fiscal unions, see, e.g. Farhi and Werning (2017), Kehoe and Pastorino (2017), and Auclert and Rognlie (2014).⁸ Where we differ from these papers is that we do not consider the same set of instruments, for example ruling out lump-sum transfers, and we study optimal policy according to aggregate productivity $A(\tau, \sigma)$ instead of a social welfare function.

Consider a group of regions, each indexed by $h \in H$, in financial autarky. There is one freely traded consumption good made using labor.

Preferences and environment. Households in region h have preferences

$$u(c_h, l_h) = \mathbb{E}_0 \sum_t \beta^t [(1 - \kappa) \log c_{ht}(s) + \kappa \log \ell_{ht}(s)],$$

where $c_{ht}(s)$ is the consumption and $\ell_{ht}(s)$ is the leisure process. The Frisch elasticity of labor supply in steady-state is $\kappa/(1 - \kappa)$. Agents in every region are in financial autarky and consume out of post-tax labor income. The budget constraint at date t and state s is

$$c_{ht}(s) = z_{ht}(s)(1 - \tau_{ht}(s))(1 - \ell_{ht}(s)),$$

where $z_{ht}(s)$ is the real wage (or labor productivity) of agents in region h , $\tau_{ht}(s)$ is a region-specific labor income tax/subsidy and $1 - \ell_{ht}(s)$ is the amount of time agents in h devote to labor. The endowment of time is normalized to be one in each region. We put no structure on each region's productivity process, so there may be aggregate risk (and the number of regions may be finite).

Aggregate feasibility requires that total consumption, summed across all regions, equal

⁸These ingredients can be added to the analysis at the cost of additional complexity and require the use of numerical methods. We leave this for future work.

total production:

$$\sum_{h \in H} c_{ht}(s) = \sum_{h \in H} z_{ht}(s)(1 - \ell_{ht}(s)).$$

The policymaker can set region-specific, state-dependent linear labor income taxes/subsidies, $\tau_{ht}(s)$. These taxes can provide insurance across regions, but they also distort labor supply.

Point of approximation. Let σ denote the standard deviation of productivity shocks. We approximate optimal policy around $\sigma = 0$, where there is no risk and regions have a constant income realization in every period and state $z_{ht}(s) = \bar{z}_h$ for all t and s . Thus the approximation is around the no-risk limiting economy, but the economy of interest has $\sigma > 0$, incomplete-insurance wedges, and potentially distorted steady-state allocations. To keep the linear-quadratic problem well-defined, we assume that $\kappa < 1$.⁹

Walrasian implementation via wedges. It is convenient to define a within-period composite consumption-leisure good for region h in period t by

$$x_{ht}(s) \equiv c_{ht}(s)^{1-\kappa} \ell_{ht}(s)^\kappa.$$

This is the within-period composite of consumption and leisure. In a Walrasian implementation, the household's maximizes utility $u(c_h, l_h)$ subject to an intertemporal budget constraint:

$$\sum_{t,s} q_t(s) \mu_{ht}^x(s) \left(p_t(s) c_{ht}(s) + \mu_{ht}^\ell(s) \ell_{ht}(s) \right) = I_h,$$

where the two wedges $\mu_{ht}^x(s)$, $\mu_{ht}^\ell(s)$, and total wealth, I_h , are chosen to replicate the allocations in the hand-to-mouth financial autarky equilibrium. The Arrow security price in the Walrasian implementation is $q_t(s)$.

In the appendix, we show that the hand-to-mouth allocations can be replicated by setting the wedge on leisure to be

$$\mu_{ht}^\ell(s) = (1 - \tau_{ht}(s)), \tag{21}$$

which is directly controlled by policy, and setting the wedge on composite consumption

⁹Throughout, we ignore the feasibility constraint that leisure is bounded above by the time constraint: $l_{ht}(s) \leq 1$ for every h, t , and s . This is permissible as long as, at $\sigma = 0$, the status-quo allocation is interior, since then the optimal policy is also locally interior. Setting $\kappa < 1$ guarantees this.

$\mu_{ht}^x(s)$ to be

$$\mu_{ht}^x(s) = \left[\frac{c_{ht}(s) / (\sum_{h'} c_{h't}(s))}{c_{h0} / (\sum_{h'} c_{h'0})} \right]^{-1}, \quad (22)$$

where $c_{ht}(s) = z_{ht}(1 - \tau_{ht}(s))$. Hence, $\mu_{ht}^x(s)$ depends on both endogenous policy and exogenous productivity shocks.

Intuitively, the wedge on composite consumption in a date and state is relatively high if households are consuming a relatively low share of aggregate consumption in that date or state compared to in period 0. Hence, temporarily lowering the labor tax, $\tau_{ht}(s)$, lowers the wedge on composite consumption. However, this comes at the expense of increasing the wedge on leisure.

Harberger triangles. The Harberger triangles objective function, and the distance to the Pareto efficient frontier, is

$$-\mathcal{H}(\Delta\sigma, \Delta\tau) = \frac{1}{2} \sum_t \beta^t \sum_h \chi_h \sum_s \pi(s) \left(\Delta \log \mu_{ht}^x(s) \Delta \log x_{ht} + \kappa \Delta \log \mu_{ht}^\ell(s) \Delta \log \ell_{ht} \right), \quad (23)$$

which are the deadweight-loss triangles associated with the incomplete-markets wedges μ_{ht}^x and the leisure wedges μ_{ht}^ℓ . Each triangle is weighted by its expenditure share, out of total wealth, in the Walrasian implementation. In the Walrasian implementation, total wealth is the discounted value of time, and the price of leisure is the opportunity cost of time (i.e. the region-state-specific productivity times the price of consumption in that date-state).

We now discuss how to populate the terms in the Harberger triangles, starting with $\Delta \log \mu_{ht}^x(s)$ and $\Delta \log \mu_{ht}^\ell(s)$. First, $\Delta\sigma$ perturbs productivities in every region, date, and state:

$$\Delta \log z_{ht}(s) \equiv \left. \frac{d \log z_{ht}(s)}{d\sigma} \right|_{\sigma=0} \Delta\sigma, \quad \Delta \log \bar{z}_h \equiv (1 - \beta) \sum_{t,s} \beta^t \pi(s) \Delta \log z_{ht}(s), \quad (24)$$

where $\Delta \log \bar{z}_h$ is the discounted average value of the productivity shock to h .

These income shocks together with the taxes on labor generate incomplete-markets wedges on composite goods $x_{ht}(s)$:

$$\Delta \log \mu_{ht}^x(s) = - \left[\Delta \log \frac{z_{ht}(s)}{z_{h0}} - \mathbb{E}_{\chi_{h'}} \left[\Delta \log \frac{z_{h't}}{z_{h'0}} \right] \right] - \left[\Delta \log \frac{(1 - \tau_{ht}(s))}{(1 - \tau_{h0})} \right] \quad (25)$$

where $\mathbb{E}_{\chi_{h'}}[\cdot]$ denotes a cross-sectional expenditure-share-weighted average. The wedge

on leisure is given by

$$\Delta \log \mu_{ht}^\ell(s) = -\Delta \tau_{ht}(s). \quad (26)$$

Equations (25) and (26) are simply loglinearizations of (21) and (22).

Now we discuss how quantities are determined. By applying Theorem 3, we can write

$$\Delta \log x_{ht}^{\text{comp}}(s) = (\Delta \log z_{ht}(s)/\bar{z}_h - \mathbb{E}_{\chi_{h'}} [\Delta \log z_{h't}(s)/\bar{z}_{h'}]) + (1 - \kappa)\Delta \log(1 - \tau_{ht}(s)), \quad (27)$$

and

$$\Delta \log \ell_{ht}^{\text{comp}}(s) = \Delta \log x_{ht}^{\text{comp}}(s) - (1 - \kappa)\Delta \log(1 - \tau_{ht}(s)). \quad (28)$$

Next, because lump-sum transfers are unavailable, the policy must satisfy the no-winners-no-losers constraints:

$$(1 - \beta)\chi_h \sum_{t,s} \beta^t \pi(s) \kappa \Delta \log(1 - \tau_{ht}(s)) = 0. \quad (29)$$

These constraints ensure that each region remains on its initial indifference curve to a first order in response to a change in policy. Finally, because the government runs a balanced budget in every period and state, we also must have that

$$\mathbb{E}_{\chi_h} [\Delta \log(1 - \tau_{ht}(s))] = 0, \quad (30)$$

where $\mathbb{E}_{\chi_h}[\cdot]$ is a cross-sectional average using household wealth shares χ_h .

Optimal policy. By Theorem 3, maximizing (23) subject to equations (25) to (30) yields optimal policy:

$$\Delta \log(1 - \tau_{ht}^{LQ}(s)) = - [\Delta \log [z_{ht}(s)/\bar{z}_h] - \mathbb{E}_{\chi_{h'}} [d \log z_{h't}(s)/\bar{z}_{h'}]].$$

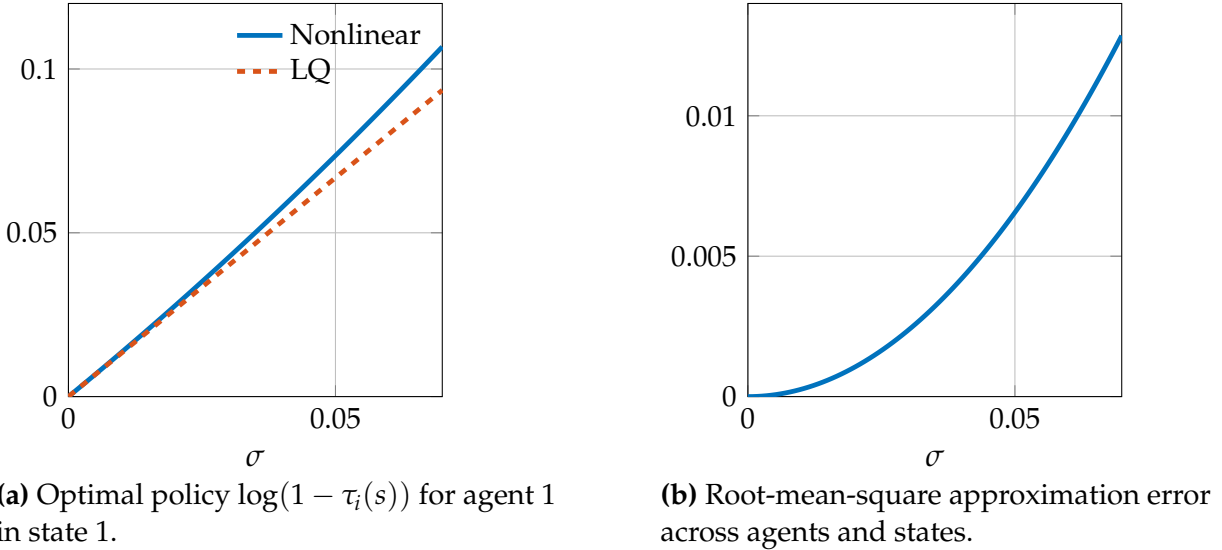
The solution therefore takes the form of an automatic fiscal stabilizer across regions: the policymaker lowers taxes on region h when its income is low relative to its discounted lifetime average *and* relative to the contemporaneous cross-sectional average across regions, and raises taxes when its income is high on both dimensions. In particular, what matters is not whether region h has a bad draw in isolation, but whether it has a bad draw relative to other regions. Thus, if all regions have low income relative to their own discounted lifetime averages, region h does not necessarily receive a subsidy: what matters for policy is its income shortfall relative to the average shortfall across regions.

Although κ , which controls the elasticity of labor supply, does not show up in the optimal policy, it does affect the objective function. Namely, substituting the optimal policy into the objective function shows that

$$\Delta \log A(\tau^*(\sigma); \sigma) \approx (1 - \kappa) \mathbb{E}_\chi [\text{Var}_{s,t} [d \log z_h(s) - \mathbb{E}_\chi [d \log z_{h'}(s)] | h]],$$

where $\text{Var}_{s,t}(\cdot | h)$ is lifetime variance, given h , in present value terms.¹⁰ Intuitively, the gain in productivity from implementing the second-best policy is proportional to the cross-sectional average of the volatility in each region's share of consumption multiplied by the labor share of time $1 - \kappa$. If households do not value labor, $\kappa = 1$, then the gain to aggregate productivity from optimal policy is zero. Productivity gains are greatest when labor is inelastically supplied, $\kappa = 0$, and the policymaker can provide perfect insurance without any labor-leisure distortions.

Figure 2: Comparison between the nonlinear policy and its approximation.



Notes: The calibration uses 50 agents, 200 states, and 100 time periods. Each agent draws an i.i.d. AR(1) productivity process with persistence parameter 0.9 and lognormal innovations with standard deviation σ . The discount factor is 0.975. The steady state Frisch elasticity of labor supply is one (so $\kappa = 1/2$).

Figure 2 provides a comparison of the nonlinear optimal policy and the approximate optimal policy. As $\sigma \rightarrow 0$, the approximate policy coincides with the optimum. In this example, σ gives the standard deviation of log productivity shocks. Hence, for a regional annual calibration, a reasonable value of σ is around 0.02, where the approximation per-

¹⁰Formally, for some process $X_{ht}(s)$, we define $\text{Var}_{s,t}(X_{ht}(s)) = \sum_{t,s} (1 - \beta) \beta^t \pi(s) (X_{ht}(s) - (1 - \beta) \sum_{t,s} \beta^t \pi(s) X_{ht}(s))^2$.

forms quite well. Of course, the performance of the approximation does deteriorate as the shocks get larger. See Appendix D for a derivation of the nonlinear optimal policy.

7 Conclusion

The general theory of second best teaches an important but mostly negative lesson: in a distorted economy, removing one wedge or correcting one margin in isolation need not improve welfare. What it does not typically provide is a constructive characterization of what optimal second-best policy should look like. This paper provides such a characterization locally. For economies where pre-existing wedges are sufficiently small, we show that optimal second-best policy can be approximated to a first order by solving a linear-quadratic problem. The resulting problem has a simple economic interpretation: optimal policy minimizes the sum of deadweight-loss triangles subject to linearized equilibrium restrictions.

We first establish this result in a representative-agent environment, where maximizing utility is locally equivalent to minimizing Harberger triangles. Those triangles can be written in terms of expenditure shares and compensated elasticities, which makes the approximate policy problem transparent and tractable. We then extend the same logic to heterogeneous-agent economies by replacing utility with an aggregate-productivity criterion measured in TFP-equivalent units. With lump-sum redistribution, this criterion captures the surplus left over after winners compensate losers. Without lump-sum redistribution, or when redistribution is costly, the policymaker faces additional linear constraints that ensure policy changes do not create first-order winners and losers. Even in that case, the same basic logic survives: to a first-order approximation, optimal policy is characterized by minimizing deadweight-loss triangles subject to linearized feasibility and redistribution constraints. A further implication of the analysis is that, both with and without redistribution, the approximate optimal policy depends on expenditure shares and substitution elasticities, but not on income elasticities.

Taken together, these results suggest that there is a useful approximate general theory of second best. Although exact optimal policy in distorted economies is often too complex to characterize directly, the local problem has a simple structure: identify the wedges that policy can affect, measure the deadweight-losses they create, and choose instruments subject to the linearized constraints imposed by equilibrium and redistribution. The pen-and-paper examples in the paper illustrate how this logic can be applied in practice and how pre-existing distortions shape the optimal response. In companion work, Baqaee et al. (2026), we apply this framework to characterize optimal monetary policy in a quan-

titative HANK model. More broadly, we hope the approach developed here proves useful for studying optimal tax, monetary, and regulatory policy in economies with many distortions, constrained instruments, and heterogeneous agents.

References

- Auclert, A. and M. Rognlie (2014). Monetary union begets fiscal union. In *Working paper*.
- Baqae, D. and A. Burstein (2025a). Aggregate productivity with heterogeneous agents. Technical report.
- Baqae, D. and A. Burstein (2025b, August). Aggregate productivity with heterogeneous agents. Working Paper 34176, National Bureau of Economic Research.
- Baqae, D., A. Burstein, and J. Guerreiro (2026). Monetary policy without redistributive concerns. Technical report.
- Baqae, D. and E. Farhi (2019). A short note on aggregating productivity. Technical report, National Bureau of Economic Research.
- Baqae, D. R. and A. Burstein (2025c, September). Efficiency costs of incomplete markets. Working Paper 34233, National Bureau of Economic Research.
- Baqae, D. R. and E. Farhi (2020). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics* 135(1), 105–163.
- Baqae, D. R., E. Farhi, and K. Sangani (2024). The darwinian returns to scale. *Review of Economic Studies* 91(3), 1373–1405.
- Barro, R. J. (1979). On the determination of the public debt. *Journal of political Economy* 87(5, Part 1), 940–971.
- Benigno, P. and M. Woodford (2006). Optimal taxation in an rbc model: A linear-quadratic approach. *Journal of Economic Dynamics and Control* 30(9-10), 1445–1489.
- Benigno, P. and M. Woodford (2012). Linear-quadratic approximation of optimal policy problems. *Journal of Economic Theory* 147(1), 1–42.
- Bigio, S. and J. La’O (2016). Financial frictions in production networks. Technical report.
- Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2007). Business cycle accounting. *Econometrica* 75(3), 781–836.
- Debreu, G. (1951). The coefficient of resource utilization. *Econometrica: Journal of the Econometric Society*, 273–292.
- Edmond, C., V. Midrigan, and D. Y. Xu (2018). How costly are markups? Technical report, National Bureau of Economic Research.
- Farhi, E. and I. Werning (2017). Fiscal unions. *American Economic Review* 107(12), 3788–3834.

- Hansen, L. P. and T. J. Sargent (2013). *Recursive models of dynamic linear economies*. Princeton University Press.
- Harberger, A. C. (1964). The measurement of waste. *The American Economic Review* 54(3), 58–76.
- Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing TFP in China and India. *The quarterly journal of economics* 124(4), 1403–1448.
- Kehoe, P. J. and E. Pastorino (2017). Fiscal unions redux. *Economic Theory* 64(4), 741–776.
- Kim, J. and S. H. Kim (2003). Spurious welfare reversals in international business cycle models. *Journal of International Economics* 60(2), 471–500.
- La’O, J. and A. Tahbaz-Salehi (2022). Optimal monetary policy in production networks. *Econometrica* 90(3), 1295–1336.
- Lipsey, R. G. and K. Lancaster (1956). The general theory of second best. *The review of economic studies* 24(1), 11–32.
- Negishi, T. (1960). Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica* 12(2-3), 92–97.
- Restuccia, D. and R. Rogerson (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics* 11(4), 707–720.
- Rotemberg, J. J. and M. Woodford (1997). An optimization-based econometric framework for the evaluation of monetary policy. *NBER macroeconomics annual* 12, 297–346.
- Rubbo, E. (2020). Networks, phillips curves and monetary policy. Technical report, mimeo, Harvard University.
- Schmitt-Grohé, S. and M. Uribe (2004). Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of economic dynamics and control* 28(4), 755–775.
- Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.
- Woodford, M. and P. Benigno (2004). *Inflation Stabilization and Welfare: The Case of a Distorted Steady State*. JSTOR.

Appendix

A	Relation to Social-Welfare-Based Policy	A1
B	Extension to Biased Productivity Variation	A4
C	Proofs	A6
D	More Details on the Fiscal-Union Example in Section 6	A20

Appendix A Relation to Social-Welfare-Based Policy

In this appendix, we compare our approach to one based on social welfare functions. First, we show how the linear-quadratic methodology can be deployed to study policy problems where the objective is a social welfare function. Then we contrast the efficiency-mandate optimal policy with optimal policy that maximizes a social welfare function with Negishi-type weights. We show, by way of an example, that these two approaches are not the same. Specifically, unlike the efficiency-mandate, the solution to maximizing an SWF with Negishi weights is not invariant to monotone transformations of utility.

A.1 Using LQ Methods to find SWF-Optimal Policy

Let $W(c; \sigma)$ be a Bergson-Samuelson social welfare function. Denote the consumption allocation given wedges μ and lump-sum transfers T by $c(\mu, T; \sigma)$. Since consumption allocations are a function of the vector of wedges, we can write social welfare as a function of the choice of wedges:

$$V(\mu; \sigma) \equiv \max_T W(c(\mu, T; \sigma); \sigma),$$

where the function may depend directly on σ because σ can affect preferences and technologies and transfers are constrained to be feasible. Define the SWF-optimal policy to be the policy that maximizes the SWF:

$$\tau^{SWF}(\sigma) \in \arg \max_{\tau} V(\mu(\tau, \sigma); \sigma).$$

Given this set up, we have the following extension of Theorem 2.

Proposition 7. *Suppose at some $\sigma = 0$, the planner can set all wedges to one $\mu(\tau^{FB}; 0) = \mathbf{1}$. Consider the LQ problem*

$$\Delta\tau^{LQ} \in \arg \max_{\Delta\tau} \frac{1}{2} \sum_{i,j} \Delta \log \mu_j \frac{\partial^2 V}{\partial \log \mu_j \partial \log \mu_i} \Delta \log \mu_i,$$

subject to linear constraints

$$\Delta \log \mu = \left. \frac{\partial \log \mu}{\partial \sigma} \right|_{(0, \tau^{SWF}(0))} \Delta \sigma + \left. \frac{\partial \log \mu}{\partial \tau} \right|_{(0, \tau^{SWF}(0))} \Delta \tau.$$

Then the LQ optimal policy approximates the optimal policy to a first order at $\sigma = 0$:

$$\left. \frac{d\tau^{SWF}}{d\sigma} \right|_{\sigma=0} = \frac{\Delta\tau^{LQ}}{\Delta\sigma}.$$

However, the second order approximation in the proposition above is not the sum of deadweight-loss triangles $\mathcal{H}(\Delta\tau, \Delta\sigma)$, and accordingly, the SWF-optimal policy generically differs from the policy that maximizes $A(\tau, \sigma)$, even to a first order:

$$\frac{d \log \tau^{SWF}}{d\sigma} \neq \frac{d \log \tau^*}{d\sigma}.$$

The example below demonstrates.

A.2 Contrast between Efficiency Mandate and Negishi-weighted SWF

One approach towards purging redistributive motives from an SWF is to consider SWFs of the form:

$$W(\mathbf{c}) = \sum_h \omega_h u_h(\mathbf{c}_h),$$

where ω_h are the Pareto weights and to calibrate the Pareto weights ω_h in such a way that the social welfare function is maximized by some initial observed decentralized allocation, given prices and aggregate income. That is, holding prices and aggregate income constant, the Pareto weights are picked to ensure that the social planner's optimal allocation of income across households coincides with the observed initial equilibrium. Pareto weights derived in this way are sometimes called the Negishi (1960) weights. We show below, by way of an example, that our approach is not the same choosing the Pareto weights in this way. Moreover, this alternative approach is not invariant to monotone transformations of utility.

To a second-order approximation, the change in social welfare due to wedges is

$$\sum_h \omega_h \sum_{i,j} \left[\frac{\partial^2 u_h}{\partial \log c_{hj} \partial \log c_{hi}} \frac{\partial \log c_{hi}}{\partial \log \mu_i} \frac{\partial \log c_{hj}}{\partial \log \mu_j} + \sum_h \omega_h \sum_i \frac{\partial u_h}{\partial \log c_{hi}} \frac{\partial^2 \log c_{hi}}{\partial \log \mu_j \partial \log \mu_i} \right] \Delta \log \mu_i \Delta \log \mu_j.$$

By inspection, the policy that maximizes this expression is not, in general, invariant to monotone transformations of the individual utility functions. Furthermore, since households are not being kept on their initial indifference curve, changes in quantities generically depend on income effects as well as substitution effects (in contrast to the representative agent models and efficiency-mandate optimal policies, where income effects are absent to a first-order).

To see the difference between the SWF optimal policy τ^{SWF} and the efficiency-mandate optimal policy τ^* , consider the following example.

Example 7. Consider the economy in Example 5 and a social welfare function

$$W(\mathbf{c}, \boldsymbol{\omega}) = \sum_h \omega_h u_h(\mathbf{c}_h, \ell_h).$$

Individual h 's utility function is

$$u(\{c_{h,i}\}_i, \ell_h) = \left[\sum_{i \in N_h} c_{h,i}^{1-\frac{1}{\eta}} + \ell_h^{1-\frac{1}{\eta}} \right]^{\frac{\rho\eta}{\eta-1}},$$

increasing in consumption and leisure. The parameter $\rho > 0$ is an arbitrary monotone transformation of the utility function, not pinned down by any primitive. Suppose we start at a first-best where all wedges are equal to one and the Pareto weights $\boldsymbol{\omega}$ are calibrated to rationalize the first-best allocation. Let χ_h denote the expenditure share of household h in this initial equilibrium.

Index the consumption tax on each $c_{h,i}$ by $\log \mu_{h,i} = \sigma \log t_{h,i}$. As σ gets larger, consumption taxes on different goods and households increase. The initially optimal policy, given the Pareto weights are equal to the Negishi weights, is to set the tax on leisure, or equivalently the subsidy on labor, $\mu_\ell = 1/\tau_\ell = 1$. We now consider the optimal labor subsidy that maximizes the SWF as σ increases from zero by $\Delta\sigma$.

Define the average wedge paid by household h to be μ_h :

$$\log \mu_h = \sum_{i \in N_h} \frac{l_i}{\chi_h} \Delta \log \mu_{h,i} + \frac{\ell_h}{\chi_h} \Delta \log \mu_\ell,$$

and the dispersion in the average wedge paid by h to be

$$\sigma_h^2 = \left[\sum_{i \in N_h} \frac{l_i}{\chi_h} (\Delta \log \mu_{h,i})^2 + \frac{\ell_h}{\chi_h} (\Delta \log \mu_\ell)^2 - (\log \mu_h)^2 \right].$$

One can show that the change in social welfare is

$$\Delta \log V \approx \text{Constant} - \rho \left[\frac{1}{1-\rho} \text{Var}_\chi(\log \mu_h) + \eta \mathbb{E}_\chi[\sigma_h^2] \right],$$

where the variance and expectation use χ_h as the weights. Clearly, despite the use of Negishi weights, the change in the SWF depends on the monotone transformation of the utility function we use (here controlled by ρ). The choice of ρ affects the penalty on the cross-sectional variance of average wedges faced by each $\log \mu_h$. It follows that the optimal policy according to this measure depends on the choice of ρ and therefore on the monotone transformation of utility we use.

By Proposition 7, the optimal tax on leisure (equivalently optimal subsidy for labor) is given by

$$\frac{\Delta \log \tau^{SWF}}{\Delta \sigma} \approx \frac{\frac{1}{1-\rho} \text{Cov}_{\chi_h} \left[\frac{\ell_h}{\chi_h}, \sum_{i \in N_h} \frac{l_i}{\chi_h} \log t_i \right] - \left[\eta \mathbb{E}_\chi \left(\sum_{i \in N_h} \left[1 - \frac{\ell_h}{\chi_h} \right] \frac{\ell_h}{\chi_h} \frac{l_i}{\chi_h - \ell_h} \log t_i \right) \right]}{\left[\frac{1}{1-\rho} + \eta \right] \left[\mathbb{E}_\chi \left[\frac{\ell_h}{\chi_h} \left[1 - \frac{\ell_h}{\chi_h} \right] \right] \right]},$$

to a first order approximation. This example shows that, although the Pareto weights are always chosen to rationalize the initial allocation for any ρ , the choice of ρ , which affects which monotone transformation of utility is being used, affects the optimal policy and that generically, $\Delta \log \tau^{SWF} \neq \Delta \log \tau^*$.

Appendix B Extension to Biased Productivity Variation

Consider an alternative definition of aggregate productivity where, instead of using a uniform TFP-equivalent variation, we measure the value of a shock in terms of equivalent productivity change for some subset of factor endowments. For concreteness, suppose that we scale the productivity of a single factor endowment, indexed by $f^* \in F$, instead of all factor endowments. That is, the resource constraints for factors are now

$$\sum_i l_{if} \leq \sum_h L_{hf}, \quad (f \neq f^*)$$

and

$$\sum_i l_{if^*} \leq Z \sum_h L_h f^*.$$

Relative to the baseline model in Section 2, the productivity shifter Z now only affects factor endowment of f^* instead of all factor endowments.

We then define an f^* -biased productivity measure as follows.

$$\tilde{A}(\tau; \sigma) \equiv \max \left\{ \begin{array}{l} \text{there is a feasible } c \text{ given } \mu(\tau, \sigma), \text{ one-time lump-sum transfers,} \\ Z \in \mathbb{R} : \text{ and aggregate productivity shifter } 1/Z, \\ \text{such that } c_h \succeq_h c_h^0 \text{ for every } h \end{array} \right\}. \quad (31)$$

The next proposition shows that the same LQ optimal policy as defined in Section 4 is also approximately optimal for a policymaker seeking to choose policy to maximize $\tilde{A}(\tau; \sigma)$.

Proposition 8. *Define optimal policy to be*

$$\tilde{\tau}(\sigma) \in \arg \max_{\tau} \tilde{A}(\tau; \sigma).$$

Suppose that at $\sigma = 0$, optimal policy can attain first-best. Consider the LQ problem that minimizes deadweight-loss triangles:

$$\Delta \tau^{LQ} \in \arg \max_{\Delta \tau} \mathcal{H}(\Delta \sigma, \Delta \tau)$$

The LQ optimal policy associated with this problem approximates both optimal policies $\tau^ \sigma$ and $\tilde{\tau}(\sigma)$ to a first order at $\sigma = 0$:*

$$\left. \frac{d\tilde{\tau}}{d\sigma} \right|_{\sigma=0} = \left. \frac{d\tau^*}{d\sigma} \right|_{\sigma=0} = \frac{\Delta \tau^{LQ}}{\Delta \sigma}.$$

This is exactly like Theorem 2 except $\tau^*(\sigma)$ has been replaced by $\tilde{\tau}(\sigma)$ and $A(\tau; \sigma)$ has been replaced by $\tilde{A}(\tau; \sigma)$.

It follows that, to a first order approximation, the optimal policy does not depend on the choice of factor endowment we use to measure TFP-equivalent variation. The same LQ policy is approximately optimal regardless of whether we use a uniform TFP-equivalent variation or a biased factor-productivity equivalent variation. So, the approximate optimal policy is invariant to the choice of which factors we use to measure productivity variation.

Appendix C Proofs

Proof of Proposition 1. Equations (6) and (11) loglinearize the resource constraints in (4). Equations (7) to (9) loglinearize the input demand and consumer demand functions, using the definitions of (compensated or constant-output) price elasticities, the fact that input demand is unit elastic with respect to output (due to constant returns to scale in production), and the definition of income elasticities. Equation (10) uses the fact that, due to constant returns to scale, the elasticity of marginal cost with respect to input prices is the same as the elasticity of the unit cost function. The elasticity of the unit cost function with respect to each input is, in turn, related to the expenditure share on that input via Shephard's lemma. Equation (12) loglinearizes the budget constraint. Equation (13) loglinearizes the budget balance condition on transfers. \square

Proof of Proposition 2. This is a straightforward consequence of the first welfare theorem. \square

Proof of Theorem 1. The optimal policy satisfies the following first order condition:

$$\frac{\partial \log W(\boldsymbol{\mu}(\sigma, \boldsymbol{\tau}^*(\sigma)); \sigma)}{\partial \log \boldsymbol{\mu}} \cdot \frac{d \log \boldsymbol{\mu}}{d \tau_i} = 0$$

for every component of $\boldsymbol{\tau}$ and every σ . Differentiate this expression with respect to σ to get

$$\begin{aligned} d\sigma \left[\frac{\partial \log \boldsymbol{\mu}}{\partial \sigma} + \frac{\partial \log \boldsymbol{\mu}}{\partial \boldsymbol{\tau}} \frac{d \boldsymbol{\tau}}{d \sigma} \right]^T & \frac{\partial^2 \log W(\boldsymbol{\mu}(\sigma, \boldsymbol{\tau}^*(\sigma)); \sigma)}{\partial \log \boldsymbol{\mu}^2} \frac{d \log \boldsymbol{\mu}}{d \tau_i} \\ & + d\sigma \frac{\partial^2 \log W(\boldsymbol{\mu}(\sigma, \boldsymbol{\tau}^*(\sigma)); \sigma)}{\partial \sigma \partial \log \boldsymbol{\mu}} \frac{d \log \boldsymbol{\mu}}{d \tau_i} \\ & + \frac{\partial \log W(\boldsymbol{\mu}(\sigma, \boldsymbol{\tau}^*(\sigma)); \sigma)}{\partial \log \boldsymbol{\mu}} \left[\frac{d^2 \log \boldsymbol{\mu}}{d \sigma d \tau_i} d\sigma + \frac{d^2 \log \boldsymbol{\mu}}{d \tau_i^2} \frac{d \boldsymbol{\tau}}{d \sigma} d\sigma \right] = 0. \end{aligned}$$

Evaluate this derivative at $(\sigma, \boldsymbol{\tau}^*(\sigma)) = (0, \boldsymbol{\tau}^*(0))$ to get

$$\left[\frac{\partial \log \boldsymbol{\mu}}{\partial \sigma} + \frac{\partial \log \boldsymbol{\mu}}{\partial \boldsymbol{\tau}} \frac{d \boldsymbol{\tau}}{d \sigma} \right]^T \frac{\partial^2 \log W(\boldsymbol{\mu}(\sigma, \boldsymbol{\tau}^*(\sigma)); \sigma)}{\partial \log \boldsymbol{\mu}^2} \frac{d \log \boldsymbol{\mu}}{d \tau_i} = 0, \quad (32)$$

which follows from the fact that at $(\sigma, \boldsymbol{\tau}^*(\sigma)) = (0, \boldsymbol{\tau}^*(0))$, the first welfare theorem implies that $\partial \log W(\boldsymbol{\mu}(\sigma, \boldsymbol{\tau}^*(\sigma)); \sigma) / \partial \log \mu_i = 0$ for every i and $\partial \log W(\mathbf{1}; \sigma) / \partial \log \mu_i = 0$ for every σ . By the inverse function theorem, inverting (32) pins down $d \boldsymbol{\tau}^* / d \sigma$.

Now consider $\Delta\tau^{LQ}$ defined by

$$\Delta\tau^{LQ} \in \arg \max_{\Delta\tau} \left[\frac{\partial \log \mu}{\partial \sigma} \Delta\sigma + \frac{\partial \log \mu}{\partial \tau} \Delta\tau \right]^T \frac{\partial^2 \log W}{\partial \log \mu^2} \Big|_{(0, \tau^*(0))} \left[\frac{\partial \log \mu}{\partial \sigma} \Delta\sigma + \frac{\partial \log \mu}{\partial \tau} \Delta\tau \right].$$

The first-order condition with respect to $\Delta\tau_i$ is

$$\frac{\partial \log \mu^T}{\partial \tau_i} \frac{\partial^2 \log W}{\partial \log \mu^2} \Big|_{(0, \tau^*(0))} \left[\frac{\partial \log \mu}{\partial \sigma} \Delta\sigma + \frac{\partial \log \mu}{\partial \tau} \Delta\tau \right] = 0.$$

Because the Hessian is symmetric, this is identical to (32), and hence $\Delta\tau^{LQ} = (d\tau^*/d\sigma)\Delta\sigma$. \square

Proof of Proposition 3. Define

$$e(\mathbf{p}, u) = \min \{ \mathbf{p} \cdot \mathbf{c} : U(\mathbf{c}) \geq u \}.$$

Take a total derivative with respect to wedges. We know that

$$\begin{aligned} d \log e &= \sum_i \frac{\partial \log e}{\partial \log p_i} d \log p_i + \frac{\partial \log e}{\partial \log u} d \log u, \\ &= \sum_i b_i d \log p_i + \frac{\partial \log e}{\partial \log u} d \log u, \end{aligned}$$

where we use Shephard's lemma and use b_i to denote the budget share on good i . Equivalently,

$$d \log W = \frac{1}{\frac{\partial \log e}{\partial \log u}} \left[d \log e - \sum_i b_i d \log p_i \right].$$

Differentiate this identity a second time to get

$$d^2 \log W = d \left[\frac{1}{\frac{\partial \log e}{\partial \log u}} \right] \left[d \log e - \sum_i b_i d \log p_i \right] + \frac{1}{\frac{\partial \log e}{\partial \log u}} d \left[d \log e - \sum_i b_i d \log p_i \right].$$

By the first welfare theorem, we know that $[d \log e - \sum_i b_i d \log p_i] = 0$ in response to a change in wedges. Hence, to a second order, we can write

$$\Delta \log W \approx \frac{1}{2} \frac{1}{\frac{\partial \log e}{\partial \log u}} \Delta \log \boldsymbol{\mu}^T \frac{d}{d \log \boldsymbol{\mu}} \left[\frac{d \log e}{d \log \boldsymbol{\mu}} \Delta \log \boldsymbol{\mu} - \sum_i b_i \frac{d \log p_i}{d \log \boldsymbol{\mu}} \Delta \log \boldsymbol{\mu} \right].$$

The proof is complete if we can show that

$$\frac{d}{d \log \boldsymbol{\mu}} \left[\frac{d \log e}{d \log \boldsymbol{\mu}} - \sum_i b_i \frac{d \log p_i}{d \log \boldsymbol{\mu}} \right] = \mathcal{H}(\Delta\sigma, \Delta\tau).$$

For the next part of the proof, we follow the steps in Baqaee and Farhi (2019). The total differentials with respect to wedges is

$$\begin{aligned} d \log e - \sum_i b_i d \log p_i &= \sum_i b_i [d \log p_i + d \log c_i] - \sum_i b_i d \log p_i \\ &= \sum_i b_i d \log c_i \\ &= \sum_i b_i \left[\frac{y_i}{c_i} d \log y_i - \frac{y_i}{c_i} \sum_j \frac{y_{ji}}{y_i} d \log y_{ji} \right] \\ &= \sum_i b_i \left[\frac{p_i y_i}{p_i c_i} d \log y_i - \frac{p_i y_i}{p_i c_i} \sum_j \frac{p_i y_{ji}}{p_i y_i} d \log y_{ji} \right] \\ &= \sum_i \frac{p_i y_i}{e} \left[d \log y_i - \sum_j \frac{p_i y_{ji}}{p_i y_i} d \log y_{ji} \right] \\ &= \frac{1}{e} \left[\sum_i p_i y_i d \log y_i - \sum_i \sum_j p_i y_{ji} d \log y_{ji} \right] \end{aligned}$$

Recall that F_i is the production function for good i . We use the fact that

$$\frac{\partial \log F_j}{\partial \log y_{ji}} = \mu_j \frac{p_i y_{ji}}{p_j y_j}$$

to write

$$\begin{aligned} d \log y_j &= \sum_j \frac{\partial \log F_j}{\partial \log y_{ji}} d \log y_{ji} + \sum_f \frac{\partial \log F_j}{\partial \log l_{jf}} d \log l_{jf}, \\ &= \sum_i \mu_j \frac{p_i y_{ji}}{p_j y_j} d \log y_{ji} + \sum_f \mu_j \frac{w_f l_{jf}}{p_j y_j} d \log l_{jf}, \\ \mu_j^{-1} \left[d \log y_j - \sum_f \mu_j \frac{w_f l_{jf}}{p_j y_j} d \log l_{jf} \right] &= \sum_i \frac{p_i y_{ji}}{p_j y_j} d \log y_{ji} \end{aligned}$$

Hence,

$$\begin{aligned}
d \log e - \sum_i b_i d \log p_i &= \frac{1}{e} \left[\sum_i p_i y_i d \log y_i - \sum_i \sum_j p_i y_{ji} d \log y_{ji} \right] \\
&= \frac{1}{e} \left[\sum_i p_i y_i d \log y_i - \sum_j p_j y_j \sum_i \frac{p_i y_{ji}}{p_j y_j} d \log y_{ji} \right] \\
&= \frac{1}{e} \left[\sum_i p_i y_i d \log y_i - \sum_j p_j y_j \mu_j^{-1} \left[d \log y_j - \sum_f \mu_j \frac{w_f l_{jf}}{p_j y_j} d \log l_{jf} \right] \right] \\
&= \frac{1}{e} \left[\sum_i p_i y_i d \log y_i - \sum_i p_i y_i \mu_i^{-1} \left[d \log y_i - \sum_f \mu_i \frac{w_f l_{if}}{p_i y_i} d \log l_{if} \right] \right] \\
&= \left[\sum_i \frac{p_i y_i}{e} \left[1 - \frac{1}{\mu_i} \right] d \log y_i + \sum_f w_f \sum_i [l_{if} d \log l_{if}] \right] \\
&= \sum_i \frac{p_i y_i}{e} \left[1 - \frac{1}{\mu_i} \right] d \log y_i.
\end{aligned}$$

Given this expression, we can now write

$$\begin{aligned}
d \left[d \log e - \sum_i b_i d \log p_i \right] &= d \left[\sum_i \frac{p_i y_i}{e} \left[1 - \frac{1}{\mu_i} \right] d \log y_i \right] \\
&= \left[\sum_i d \left[\frac{p_i y_i}{e} \right] \left[1 - \frac{1}{\mu_i} \right] d \log y_i \right] + \left[\sum_i \frac{p_i y_i}{e} d \left[1 - \frac{1}{\mu_i} \right] d \log y_i \right] \\
&\quad + \left[\sum_i \frac{p_i y_i}{e} \left[1 - \frac{1}{\mu_i} \right] d^2 \log y_i \right] \\
&= \sum_i \frac{p_i y_i}{e} d \left[1 - \frac{1}{\mu_i} \right] d \log y_i \\
&= \sum_i \frac{p_i y_i}{e} d \log \mu_i d \log y_i,
\end{aligned}$$

where we use the fact that, at the point of approximation, $\mu = 1$. Putting this altogether, we get

$$\Delta \log W \approx \frac{1}{2} \frac{1}{\frac{\partial \log e}{\partial \log u}} \sum_i \frac{p_i y_i}{e} \Delta \log \mu_i \Delta \log y_i.$$

The proof is complete if we can show that $1/(\partial \log e / \partial \log u)$ is the elasticity of utility with respect to total wealth. To see this, observe that, if v is the indirect utility function,

then $v(\mathbf{p}, e(\mathbf{p}, u)) = u$ is an identity. Log differentiate this identity with respect to u to get:

$$\frac{\partial \log v}{\partial \log e} \frac{\partial \log e}{\partial \log u} = 1.$$

From this it follows that $1/(\partial \log e / \partial \log u)$ is the elasticity of utility with respect to wealth $\partial \log v / \partial \log e$. \square

Proof of Proposition 4. From the proof of Proposition 3, we know that, at the point of approximation where $\boldsymbol{\mu} = \mathbf{1}$, we have $[d \log I - \sum_i b_i d \log p_i] = 0$, where I is the income (or expenditures) of the representative agent. Since the income elasticities only enter the equilibrium equations in Proposition 1 multiplied by $[d \log I - \sum_i b_i d \log p_i]$, this implies that the value of the income elasticities disappear from the LQ problem altogether. \square

Proof of Proposition 5. Let $\mathcal{C}(\boldsymbol{\mu}, Z)$ be the consumption possibility set given wedges $\boldsymbol{\mu}$, aggregate productivity shifter Z , and lump-sum transfers. By the second welfare theorem, $\mathcal{C}(\mathbf{1}, Z)$ is the Pareto frontier. By definition, there exists some allocation $\mathbf{c}^* \in \mathcal{C}(\boldsymbol{\mu}, 1/A(\sigma, \tau))$ such that $\mathbf{c}_h^* \succeq_h \mathbf{c}^0$ for every h . By definition of the Pareto frontier, there exists some allocation $\mathbf{c}^{**} \in \mathcal{C}(\mathbf{1}, 1/A(\sigma, \tau))$ that weakly Pareto dominates \mathbf{c}^* . Hence, $\mathcal{D}(\sigma) \geq A(\sigma, \tau)$. \square

Proof of Corollary 2. Substituting the expression for dT_h into the expression for $d \log I_h$ in Proposition 1 yields:

$$d \log I_h = \sum_i \frac{p_i c_{hi}}{I_h} \frac{d \log p_i}{d \log \boldsymbol{\mu}} d \log \boldsymbol{\mu}.$$

Substituting this into (9) and dropping unnecessary equations yields the system of equations in (2).

To see that the system of equations in Corollary 2 implies that

$$\sum_i p_i c_{hi} \frac{d \log c_{hi}}{d \log \boldsymbol{\mu}} d \log \boldsymbol{\mu} = 0,$$

observe the following. Using (9), and writing total differentials with respect to $d \log \boldsymbol{\mu}$, we have that

$$\begin{aligned} \sum_i p_i c_{hi} d \log c_{hi} &= \sum_i p_i d c_{hi} \\ &= \sum_i p_i \sum_k \frac{\partial c_{hi}}{\partial p_k} d p_k \end{aligned}$$

using symmetry of cross-partials of compensated demand

$$= \sum_i p_i \sum_k \frac{\partial c_{hk}}{\partial p_i} dp_k = \sum_k \underbrace{\left[\sum_i p_i \frac{\partial c_{hk}}{\partial p_i} \right]}_{=0} dp_k = 0.$$

The final line uses the fact that the compensated demand for each good is homogeneous of degree zero in all prices. \square

Proof of Theorem 2 and Proposition 6. The optimal policy is

$$\tau^*(\sigma) \in \arg \max_{\tau} A(\mu(\tau, \sigma); \sigma).$$

An interior optimum satisfies

$$\frac{\partial \log A}{\partial \log \mu} \cdot \frac{\partial \log \mu}{\partial \tau} = 0.$$

Suppose that an interior optimum exists in a neighborhood of $\sigma = 0$, so we can differentiate this in σ to get

$$\begin{aligned} & \left[\frac{\partial \log \mu}{\partial \sigma} + \frac{\partial \log \mu}{\partial \tau} \frac{d\tau^*}{d\sigma} \right]^T \frac{\partial^2 \log A}{\partial \log \mu^2} \cdot \frac{\partial \log \mu}{\partial \tau} \\ & + \frac{\partial \log A}{\partial \log \mu} \cdot \left[\frac{\partial^2 \log \mu}{\partial \tau^2} \frac{d\tau^*}{d\sigma} + \frac{\partial^2 \log \mu}{\partial \tau \partial \sigma} \right] + \frac{\partial^2 \log A}{\partial \log \mu \partial \sigma} \cdot \frac{\partial \log \mu}{\partial \tau} = 0. \end{aligned}$$

Since $\partial A(1; \sigma) / \partial \mu = \mathbf{0}$ for every σ in a neighborhood of $\sigma = 0$ due to the first-welfare theorem, we have that

$$\frac{\partial^2 \log A}{\partial \log \mu \partial \sigma} = 0.$$

Hence, this simplifies to

$$\left[\frac{\partial \log \mu}{\partial \sigma} + \frac{\partial \log \mu}{\partial \tau} \frac{d\tau^*}{d\sigma} \right]^T \frac{\partial^2 \log A}{\partial \log \mu^2} \cdot \frac{\partial \log \mu}{\partial \tau} = 0.$$

Note that this is exactly the first-order condition we get if we solve the problem

$$\Delta\tau^{LQ} = \arg \max_{\Delta\tau} \left[\frac{\partial \log \boldsymbol{\mu}}{\partial \sigma} \Delta\sigma + \frac{\partial \log \boldsymbol{\mu}}{\partial \tau} \Delta\tau \right]^T \frac{\partial^2 \log A}{\partial \log \boldsymbol{\mu}^2} \times \left[\frac{\partial \log \boldsymbol{\mu}}{\partial \sigma} \Delta\sigma + \frac{\partial \log \boldsymbol{\mu}}{\partial \tau} \Delta\tau \right],$$

so that

$$\frac{d\tau^*}{d\sigma} = \frac{\Delta\tau^{LQ}}{\Delta\sigma}.$$

All it remains is to characterize

$$\frac{\partial^2 \log A(\boldsymbol{\mu}; 0)}{\partial \log \boldsymbol{\mu}^2} \Big|_{\boldsymbol{\mu}=1}.$$

To get this Hessian, let $e_h(\mathbf{p}, u)$ be the expenditure function of household h . Index all equilibrium variables by wedges and TFP. So, for example, $\mathbf{p}(\boldsymbol{\mu}, -\log A)$ are prices that prevail given wedges $\boldsymbol{\mu}$ and aggregate productivity shifter $1/A$ (conditional on some lump-sum transfers).

Note that, in an interior optimum, the transfers, if they are to compensate households, must hold for every h :

$$T_h = e_h(\mathbf{p}(\boldsymbol{\mu}, -\log A), u_h^0) - \left[\sum_f \omega_{hf} w_f(\boldsymbol{\mu}, -\log A) \frac{L_f}{A} + \sum_i \omega_{hi} \left(1 - \frac{1}{\mu_i} \right) p_i(\boldsymbol{\mu}, -\log A) y_i(\boldsymbol{\mu}, -\log A) \right].$$

Let nominal GDP be the numeraire and let $\lambda_f = w_f L_f / A$, $\lambda_i = p_i y_i$, $b_{hi} = p_i c_{hi}$ and $\chi_h = \sum_i p_i c_{hi}$. Then, the transfers must satisfy:

$$T_h = e_h(\mathbf{p}(\boldsymbol{\mu}, -\log A), u_h^0) - \left[\sum_f \omega_{hf} \lambda_f(\boldsymbol{\mu}, -\log A) + \sum_i \omega_{hi} \left(1 - \frac{1}{\mu_i} \right) \lambda_i(\boldsymbol{\mu}, -\log A) \right],$$

We log differentiate this expression with respect to $\boldsymbol{\mu}$ to get:

$$\begin{aligned}
dT_h &= \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log A), u_h^0) \left[\frac{d \log p_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log A} d \log A \right] \\
&\quad - \chi_h \sum_f \omega_{hf} \lambda_f(\boldsymbol{\mu}, -\log A) \left[\frac{d \log \lambda_f}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log \lambda_f}{d \log A} d \log A \right] \\
&\quad - \chi_h \sum_i \omega_{hi} \left(1 - \frac{1}{\mu_i} \right) \lambda_i(\boldsymbol{\mu}, -\log A) \left[\frac{d \log \lambda_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log \lambda_i}{d \log A} d \log A \right] \\
&\quad - \chi_h \sum_i \omega_{hi} \frac{1}{\mu_i} d \log \mu_i \lambda_i(\boldsymbol{\mu}, -\log A).
\end{aligned}$$

Adding up and imposing $\sum_h dT_h = 0$, we get:

$$\begin{aligned}
\sum_h dT_h &= \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log A), u_h^0) \left[\frac{d \log p_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log A} d \log A \right] \\
&\quad - \sum_h \chi_h \sum_f \omega_{hf} \lambda_f(\boldsymbol{\mu}, -\log A) \left[\frac{d \log \lambda_f}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log \lambda_f}{d \log A} d \log A \right] \\
&\quad - \sum_h \chi_h \sum_i \omega_{hi} \left(1 - \frac{1}{\mu_i} \right) \lambda_i(\boldsymbol{\mu}, -\log A) \left[\frac{d \log \lambda_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log \lambda_i}{d \log A} d \log A \right] \\
&\quad - \sum_h \chi_h \sum_i \omega_{hi} \frac{1}{\mu_i} d \log \mu_i \lambda_i(\boldsymbol{\mu}, -\log A) = 0.
\end{aligned}$$

Since nominal GDP is the numeraire, we can write

$$\sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log A), u_h^0) \left[\frac{d \log p_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log A} d \log A \right] = 0.$$

Hence,

$$\begin{aligned}
&\sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log A), u_h^0) \left[\frac{d \log p_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log A} d \log A \right] \\
&= - \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log A), u_h^0) \left[\frac{d \log c_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log c_i}{d \log A} d \log A \right] = 0.
\end{aligned}$$

Hence, suppressing the derivatives with respect to $\boldsymbol{\mu}$ and remembering everything is a

differential in μ :

$$\begin{aligned} \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\mu, -\log A), u_h^0) \frac{d \log c_i}{d \log A} d \log A &= \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\mu, -\log A), u_h^0) d \log c_i \\ &= \sum_i \lambda_i(\mu, -\log A; u_h^0) \left[1 - \frac{1}{\mu_i}\right] d \log y_i. \end{aligned}$$

To get from the second to last line to the last line mimic the proof of Proposition 3.

Differentiate this a second time and evaluate at $\mu = 1$ to get

$$\underbrace{\sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(1, -\log A), u_h^0) \frac{d \log c_i}{d \log A}}_{=1} d^2 \log A = \sum_i \lambda_i d \log \mu_i d \log y_i,$$

where we use the fact that $d \log A / d \log \mu = 0$ at the point of approximation. Hence,

$$d^2 \log A = \sum_i \lambda_i d \log \mu_i d \log y_i.$$

Next, we show that

$$\log A(\mu(\tau^*(\sigma), \sigma); \sigma) \approx \mathcal{H}(\Delta\tau^*, \Delta\sigma) - \mathcal{H}(0, \Delta\sigma).$$

To see, this note that

$$\begin{aligned} \log A(\mu(\tau^*(\sigma), \sigma); \sigma) &\approx \log A(\mu(\tau^*(0), 0); 0) \\ &+ \frac{\partial \log A}{\partial \log \mu} \left[\frac{\partial \log \mu}{\partial \tau^*} \frac{d\tau^*}{d\sigma} + \frac{\partial \log \mu}{\partial \sigma} \right] \Delta\sigma + \frac{\partial \log A}{\partial \sigma} \Delta\sigma \\ &+ \frac{1}{2} \left[\frac{\partial \log \mu}{\partial \tau^*} \frac{d\tau^*}{d\sigma} + \frac{\partial \log \mu}{\partial \sigma} \right]^2 \frac{\partial^2 \log A}{\partial \log \mu^2} \left[\frac{\partial \log \mu}{\partial \tau^*} \frac{d\tau^*}{d\sigma} + \frac{\partial \log \mu}{\partial \sigma} \right] (\Delta\sigma)^2 \\ &+ \frac{1}{2} \frac{\partial^2 \log A}{\partial \sigma \partial \log \mu} \left[\frac{\partial \log \mu}{\partial \tau^*} \frac{d\tau^*}{d\sigma} + \frac{\partial \log \mu}{\partial \sigma} \right] (\Delta\sigma)^2 \\ &+ \frac{1}{2} \frac{\partial \log A}{\partial \log \mu} \frac{d}{d\sigma} \left[\frac{\partial \log \mu}{\partial \tau^*} \frac{d\tau^*}{d\sigma} + \frac{\partial \log \mu}{\partial \sigma} \right] (\Delta\sigma)^2 + \frac{1}{2} \frac{\partial^2 \log A}{\partial \sigma^2} (\Delta\sigma)^2. \end{aligned}$$

Note that, by construction,

$$\log A(\mu(\tau^0(\sigma), \sigma); \sigma) = 0$$

hence, differentiating this once with respect to σ , we get

$$\frac{\partial \log A}{\partial \log \mu} \left[\frac{\partial \log \mu}{\partial \sigma} \right] \Delta \sigma + \frac{\partial \log A}{\partial \sigma} \Delta \sigma = 0,$$

where recall that $\frac{\partial \log \mu}{\partial \sigma}$ includes changes in wedges due to changes in status quo policy. We differentiate it again with respect to σ , we get

$$\begin{aligned} & \left[\frac{\partial \log \mu}{\partial \sigma} \right]^T \frac{\partial^2 \log A}{\partial \log \mu^2} \left[\frac{\partial \log \mu}{\partial \sigma} \right] (\Delta \sigma)^2 \\ & + \frac{\partial \log A}{\partial \sigma \partial \log \mu} \left[\frac{\partial \log \mu}{\partial \sigma} \right] (\Delta \sigma)^2 + \frac{\partial \log A}{\partial \log \mu} \frac{d}{d\sigma} \left[\frac{\partial \log \mu}{\partial \sigma} \right] \Delta \sigma + \frac{\partial^2 \log A}{\partial \sigma^2} (\Delta \sigma)^2 = 0. \end{aligned}$$

At the point of approximation, we know that

$$\frac{\partial \log A}{\partial \log \mu} = 0.$$

Hence, from the equation above,

$$\frac{\partial \log A}{\partial \sigma} = 0.$$

We also know that

$$\left. \frac{\partial^2 \log A(\mu; \sigma)}{\partial \sigma \partial \log \mu} \right|_{\mu=1} = 0$$

because $dA(\mu; \sigma)/d\mu = 0$ at $\mu = 1$ for every σ . Hence, we can conclude that

$$\left[\frac{\partial \log \mu}{\partial \sigma} \right]^T \frac{\partial^2 \log A}{\partial \log \mu^2} \left[\frac{\partial \log \mu}{\partial \sigma} \right] (\Delta \sigma)^2 = -\frac{\partial^2 \log A}{\partial \sigma^2}.$$

Going back to our approximation of the objective function at the optimum, we get

$$\begin{aligned} \log A(\mu(\tau^*(\sigma), \sigma); \sigma) & \approx \frac{1}{2} \left[\frac{\partial \log \mu}{\partial \tau^*} \frac{d\tau^*}{d\sigma} + \frac{\partial \log \mu}{\partial \sigma} \right]^2 \frac{\partial^2 \log A}{\partial \log \mu^2} \left[\frac{\partial \log \mu}{\partial \tau^*} \frac{d\tau^*}{d\sigma} + \frac{\partial \log \mu}{\partial \sigma} \right] (\Delta \sigma)^2 \\ & - \frac{1}{2} \left[\frac{\partial \log \mu}{\partial \sigma} \right]^T \frac{\partial^2 \log A}{\partial \log \mu^2} \left[\frac{\partial \log \mu}{\partial \sigma} \right] (\Delta \sigma)^2 \\ & = \mathcal{H}(\Delta \sigma, \Delta \tau^*) - \mathcal{H}(\Delta \sigma, 0). \end{aligned}$$

□

Proof of Corollary 4. At the point of approximation where $\mu = 1$, we have $[d \log I_h - \sum_i b_{h,i} d \log p_i] = 0$, where I_h is the income (or expenditures) of agent h . Since the income elasticities only

enter the equilibrium equations in Proposition 1 multiplied by $[d \log I_h - \sum_i b_{h,i} d \log p_i]$, this implies that the value of the income elasticities disappear from the LQ problem altogether. \square

Proof of Theorem 3. To prove this result, we use the notion of a compensated representative agent. Consider some policy τ that, given the factor-augmenting contraction $1/A^{\text{costly}}(\tau, \sigma)$, keeps every agent indifferent to the status quo. That is, there exists a feasible allocation \mathbf{c}^* such that $u_h(\mathbf{c}_h^*) = u_h(\mathbf{c}_h^0)$ with factor productivity is $1/A^{\text{costly}}(\tau, \sigma)$. At this point, there are some equilibrium prices, denoted by \mathbf{p}^* .

Following Baqaee and Burstein (2025b), we show that this allocation \mathbf{c}^* solves a utility maximization problem for a fictitious representative agent. To show this, define the individual consumption equivalent variation function to be the scalar-valued function \tilde{u}_h that solves

$$u_h(\mathbf{c}_h / \tilde{u}_h) = u_h(\mathbf{c}_h^0).$$

Define the utility of a fictitious agent to be:

$$U(\mathbf{c}) = \min_h \{ \tilde{u}_h(\mathbf{c}_h) \}.$$

Following Lemma 1 in the Appendix of Baqaee and Burstein (2025b), given the supporting prices \mathbf{p}^* and aggregate income $\sum_h \mathbf{p}^* \cdot \mathbf{c}_h^*$, this fictitious agent picks the same allocation \mathbf{c}^* :

$$\mathbf{c}^* \in \arg \max \left[U(\mathbf{c}) : \mathbf{p}^* \cdot \sum_h \mathbf{c}_h \leq \mathbf{p}^* \cdot \sum_h \mathbf{c}_h^* \right]$$

and

$$U(\mathbf{c}^*) = 1.$$

Hence, the prices \mathbf{p}^* and quantities \mathbf{c}^* also form the equilibrium of an alternative economy with the same wedges and technologies but with a fictitious representative agent. Furthermore, the equilibrium consumption allocation of the economy with the representative agent is homogeneous of degree one in factor-productivity shifter (since $U(\mathbf{c})$ is homogeneous of degree one by construction, production has constant returns to scale, and the wedges $\boldsymbol{\mu}(\tau^*, \sigma)$ are not a function of factor-augmenting productivity shifters in the equilibrium with the compensated representative agent). Hence, we can write

$$\mathbf{c}^* = \frac{\mathbf{c}^{**}}{A^{\text{costly}}(\tau, \sigma)},$$

where \mathbf{c}^{**} is the allocation in the economy with the fictitious representative agent given

the same wedges but keeping aggregate factor-augmenting productivity at $Z = 1$. Since $U(c)$ is homogeneous of degree one this implies that

$$U(c^{**}) = A^{\text{costly}}(\tau, \sigma).$$

This shows that, given some wedges μ , we can solve for $A^{\text{costly}}(\tau, \sigma)$ by solving for the utility of the compensated representative agent given those same wedges μ , as long as the associated c^* allocation satisfies $u_h(c_h^*) = u_h(c_h^0)$ for every h . That is, we can write

$$U(c(\mu(\tau, \sigma)); \sigma) = A^{\text{costly}}(\tau, \sigma).$$

The left-hand side is just a standard representative agent problem. Hence, we can now apply Theorem 1 and Proposition 3 to this economy to get the desired result. To apply these results, we require the mild technical assumption that $U(c(\mu(\tau, \sigma)); \sigma)$ be differentiable in μ . This typically holds as long as supply of goods is not inelastic at the individual household level.

The changes in quantities due to changes in wedges in this fictional representative agent problem are the same as the compensated changes in Corollary 2. Furthermore, there is a restriction that the allocation has to satisfy $u_h(c_h^*) = u_h(c_h^0)$, for every h (otherwise it is not an interior optimum). Differentiating $u_h(c_h^*) = u_h(c_h^0)$ with respect to τ at $\sigma = 0$ yields the additional constraint that (20). \square

Proof of Proposition 8. First, mimicking the proof of Theorem 2 shows that we can solve for $d\tilde{\tau}/d\sigma$ by solving the LQ problem given $\partial^2 \log \tilde{A} / \partial \log \mu^2$. The proof is complete if we can show that $\partial^2 \log \tilde{A} / \partial \log \mu^2$ is proportional to $\partial^2 \log A / \partial \log \mu^2$ at $\sigma = 0$.

Below we show that, at $\sigma = 0$, we have

$$\frac{d^2 \log \tilde{A}}{d \log \mu^2} = \frac{1}{\lambda_{f^*}} \sum_i \lambda_i d \log \mu_i d \log y_i = \frac{1}{\lambda_{f^*}} \frac{d^2 \log A}{d \log \mu^2}.$$

In words, since $\frac{d^2 \log \tilde{A}}{d \log \mu^2}$ is proportional to $\frac{d^2 \log A}{d \log \mu^2}$, where the constant of proportionality is the reciprocal of the factor share that is being contracted $1/\lambda_{f^*}$. Given this, it follows that the optimal policy that maximizes \tilde{A} is the same as the one that maximizes A to a first order.

To see this, let $e_h(p, u)$ be the expenditure function of household h . Index all equilibrium variables by wedges and the vector of factor-augmenting productivities. So, for example, $p(\mu, -\log \tilde{A})$ are prices that prevail given wedges μ and biased-productivity shifter $1/\tilde{A}_f$ on factor f (conditional on some lump-sum transfers).

Note that, in an interior optimum, the transfers, if they are to compensate households, must hold for every h :

$$T_h = e_h(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}), u_h^0) - \left[\sum_f \omega_{hf} w_f(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) \frac{L_f}{\tilde{A}_f} + \sum_i \omega_{hi} \left(1 - \frac{1}{\mu_i}\right) p_i(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) y_i(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) \right].$$

Let nominal GDP be the numeraire and let $\lambda_f = w_f L_f / \tilde{A}_f$, $\lambda_i = p_i y_i$, $b_{hi} = p_i c_{hi}$ and $\chi_h = \sum_i p_i c_{hi}$. Then, the transfers must satisfy:

$$T_h = e_h(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}), u_h^0) - \left[\sum_f \omega_{hf} \lambda_f(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) + \sum_i \omega_{hi} \left(1 - \frac{1}{\mu_i}\right) \lambda_i(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) \right],$$

We log differentiate this expression with respect to $\boldsymbol{\mu}$ to get:

$$\begin{aligned} dT_h &= \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}), u_h^0) \left[\frac{d \log p_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log \tilde{\mathbf{A}}} \cdot d \log \tilde{\mathbf{A}} \right] \\ &\quad - \chi_h \sum_f \omega_{hf} \lambda_f(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) \left[\frac{d \log \lambda_f}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log \lambda_f}{d \log \tilde{\mathbf{A}}} \cdot d \log \tilde{\mathbf{A}} \right] \\ &\quad - \chi_h \sum_i \omega_{hi} \left(1 - \frac{1}{\mu_i}\right) \lambda_i(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) \left[\frac{d \log \lambda_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log \lambda_i}{d \log \tilde{\mathbf{A}}} \cdot d \log \tilde{\mathbf{A}} \right] \\ &\quad - \chi_h \sum_i \omega_{hi} \frac{1}{\mu_i} d \log \mu_i \lambda_i(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}). \end{aligned}$$

Adding up and imposing $\sum_h dT_h = 0$, we get:

$$\begin{aligned} \sum_h dT_h &= \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}), u_h^0) \left[\frac{d \log p_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log \tilde{\mathbf{A}}} \cdot d \log \tilde{\mathbf{A}} \right] \\ &\quad - \sum_h \chi_h \sum_f \omega_{hf} \lambda_f(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) \left[\frac{d \log \lambda_f}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log \lambda_f}{d \log \tilde{\mathbf{A}}} \cdot d \log \tilde{\mathbf{A}} \right] \\ &\quad - \sum_h \chi_h \sum_i \omega_{hi} \left(1 - \frac{1}{\mu_i}\right) \lambda_i(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) \left[\frac{d \log \lambda_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log \lambda_i}{d \log \tilde{\mathbf{A}}} \cdot d \log \tilde{\mathbf{A}} \right] \\ &\quad - \sum_h \chi_h \sum_i \omega_{hi} \frac{1}{\mu_i} d \log \mu_i \lambda_i(\boldsymbol{\mu}, -\log \tilde{\mathbf{A}}) = 0. \end{aligned}$$

Since nominal GDP is the numeraire, we can write

$$\sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{A}), u_h^0) \left[\frac{d \log p_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log \tilde{A}} \cdot d \log \tilde{A} \right] = 0.$$

Hence,

$$\begin{aligned} & \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{A}), u_h^0) \left[\frac{d \log p_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log \tilde{A}} \cdot d \log \tilde{A} \right] \\ &= - \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{A}), u_h^0) \left[\frac{d \log c_i}{d \log \boldsymbol{\mu}} \cdot d \log \boldsymbol{\mu} - \frac{d \log p_i}{d \log \tilde{A}} \cdot d \log \tilde{A} \right] = 0. \end{aligned}$$

Hence, suppressing the derivatives with respect to $\boldsymbol{\mu}$ and remembering everything is a differential in $\boldsymbol{\mu}$:

$$\begin{aligned} \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{A}), u_h^0) \frac{d \log c_i}{d \log \tilde{A}} \cdot d \log \tilde{A} &= \sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\boldsymbol{\mu}, -\log \tilde{A}), u_h^0) \cdot d \log c_i \\ &= \sum_i \lambda_i(\boldsymbol{\mu}, -\log \tilde{A}; u_h^0) \left[1 - \frac{1}{\mu_i} \right] d \log y_i. \end{aligned}$$

To get from the second to last line to the last line mimic the proof of Proposition 3.

Differentiate this a second time and evaluate at $\boldsymbol{\mu} = \mathbf{1}$ to get

$$\underbrace{\sum_h \chi_h \sum_i b_{hi}(\mathbf{p}(\mathbf{1}, -\log \tilde{A}), u_h^0) \frac{d \log c_i}{d \log \tilde{A}} \cdot d^2 \log A}_{=\lambda_{f^*}} = \sum_i \lambda_i d \log \mu_i d \log y_i,$$

where we use the fact that $d \log A / d \log \boldsymbol{\mu} = 0$ at the point of approximation, and the fact that aggregate productivity obeys Hulten's theorem when $\boldsymbol{\mu} = \mathbf{1}$ (see Baqaee and Burstein, 2025a).

Hence,

$$\frac{d^2 \log \tilde{A}}{d \log \boldsymbol{\mu}^2} = \frac{1}{\lambda_{f^*}} \sum_i \lambda_i d \log \mu_i d \log y_i = \frac{1}{\lambda_{f^*}} \frac{d^2 \log A}{d \log \boldsymbol{\mu}^2}.$$

□

Proof of Proposition 7. This can be proved by mimicking the proof of Theorem 1. □

Appendix D More Details on the Fiscal-Union Example in Section 6

D.1 Derivation of LQ Problem

D.1.1 Wedges that Decentralize the Hand-to-Mouth Allocation

We begin from the decentralized hand-to-mouth allocation in Section 6. The household's intratemporal problem implies

$$\ell_{ht}(s) = \kappa \quad \text{and} \quad c_{ht}(s) = (1 - \kappa)z_{ht}(s)(1 - \tau_{ht}(s)).$$

Hence the within-period composite good

$$x_{ht}(s) \equiv c_{ht}(s)^{1-\kappa} \ell_{ht}(s)^\kappa$$

satisfies

$$\Delta \log x_{ht}(s) = (1 - \kappa) [\Delta \log z_{ht}(s) + \Delta \log(1 - \tau_{ht}(s))]. \quad (33)$$

We now construct wedges $\mu_{ht}^x(s)$ and $\mu_{ht}^\ell(s)$ that decentralize this same allocation in a complete-markets Walrasian implementation. The wedge on leisure is pinned down by the intratemporal optimality condition. In the Walrasian implementation, the within-period first-order condition requires

$$\frac{\kappa}{1 - \kappa} \frac{c_{ht}(s)}{\ell_{ht}(s)} = \mu_{ht}^\ell(s) z_{ht}(s).$$

Substituting the hand-to-mouth allocation gives

$$\mu_{ht}^\ell(s) = 1 - \tau_{ht}(s), \quad (34)$$

which is equation (21) in the text.

To characterize the wedge on the composite good, note that in the complete-markets implementation the household chooses the date-state profile of the composite good subject to a single intertemporal budget constraint. Therefore, for a given household h , relative allocations of $x_{ht}(s)$ across date-state pairs are governed by the relative shadow prices of the composite good. Since $\ell_{ht}(s) = \kappa$ is constant across date-state pairs in the hand-to-mouth allocation, matching the relative allocation of the composite good is equivalent to matching the relative allocation of consumption. Normalizing relative to date 0, the

wedge required to implement the hand-to-mouth allocation is therefore

$$\mu_{ht}^x(s) = \left[\frac{c_{ht}(s) / \sum_{h'} c_{h't}(s)}{c_{h0} / \sum_{h'} c_{h'0}} \right]^{-1}, \quad (35)$$

which is equation (22) in the text.

Using $c_{ht}(s) = (1 - \kappa)z_{ht}(s)(1 - \tau_{ht}(s))$, equation (35) becomes

$$\mu_{ht}^x(s) = \left[\frac{z_{ht}(s)(1 - \tau_{ht}(s)) / \sum_{h'} z_{h't}(s)(1 - \tau_{h't}(s))}{z_{h0}(1 - \tau_{h0}) / \sum_{h'} z_{h'0}(1 - \tau_{h'0})} \right]^{-1}.$$

Loglinearizing around $\sigma = 0$ yields

$$\Delta \log \mu_{ht}^x(s) = - \left[\Delta \log \frac{z_{ht}(s)}{z_{h0}} - \mathbb{E}_{\chi_{h'}} \left[\Delta \log \frac{z_{h't}(s)}{z_{h'0}} \right] \right] - \left[\Delta \log \frac{1 - \tau_{ht}(s)}{1 - \tau_{h0}} \right], \quad (36)$$

which is equation (25).

D.1.2 Quantity Responses and Harberger Triangles

Applying Corollary 2, given the wedges above, gives the compensated quantity distortion:

$$\Delta \log x_{ht}^{\text{comp}}(s) = \left[\Delta \log \frac{z_{ht}(s)}{\bar{z}_h} - \mathbb{E}_{\chi_{h'}} \left[\Delta \log \frac{z_{h't}(s)}{\bar{z}_{h'}} \right] \right] + (1 - \kappa) \Delta \log(1 - \tau_{ht}(s)), \quad (37)$$

which is equation (27).

An analogous calculation for leisure yields

$$\Delta \log \ell_{ht}^{\text{comp}}(s) = - \left[\Delta \log \frac{z_{ht}(s)}{\bar{z}_h} - \mathbb{E}_{\chi_{h'}} \left[\Delta \log \frac{z_{h't}(s)}{\bar{z}_{h'}} \right] \right]. \quad (38)$$

Using (37), this can be written as

$$\Delta \log \ell_{ht}^{\text{comp}}(s) = \Delta \log x_{ht}^{\text{comp}}(s) - (1 - \kappa) \Delta \log(1 - \tau_{ht}(s)), \quad (39)$$

which is equation (28).

Equations (34), (36), (37), and (39) are the four ingredients needed to specialize the general Harberger-triangles formula to this environment.

D.1.3 Nonlinear Solution

This appendix derives the exact nonlinear optimal policy for the infinite-horizon fiscal union environment in Section 6.

Households in region h have preferences

$$(1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) [\log c_{ht}(s) + \log \ell_{ht}(s)]$$

and budget constraints

$$c_{ht}(s) = z_{ht}(s)(1 - \tau_{ht}(s))(1 - \ell_{ht}(s)).$$

Since the problem is separable across date-state pairs, the household's intratemporal choice satisfies

$$\ell_{ht}(s) = \frac{1}{2}, \quad c_{ht}(s) = \frac{z_{ht}(s)(1 - \tau_{ht}(s))}{2}.$$

Under the status quo, $\tau_{ht}(s) = 0$ for all h, t, s . We solve for the optimal policy using techniques from Baqaee and Burstein (2025a). In particular, define the consumption-equivalent variation of region h from policy τ by

$$\tilde{u}_h = \exp \left\{ (1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \log(1 - \tau_{ht}(s)) \right\}.$$

Following Lemma 1 in the appendix of Baqaee and Burstein (2025a), aggregate productivity can be written as

$$A = \min_{h \in H} \left\{ (1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \log(1 - \tau_{ht}(s)) \right\}.$$

The government chooses $\tau_{ht}(s)$ subject to balanced budgets in every date-state:

$$\sum_{h \in H} \tau_{ht}(s) z_{ht}(s) = 0 \quad \text{for every } (t, s).$$

At the optimum, the policymaker equalizes consumption-equivalent gains across regions. Let the common value be denoted by $\log(1 - \tau)$. Then the problem can be written as

$$\max \log(1 - \tau)$$

subject to

$$(1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \log(1 - \tau_{ht}(s)) = \log(1 - \tau) \quad \text{for all } h,$$

together with the date-state budget constraints.

The Lagrangian is

$$\begin{aligned} \mathcal{L} = \log(1 - \tau) + \sum_{h \in H} \lambda_h \left[(1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \log(1 - \tau_{ht}(s)) - \log(1 - \tau) \right] \\ - (1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \mu_t(s) \sum_{h \in H} \tau_{ht}(s) z_{ht}(s). \end{aligned}$$

The first-order conditions imply

$$\sum_{h \in H} \lambda_h = 1$$

and

$$\lambda_h \frac{1}{1 - \tau_{ht}(s)} = \mu_t(s) z_{ht}(s).$$

Hence

$$(1 - \tau_{ht}(s)) z_{ht}(s) = \frac{\lambda_h}{\mu_t(s)}.$$

Summing across h and using budget balance yields

$$\sum_{h \in H} z_{ht}(s) = \frac{1}{\mu_t(s)}.$$

Define

$$Z_t(s) \equiv \frac{1}{|H|} \sum_{h \in H} z_{ht}(s).$$

Then

$$\mu_t(s) = \frac{1}{|H| Z_t(s)},$$

so

$$\log(1 - \tau_{ht}(s)) = \log(|H| \lambda_h) - \log\left(\frac{z_{ht}(s)}{Z_t(s)}\right).$$

The constant λ_h is determined by the requirement that all regions receive the same

consumption-equivalent gain. Imposing

$$(1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \log(1 - \tau_{ht}(s)) = \log(1 - \tau)$$

and using $\sum_{h \in H} \lambda_h = 1$ gives

$$\lambda_h = \frac{\exp \left\{ (1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \log \left(\frac{z_{ht}(s)}{Z_t(s)} \right) \right\}}{\sum_{h' \in H} \frac{1}{|H|} \exp \left\{ (1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \log \left(\frac{z_{h't}(s)}{Z_t(s)} \right) \right\}}.$$

Substituting back yields the exact nonlinear optimal policy:

$$\log(1 - \tau_{ht}(s)) = \log \left(\frac{\exp \left\{ (1 - \beta) \sum_{t'=0}^{\infty} \beta^{t'} \sum_{s'} \pi(s') \log \left(\frac{z_{h't'}(s')}{Z_{t'}(s')} \right) \right\}}{\sum_{h' \in H} \frac{1}{|H|} \exp \left\{ (1 - \beta) \sum_{t'=0}^{\infty} \beta^{t'} \sum_{s'} \pi(s') \log \left(\frac{z_{h't'}(s')}{Z_{t'}(s')} \right) \right\}} \right) - \log \left(\frac{z_{ht}(s)}{Z_t(s)} \right). \quad (40)$$

Equation (40) is the exact nonlinear counterpart to the policy rule in Section 6. The government assigns each region a constant term, common across all date-state pairs for that region, and then varies labor taxes with region h 's productivity relative to the contemporaneous cross-sectional average. Regions are subsidized in date-state pairs in which their productivity is low relative to $Z_t(s)$ and taxed in date-state pairs in which their productivity is high relative to $Z_t(s)$, with the constant term chosen so that discounted expected log utility gains are equalized across regions.

The connection with the linear-quadratic formula in the main text is immediate. Suppose

$$\log z_{ht}(s) = \sigma x_{ht}(s),$$

and expand (40) around $\sigma = 0$. Since

$$\log Z_t(s) = \log \left(\frac{1}{|H|} \sum_{h' \in H} e^{\sigma x_{h't}(s)} \right) = \sigma \mathbb{E}_{\chi_{h'}} [x_{h't}(s)] + o(\sigma),$$

the first-order approximation to (40) is

$$\Delta \log(1 - \tau_{ht}(s)) = - \left[\Delta \log \left(\frac{z_{ht}(s)}{\bar{z}_h} \right) - \mathbb{E}_{\chi_{h'}} \left[\Delta \log \left(\frac{z_{h't}(s)}{\bar{z}_{h'}} \right) \right] \right],$$

where

$$\Delta \log \bar{z}_h \equiv (1 - \beta) \sum_{t=0}^{\infty} \beta^t \sum_s \pi(s) \Delta \log z_{ht}(s).$$

This is exactly the policy rule reported in Section 6, now interpreted as the first-order approximation to the exact nonlinear dynamic allocation rule in (40).