Efficiency, Insurance, and Redistribution Effects of Government Policies*

Anmol Bhandari
University of Minnesota

David Evans
University of Oregon

Mikhail Golosov
University of Chicago

Thomas Sargent
NYU

February 2026

Abstract

This paper decomposes welfare measures of policy reforms into parts attributable to redistribution and parts due to efficiency. We further decompose efficiency into subcomponents such as gains from better insurance against idiosyncratic and aggregate risk. Our decomposition of welfare measures associated with alternative feasible allocations is cast in terms of a coordinate system that uses generalized Pareto–Negishi weights to capture inequality and production and consumption wedges to capture distortions. Our decomposition has several desirable properties. It attributes welfare changes from movements along a Pareto frontier to redistribution; it attributes negative efficiency changes to movements away from the Pareto frontier; and it produces subcomponent shares of welfare changes that are numeraire-invariant and symmetric with respect to the direction of the reform. Our decomposition can be explained in terms of an implicit tax-and-transfer system in which redistribution captures real income changes, efficiency captures deadweight losses and output costs.

*We thank Casey Mulligan, Eduardo Dávila, Axelle Ferriere, Harald Uhlig, Nicolas Werquin, and especially Guido Lorenzoni and Elisa Rubbo, as well as participants at various conferences for discussions and comments, and Leo Aparisi De Lannoy, Jiawei Fan and Karan Jain for excellent research assistance. Bhandari, Evans, and Golosov thank the NSF for support (grant #36354.00.00.00).

1 Introduction

Government policies affect social welfare through various channels. They redistribute resources among agents, distort the production and allocation of goods, and provide insurance against stochastic shocks. A quantitative analyst often wants to decompose welfare effects of policy reforms into parts attributable to these channels.[1] This paper proposes a decomposition that can be applied to a broad class of economies and that satisfies desirable normative properties.

We consider a setting with heterogeneous agents who have preferences over multiple final consumption goods. Our specification is general enough to include dynamic stochastic models with aggregate and idiosyncratic shocks as well as flexible production networks. Allocations may be distorted due to incomplete markets, monopoly power, borrowing constraints, or various government policies such as taxes.

A welfare measure aggregates the preferences of heterogeneous consumers, so welfare changes can come partly from redistribution across agents and partly from changes in efficiency. Our first task is to decompose welfare changes into these two components. To distinguish redistribution from efficiency, we want a notion of efficiency that does not incorporate attitudes about redistribution. Pareto efficiency serves our purpose: at a Pareto-efficient allocation, it is impossible to make an agent better off without making another agent worse off. Our attachment to Pareto efficiency motivates us to impose two minimal consistency requirements on our decomposition. First, a welfare difference between two Pareto-efficient allocations must be attributed entirely to redistribution. Second, policies that move an allocation off the Pareto frontier (or back onto it) must be evaluated as lowering efficiency (or improving it).

To create a welfare decomposition with these properties, we represent allocations with a coordinate system that measures both implicit inequality and distortions. Our approach is illustrated well by first considering points on the Pareto frontier. Here we can take advantage of classic insights of Negishi (1960), who mapped a point on the Pareto frontier to a vector of implicit welfare weights, often referred to as Pareto–Negishi (PN) weights. PN weights provide a natural coordinate system for when we want to interpret movements along the Pareto frontier as redistribution.

To extend the coordinate system to all feasible allocations, we generalize the notion

[1] See, for example, Abbott et al. (2019), Cho et al. (2015), Conesa et al. (2009), Dyrda and Pedroni (2021), Guvenen et al. (2019), Heathcote et al. (2017), Koehne and Kuhn (2015), Nakajima and Takahashi (2020), Seshadri and Yuki (2004).

of PN weights to measure implicit inequality associated with inefficient allocations. So in addition to PN weights, we use a vector of "wedges" that capture departures from efficiency. These wedges measure shortfalls in production and deviations from an optimal allocation of resources across goods for each individual. Any change in allocations alters both the generalized PN weights and the wedges; by inverting the relationship between allocations and coordinates, we can assess how much of the welfare difference is accounted for by changes in generalized PN weights, to be attributed to redistribution, and how much is accounted for by changes in wedges, to be attributed to a change in efficiency.

Our approach naturally extends to finer divisions of efficiency. Since our vector of wedges consists of two components—one measuring distortions in the production of aggregate resources and the other measuring distortions in the allocation of individual resources across specific goods—we can immediately separate the efficiency component into changes in production efficiency and changes in allocative efficiency. By focusing on wedges that measure distortions in how physical goods are allocated across different stochastic states of nature, we can also isolate an insurance component of government policies, then refine it into parts attributable to insurance against aggregate risks and parts attributable to insurance against idiosyncratic shocks.

By construction our decomposition satisfies the Pareto-consistency requirements that we want. It is numeraire-invariant and symmetric with respect to the direction of the reform. It also has an income-based redistribution property: a redistribution is associated with a change in the distribution of real disposable income across agents. A final feature of our decomposition is that the efficiency component depends on social welfare weights. While there exist measures of efficiency that are agnostic to social welfare weights and measures of redistribution that are not income-based, we show that using such measures leads to misleading attributions of welfare changes to redistribution and efficiency.

Our decomposition has a natural economic interpretation. A feasible allocation can be regarded as a competitive equilibrium allocation for a complete-markets economy with fictitious distortionary taxes and person-specific lump-sum transfers. A change in allocation can be interpreted as having come from a reform of this fictitious tax system. Viewed in this way, the allocative efficiency term is determined by the deadweight losses from changes in distortionary taxes faced by consumers; the production efficiency term is determined by changes in distortionary taxes faced by producers as well as wasteful government spending; and the redistributive component is determined by changes in the net transfer payments that consumers receive through the tax-and-transfer system. These components can be described with concepts appearing in the public finance literature (Diamond and Mirrlees, 1971b; Diamond, 1975).

We compare our decomposition to two prominent alternatives, one developed by Benabou (2002) and Floden (2001) and another by Dávila and Schaab (2022). We show that these decompositions do not satisfy the two Pareto-consistency properties that ours incorporates. For example, they can find that Pareto-improving policies reduce efficiency, that measured efficiency gains increase as distortions grow, or that distortionary taxes are more efficient than a lump-sum transfer that achieves the same redistribution. In stochastic environments, the insurance components in their decompositions inherit similar problems.

We apply our decomposition to two calibrated incomplete-markets economies. First, we revisit a standard permanent income-tax reform in which a higher proportional tax on household income is financed by adjusting uniform lump-sum transfers. We show how our decomposition separates the welfare effect of the reform into redistribution, insurance, and several subcomponents of efficiency—most notably intertemporal smoothing and intratemporal labor distortions. In this setting, aggregate utilitarian welfare is maximized at a substantially higher tax rate than in a status quo that is calibrated to the observed US fiscal policy. Our decomposition indicates that welfare gains are driven mostly by redistribution and improved insurance. Various efficiency channels have offsetting effects but overall they are negative and dominated by distortions to labor supply. At a utilitarian optimum, the reform generates sizable welfare gains for low- and middle-wealth households and welfare losses for high-wealth households. Cross-sectional gradients of welfare effects are almost entirely attributable to the redistribution component. Insurance contributes positively for all groups, while efficiency losses are concentrated among high-wealth households through the labor wedge.

We also study a public-debt reform of Aguiar et al. (2024) that raises government debt and delivers a Pareto improvement. Our decomposition shows that almost all welfare gains come from improved insurance and intertemporal smoothing, with only a minor role being played by redistribution.

The rest of the paper is organized as follows. Section 2 introduces the environment. Section 3 describes our coordinate system, our proposed welfare decomposition, and discusses its properties. Section 4 presents an alternative characterization of our decomposition in terms of a fictitious tax-and-transfer system in an Arrow–Debreu economy. Section 5 compares our approach to other decompositions that have been proposed. Section 6 provides two quantitative applications of our decomposition. Section 7 concludes. Proofs of claims in

Sections 3 and 4 are in Appendix A, and details of numerical implementation and additional results for Sections 5 and 6 are in Appendix B.

2 Preliminaries

I consumers derive utility from K goods. Consumer i's utility function $u_i : \mathbb{R}_+^K \to \mathbb{R}$ is strictly increasing in all arguments, strictly concave, and twice differentiable, with the marginal utility of good k denoted by $u_{i,k}$.

An allocation is an array $x \in \mathbb{R}_+^{I\times K}$. A typical entry $x_{i,k}$ denotes consumption of good k by consumer i; $x_i = \{x_{i,k}\}_k$ denotes the consumption bundle of consumer i. An allocation x determines aggregate consumption $X = \sum_i x_i \in \mathbb{R}^K$. A convex and compact set $\mathcal{Y} \subset \mathbb{R}^K$ satisfying free disposal describes production possibilities. An allocation x is feasible if $X \in \mathcal{Y}$.

As noted by Arrow (1953) and Debreu (1959), such an environment encompasses both deterministic and stochastic economies, since there is no conceptual distinction between consumption of different physical goods and consumption of the same physical good in different states of nature. The K goods from which consumers derive utility are often referred to as final goods. Our specification of $\mathcal{Y}$ is general enough to allow for production processes that include multiple sectors or firms, additional intermediate goods, and production networks. We list a few cases to illustrate how such sets arise in specific macroeconomic environments. These examples will be used in later sections.

Illustrations of production possibility sets. We first review how to deduce production sets from production functions.[2] Suppose that a production technology F converts initial endowments of goods into final consumption goods. Let $a \in \mathbb{R}^{I\times L}$ denote a matrix of individual endowments, and $A \in \mathbb{R}^L$ the vector of aggregate endowments with $A_l = \sum_i a_{i,l}$ being the total endowment of good l for $l = 1, \dots, L$ and $a_{i,l}$ being person i's endowment of good l. We can write feasibility $F(X, A) \leq 0$. Here the production set $\mathcal{Y}$ consists of all aggregate bundles X that satisfy $F(X, A) \leq 0$. In some applications, it is convenient to assume that endowments are in final goods, in which case the aggregate technology can be represented as the constraint $F(X - A) \leq 0$.

We say that a technology is *linear* if F takes the form $F(X - A) = \sum_k b_k(X_k - A_k)$ for some parameters $b = \{b_k\}_k$. We have an *endowment economy* if F takes the form $F(X - A) = \max_k\{X_k - A_k\}$. In an endowment economy, F is not differentiable because

[2] See also Mas-Colell et al. (1995) for a textbook treatment of production functions.

goods cannot be transformed into one another, so the boundary of the feasible set exhibits kinks. Likewise, in stochastic environments with aggregate states, F is non-differentiable across states of nature: units of a physical good in one state cannot be converted into units of the same good in another state.

An example of linear technology are consumption–leisure models in which agents' labor supplies are converted into outputs of the final consumption good linearly. In such models, agents have a fixed endowment of time that they allocate between leisure ℓ and labor supply $l = a_\ell - \ell$. Let $\bar{A}_\ell$ denote the aggregate time endowment, so that aggregate leisure $\mathcal{L}$ and aggregate labor supply L satisfy $\mathcal{L} + L = \bar{A}_\ell$. The aggregate technology can then be represented in terms of either aggregate labor or aggregate leisure as $F(C, \mathcal{L}) = F(C, \bar{A}_\ell - L)$. We shall use such leisure and labor representations interchangeably.

Finally, various types of production networks or intermediate inputs are subsumed in F. We illustrate this using a simple two-tier production network. Consider a consumption–leisure model in which the final consumption good C is produced by a final-good sector that combines labor L_1 and intermediate goods Z using the technology $C = Z^\theta L_1^{1-\theta}$. Intermediate goods are produced from labor according to $Z = L_2$, and the aggregate labor supply is $L = L_1 + L_2$. It is easy to show that the resulting aggregate production technology is $C \leq \theta^\theta (1-\theta)^{1-\theta} L$. Equivalently, in terms of a consumption–leisure representation, feasibility can be written as $F(C, \mathcal{L}) \leq 0$ with $F(C, \mathcal{L}) = C + \theta^\theta (1-\theta)^{1-\theta} (\mathcal{L} - \bar{T})$, which is a special case of the linear technologies defined above.

In many interesting economies, the aggregate allocation X implied by x lies inside the production possibility frontier (PPF), i.e., in the interior of $\mathcal{Y}$. In the two-tier example above, X is on the PPF if and only if the allocation of labor between the final- and intermediate-good sectors satisfies $L_1/L_2 = (1-\theta)/\theta$. This condition would hold in a classical undistorted competitive equilibrium but would typically fail in the presence of monopoly power or tax distortions. This mirrors a more general property that in network economies, such as the one considered by Liu (2019) or Baqaee and Farhi (2019), the equilibrium allocation of final goods typically lies in the interior of the PPF when intermediate-goods production is distorted.

Social welfare. We assume that a researcher evaluates social welfare using the function $W(x) = \sum_i \bar{\alpha}_i u_i(x_i)$, where $\bar{\alpha}_i \geq 0$ for all i.[3] Let x^* be the equilibrium allocation under

[3] We use a welfare function that is linear in utility for expositional simplicity. The welfare index can be any function of the allocation as long as it preserves Pareto improvements, that is, if x^{**} Pareto dominates

status-quo government policies, and x^{**} be the allocation under an alternative policy. The welfare effect of a policy reform is $W(x^{**})-W(x^*)$. Our goal is to decompose $W(x^{**})-W(x^*)$ into two components: a redistribution component that measures how welfare changes due to reallocation of resources across agents, and an efficiency component that measures gains or losses associated with distortions in the production and allocation of resources. With such a decomposition in hand, we can decompose the efficiency term into additional subcomponents representing, for instance, efficiency changes due to improved insurance.

3 Welfare decompositions

We construct our decomposition through a coordinate system that allows us to categorize and understand the welfare effects of policy reforms. We begin with points on the Pareto frontier, where welfare effects can be classified exclusively as redistribution. Standard economic theory tells us that movements along the frontier can be indexed by Pareto–Negishi (PN) weights, which we interpret as a coordinate system for the frontier. We extend these ideas to all feasible allocations. A reform that moves inside the Pareto frontier should require a reduction in efficiency. Motivated by this property, we associate each inefficient allocation with a point on the Pareto frontier and a vector of economic distortions relative to that point—yielding a coordinate system consisting of PN weights and efficiency wedges. We develop a decomposition using this coordinate system by associating changes in PN weights with redistribution and changes in efficiency wedges with efficiency gains or losses. We conclude by discussing the intuitive properties of this decomposition and then extend the analysis to stochastic environments. Throughout, we focus on interior allocations in the main text, but include analysis of all feasible allocations in Appendix A.

3.1 Efficient allocations

Let x be a point on the Pareto frontier that attains the maximum of the following problem

$$\max_{\tilde{x},\tilde{X}} \sum_i \alpha_i u_i\left(\tilde{x}_i\right) \quad \text{s.t.} \quad \sum_i \tilde{x}_i \leq \tilde{X}, \quad \tilde{X} \in \mathcal{Y}, \tag{1}$$

for some Pareto–Negishi weights α. At an interior allocation, the PN weight of consumer i is inversely proportional to the marginal utility of consumption: $\alpha_i \propto \frac{1}{u_{i,k}(x_i)}$. Given an

x^* then $W(x^{**}) \geq W(x^*)$. For example, any welfare aggregator $W(x) = \mathcal{G}(\{u_i(x_i)\}_i)$ that is increasing in all arguments. This allows $W(x)$ to be evaluated in any units, e.g. % of consumption.

allocation, any good k can be used to construct the PN weights, since on the Pareto frontier ratios of marginal utilities are equalized across agents for all goods.

Let $\mathsf{P} \in \mathbb{R}^K$ be the vector of Lagrange multipliers associated with feasibility constraints $\sum_i \tilde{x}_i \leq \tilde{X}$ in maximization problem (1). Rewrite problem (1) as

$$\max_{\tilde{X} \in \mathcal{Y}} \sum_k \mathsf{P}_k \tilde{X}_k + \sum_i \alpha_i \max_{\tilde{x}_i} \Big\{ u_i(\tilde{x}_i) - \alpha_i^{-1} \sum_k \mathsf{P}_k \tilde{x}_{i,k} \Big\}. \tag{2}$$

The maximizers (x, X) satisfy

$$\sum_{i,k} \mathsf{P}_k x_{i,k} = \sum_k \mathsf{P}_k X_k. \tag{3}$$

Lagrange multipliers P act as implicit prices that convert quantities of physical goods into a common notion of income. In general, the multipliers P depend on the underlying PN weights α, so let $\mathsf{P}(\alpha)$ denote the Lagrange multipliers associated with the feasibility constraints in problem (1) given weights α.

The first maximization in problem (2) shows that at a Pareto optimum the planner chooses aggregate quantities to maximize aggregate income. The collection of maximization problems over $\tilde{x}_i, i \in I$, in problem (2) determines the optimal allocation of total output among individual agents. Equation (3) shows that the sum of individual spending or disposable income equals aggregate income.

Define the indirect utility function V_i for consumer i at a price vector $p \in \mathbb{R}^K_+$ and income y_i as

$$V_i(p, y_i) := \max_{\tilde{x}_i} u_i(\tilde{x}_i) \quad \text{s.t.} \quad \sum_k p_k \tilde{x}_{i,k} \leq y_i. \tag{4}$$

Using this function, we characterize Pareto-efficient allocations as follows:

Lemma 1. *Let x be any interior Pareto-efficient allocation and $\alpha \in \mathbb{R}^I_+$ be the associated Pareto–Negishi weights and $p \in \mathbb{R}^K_+$ be any supporting price vector such that $p \propto \mathsf{P}(\alpha)$. Define individual incomes $y_i := \sum_k p_k x_{i,k}$ and aggregate income $Y := \sum_i y_i$. The triple (Y, y, x) satisfies:*

(i) ***Aggregate income.*** *Given p, aggregate income Y is the maximum attainable income:*

$$Y = Y^{\max}(p) := \max_{\tilde{X} \in \mathcal{Y}} \sum_k p_k \tilde{X}_k. \tag{5}$$

(ii) ***Income distribution.*** *Given (p, Y), the income distribution y solves*

$$\max_{\tilde{y} \geq 0} \sum_i \alpha_i V_i(p, \tilde{y}_i) \quad s.t. \quad \sum_i \tilde{y}_i \leq Y. \tag{6}$$

(iii) ***Individual allocation.*** *The consumption bundle x_i of consumer i solves the individual problem (4) given (p, y_i), i.e., x_i satisfies*

$$1 = \frac{p_1}{p_k} \frac{u_{i,k}(x_i)}{u_{i,1}(x_i)} \quad \textit{for all } k > 1. \tag{7}$$

Proof. See Appendix A.1.1. □

Lemma 1 computes an efficient allocation in three steps. First, the planner chooses aggregate quantities by maximizing aggregate income as in (5). Second, the planner allocates aggregate income across consumers by solving (6). Finally, each consumer allocates income across goods by solving (4).

At an interior solution to (6), the optimality condition implies $\alpha_i \propto 1/V_{i,y}(p, y_i)$, where $y_i = \sum_k p_k x_{i,k}$ and since any supporting price vector satisfies $p \propto \mathsf{P}(\alpha)$, the PN weights satisfy a fixed point given by the following lemma.

Lemma 2. *Let x be any interior Pareto-efficient allocation and α be the associated Pareto–Negishi weights normalized to satisfy $\sum_i \alpha_i = 1$. Then α is the unique solution to*

$$\alpha_i = \frac{\dfrac{1}{V_{i,y}(\mathsf{P}(\alpha), \sum_k \mathsf{P}_k(\alpha) x_{i,k})}}{\sum_{i'} \dfrac{1}{V_{i',y}\big(\mathsf{P}(\alpha), \sum_k \mathsf{P}_k(\alpha) x_{i',k}\big)}}, \tag{8}$$

where $V_{i,y}$ denotes the derivative of V_i with respect to income.

Proof. See Appendix A.1.2. □

In general, (8) is nonlinear because the supporting multipliers $\mathsf{P}(\alpha)$ depend on α. When Gorman aggregation holds, relative multipliers $\mathsf{P}_k(\alpha)/\mathsf{P}_1(\alpha)$ are independent of α, so supporting prices can be chosen to be constant along the Pareto frontier.[4] In such cases, one can find prices by solving $X \in \arg\max_{\tilde{X} \in \mathcal{Y}} u(\tilde{X})$ for a common utility function u, and then setting $\mathsf{P}_k(\alpha) \propto u_k(X)$. For identical separable CRRA preferences $u_i(x_i) = \frac{1}{1-\sigma} \sum_k d_k x_{i,k}^{1-\sigma}$,

[4] Prices are also constant if the technology is linear and in that case they equal the marginal rates of transformation.

we have $V_{i,y}(p,y) \propto y^{-\sigma}$ and hence $\alpha_i(x) \propto y_i(x)^\sigma$. For identical separable CARA preferences $u_i(x_i) = -\frac{1}{\gamma}\sum_k d_k \exp(-\gamma x_{i,k})$, $V_{i,y}(p,y)$ is exponential in income, implying $\alpha_i(x) \propto \exp\bigl(\gamma y_i(x)/\sum_k p_k\bigr)$; the normalization by $\sum_k p_k$ ensures invariance to rescaling of prices. Even when conditions for exact Gorman aggregation are not met, many models—for instance, Bewley–Huggett–Aiyagari-style settings—feature an approximate aggregation that leads to relative prices that are nearly constant along the Pareto frontier. We illustrate this in our quantitative application in Section 6.

We focused first on allocations on the Pareto frontier deliberately as it allows us to unambiguously separate redistribution from efficiency. By definition, at a Pareto-efficient allocation it is impossible to make any agent strictly better off without making someone else worse off. Thus, the welfare effects associated with movements along the frontier must be due only to redistribution. We refer to this property as *weak Pareto-consistency.*

Conditional on the desired level of redistribution, summarized by the Pareto–Negishi weights α, prices serve as a sufficient statistic for efficient production and efficient allocation of income across goods for each consumer. This is reflected in Part (i) and Part (iii) of Lemma 1 where the equations (5) and (7) do not depend on α. The supporting prices also determine who benefits from redistribution along the Pareto frontier. Let $\hat{x}_{i,k}$ denote an infinitesimal reallocation along the frontier. This change is associated with adjustments in prices $\hat{p}_k$ and nominal incomes $\hat{y}_i = \sum_k(\hat{p}_k x_{i,k} + p_k \hat{x}_{i,k})$. The welfare change for consumer i is given by $\hat{U}_i = V_{i,y}\hat{y}_i + V_{i,k}\hat{p}_k$, where $V_{i,k}$ denotes the derivative of $V_i(p,y)$ with respect to p_k. Applying Roy's identity, this simplifies to

$$\hat{U}_i = V_{i,y} \sum_k p_k \hat{x}_{i,k}. \tag{9}$$

The term $\sum_k p_k \hat{x}_{i,k}$ is the real change in disposable income. Redistribution along the Pareto frontier requires changes in real disposable income—we term this *income-based redistribution.* In the next section, we show how to generalize this property to all feasible allocations.

3.2 Representation of distorted allocations

We now generalize this approach to all feasible allocations, not just those on the Pareto frontier. Because the Pareto frontier represents the set of maximally efficient allocations, movements inside the frontier necessarily involve reductions in efficiency. We label this property *strong Pareto-consistency.*

We use strong Pareto-consistency to inform our coordinate system by associating each inefficient allocation with a point on the Pareto frontier, indexed by Pareto–Negishi weights α. Associated with that point is a vector of prices p_k that summarize how to efficiently produce and allocate goods. We represent inefficiency through a vector of distortions $t = (\xi, \tau)$, where ξ captures distortions in production and τ captures distortions in the allocation of income across goods. This representation extends Lemma 1 to all feasible allocations.

Lemma 3. *Let x be any interior feasible allocation and $p \propto \mathsf{P}(\alpha)$ be any supporting price vector. Define individual incomes $y_i := \sum_k p_k x_{i,k}$, and let $Y := \sum_i y_i$ denote aggregate income. There exist Pareto–Negishi weights $\alpha \in \mathbb{R}^I_+$, allocation wedges $\tau = \{\tau_{i,k}\}_{i,k}$ with $\tau_{i,1} = 0$, and a production wedge $\xi \in [0, 1]$ for which the triple (Y, y, x) satisfies*

(i) ***Aggregate income.*** *Aggregate income equals the maximum attainable income reduced by the production wedge:*

$$Y = (1 - \xi)\, Y^{\max}(p), \tag{10}$$

where $Y^{\max}(p)$ is defined in equation (5).

(ii) ***Income distribution.*** *Given p and Y, the income distribution y solves*

$$\max_{\widetilde{y} \geq 0} \sum_i \alpha_i V_i(p, \widetilde{y}_i) \quad s.t. \quad \sum_i \widetilde{y}_i \leq Y, \tag{11}$$

where the indirect utility function $V_i(p, y)$ is defined in equation (4).

(iii) ***Individual allocation.*** *The consumption bundle x_i of each household satisfies*

$$1 + \tau_{i,k} = \frac{p_1}{p_k} \frac{u_{i,k}(x_i)}{u_{i,1}(x_i)} \quad \textit{for all } k > 1. \tag{12}$$

Proof. See Appendix A.1.3. □

Lemma 3 shows that any feasible allocation can be represented by a triple (α, ξ, τ) consisting of Pareto–Negishi weights α and two types of wedges: a production wedge ξ, and allocation wedges τ that capture inefficiencies in production and consumption, respectively. Central to this representation is the vector of PN weights $\alpha(x)$ implied by the allocation x. These weights solve the same system of nonlinear equations as in Lemma 2.[5] Once $\alpha(x)$

[5] In the proof of Lemma 3, we apply Brouwer's fixed-point theorem to the mapping defined by (8) to show the existence of $\alpha(x)$. In standard models and the examples we study, the weak dependence of P on α implies that the fixed point is unique. In general, there can be multiplicity, and for the purpose of constructing our decomposition, we use the following selection procedure. Consider comparing a status quo allocation x^* to a

is computed, shadow prices are $p(x) \propto \mathsf{P}(\alpha(x))$, and the wedges $\xi(x)$ and $\tau(x)$ are then determined by conditions (10) and (12). We use a vector $t(x) = (\xi(x), \tau(x))$ to summarize distortions associated with allocation x.

Defining α in terms of the marginal utility of income requires computing prices. On the Pareto frontier, the Lagrange multipliers $\mathsf{P}(\alpha)$ served as supporting prices that could be used to decentralize any point on the frontier with lump-sum transfers. This allowed us to interpret welfare changes from reforms along the Pareto frontier as redistribution, reflecting the tight link between changes in utilities and changes in real disposable income. This link necessarily breaks down for allocations inside the Pareto frontier. When marginal rates of substitution are not equalized across agents, there is no single price vector that can consistently aggregate consumption into a measure of disposable income for all consumers.

Our coordinate system gives a natural solution to this problem. A by-product of the fixed point for α are supporting prices associated with the generalized Pareto–Negishi weights. Because the Pareto–Negishi weights are determined by the income distribution, an infinitesimal change $\hat{x}_{i,k}$ can only result in a change in these weights if there is an associated change in real income $\sum_k p_k \hat{x}_{i,k}$. This extends the notion of *income-based redistribution* used on the Pareto frontier to all feasible allocations.

In our coordinate system, the two measures of efficiency—the scalar ξ and the vector of wedges τ—play distinct roles. The production wedge ξ summarizes how far aggregate income Y falls short of the maximum attainable income $Y^{\max}(p)$ at supporting prices p, so $1 - \xi$ measures inefficiencies in the aggregate production plan X.[6] The allocation wedges $\tau = \{\tau_{i,k}\}_{i,k}$ measure deviations of individuals' marginal rates of substitution from socially optimal rates of transformation. They can be interpreted as the implicit taxes and subsidies that distort households' allocation of resources across goods.

The vectors α and t have $I - 1$ and $I(K-1)+1$ degrees of freedom, respectively. Under some mild auxiliary conditions on preferences,[7] the pair (α, t) describes a unique allocation $\mathcal{X}(\alpha, t)$ as follows. Let $p \propto \mathsf{P}(\alpha)$ be any supporting price vector. Total private income Y is

post-reform allocation x^{**}. In the case of multiple solutions to (8), we select α^* at the status quo allocation x^* to minimize the welfare cost of distortions $W(x^{PF}(\alpha^*)) - W(x^*)$, and select α^{**} at the reform allocation x^{**} to minimize $\|\alpha^{**} - \alpha^*\|^2$.

[6] Our production wedge $1 - \xi$ may appear similar to the *coefficient of resource utilization* of Debreu (1951), but the two are generically different. Debreu's coefficient aggregates all deviations from the Pareto frontier into a single scalar index. Our representation instead separates underutilization and technological inefficiency into a scalar ξ and allocation inefficiencies into a vector τ. The two measures coincide only when technology is CRS, supporting prices $p \propto \mathsf{P}(\alpha)$ are independent of α, and allocation wedges vanish ($\tau = 0$). See Appendix A.1.7 for details.

[7] We give these conditions in Appendix A.1.4. A sufficient condition, which is much stronger than needed for the result, is that all goods are normal.

determined by the production wedge ξ via $Y = (1-\xi)\,Y^{\max}(p)$. Given income Y and PN weights α, the distribution of income y solves the planner's problem (6) in condition (ii) of Lemma 1. This ensures that the allocation of realized income Y is consistent with the social preferences embedded in α even when $Y < Y^{\max}(p)$. The consumption bundles $x_{i,k}$ then satisfy the distorted first-order conditions (12) and the budget constraint $\sum_k p_k x_{i,k} = y_i$.

The supporting price vector p continues to be defined only up to a positive normalization of the Lagrange multipliers $\mathsf{P}(\alpha)$. If p is a supporting price vector for a given α, then so is $\tilde{p} = \lambda p$ for any $\lambda > 0$, with incomes and $Y^{\max}(p)$ scaling proportionally. All of the objects we have constructed from p are invariant to such rescalings. Thus, the allocation wedges τ depend on p only through relative prices p_1/p_k, and the production wedge ξ is defined from the ratio $Y/Y^{\max}(p)$, which is unchanged when both Y and $Y^{\max}(p)$ are multiplied by the same factor. Likewise, the PN weights $\alpha(x)$ are constructed from marginal utilities of income evaluated at p-implied incomes, and rescaling p rescales all incomes proportionally without affecting the fixed point (8) that defines $\alpha(x)$. Consequently, our (α, t) representation and the resulting welfare decomposition do not depend on how we normalize p. Later, we shall use this property when we construct a decomposition that is numeraire-independent.

Welfare can be written in the (α, t) coordinates as $\mathcal{W}(\alpha, t) := \sum_i \bar{\alpha}_i \, u_i(\mathcal{X}_i(\alpha, t))$. By construction, $(\alpha, 0)$ represents an allocation on the Pareto frontier. Consistent with interpreting t as a measure of distortions, the following lemma asserts that distortions reduce welfare relative to the efficient benchmark, and is a key ingredient in guaranteeing that our welfare decomposition satisfies strong Pareto-consistency.

Lemma 4. *For all feasible x with corresponding α and t,* $\mathcal{W}(\alpha, 0) \geq \mathcal{W}(\alpha, t)$.

Proof. See Appendix A.1.5. □

3.3 Marginal and Shapley-value welfare decompositions

We apply our coordinate system (α, t) to construct a welfare decomposition for a policy reform that moves the economy from allocation x^* to allocation x^{**}. Using our representation $W(x) = \mathcal{W}(\alpha(x), t(x))$, we can write the welfare effect of this reform as

$$W(x^{**}) - W(x^*) = \mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^*, t^*), \tag{13}$$

where, for brevity, we write $\alpha^* := \alpha(x^*)$ and $t^* := t(x^*)$, and similarly $\alpha^{**} := \alpha(x^{**})$ and $t^{**} := t(x^{**})$.

We begin with a marginal reform that changes x^* by a small amount. Fix x^* and set $x^{**} = x^* + \varepsilon \hat{x}$, where ε is a scalar and $\hat{x} \in \mathbb{R}^{I \times K}$ is the direction of change. Use hats to denote marginal changes as $\varepsilon \to 0$, e.g., $\hat{W} := \lim_{\varepsilon \to 0} \frac{W(x^* + \varepsilon \hat{x}) - W(x^*)}{\varepsilon}$. Dividing (13) by ε and taking a limit as $\varepsilon \to 0$, we obtain

$$\hat{W} = \underbrace{\mathcal{W}_\alpha \hat{\alpha}}_{\hat{\mathsf{R}}} + \underbrace{\mathcal{W}_t \hat{t}}_{\hat{\mathsf{E}}}, \tag{14}$$

where $\mathcal{W}_\alpha$ and $\mathcal{W}_t$ denote the gradients of $\mathcal{W}$ with respect to α and t, evaluated at (α^*, t^*), and where $\hat{\alpha}$ and $\hat{t}$ are corresponding marginal changes in coordinates induced by the perturbation $\hat{x}$. The Gateaux derivative $\hat{\mathsf{R}} := \mathcal{W}_\alpha \hat{\alpha}$ captures the effect of changes in inequality on welfare holding distortions fixed, while $\hat{\mathsf{E}} := \mathcal{W}_t \hat{t}$ captures the effect of changes in distortions on welfare holding inequality fixed. Whenever $\hat{W} \neq 0$, the ratios $\hat{\mathsf{R}}/\hat{W}$ and $\hat{\mathsf{E}}/\hat{W}$ sum to one and can be interpreted as the fractions of the marginal welfare change attributable to redistribution and to efficiency. These fractions can be negative or exceed one in absolute value when redistributive and efficiency effects work in opposite directions.

For non-marginal reforms, the attribution of the total welfare difference $\mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^*, t^*)$ to changes in α and t is not unique because $\mathcal{W}$ is generally non-separable in its arguments. A common approach in such settings—see Shorrocks (2013) and Lundberg and Lee (2017)—is to use *Shapley values* to assign contributions to each argument.[8] In our setting, Shapley values associated with α and t are

$$\mathsf{R}(x^*, x^{**}) = \tfrac{1}{2}\big[\mathcal{W}(\alpha^{**}, t^*) - \mathcal{W}(\alpha^*, t^*)\big] + \tfrac{1}{2}\big[\mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^*, t^{**})\big],$$

$$\mathsf{E}(x^*, x^{**}) = \tfrac{1}{2}\big[\mathcal{W}(\alpha^*, t^{**}) - \mathcal{W}(\alpha^*, t^*)\big] + \tfrac{1}{2}\big[\mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^{**}, t^*)\big].$$

This yields the global decomposition

$$W(x^{**}) - W(x^*) = \mathsf{R}(x^*, x^{**}) + \mathsf{E}(x^*, x^{**}). \tag{15}$$

We use R and E as our measures of redistribution and efficiency for non-marginal reforms.

[8] Let $f(z)$ be a function of $z \in \mathbb{R}^N$. The Shapley value of the nth component z_n in the difference $f(z^{**}) - f(z^*)$ is defined as the average marginal contribution of z_n across all possible orders in which $z^* = (z_1^*, \ldots, z_N^*)$ can be transformed into $z^{**} = (z_1^{**}, \ldots, z_N^{**})$ by changing one component at a time. We experimented with several alternative ways of attributing the change in $\mathcal{W}$ to its arguments and found that they perform similarly in the examples and calibrated economies we consider below. We derive them in Appendices A.2.1 and A.2.2. For quantitative applications, we use the Shapley-value-based decomposition because it is easy to implement and widely used, and we report outcomes under alternative schemes in Appendix B.7.

It is easy to verify that for marginal reforms R and E converge to $\hat{\mathsf{R}}$ and $\hat{\mathsf{E}}$, so that (15) is a global extension of (14).

3.4 Properties of the decomposition

We now collect several properties of the Shapley-value based welfare decomposition.

Proposition 1. *Decomposition* (15) *satisfies:*

(i) Weak Pareto-consistency: If x^* and x^{**} are on the Pareto frontier, then $\mathsf{E}(x^*, x^{**}) = 0$.

(ii) Strong Pareto-consistency: If x^* is on the Pareto frontier but x^{**} is not, then $\mathsf{E}(x^*, x^{**}) \leq 0$; if x^{**} is on the Pareto frontier but x^* is not, then $\mathsf{E}(x^*, x^{**}) \geq 0$. When social welfare weights $\bar{\alpha}$ are strictly positive, these inequalities are strict.

(iii) Numeraire invariance: The ratios $\mathsf{R}(x^*, x^{**})/(W(x^{**})-W(x^*))$ and $\mathsf{E}(x^*, x^{**})/(W(x^{**})-W(x^*))$ are independent of the choice of numeraire good (the normalization of the price vector) and of the labeling of goods (which good is indexed as good 1).

(iv) Direction invariance: The decomposition is reflexive in the sense that

$$\frac{\mathsf{R}(x^*, x^{**})}{W(x^{**}) - W(x^*)} = \frac{\mathsf{R}(x^{**}, x^*)}{W(x^*) - W(x^{**})}, \qquad \frac{\mathsf{E}(x^*, x^{**})}{W(x^{**}) - W(x^*)} = \frac{\mathsf{E}(x^{**}, x^*)}{W(x^*) - W(x^{**})}.$$

(v) Income-based redistribution: If $\sum_k p_k^* x_{i,k}^* = \sum_k p_k^* x_{i,k}^{**}$ for all i then $\mathsf{R}(x^*, x^{**}) = 0$.

Proof. See Appendix A.1.6. □

Properties (i)–(iv) are natural desiderata for any decomposition that aims to separate redistribution from efficiency. The first two properties formalize a Pareto-based notion of efficiency that is purged of distributional judgments. If both allocations are Pareto efficient, then any change in social welfare must reflect a change in the distribution of resources, and not a change in the feasible set; Weak Pareto-consistency therefore requires $\mathsf{E} = 0$. Conversely, when one allocation lies on the frontier and the other does not, Strong Pareto-consistency requires that the efficiency component has the correct sign: moving away from the frontier lowers efficiency, while moving toward the frontier raises it (strictly so when all welfare weights are strictly positive).

Properties (iii) and (iv) ensure that the *shares* of the welfare change attributed to redistribution and efficiency are not artifacts of normalization choices. Numeraire invariance guarantees that changing units (or relabeling goods) does not affect the decomposition. Direction invariance guarantees that the relative contributions of redistribution and efficiency are the same whether we view the comparison as a reform from x^* to x^{**} or as the reverse reform.

Property (v) is substantive: it extends the income-based redistribution insight on the Pareto-frontier from Section 3.1 to distorted allocations. Along the Pareto frontier, supporting prices provide a common metric for measuring each agent's real disposable income, $y_i = \sum_k p_k x_{i,k}$, and changes in welfare can be interpreted as redistribution precisely because welfare changes are tightly linked to changes in these real incomes. Property (v) generalizes this idea by requiring that, even away from the frontier, a policy that leaves all agents' real disposable incomes (evaluated at the baseline supporting prices p^*) unchanged has *no* redistribution component. To see why this criterion is appealing, consider the following example.

Example 1 (A pure distortion with no redistribution)**.** Consider an economy with two consumers and two goods, consumption and leisure. Technology is linear: time can be converted into consumption at a constant marginal rate. Consumer 2's labor earnings are taxed at a proportional rate τ, and tax revenues are rebated lump-sum to consumer 2. Consider a reform that changes the tax rate from τ^* to τ^{**}.

In Example 1, the reform alters consumer 2's marginal incentives by distorting their labor supply, but it does not transfer any resources across consumers—consumer 1's allocation is entirely untouched and all taxes paid by agent 2 are rebated lump-sum back to them. Intuitively, any welfare change in this example should arise entirely from changes in how efficiently agent 2 allocates their time between labor and leisure, not from redistribution across consumers. This intuition is captured by Property (v): since each consumer's real disposable income, evaluated at the pre-reform supporting prices p^*, is unchanged (i.e., $\sum_k p_k^* x_{i,k}^* = \sum_k p_k^* x_{i,k}^{**}$ for $i = 1, 2$), our decomposition correctly identifies zero contribution of redistribution to the welfare change.

The reform in Example 1 affects the distribution of utilities even if the distribution of income is unchanged, as consumer 2 is made worse off by the distortion (while consumer 1 is unaffected). It is possible to construct a decomposition that connects redistribution to changes in utilities rather than real incomes. Representing allocations by an I-dimensional vector of utilities requires using a scalar to represent efficiency since capturing redistribution requires $I - 1$ degrees of freedom. In Appendix A.2.3, we use the coefficient of resource utilization suggested by Debreu (1959) as that scalar and construct a decomposition that features a utility-based notion of redistribution.

We show that this alternative decomposition satisfies properties (i)–(iv) but not Property (v). Applying this decomposition to Example 1, we find that it can attribute anywhere between $-\infty\%$ and 100% of the welfare gain to redistribution and thus fail to capture the economic logic underlying the reform.

3.5 Extensions

The efficiency term can be decomposed to highlight distinct channels. First, the array $t = (\xi, \tau)$ consists of a production wedge ξ, which captures distortions in the production of aggregate resources, and allocation wedges τ, which capture distortions in the allocation of goods across consumers. We can represent welfare as $W(x) = \mathcal{W}(\alpha, \xi, \tau)$ and use the Shapley value contributions of ξ and τ to the efficiency component E, denoted by E^{pr} and E^{al}, to measure productive and allocative distortions. We thus obtain

$$W(x^{**}) - W(x^{*}) = \mathsf{R}(x^{*}, x^{**}) + \mathsf{E}^{pr}(x^{*}, x^{**}) + \mathsf{E}^{al}(x^{*}, x^{**}), \tag{16}$$

so E^{pr} and E^{al} are two subcomponents of the efficiency term E. A marginal counterpart is

$$\hat{W} = \underbrace{\mathcal{W}_{\alpha}\hat{\alpha}}_{=\hat{\mathsf{R}}} + \underbrace{\mathcal{W}_{\xi}\hat{\xi}}_{=\hat{\mathsf{E}}^{pr}} + \underbrace{\mathcal{W}_{\tau}\hat{\tau}}_{=\hat{\mathsf{E}}^{al}}. \tag{17}$$

Second, we can single out the role of insurance in stochastic economies. In incomplete-markets heterogeneous-agent models like the one studied by Aiyagari (1994), taxation of labor or capital income affects both the allocation of physical goods and the volatility of consumption. In quantitative work, researchers often want to measure welfare consequences of these two channels separately.[9] An Arrow–Debreu representation that indexes goods by their physical characteristics and the stochastic state of nature in which they arrive provides a natural framework for such decompositions. Our distortion vector τ captures misallocations both across physical goods and across states of nature. By singling out the part of τ that describes distortions across states of nature, we obtain a measure of efficiency losses due to imperfect insurance.

To be concrete, suppose that there are N physical goods and S states of nature, so that $N \times S = K$. To distinguish between physical goods and states explicitly, we denote

[9] See, e.g., Heathcote et al. (2017), Dyrda and Pedroni (2021), or Bhandari et al. (2021).

a typical element of x by $x_{i,n}(s)$. In the same way, typical elements of vectors p and τ are written as $p_n(s)$ and $\tau_{i,n}(s)$. Given the numeraire independence property, we use the pair $(n, s) = (1, 1)$ as the numeraire without loss of generality. Let $\Pr(s)$ be the probability that state s is realized and $u_i(x_i) = \sum_s \Pr(s)\, v_i(x_i(s), s)$ be the expected utility of agent i. Let x_i^{ins} solve

$$\max_{\tilde{x}_i} u_i(\tilde{x}_i) \quad \text{s.t.} \quad \sum_s p_n(s)\tilde{x}_{i,n}(s) \leq \sum_s p_n(s)x_{i,n}(s) \quad \text{for all } n. \tag{18}$$

Allocation x_i^{ins} attains the highest utility that consumer i could attain if she were able to freely reallocate physical goods across states of nature using the Arrow–Debreu securities. A natural interpretation of x_i^{ins} is the best insurance that agent i could obtain, given the existing allocation of physical goods she receives. Motivated by this observation, we decompose τ into two components, τ^g and τ^{ins}, as follows:

$$1 + \tau_{i,n}(s) = \underbrace{\frac{p_1(1)}{p_n(s)} \frac{\Pr(s)\, v_{i,n}(x_i^{ins}(s), s)}{\Pr(1)\, v_{i,1}(x_i^{ins}(1), 1)}}_{:=1+\tau_{i,n}^g} \underbrace{\frac{v_{i,n}(x_i(s), s)/v_{i,1}(x_i(1), 1)}{v_{i,n}(x_i^{ins}(s), s)/v_{i,1}(x_i^{ins}(1), 1)}}_{:=1+\tau_{i,n}^{ins}(s)}.$$

In this construction, τ^{ins} captures insurance imperfections, measured by deviations in the marginal rates of substitution relative to the numeraire across states from full insurance, while τ^g captures distortions in the trade of physical goods.

This approach expresses welfare $W(x)$ in coordinates $(\alpha, \xi, \tau^g, \tau^{ins})$. By computing the Shapley value contributions of τ^{ins} and τ^g to the allocative efficiency component E^{al} we obtain two terms: E^g that captures inefficiencies in the allocation of physical goods and E^{ins} that captures inefficiencies in providing insurance. This extends (16) to

$$W(x^{**}) - W(x^*) = \mathsf{R}(x^*, x^{**}) + \mathsf{E}^{pr}(x^*, x^{**}) + \mathsf{E}^g(x^*, x^{**}) + \mathsf{E}^{ins}(x^*, x^{**}). \tag{19}$$

Proposition 1 still holds for this decomposition, and its Properties (iii) and (iv) also apply to all subcomponents of E. Following analogous steps, one can further decompose the insurance term E^{ins} into separate components capturing insurance against aggregate and idiosyncratic shocks. We do this in Appendix A.3.

4 Decomposition of policy reforms

Our Section 3 decomposition is cast in terms of allocations and is silent about policies that generate those allocations. In this section, we construct an alternative representation of (17) that is cast in terms of a rich set of government policies.

We restrict attention to the case in which the production set $\mathcal{Y}$ is generated by a constant-returns-to-scale technology $F(X - A) \leq 0$, where A is the vector of aggregate endowments constructed from individual endowments $a \in \mathbb{R}^{I \times K}$. This assumption simplifies our exposition, but is not essential to our analysis.[10]

We consider government policies that can be summarized by four objects: government purchases $G \in \mathbb{R}^K_+$, consumer-specific lump-sum taxes/transfers $T \in \mathbb{R}^I$, consumer-specific consumption taxes $\iota \in \mathbb{R}^{I \times K}$ with normalization $\iota_{i,1} = 0$ for all i, and producer taxes $\varsigma \in \mathbb{R}^K$ with normalization $\varsigma_1 = 0$.

Definition. A competitive equilibrium given government policy (G, T, ι, ς) consists of consumption $x \in \mathbb{R}^{K \times I}_+$, aggregate net output $Z = X - A \in \mathbb{R}^K$, and prices $r \in \mathbb{R}^K$ such that

(i) for each i, x_i solves $\max_{\tilde{x}_i} u_i(\tilde{x}_i)$ subject to $\sum_k (1 + \iota_{i,k}) r_k (\tilde{x}_{i,k} - a_{i,k}) \leq T_i$;

(ii) Z solves $\max_{\tilde{Z}} \sum_k (1 - \varsigma_k) r_k \tilde{Z}_k$ subject to $F(\tilde{Z}) \leq 0$;

(iii) the allocation is feasible, $\sum_i x_{i,k} + G_k = Z_k + \sum_i a_{i,k}$ for all k;

(iv) the government budget constraint is satisfied,

$$\sum_i \underbrace{\left[T_i - \sum_k \iota_{i,k} r_k \left(x_{i,k} - a_{i,k}\right)\right]}_{:= T_i^{net}} = \underbrace{\sum_k \varsigma_k r_k Z_k - \sum_k r_k G_k}_{:= S} . \tag{20}$$

This class of economies includes many interesting examples. The equilibrium allocation x of any economy with frictions—for example, due to monopoly power or market incompleteness—can be decentralized *as if* it were the competitive equilibrium allocation for some vector of policies (G, T, ι, ς). In such a decentralization, G need not correspond to actual government purchases. As discussed in Section 2, distortions affecting intermediate producers in network economies place the allocation x inside the PPF. From a welfare point

[10] Any convex technology can be represented as a constant-returns-to-scale technology $F(X - A) \leq 0$ by suitably redefining the commodity space; see McKenzie (1959). Our approach extends to this expanded commodity space in a straightforward manner. To avoid carrying around extra notation for two sets of commodity spaces, we assume, without loss of generality, that the primitive F is constant-returns-to-scale in this section.

of view, such equilibria are equivalent to competitive equilibria of economies without distortions in intermediate-good production but in which some vector of resources G is simply discarded.[11]

It is convenient to rewrite the government budget constraint (20) as $\sum_i T_i^{net} = S$. Here T_i^{net} is the net transfer to consumer i, i.e., the sum of all transfers and taxes paid by that consumer, and S is the government surplus that is available to finance net transfers to consumers. This surplus includes tax revenues from producers net of non-transfer wasteful government expenditures.

Consider an allocation x associated with coordinates (α, ξ, τ) and implicit prices p. This allocation can be supported by a competitive equilibrium for some government policy (G, T, ι, ς) and equilibrium prices r. Moreover, these policies can be chosen so that implicit prices and allocative wedges coincide with equilibrium prices and consumer taxes: $p = r$ and $\tau = \iota$. Given this, we can work directly with the policy vector (G, T, τ, ς) and use p to denote producer prices. In such an equilibrium, the production wedge ξ equals $\frac{S}{\sum_k p_k A_k}$ or the ratio of government surplus (excluding net transfers to consumers) relative to GDP.

We take advantage of this insight and use it to construct an alternative interpretation of our Section 3 decomposition. The allocative efficiency component E^{al} captures the effect of changing consumer taxes; the productive efficiency component E^{pr} captures the effect of changes in wasteful government expenditures net of production taxes or subsidies; the redistributive component R captures the effect of changes in the (after-tax) income distribution, which we denote by $y = \{y_i\}_i$, where $y_i = \sum_k p_k x_{i,k} = \sum_k p_k[a_{i,k} - \tau_{i,k}(x_{i,k} - a_{i,k})] + T_i$. We can express our marginal decomposition (17) that was written using changes in (α, ξ, τ) coordinates using changes in policy instruments (G, T, τ, ς). In the process of doing this, we establish connections between our decomposition and some important objects from consumer-theory and the public-finance literature about effects of taxation.

It is helpful to start from some basic components of consumer theory. Let x be a feasible allocation, and let p and τ be corresponding vectors of producer prices and consumer taxes. The consumption vector x_i can be viewed as solving the optimum problem of a consumer who faces the price vector $q_i \in \mathbb{R}^K$ with $q_{i,k} = (1 + \tau_{i,k})p_k$ and income $m_i = \sum_k q_{i,k} x_{i,k}$. The vector x_i satisfies $x_i = \mathsf{x}_i(q_i, m_i)$, where $\mathsf{x}_i \in \mathbb{R}^K$ denotes the Marshallian demand of consumer i. The derivatives $\{\frac{\partial \mathsf{x}_{i,k}}{\partial m}\}_k$ and $\{\frac{\partial \mathsf{x}_{i,k}}{\partial q_l}\}_{k,l}$, evaluated at (q_i, m_i), are income and uncompensated substitution effects, respectively. We use $\zeta^c_{i,kl}$ to denote the

[11] This equivalence underlies the celebrated "production efficiency" result of Diamond and Mirrlees (1971a), who show that taxing intermediate goods is inferior to taxing final goods.

compensated elasticity of demand for good k with respect to the price of good l. The compensated elasticities can be constructed from income and uncompensated substitution effects via the Slutsky equation. Utility of consumer i at allocation x satisfies $u_i(x_i) = V_i(q_i, m_i)$. Let η_i and $\vartheta^c_{i,k}$ be defined by $\eta_i = \sum_k p_k \tau_{i,k} \frac{\partial \mathsf{x}_{i,k}}{\partial m}$ and $\vartheta^c_{i,k} = \sum_l \tau_{i,l} p_l x_{i,l} \zeta^c_{i,lk}$; these are weighted sums of income and compensated substitution elasticities, with weights given by tax payments.

It is useful first to focus on the case in which the relative Lagrange multipliers $\mathsf{P}(\alpha)$ are independent of α. As discussed in Section 3.1, this case includes a variety of environments: economies in which consumers have identical homothetic preferences, economies with linear technologies, and standard CARA and CRRA models. We will discuss how our results extend to the general case with $\mathsf{P}(\alpha)$ depending on α later in this section. We now state our main result. After that we will offer economic interpretations of the objects that appear in (21) and their connection to the public finance literature.

Proposition 2. *Let x^* be an allocation and define $\omega_i \in \mathbb{R}^I_+$ by $\omega_i = \frac{V_{i,y}/V_{i,yy}}{\sum_{i'} V_{i',y}/V_{i',yy}}$, where derivatives of V_i are evaluated at (p^*, y^*_i). Suppose that the relative Lagrange multipliers $\mathsf{P}(\alpha)$ are independent of α. Then (17) can be written as*

$$\hat{W} = \underbrace{\sum_i \frac{\bar{\alpha}_i u_{i,1}}{1-\eta_i} \big(\widehat{T}^{net}_i - \omega_i \widehat{S}\big)}_{=\hat{\mathsf{R}}} + \underbrace{\sum_i \frac{\bar{\alpha}_i u_{i,1}}{1-\eta_i} \omega_i \widehat{S}}_{=\hat{\mathsf{E}}^{pr}} + \underbrace{\sum_i \frac{\bar{\alpha}_i u_{i,1}}{1-\eta_i} \sum_k \vartheta^c_{i,k} \widehat{\ln(1+\tau_{i,k})}}_{=\hat{\mathsf{E}}^{al}}. \tag{21}$$

Proof. See Appendix A.4.1. □

Consistent with our earlier discussion, changes in consumer taxes $\widehat{\ln(1+\tau)}$ map into the allocative efficiency component of our decomposition, and changes in surplus $\widehat{S}$ map into the productive efficiency component. Redistribution is captured by changes in $\{\widehat{T}^{net}_i - \omega_i \widehat{S}\}_i$, that is, changes in net transfers to agents after an inequality-preserving share of changes in wasteful expenditures net of production taxes or subsidies is accounted for.

To explain the interpretation of the terms ω, η, and ϑ^c that appear in these formulas and the economics that makes our decomposition take form (21), we offer a sequence of examples.

Example 2. Interpretation of weights ω. Consider an economy with a single good. Government policies can be summarized by a vector of lump-sum taxes T and government expenditures on the single good G_1. Consider a perturbation $(\hat{T}, \hat{G}_1)$. Decomposition (21) becomes

$$\hat{W} = \underbrace{\sum_i \bar{\alpha}_i u_{i,1}(\hat{T}_i - \omega_i \widehat{S})}_{=\hat{\mathsf{R}}} + \underbrace{\sum_i \bar{\alpha}_i u_{i,1} \omega_i \widehat{S}}_{=\hat{\mathsf{E}}^{pr}}, \tag{22}$$

where $\widehat{S} = -\hat{G}_1$. Since there is no distinction between consumption and income in a one-good economy, $y_i = x_{i,1}$ and $V_i(1, y_i) = u_i(x_i)$, we can write weight ω_i as $\omega_i \propto \frac{u_{i,1}(x_i^*)}{u_{i,11}(x_i^*)}$, where $u_{i,1}$ and $u_{i,11}$ are first and second derivatives of consumer i's utility function with respect to consumption of the single good.

To interpret this expression, observe that the PN weights α^* that we use to capture the inequality in x^* satisfy $\frac{\alpha_1^*}{\alpha_i^*} = \frac{u_{i,1}(x_i^*)}{u_{1,1}(x_1^*)}$ for all $i > 1$. Inequality, as measured by ratios of marginal utilities, remains unchanged when a change in aggregate resources $\widehat{S}$ is distributed across consumers in proportion to $\frac{u_{i,1}(x_i^*)}{u_{i,11}(x_i^*)}$. Thus, $\hat{\mathsf{E}}^{pr}$ in (22) computes how much welfare would change if a change $\widehat{S} = -\hat{G}_1$ in resources were distributed in a way that keeps inequality unchanged; $\hat{\mathsf{R}}$ captures the welfare effect of departing from this inequality-preserving rule.

This interpretation extends to multi-agent and multi-good economies. There the PN weights α^* satisfy $\frac{\alpha_1^*}{\alpha_i^*} = \frac{V_{i,y}(y_i^*)}{V_{1,y}(y_1^*)}$ for all $i > 1$, so an "inequality-preserving" distribution of additional aggregate income $\widehat{S}$ is described by the weights $\omega_i \propto \frac{V_{i,y}(y_i^*)}{V_{i,yy}(y_i^*)}$ that appear in equation (21). These weights are often easy to compute in closed form. For example, with CRRA preferences the weights are $\omega_i = y_i^*/Y^*$, i.e., the income share of agent i. Thus, inequality is unchanged if additional income is distributed in a way that keeps income shares of all agents unchanged. This distribution of resources also freezes standard statistical measures of inequality including the Gini coefficient and various quantile ratios. When preferences are of CARA form we have $\omega_i = 1/I$, so that inequality is unchanged if aggregate income is distributed equally among consumers.

Example 3. Role of income elasticity η. Consider an economy with a single consumer and linear technology. Suppose that the government finances wasteful purchases G_1^* of good 1 with a combination of consumer taxes and transfers (τ^*, T^*). Suppose the policy reform increases these purchases by $\hat{G}_1$, financed by an increase in the lump-sum tax $\hat{T}_1$ and no changes in consumer taxes, so that $\hat{\tau} = 0$.

Since there is a single consumer, there is no redistribution, so $\hat{\mathsf{R}} = 0$. Since consumer taxes are unchanged, the allocative distortions are unchanged, $\hat{\mathsf{E}}^{al} = 0$, the change in welfare equals the change in the utility of the single consumer, so our decomposition assigns all welfare changes to productive efficiency:

$$\hat{W} = \hat{u}_1 = \underbrace{-\frac{u_{1,1}}{1-\eta_1}\hat{G}_1}_{=\hat{\mathsf{E}}^{pr}}. \tag{23}$$

This equation indicates that since an increase in government expenditures is wasteful, the consumer views it as a loss of efficiency. While government expenditures increase by $\hat{G}_1$, consumer income decreases by $\frac{1}{1-\eta_1}\hat{G}_1$. To interpret $\frac{1}{1-\eta_1}$, suppose for concreteness that income effects and taxes are both positive. If the consumer loses an extra dollar of transfers from the government, an income effect induces her to decrease purchases of consumption goods. Lower purchases of consumption goods translate into lower tax revenues that the government collects through consumption taxes.

By construction, η_1 captures the offsetting revenue losses from consumer 1. Because η_1 dollars are lost through lower consumption-tax receipts, a one-dollar increase in the statutory lump-sum tax raises only $1-\eta_1$ dollars of *net* government revenue. Financing one additional dollar of government spending requires an increase in the lump-sum tax of $1/(1-\eta_1)$, which is the marginal cost of public funds in this example. This is why the income adjustment in equation (23) scales $\hat{G}_1$ by $1/(1-\eta_1)$. Thus, $\hat{\mathsf{E}}^{pr} = -u_{1,1} \times \frac{1}{1-\eta_1} \times \hat{G}_1$ has a simple interpretation: the last term in the product measures the net change in resources that are being wasted; the middle term converts that waste into a change in consumer income by adjusting it for the fiscal externality that emerges if the status quo allocation is already distorted with commodity taxes; then the marginal utility $u_{1,1}$ converts these monetary values into utils.

The same observations carry over to a general multi-good economy. The production efficiency term $\hat{\mathsf{E}}^{pr}$ can be written as $\sum_i \bar{\alpha}_i \left(u_{i,1} \times \frac{1}{1-\eta_i} \times \omega_i \widehat{S} \right)$, where $\widehat{S}$ denotes the increase in surplus due to a change in total government spending, valued at producer prices. The expression inside the parentheses calculates the utility change of consumer i from the inequality-preserving distribution of resources. Consumer i gets an additional $\omega_i \widehat{S}$ of such resources from the government, but the net amount of resources needs to be adjusted for the fiscal externality η_i and converted to utils to compute welfare. Welfare weights $\bar{\alpha}$ are then used to aggregate individual utility changes into the change in social welfare. The same fiscal externality adjustment appears in the other components of our decomposition (21).

Example 4. Role of compensated elasticities ϑ^c. Suppose everything is the same as in the previous single agent example except that now the government uses some arbitrary changes in taxes and transfers $(\hat{\tau}, \hat{T})$ to finance $\hat{G}_1$. In this case, our decomposition becomes

$$\hat{W} = \hat{u}_1 = \underbrace{-\frac{u_{1,1}}{1-\eta_1}\hat{G}_1}_{=\hat{\mathsf{E}}^{pr}} + \underbrace{\frac{u_{1,1}}{1-\eta_1}\sum_k \vartheta^c_{1,k}\ln\widehat{(1+\tau_{1,k})}}_{=\hat{\mathsf{E}}^{al}}. \tag{24}$$

Changes in taxes now affect our measure of allocative efficiency. Public finance theorists call the term $\sum_k \vartheta^c_{1,k}\ln\widehat{(1+\tau_{1,k})}$ the marginal deadweight loss from tax changes or MDWL.[12] The deadweight loss measures the excess burden of distortionary taxation relative to lump-sum taxation, so it naturally appears in the allocative efficiency term. Note that the $\hat{\mathsf{E}}^{al}$ term has a very similar structure to the $\hat{\mathsf{E}}^{pr}$ term and can be written as a product of three terms, $u_{1,1} \times \frac{1}{1-\eta_1} \times MDWL$, where the marginal deadweight loss $MDWL$ is first converted to an effective loss of income for the consumer by adjusting for the fiscal externality, and then to utils by scaling with the marginal utility of consumption.

In constructing decomposition (21) we chose good 1 as the numeraire. But as Proposition 1 shows, our decomposition is independent of the choice of numeraire. Consequently, the ratios $\hat{\mathsf{R}}/\hat{W}$, $\hat{\mathsf{E}}^{pr}/\hat{W}$, and $\hat{\mathsf{E}}^{al}/\hat{W}$ in decomposition (21) are all independent of which good is labeled as $k = 1$.[13]

We derived equation (21) under the assumption that the relative Lagrange multipliers $\mathsf{P}(\alpha)$ are independent of α. When efficient prices $\mathsf{P}(\alpha)$ vary with α, the structure of decomposition (21) is similar, except that we must distinguish between changes in real incomes (i.e., income computed at constant prices) and changes in nominal incomes, then make some adjustments for redistributive effects of tax changes. We discuss such adjustments in Appendix A.4.1.

4.1 On aggregation of efficiency losses

Equation (21) decomposes the utility change $\hat{u}_i$ of each consumer into redistribution and efficiency components, and then aggregates these individual decompositions into a welfare decomposition using social welfare weights $\bar{\alpha}$. This approach contrasts with standard measures of efficiency—such as Kaldor (1939) and Hicks (1939)'s sum of compensated gains,

[12]For example, see Auerbach and Hines (2002). Their expression for the marginal deadweight loss, equation (2.5), equals our term $\sum_k \vartheta^c_{1,k}\ln\widehat{(1+\tau_{1,k})}$. The fiscal externality term $1-\eta_i$ and the marginal deadweight loss term $\sum_k \vartheta^c_{i,k}\ln\widehat{(1+\tau_{i,k})}$ both appear often in the public finance literature, for example in formulas of Diamond and Mirrlees (1971b) and Diamond (1975) that characterize optimal taxes.

[13]While $u_{1,1}$ appears to depend on the numeraire, the product $u_{1,1} \times \frac{1}{1-\eta_1}$ is independent of the numeraire.

Harberger (1971)'s sum of deadweight losses, Debreu (1951)'s coefficient of resource utilization, and more recently Baqaee and Burstein (2025)'s measure of aggregate efficiency—which are constructed without reference to social welfare weights. While such measures may have independent applied interest, we argue in this section that ignoring social weights leads to a misleading aggregation of the different components of the welfare change.

We illustrate this point using Harberger's measure, although the argument applies more broadly. Harberger proposed measuring aggregate efficiency by summing individual deadweight losses using equal money-metric weights.[14] Such deadweight losses are often depicted as "Harberger triangles." In our context, Harberger's measure can be expressed as $\hat{\mathsf{E}}^{harb} = \sum_i \sum_k \vartheta^c_{i,k} \ln \widehat{(1+\tau_{i,k})}$ where recall that $\vartheta^c_{i,k}$ are expenditure-weighted compensated elasticities. The counterpart of this term in our decomposition is $\hat{\mathsf{E}}^{al} = \sum_i \frac{\bar{\alpha}_i u_{i,1}}{1-\eta_i} \sum_k \vartheta^c_{i,k} \ln \widehat{(1+\tau_{i,k})}$. Although both aggregators use $\{MDWL_i\}_i$ as measures of inefficiency, they use them differently. The following example emphasizes advantages of aggregating deadweight losses using welfare weights when constructing welfare decompositions.

Consider again Example 1 from Section 3.4, in which we studied a reform that taxes one agent and rebates the revenue to that same agent. By design, there are no changes in the resources of consumer 1, so $\hat{u}_1 = 0$ and the welfare change is $\hat{W} = \bar{\alpha}_2 \hat{u}_2 < 0$. As discussed in Section 3.4, there is no redistribution of resources across consumers, so from Property (v) of Proposition 1 it follows that $\hat{\mathsf{R}} = 0$. Furthermore, since tax revenues are returned back to households there is no waste, and thus the production-efficiency component $\hat{\mathsf{E}}^{pr} = 0$. It follows that $\hat{\mathsf{E}}^{al} = \bar{\alpha}_2 \hat{u}_2$. Therefore, our decomposition interprets the welfare loss resulting from this tax increase as coming 100% from a reduction in efficiency, regardless of social weights: $\hat{\mathsf{E}}/\hat{W} = 1$ and $\hat{\mathsf{R}}/\hat{W} = 0$ for all $\bar{\alpha}$.

Because Harberger's efficiency measure aggregates without reference to the social welfare function, his measure in a welfare decomposition necessarily makes the contribution of efficiency sensitive to the choice of social weights. To see how this manifests in the present example, define $\hat{W}^{\text{norm}} = \hat{W} \big/ \sum_i \bar{\alpha}_i u_{i,1}$, which expresses the change in welfare in units of the numeraire good. This aligns with Harberger's practice of measuring efficiency in monetary units. Efficiency and redistribution components become $\hat{\mathsf{E}}^{harb}/\hat{W}^{\text{norm}}$

[14]Harberger was aware of the ad hoc nature of this aggregation method. He explicitly writes: "... Giving equal weight to all dollars of income is mathematically the simplest rule, and our data come that way in any event. In a sense, the second obstacle imposes, rather arbitrarily to be sure, a solution to the perplexing difficulties posed by the first. This solution is obviously a far-from-perfect measure of national welfare—indeed it is surprising how little dissatisfaction has been expressed (until quite recently) with its use as such" (Harberger, 1971, pp. 788–789). See also Auerbach (1985) for a discussion of issues related to aggregating deadweight losses with heterogeneous consumers.

and $\hat{\mathsf{R}}^{harb}/\hat{W}^{\text{norm}} = 1 - \hat{\mathsf{E}}^{harb}/\hat{W}^{\text{norm}}$, where $\hat{\mathsf{R}}^{harb}$ is defined residually to make the two components sum to total welfare. By choosing $\bar{\alpha}_2$ sufficiently small, we can make $\hat{\mathsf{E}}^{harb}/\hat{W}^{\text{norm}}$ arbitrarily large and positive (and the redistribution component $\hat{\mathsf{R}}^{harb}/\hat{W}^{\text{norm}}$ arbitrarily large in magnitude and negative). Mechanics like these are not confined to Harberger's measure of efficiency. They apply to any efficiency measure that is constructed without using social weights.[15] Ultimately, the problem with using any such measure is that the resulting decomposition fails to satisfy Property (v) of Proposition 1 which isolates redistribution with changes in real disposable incomes.

We conclude this section by discussing the role that the cardinality of preferences plays in our decomposition. Welfare measures require inter-person comparisons of utilities and consequently depend on cardinality (see, e.g., Sen (1970)). As a result, cardinal measures $\{\bar{\alpha}_i u_{i,1}\}_i$ appear in our welfare decomposition. The cardinality of u_i and the Pareto weights $\bar{\alpha}_i$ are tightly connected: $\bar{\alpha}$ captures societal views about inequality for specific cardinalizations of utility functions. Re-cardinalizing utility functions requires re-cardinalizing the Pareto weights as well.[16]

4.2 Insurance

Proposition 2 applies to both deterministic and stochastic economies, with insurance representing a subcomponent of $\hat{\mathsf{E}}^{al}$. In principle, we can analyze risk by treating consumption in different states of the world as different goods, but for applications this approach is not attractive. Most macroeconomic and finance applications want to frame risk and insurance in terms of statistical measures—variances, covariances, and coefficients of risk aversion. We now discuss our decomposition from that perspective.

To keep our exposition transparent, we focus on an economy with a single physical good.

[15] The Kaldor–Hicks compensation test of Kaldor (1939) and Hicks (1939) involves another example of an efficiency measure that fits into our discussion. It evaluates the change in efficiency between allocations x^* and x^{**} by summing their money-metric willingness to pay to move from bundle x_i^* to x_i^{**} , across individuals i. In addition to the limitations discussed here, this measure has a more fundamental problem: it does not, in general, yield an internally consistent social ordering and its application can lead to inconsistent rankings. See Scitovsky (1941) and Kuznets (1948) for early discussions and Chipman and Moore (1978) for a comprehensive review of this approach.

[16] It is common in the normative literature to assume utilitarian Pareto weights, $\bar{\alpha}_i = 1$ for all i, appealing to a *veil-of-ignorance* argument. Implicitly, these papers also take a stance on the cardinality of the utility functions. For example, consider a commonly used preference specification over consumption c and labor l of the form $v(c) - \nu(l)/\varepsilon_i$, where ε_i is an exogenous parameter capturing heterogeneity. This preference specification is equivalent, in the ordinal sense, to preferences of the form $\varepsilon_i v(c) - \nu(l)$, as these two utility functions capture the same preference orderings over all possible bundles (c, l) for each consumer i. Despite their observational equivalence, these two preference specifications lead to very different normative conclusions when welfare is evaluated using equal weights $\bar{\alpha}_i = 1$ for all i.

Throughout the main text we assume CRRA expected utility $u_i(x_i) = \sum_s \Pr(s)\, v(x_i(s))$ with $v(c) = \frac{c^{1-\sigma}}{1-\sigma}$; we prove a parallel CARA analysis in Appendix A.4.2.

Idiosyncratic risk. Using the notation of Section 3.5, in a single-good environment, an allocation takes the form $x = \{x_{i,1}(s)\}_{i,s}$. To model idiosyncratic shocks, interpret each i as a group of ex-ante identical consumers in which, ex post, a fraction $\Pr(s)$ of members of group i receive endowment $a_i(s)$.[17] Since shocks are purely idiosyncratic, the aggregate endowment $A_1 = \sum_i \mathbb{E}[a_i]$ of the single good is constant. We study a policy that reallocates these endowments across groups to produce an allocation $x = \{x_i(s)\}_s$ with $\sum_i \mathbb{E}x_{i,1} = A_1$. Let $\hat{x}$ denote a marginal reform that satisfies aggregate feasibility $\sum_i \mathbb{E}\hat{x}_{i,1} = 0$.

For CRRA preferences, supporting prices $p = \{p_1(s)\}_s$ can be set equal to state probabilities, so that $p_1(s) = \Pr(s)$; see Section 3.1. Hence, the income of consumer i equals her expected consumption, $y_i = \mathbb{E}x_{i,1}$, and aggregate income equals aggregate consumption, $Y = X_1 = \sum_i \mathbb{E}x_{i,1}$.

The PN weights satisfy $\frac{\alpha_i}{\alpha_1} = \left(\frac{\mathbb{E}x_{i,1}/X_1}{\mathbb{E}x_{1,1}/X_1}\right)^{\sigma}$, so inequality is determined entirely by the dispersion of expected consumption shares. Allocative wedges are $\frac{1+\tau_{i,1}(s)}{1+\tau_{i,1}(1)} = \left(\frac{x_{i,1}(s)/\mathbb{E}x_{i,1}}{x_{i,1}(1)/\mathbb{E}x_{i,1}}\right)^{-\sigma}$, where we normalize state $s = 1$ as the reference state. These wedges measure deviations of realized consumption from expected consumption, thereby capturing incomplete insurance.

There are no production distortions, so welfare changes operate solely through redistribution and insurance. Applying our decomposition (see Appendix A.4.2 for detailed derivation) yields

$$\hat{W} = \underbrace{\sum_i \bar{\alpha}_i\, \mathbb{E}\left[v'(x^*_{i,1})\, x^*_{i,1} \ln\widehat{\left(\frac{\mathbb{E}x_{i,1}}{X_1}\right)}\right]}_{=\hat{\mathsf{R}}} + \underbrace{\sum_i \bar{\alpha}_i\, \mathbb{E}\left[v'(x^*_{i,1})\, x^*_{i,1} \ln\widehat{\left(\frac{x_{i,1}}{\mathbb{E}x_{i,1}}\right)}\right]}_{=\hat{\mathsf{E}}^{ins}}. \tag{25}$$

The term in the first braces labeled $\hat{\mathsf{R}}$ captures changes in expected consumption shares holding the pattern of state-by-state deviations fixed. The insurance term $\hat{\mathsf{E}}^{ins}$ captures changes in the state-by-state profile of consumption risk, holding each $\mathbb{E}x_{i,1}$ fixed. Thus, mean-preserving spreads of consumption are interpreted as 100% efficiency losses. This reflects Part (v) of Proposition 1, since the value of the consumption bundle here is simply expected consumption.

[17] A formal treatment of idiosyncratic shocks uses a continuum of agents and a well-defined stochastic process; see Uhlig (1996). Our discrete-type representation produces identical allocations and decompositions.

A second-order approximation of (25) yields

$$\hat{W} \approx \underbrace{\sum_i \bar{\alpha}_i \left(\mathbb{E}x^*_{i,1}\right)^{1-\sigma} \widehat{\ln \mathbb{E}x_{i,1}}}_{=\hat{\mathsf{R}}} \; - \; \underbrace{\frac{\sigma}{2} \sum_i \bar{\alpha}_i \left(\mathbb{E}x^*_{i,1}\right)^{1-\sigma} \widehat{\mathrm{var}(\ln x_{i,1})}}_{=\hat{\mathsf{E}}^{ins}}, \tag{26}$$

highlighting that insurance for consumer i is measured by the change in the variance of $\ln x_{i,1}$ scaled by the coefficient of relative risk aversion σ and the level of expected consumption.

Aggregate risk. Contrast the preceding idiosyncratic case with an economy exposed to aggregate shocks. The feasibility constraint now requires $\sum_i x_{i,1}(s) = A_1(s)$ for every aggregate state s, where $A_1(s)$ is the aggregate endowment of the single physical good in state s. Following Section 3.1, the supporting price vector satisfies $p_1(s) \propto \Pr(s)\, A_1(s)^{-\sigma}$. Without loss of generality we can normalize these prices so that $\sum_s p_1(s) = 1$. With complete markets, these prices correspond to the canonical risk-neutral probability measure. Let $\widetilde{\mathbb{E}}$ denote a mathematical expectation computed with respect to this measure.

The previous analysis, which focused on idiosyncratic risk, extends once one replaces physical probabilities $\Pr(s)$ with the risk-neutral weights $p_1(s)$. Thus, the income of household i is $y_i = \sum_s p_1(s) x_{i,1}(s) = \widetilde{\mathbb{E}}\, x_{i,1}$. Substituting risk-neutral expectations into the decomposition (25) yields the following decomposition for aggregate shocks:

$$\hat{W} = \underbrace{\sum_i \bar{\alpha}_i \, \mathbb{E}\left[v'(x^*_{i,1})\, x^*_{i,1} \ln \widehat{\left(\frac{\widetilde{\mathbb{E}} x_{i,1}}{\widetilde{\mathbb{E}} X_1} \right)} \right]}_{=\hat{\mathsf{R}}} \; + \; \underbrace{\sum_i \bar{\alpha}_i \, \mathbb{E}\left[v'(x^*_{i,1})\, x^*_{i,1} \ln \widehat{\left(\frac{x_{i,1}}{\widetilde{\mathbb{E}} x_{i,1}} \right)} \right]}_{=\hat{\mathsf{E}}^{ins}}. \tag{27}$$

See Appendix A.4.2 for a detailed derivation.

To express this decomposition in terms of physical probabilities and covariances, use the approximation $\widetilde{\mathbb{E}} Z \; \approx \; \mathbb{E} Z \; - \; \sigma \,\mathrm{cov}(Z, \ln X_1)$, which follows from the CRRA pricing kernel. Substituting this into (27) yields the second-order approximation

$$\begin{aligned} \hat{W} \approx & \underbrace{\sum_i \bar{\alpha}_i \left(\mathbb{E}x^*_{i,1}\right)^{1-\sigma} \left[\widehat{\ln \mathbb{E}x_{i,1}} - \sigma \,\mathrm{cov}\Big(\widehat{\ln x_{i,1}}, \ln X_1\Big) \right]}_{=\hat{\mathsf{R}}} \\ & - \underbrace{\frac{\sigma}{2} \sum_i \bar{\alpha}_i \left(\mathbb{E}x^*_{i,1}\right)^{1-\sigma} \mathrm{var}(\ln \widehat{x_{i,1} - \ln} X_1)}_{=\hat{\mathsf{E}}^{ins}}. \end{aligned}$$

Relative to decomposition (26) for idiosyncratic risk, the key difference is the appearance of covariance terms $\text{cov}\left(\widehat{\ln x_{i,1}}, \ln X_1\right)$ that measure how reforms tilt consumption toward "good" or "bad" aggregate states. A reform that increases consumption precisely in states where aggregate consumption $X_1(s)$ is low provides a valuable hedge because marginal utility is highest in those states. The covariance terms measure this exposure and appear scaled by σ, so they can be interpreted as an *effective risk premium* associated with shifting consumption toward bad aggregate states. When a reform reduces one agent's exposure to aggregate risk while increasing another's, the decomposition records it as redistribution: a valuable hedging opportunity is being transferred across agents.

The insurance term here reflects changes in the variance of $\ln x_{i,1}(s) - \ln X_1(s)$, i.e., the variance of individual consumption *relative* to aggregate consumption. This reflects that with CRRA preferences and aggregate shocks, "perfect insurance" means that individual consumption co-moves one-for-one with aggregate consumption, making relative consumption constant across states.

This analysis can be extended to apply when the aggregate feasibility condition includes wasteful government purchases $G = \{G_1(s)\}_s$. Such expenditures create a wedge between $A_1(s)$ and $X_1(s)$ and enter through the production-efficiency term $\hat{\mathsf{E}}^{pr} = \sum_i \bar{\alpha}_i \, \mathbb{E}\left[v'(x^*_{i,1})\, x^*_{i,1} \, \widehat{\ln(\widetilde{\mathbb{E}} X_1)}\right]$. The expressions for $\hat{\mathsf{R}}$ and $\hat{\mathsf{E}}^{ins}$ in (27) remain unchanged, yielding $\hat{W} = \hat{\mathsf{R}} + \hat{\mathsf{E}}^{pr} + \hat{\mathsf{E}}^{ins}$. This decomposition also follows directly from the identity $x_{i,1}(s) = \frac{\widetilde{\mathbb{E}} x_{i,1}}{\widetilde{\mathbb{E}} X_1} \times \widetilde{\mathbb{E}} X_1 \times \frac{x_{i,1}(s)}{\widetilde{\mathbb{E}} x_{i,1}}$, which in terms of log differences implies

$$\widehat{\ln x_{i,1}(s)} = \widehat{\ln\left(\frac{\widetilde{\mathbb{E}} x_{i,1}}{\widetilde{\mathbb{E}} X_1}\right)} + \widehat{\ln \widetilde{\mathbb{E}} X_1} + \widehat{\ln\left(\frac{x_{i,1}(s)}{\widetilde{\mathbb{E}} x_{i,1}}\right)},$$

the counterpart of the idiosyncratic expansion. The CARA case follows similar steps, except that inequality and insurance are expressed in levels rather than logs, and are scaled by the coefficient of absolute risk aversion. Details are provided in Appendix A.4.2.

5 Comparisons with alternative decompositions

An early welfare decomposition was developed by Benabou (2002), who wanted to understand the sources of welfare gains from changes in tax policy in a particular economic environment. Floden (2001) extended Benabou's decomposition to more general settings. Floden (2001)'s version has subsequently been used by various researchers. We refer to it as the BF decomposition. More recently, Dávila and Schaab (2022) developed an alternative that we refer to as the DS decomposition. To compare these decompositions, we focus on an economy in which consumers have identical preferences $u(c_i, \ell_i)$ over consumption and leisure.[18] For now, we assume that there is no uncertainty. We add uncertainty later in this section.

Let $\bar{c} = \frac{1}{I}\sum_i c_i$ and $\bar{\ell} = \frac{1}{I}\sum_i \ell_i$ denote average consumption and leisure per capita. Floden defines his measure of efficiency ω_{E} as a solution of $u\big((1+\omega_{\mathsf{E}})\bar{c}^*, \bar{\ell}^*\big) = u\big(\bar{c}^{**}, \bar{\ell}^{**}\big)$. Thus, $1 + \omega_{\mathsf{E}}$ measures how many units of the numeraire consumption good a consumer would have to be compensated to be willing to hold the aggregate per-capita bundle $(\bar{c}^{**}, \bar{\ell}^{**})$ instead of $(\bar{c}^*, \bar{\ell}^*)$. Since ω_{E} is measured in units of the numeraire good, Floden also defines the total change in welfare ω_W in units of the numeraire good as a solution of $\sum_i \bar{\alpha}_i u((1+\omega_W)c_i^*, \ell_i^*) = \sum_i \bar{\alpha}_i u(c_i^{**}, \ell_i^{**})$. He defines his efficiency component, the equivalent of our ratio $\frac{\mathsf{E}(x^*, x^{**})}{W(x^{**}) - W(x^*)}$, as $\frac{\ln(1+\omega_{\mathsf{E}})}{\ln(1+\omega_W)}$. In deterministic economies, Floden's redistribution component is the residual $1 - \frac{\ln(1+\omega_{\mathsf{E}})}{\ln(1+\omega_W)}$.

Dávila and Schaab start by constructing a marginal decomposition. For allocation x and its marginal change $\hat{x}$, they define $dc_i := \hat{c}_i + \frac{u_\ell(c_i,\ell_i)}{u_c(c_i,\ell_i)}\hat{\ell}_i$. The term dc_i expresses the change $\hat{u}_i$ in agent i's utility in units of the numeraire consumption good. DS define their measure of efficiency $d\bar{c}$ as the sum of these valuations across consumers, $d\bar{c} := \frac{1}{I}\sum_i dc_i$. Because this object is measured in units of the numeraire, DS also express the marginal change dW in social welfare in the same units, so $dW := \hat{W}/(\sum_i \bar{\alpha}_i u_{i,c})$. DS's marginal efficiency component is then $d\bar{c}/dW$, and the marginal redistributive component is $1 - d\bar{c}/dW$. To construct a global decomposition, DS parametrize perturbations using a scalar $\theta \in [0, 1]$, where $\theta = 0$ corresponds to the status quo and $\theta = 1$ to the alternative allocation; they then accumulate their marginal assessments using a line integral along this path.

The BF and DS decompositions both define efficiency in terms of aggregates, ω_{E} and $d\bar{c}$, that use no social weights and are measured in consumption units. Consequently, both approaches are subject to our Section 4.1 criticism of Harberger's approach. Moreover, as we show next, the BF and DS decompositions do not satisfy our Section 2 consistency requirements and therefore have difficulty in identifying efficiency losses from introducing distortions. We illustrate this in the following two-agent "worker–capitalist" economy.

Example 5. Consider an economy with two types of households that live for one period. A representative firm uses a linear production technology that converts labor into consumption.

[18]Both Benabou (2002) and Floden (2001) require identical preferences for all agents. In addition, Floden (2001) requires preferences to satisfy certain auxiliary conditions, which we assume to hold throughout this section.

Each household is endowed with one unit of time that can be allocated between labor and leisure. Households of type 1 have zero labor productivity, while households of type 2 have unit labor productivity. In addition, each household $i \in \{1, 2\}$ has an initial endowment a_i of the consumption good. We interpret $a_i > 0$ as net savings and $a_i < 0$ as net debt. To ensure that the utility of the unproductive household (type 1) is well-defined, we assume that $a_1 > 0$. We refer to type 1 households as "capitalists," since their income derives entirely from their asset holdings, and to type 2 households as "workers," since their income comes from supplying labor.

Let allocation x^* emerge from the *laissez-faire* competitive equilibrium of this economy. We study policies that impose a lump-sum tax $T < a_1$ on the capitalist and redistribute its proceeds to the worker. We consider two redistribution schemes called Policy I and Policy II. Under Policy I, the worker receives a lump-sum transfer. Under Policy II, the worker receives the same amount, but now via a proportional labor subsidy.[19] Let x_T^I and x_T^{II} denote equilibrium allocations under these two policies. By construction, the baseline *laissez-faire* allocation with $T = 0$ satisfies $x^* = x_0^I = x_0^{II}$. For each T, the allocation x_T^I lies on the Pareto frontier. At a given T, the capitalist attains the same utility under both policies, whereas the worker is strictly better off under Policy I. Hence x_T^I Pareto-dominates x_T^{II} for all $T \neq 0$.

Though highly stylized, Example 5 captures tax policy trade-offs that arise in canonical heterogeneous-agent economies. Households with substantial wealth (capitalists) supply little labor and finance consumption from asset incomes, while others (workers) supply a lot of labor, hold little wealth, and are often in debt. Motivated by concerns about inequality, a researcher may want to tax the wealth or capital income of the rich, since these taxes are essentially non-distortionary in the short run. The revenue gathered can be used to support poorer households through transfer payments or reductions in marginal labor-income taxes.

We begin with Policy I, which redistributes using lump-sum transfers. By construction, this policy reform is efficient and moves the equilibrium allocation along the Pareto frontier. Hence, by Weak Pareto-consistency, our decomposition gives $E(x^*, x_T^I) = 0$ for all T, so all welfare changes are attributed to redistribution.

The DS decomposition reaches the same conclusion as long as the initial allocation x^* is an interior solution. In that case, DS's statistic dc_i correctly captures the marginal transfer

[19] Our setup allows T to take any sign. If $T < 0$, these policies transfer resources to the capitalist, either through a lump-sum tax on the worker or through a proportional tax on their labor income. The analysis is symmetric in the sign of T.

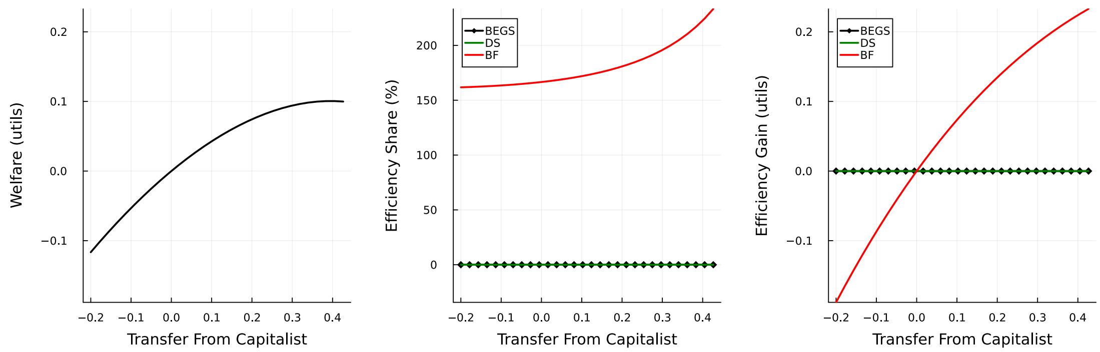

Figure 1: Welfare decomposition in Example 5 for Policy I (redistribution using lump-sum transfers). Type $i = 1$ households are the "capitalists". The x-axis is the transfer T from the capitalists and the y-axis is (a) social welfare (utils), (b) efficiency share as a percentage of the total welfare change, and (c) efficiency gains (utils). To construct this decomposition, we assumed logarithmic preferences, utilitarian social welfare weights, unit time endowment, and initial assets $a_1 = 1.1$ and $a_2 = -0.4$.

received by household type i, and since net transfers sum to zero, their efficiency term $d\bar{c}$ is also equal to zero.

In contrast, there is no reason to expect that the BF approach will arrive at the same conclusion. The BF efficiency component ω_{E} is constructed by applying the common utility function to hypothetical allocations and can be zero, positive or negative for these reforms that move along the Pareto frontier. To see this, consider preferences of the form used by Benabou (2002):

$$\ln c - \frac{2}{1+1/\gamma}(1-\ell)^{1+1/\gamma}, \tag{28}$$

For these preferences, it can be shown analytically that $\omega_{\mathsf{E}} > 0$ ($\omega_{\mathsf{E}} < 0$) for all $T > 0$ ($T < 0$) if the worker has debt. Appendix A.5 provides the derivations. We summarize this discussion using Figure 1 for an illustrative set of parameters. In panel (a) we plot welfare change $W(x_T^I) - W(x^*)$ from Policy I for various levels of T. In panel (b) we plot the efficiency component $\frac{\mathsf{E}(x^*, x_T^I)}{W(x_T^I) - W(x^*)}$ for our decomposition, the BF decomposition, and the DS decomposition, and in panel (c) we plot the $\mathsf{E}(x^*, x_T^I)$ term for all three decompositions expressed for comparability in the same units (utils).[20] Our decomposition and the DS decomposition coincide, while the BF decomposition attributes a large fraction of the welfare change to efficiency changes and violates Weak Pareto-consistency.

We now turn to Policy II. Unlike Policy I, this reform introduces a labor wedge and there-

[20]The assumption of log utility implies that the BF decomposition is already expressed in utils. To convert the DS decomposition to utils, we use their global procedure to compute the fraction of welfare gain attributed to redistribution and efficiency, then multiply by the total welfare gain.

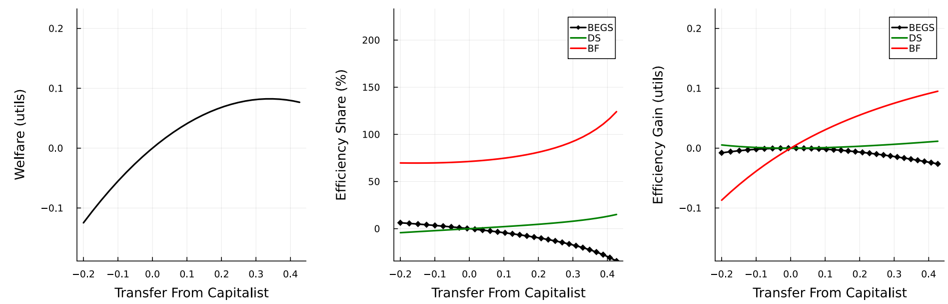

Figure 2: Welfare decomposition in Example 5 for Policy II (redistribution using proportional subsidy/tax). Type $i = 1$ households are the "capitalists". The x-axis is the transfer T from the capitalists and the y-axis is (a) social welfare (utils), (b) efficiency share as a percentage of the total welfare change, and (c) efficiency gains (utils). All parameters are the same as in Figure 1.

fore moves the equilibrium allocation from the Pareto frontier into its interior. By Strong Pareto-consistency, our decomposition correctly identifies an efficiency loss, so $E(x^*, x_T^{II}) < 0$ for all $T \neq 0$. In contrast, neither the BF nor the DS decomposition is guaranteed to detect efficiency losses in this setting.

To illustrate, suppose preferences are given by (28) and the worker is initially in debt. In this case, the BF decomposition finds efficiency losses for $T < 0$ but efficiency gains for $T > 0$, while the DS decomposition finds efficiency gains for all T. See Appendix A.5 for derivations. Figure 2, the analogue of Figure 1 for Policy II, shows these patterns. Panel (c) demonstrates that our efficiency term is always negative, DS's measure is always positive, and BF's measure switches sign at $T = 0$.

The shapes in panel (c) are also revealing. Our efficiency component is inverse-U-shaped: marginal efficiency losses become larger when the existing transfer T, and therefore the underlying labor subsidy (or tax), is already large. This reflects the standard result that deadweight losses of taxation are convex in the tax rate. Formally, the local decomposition in (21) shows that the efficiency term is proportional to $\zeta_2^c \tau^* \widehat{\ln(1+\tau)}$, where τ^* is the labor-subsidy rate associated with T^*, $\widehat{\ln(1+\tau)}$ is the marginal change in that rate, and $\zeta_2^c < 0$ is the compensated labor-supply elasticity of agent 2. By contrast, the DS efficiency component is U-shaped, implying not only that distortionary taxes generate efficiency gains, but also that these gains increase with the size of the existing distortion.[21]

[21] In the example we discuss, the global DS decomposition violates Strong Pareto-consistency. DS primarily focus on marginal decompositions in their analysis. One can formulate a local version of Strong Pareto-consistency that is appropriate for marginal decompositions: the derivative of the marginal efficiency in any feasible direction is non-positive at allocations on the Pareto frontier. The marginal DS decomposition fails this local version of Strong Pareto-consistency. This failure is reflected in their U-shaped efficiency component, which attains its minimum on the frontier—implying that, locally, the marginal efficiency component

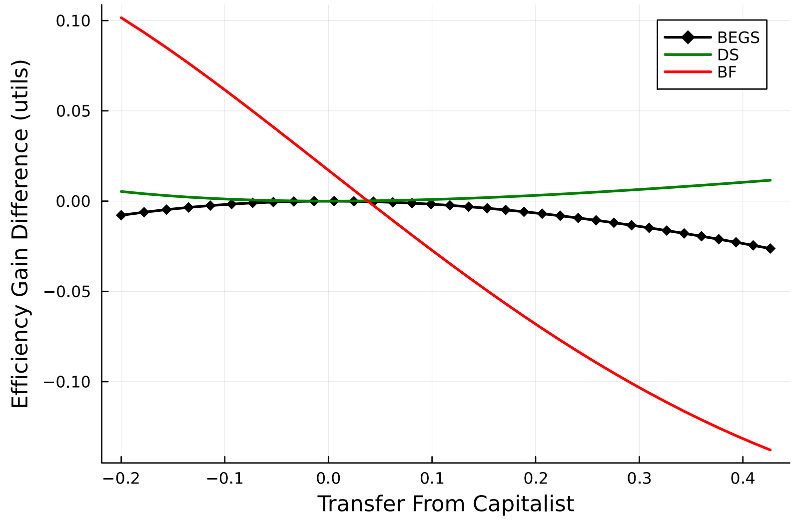

Figure 3: Efficiency comparison of Policy I (lump-sum transfers) and Policy II (proportional subsidy/tax) in Example 5. Type $i = 1$ households are the "capitalists". The x-axis is the transfer T from the capitalists and the y-axis is $\mathsf{E}(x^*, x_T^{II}) - \mathsf{E}(x^*, x_T^I)$. All parameters are the same as in Figure 1.

These differences matter for policy evaluation. Suppose that we compare several redistribution schemes that transfer the same amount of resources from capitalists to workers and want to use a welfare decomposition to indicate the most efficient one. A natural approach is to compare efficiency components E across reforms and choose the policy with the smallest efficiency loss. Figure 3 plots $\mathsf{E}(x^*, x_T^{II}) - \mathsf{E}(x^*, x_T^I)$. Because both policies deliver the same redistribution but Policy I is implemented in a non-distortionary manner, our decomposition finds this difference negative for all T, indicating that Policy I is always more efficient. In contrast, the DS decomposition ranks Policy II as more efficient for all T, while the BF decomposition can rank either policy as superior depending on the value of T.

5.1 Insurance

As in our approach, both BF and DS separate welfare gains attributable to an "insurance" component from welfare gains attributable to overall efficiency. In this section, we show how the concerns we highlighted in deterministic environments carry over to insurance components in stochastic settings. We briefly review how BF and DS construct their insurance measures and then present a simple example.

To measure the contribution of insurance, Floden proposes the following procedure. For

is *increasing* in distortions.

each agent i, define the certainty equivalent c_i^{ce} by solving $u(c_i^{ce}, \bar{\ell}) = \mathbb{E}[u(c_i, \ell_i)]$. Aggregate these certainty equivalents as $\bar{c}^{ce} = \frac{1}{I}\sum_i c_i^{ce}$. Then compute the proportional consumption loss p_{unc} from $u((1-p_{unc})\bar{c}, \bar{\ell}) = u(\bar{c}^{ce}, \bar{\ell})$ for both allocations x^* and x^{**}. The insurance measure $\omega_{\mathsf{E}^{ins}}$ solves $1 + \omega_{\mathsf{E}^{ins}} = \frac{1-p_{unc}^{**}}{1-p_{unc}^{*}}$. The insurance component of the BF decomposition is $\frac{\ln(1+\omega_{\mathsf{E}^{ins}})}{\ln(1+\omega_W)}$, the analogue of our ratio $\frac{\mathsf{E}^{ins}(x^*,x^{**})}{W(x^{**})-W(x^*)}$. In the DS approach, the marginal contribution of insurance is $dIns = \frac{1}{I}\sum_i \mathrm{cov}\left(\frac{u_{i,c}}{\mathbb{E}[u_{i,c}]}, dc_i - d\bar{c}\right)$, and the insurance component is given by $dIns/dW$. The next example illustrates problems that arise with both approaches.

Example 6. Consider the capitalist–worker environment of Example 5, but now suppose that workers face idiosyncratic productivity shocks. Assume that markets are incomplete and that in the *laissez-faire* equilibrium workers cannot trade Arrow-Debreu securities to insure these shocks. Consider a policy that imposes a proportional tax $\tau > 0$ on labor earnings and that rebates revenues to workers in a lump-sum fashion. Because capitalists do not supply labor, the reform does not affect them.

This intervention is the type of imperfect insurance policy commonly studied in heterogeneous-agent models (see, e.g., Aiyagari and McGrattan (1998); Benabou (2002); Heathcote et al. (2017), and the application in Section 6). The policy reduces the volatility of workers' after-tax income but introduces labor-supply distortions, so its welfare effect reflects a trade-off between improved insurance and higher marginal tax wedges. To focus on the insurance channel, we abstract from the labor-supply distortions by taking $\gamma \to 0$, which makes labor supply inelastic. Under this assumption, taxes do not distort behavior, and any $\tau \in (0,1]$ Pareto-dominates the *laissez-faire* allocation by providing better risk-sharing among the workers. Let us evaluate the fraction of welfare gains attributable to insurance using our decomposition, the BF decomposition, and the DS decomposition.

All three decompositions classify this reform as free of production-efficiency losses and thereby attribute welfare gains to some combination of insurance and redistribution. Because the reform does not alter the allocation of resources between workers and capitalists, our decomposition assigns the entire welfare gain to insurance for any τ and any choice of social weights $\bar{\alpha}$. This is ultimately a consequence of Part (v) of Proposition 1. In contrast, the BF and DS decompositions are highly sensitive to $\bar{\alpha}$. Figure 4 illustrates that, depending on the social weights used to evaluate welfare, the BF and DS insurance components can be made arbitrarily small or arbitrarily large and close to one. Their redistributive components—which in this setting are simply one minus the insurance component—can likewise take values arbitrarily close to zero or one.

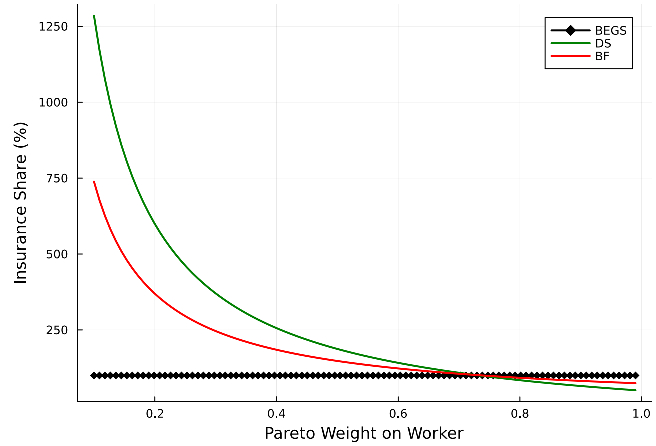

Figure 4: Insurance share of the welfare decomposition in Example 6 for a reform with $\tau = 100\%$. The "worker" is type $i = 2$ household. The x-axis shows the Pareto weight assigned to the worker in evaluating social welfare. The y-axis shows the percentage of welfare gains attributed to insurance by each decomposition. To construct this decomposition, we assumed risk aversion coefficient $\sigma = 1$, labor supply elasticity parameter $\gamma = 0.5$, a log-normal distribution for idiosyncratic productivity shocks with standard deviation of 0.1, and an initial asset distribution given by $a_1 = 10$ and $a_2 = 0$.

The issue highlighted in this example is closely related to our Section 4.1 argument that Harberger-type aggregated deadweight-loss measures are not suitable for welfare decompositions.[22] Aggregating individual insurance gains by simple summation implicitly takes a stand on social weights. However, those implicit weights are not internally consistent across insurance and redistribution components, and they do not align with the social weights used to evaluate total welfare.

5.2 Dependence on numeraire

In Proposition 1, we emphasized that our decomposition is invariant to the choice of numeraire good. In contrast, both the BF and DS decompositions depend on the choice of numeraire. That dependence affects not only the magnitudes of measured efficiency and insurance gains but also their signs, which can change when the numeraire is changed. We illustrate this point with a simple example using the DS decomposition; an analogous example can be constructed for the BF decomposition as well.

Example 7. Suppose that there are two consumers with preferences $u_i = \ln c_i + \ln \ell_i$ over consumption and leisure. Let labor productivity be one for both consumers, so that an

[22] This general property also applies to decompositions based on Debreu (1951)'s coefficient of resource utilization. See Section A.1.7 for more details

additional unit of leisure corresponds to one unit of forgone consumption. Consider an allocation x^* given by $(c_1^*, \ell_1^*) = (\frac{1}{2}, \frac{1}{2})$ and $(c_2^*, \ell_2^*) = (\frac{3}{4}, \frac{1}{2})$. Consider a reform that changes the allocation to $(\frac{1}{2} - 2\varepsilon, \frac{1}{2} + \varepsilon)$ and $(\frac{3}{4}, \frac{1}{2} + \varepsilon)$ for a small $\varepsilon > 0$. Depending on whether c or ℓ is the numeraire, the efficiency component in the DS decomposition flips sign.

This example is deliberately constructed to impose symmetries between c and ℓ in preferences and technology. Despite those symmetries, an arbitrary choice of numeraire reverses the qualitative interpretation of the same reform: using c as the numeraire makes the reform appear efficiency-enhancing, while using ℓ makes it appear efficiency-reducing.

6 Quantitative applications

In this section, we examine two fiscal reforms in calibrated incomplete-markets economies. The first is a standard income-tax experiment in which a permanent increase in distortionary taxes is financed with lump-sum transfers. The second reform implements the robust Pareto improvement of Aguiar et al. (2024). For both applications, we decompose the sources of welfare gains—both in aggregate and across the distribution. We begin by outlining the environment.

Environment. Time is discrete and infinite. A unit measure of households face idiosyncratic labor-productivity shocks. Each household i enters period $t = 0$ with assets $a_{i,0}$ and an idiosyncratic productivity state $\epsilon_{i,0} = (\epsilon_{i,0}^P, \epsilon_{i,0}^T)$ consisting of persistent and transient components that evolve according to

$$\begin{aligned} \log \epsilon_{i,t}^P &= \rho^P \log \epsilon_{i,t-1}^P + \eta_{i,t}^P, \\ \log \epsilon_{i,t}^T &= \eta_{i,t}^T, \end{aligned}$$

where $\{\eta_{i,t}^P, \eta_{i,t}^T\}$ are independent Gaussian shocks, and $\rho^P \in (0, 1)$.

Household i chooses consumption $c_{i,t}$, labor $l_{i,t}$ and next period's assets $a_{i,t+1}$ to solve

$$\max_{\{c_{i,t}, l_{i,t}, a_{i,t+1}\}_{i,t}} \mathbb{E}_0 \sum_{t \geq 0} \beta^t \left[\frac{c_{i,t}^{1-\sigma}}{1-\sigma} - \Psi \frac{l_{i,t}^{1+\gamma}}{1+\gamma} \right],$$

subject to

$$c_{i,t} + a_{i,t+1} = (1 - \tau_t) \big[r_t a_{i,t} + w_t \epsilon_{i,t}^P \epsilon_{i,t}^T l_{i,t} \big] + a_{i,t} + T_t, \qquad a_{i,t+1} \geq \underline{a},$$

given $(a_{i,0}, \epsilon_{i,0})$. Here τ_t is a proportional tax on total household income and T_t is a uniform lump-sum transfer.

A representative firm operates a Cobb–Douglas technology $Y_t = AK_t^{\theta} L_t^{1-\theta}$ and hires capital and labor at factor prices $\{(1+\varsigma_t)r_t,\, w_t\}_t$, where ς_t is a tax on capital services used by the firm.

Aggregate variables are obtained by summing across households:

$$A_t = \int a_{i,t} di, \qquad L_t = \int \epsilon_{i,t}^P \epsilon_{i,t}^T l_{i,t} di, \qquad C_t = \int c_{i,t} di.$$

Aggregate savings are allocated between capital and government debt, $K_{t+1} + B_t = A_{t+1}$, and the resource constraint is

$$Y_t = C_t + G_t + K_{t+1} - (1-\delta)K_t.$$

Given an initial distribution μ_0 over $(a_{i,0}, \epsilon_{i,0})$ and a fiscal policy $\{\tau_t, T_t, \varsigma_t, B_t, G_t\}_{t\geq 0}$, a competitive equilibrium is an allocation $\{c_{i,t}, l_{i,t}, a_{i,t+1}\}_{i,t}$ and a sequence of prices $\{r_t, w_t\}_t$ for which households and firms optimize, markets clear, and the government budget constraint holds at every t.

6.1 Income-tax reform

We begin by studying a one-time, permanent change in the income tax rate τ, financed by adjusting the lump-sum transfer $\{T_t\}_t$. We set $B_t = B$, $G_t = G$, $\varsigma_t = 0$, and $\tau_t = \tau^*$ and compute the associated invariant distribution $\mu^{ss}(\tau^*)$. Our status-quo allocation is the stationary equilibrium corresponding to a policy (τ^*, B, G). We want to study the welfare effects of once and for all moving to a new policy τ^{**} associated with a new equilibrium with initial distribution $\mu_0 = \mu^{ss}(\tau^*)$ and policies (τ^{**}, B, G), where the path of lump-sum transfers $\{T_t\}_t$ is adjusted so that the government budget constraint holds at all t.

This experiment is a stylized version of a standard experiment for evaluating income-tax reforms in an incomplete-markets model in which the policy τ^* represents the current U.S. system, and the invariant distribution $\mu^{ss}(\tau^*)$ is taken to approximate the current U.S. economy. The reform τ^{**} represents either a specific alternative or an "optimal" policy under a given welfare criterion. It is well understood that such tax reforms redistribute resources across households, provide insurance against idiosyncratic shocks, and distort incentives to work and save.[23] Our objective is to quantify the relative importance of these sources of

[23]See Huang et al. (1997), Aiyagari and McGrattan (1998), Conesa and Krueger (1999), Floden (2001),

welfare effects.

Decomposition. We apply the full decomposition provided in equation (19). Physical goods are indexed by time t and idiosyncratic states are indexed by histories s^t, so an allocation is $x = \{c_{i,t}(s^t),\, l_{i,t}(s^t)\}_{i,t,s^t}$. Implementing the decomposition requires computing coordinates $(\alpha(x),\, \xi(x),\, \tau(x))$, where $\tau(x)$ consists of goods wedges $\tau^g(x) := \{\tau^g_{i,c,t},\, \tau^g_{i,l,t}\}_{i,t}$ and insurance wedges $\tau^{ins}(x) := \{\tau^{ins}_{i,c,t}(s^t),\, \tau^{ins}_{i,l,t}(s^t)\}_{i,t,s^t}$, as well as evaluating social welfare $\mathcal{W}(\alpha, \xi, \tau)$ for counterfactual allocations in which one coordinate is varied at a time.

For given PN weights α, supporting prices $p(\alpha) := \{q^{PF}_t(\alpha)\Pr(s^t),\, w^{PF}_t(\alpha)\Pr(s^t)\}_{t,s^t}$ of date t consumption and effective-leisure are obtained by maximizing α-weighted utilities subject to aggregate feasibility, and reading off period t prices (q^{PF}_t, w^{PF}_t) from the no-arbitrage conditions of the representative firm. Let $\varepsilon_{i,t}(s^t) := \epsilon^P_{i,t}(s^t)\epsilon^T_{i,t}(s^t)$ denote efficiency units of labor, and $\bar{\ell}$ the time endowment. These prices imply lifetime incomes

$$y_i(p; x) := \sum_t \sum_{s^t} \left[q^{PF}_t \Pr(s^t)\, c_{i,t}(s^t) + w^{PF}_t \Pr(s^t)\, \varepsilon_{i,t}(s^t)(\bar{\ell} - l_{i,t}(s^t))\right]$$

and associated indirect utilities $V_i(y_i; p)$. The PN weights satisfy the fixed-point condition $\alpha_i \propto V_{i,y}\big(y_i(p(\alpha); x); p(\alpha)\big)$. In Appendix B.2, we describe an algorithm that iterates on a deterministic sequence $\{K^{PF}_t, L^{PF}_t\}_t$ to compute this fixed point.

Given $\alpha(x)$, recovering $\xi(x)$ and $\tau(x)$ is immediate. For example, the allocative components of the goods wedges are given by

$$1 + \tau^g_{i,c,t} = \frac{\beta^t}{q^{PF}_t}\left(\frac{\mathbb{E}_0 c_{i,t}}{c_{i,0}}\right)^{-\sigma}, \qquad 1 + \tau^g_{i,l,t} = \frac{\beta^t}{w^{PF}_t}\left(\frac{\chi}{c^{-\sigma}_{i,0}}\right)\left(\frac{\mathbb{E}_0[\varepsilon_{i,t} l_{i,t}]}{\mathbb{E}_0[\varepsilon^{1+1/\gamma}_{i,t}]}\right)^{\gamma}.$$

To compute the Shapley-value contribution of each coordinate, we also require the reverse mapping from (α, ξ, τ) back to allocations in order to evaluate $\mathcal{W}(\alpha, \xi, \tau)$. Under separable CRRA preferences, lifetime utility from consumption (or leisure) can be written as the product of two terms: one involving the expected time-0 allocation, which depends on (α, ξ, τ^g), and a risk-adjustment term, which depends on the insurance wedges τ^{ins}. This structure yields a computationally convenient two-step procedure for recovering welfare from coordinates. In Appendix B.2, we provide details.

Domeij and Heathcote (2004), Meh (2005), Erosa and Koreshkova (2007), Conesa et al. (2009), Krueger and Ludwig (2013), Gottardi et al. (2015), Krueger and Ludwig (2016), Heathcote et al. (2017), Rohrs and Winter (2017), McGrattan and Prescott (2017), Acikgöz et al. (2018), Hosseini and Shourideh (2019), Boar and Midrigan (2019), Boar and Midrigan (2020), Chien and Wen (2021), Bruggemann (2021), Dyrda and Pedroni (2021), Ferriere et al. (2022).

Calibration. We follow a standard approach to parameterize our economy. We set $\sigma = 1$ and $\gamma = 2$, and choose the labor disutility parameter χ to obtain average hours equal to 1/3. The subjective discount factor β is set to 0.96 to generate an after-tax return on capital (net of growth) of about 3%. We set the capital share and depreciation rate (θ, δ) to $(36\%, 10\%)$. We set baseline fiscal policy parameters (τ^*, B, G) to target a marginal income tax rate of 30%, public debt to output of 100%, and government spending (excluding transfer payments) to output equal to 15%. We adopt Krueger et al. (2016)'s choices for the idiosyncratic productivity process. This calibration generates a standard deviation of log wage earnings of 55%. A complete list of parameters and transition paths for selected aggregates are provided in Appendix B.

6.1.1 Results

We begin with a utilitarian welfare criterion and discuss the associated decomposition of welfare gains aggregated across households. We then examine its cross-sectional counterpart. Varying the size of the reform—i.e., the income tax rate τ^{**} from 10% to 80%—produces a familiar inverse-U relationship between welfare and the tax rate. With utilitarian weights, aggregate welfare is maximized at a tax rate of 48%.

Aggregate components. Applying our decomposition to the income-tax reform, we first characterize how the aggregate components of welfare vary with the tax rate. The goods-related efficiency component, $\mathsf{E}^{g}(\tau^*, \tau^{**}) + \mathsf{E}^{pr}(\tau^*, \tau^{**})$, exhibits an inverse-U shape as a function of τ^{**}, whereas the redistribution component $\mathsf{R}(\tau^*, \tau^{**})$ and the insurance component $\mathsf{E}^{ins}(\tau^*, \tau^{**})$ both rise monotonically with the tax rate; see the left panel of Figure 5. At the utilitarian optimum, redistribution and insurance account for 68% and 124% of the total welfare gain, respectively, while goods-related efficiency contributes -92%.

A monotone increase in redistribution and insurance is understandable. Higher income taxes extract more resources from asset- and income-rich households and finance larger lump-sum transfers, thereby providing additional insurance against idiosyncratic risks. The inverse-U shape of goods-related efficiency reflects competing forces. A higher marginal tax rate distorts the intratemporal labor-leisure margin, but the associated increase in transfers improves intertemporal consumption smoothing. In addition, the firm's input mix is distorted relative to the socially optimal marginal rate of transformation. Incomplete markets foster excessive precautionary savings that lower the equilibrium interest rate and lead to overaccumulation of capital relative to a complete-markets benchmark. Labor income taxation in turn depresses labor supply, raises the capital-labor ratio, and induces wages that are too high relative to ones that support an efficient allocation. As a result, the economy has too much capital and too little labor. By raising income taxes, the reform mitigates the overaccumulation distortion but exacerbates the labor-supply distortion.

Our coordinate system allows us to distinguish intratemporal and intertemporal components of E^g by using the wedges τ^g.[24] The right panel of Figure 5 shows that under our calibration the reform always improves intertemporal smoothing but does so at the expense of worsening the within-period allocation between consumption and leisure. Gains from intertemporal smoothing are modest; for sufficiently large tax increases the intratemporal distortion dominates gains from intertemporal smoothing, making the overall goods-efficiency component negative. Our decomposition further shows that the production-efficiency component $\mathsf{E}^{pr}(\tau^*, \tau^{**})$ coming from a sub-optimal mix of aggregate capital and labor is negative but quantitatively small throughout.

Finally, the additive structure of the Shapley-value decomposition provides a natural way to assess how much of the optimal income-tax rate is motivated by redistribution. To isolate the part of the welfare gain unrelated to redistribution, we evaluate each reform using

$$W(\tau^{**}) - W(\tau^*) - \mathsf{R}(\tau^*, \tau^{**}),$$

which aggregates the insurance, intratemporal, intertemporal, and production-efficiency components. Under this redistribution-neutral welfare criterion, the tax rate that maximizes welfare falls from the utilitarian optimum of 48% to 43%, or about 89% of the full optimum.

Distributional analysis. When social welfare is additive across households, our procedure delivers a decomposition that is also additive across households. Whenever welfare gains are heterogeneous, this additive individual-level decomposition is useful for diagnosing the sources of gains and losses for different groups.

As an illustration, we group households by initial wealth quintiles and plot the cross-sectional distribution of welfare components for the reform that raises the income tax rate from the calibrated U.S. steady-state value of 30% to the optimal rate of 48%. In the left panel of Figure 6, we observe that the top two wealth groups experience welfare losses, while

[24] The intratemporal wedge is given by $(1+\tau^g_{i,l,t})/(1+\tau^g_{i,c,t}) - 1$, and the intertemporal wedge is $\tau^g_{i,c,t}$. Both wedges are simple transformations of τ^g. We apply the same Shapley-value logic as in the aggregate decomposition to isolate the contribution of each wedge by varying one wedge at a time while holding the others fixed.

Figure 5: COMPONENTS OF WELFARE CHANGE

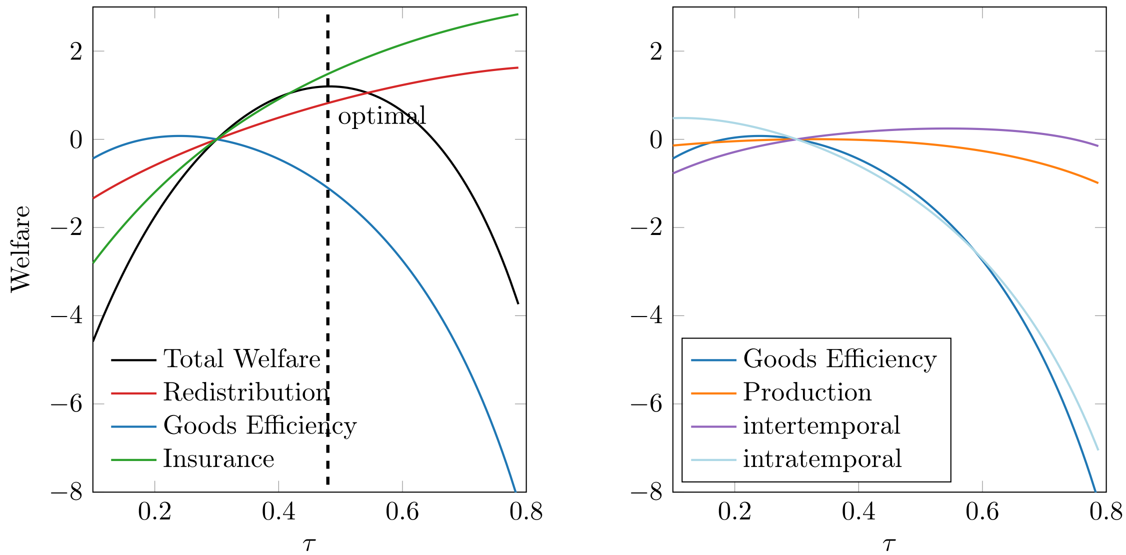

Notes: Welfare decomposition for $\tau \in [10\%, 80\%]$. In the left panel, we plot total welfare and its components: redistribution, goods-related efficiency, and insurance. In the right panel, we further decompose goods-related efficiency into production, intertemporal, and intratemporal efficiency components.

the bottom three groups experience gains. As anticipated, the most pronounced differences across groups arise from the redistribution component, which largely drives the welfare gradient: redistribution contributes 66% of total gains for the lowest wealth quintile and accounts for 69% of the losses for the highest wealth quintile.

Insurance contributes positively for all groups, though the magnitude is smaller for the highest-wealth households. Goods-related efficiency contributions are negative for all groups but are smallest in magnitude for the bottom quintile and largest for the top quintile. In the right panel of Figure 6, we decompose the aggregate goods-related efficiency component into its subparts; the same logic applies at the cross-sectional level, allowing us to examine how the intertemporal and intratemporal components of E^g, as well as the residual E^{pr} component, vary across households. We find that wealth-poor households benefit primarily through improved intertemporal smoothing, whereas the intratemporal labor wedge is most important for high-wealth households.

In Appendix B, we repeat the analysis grouping households by income rather than wealth. Given the high correlation between wealth and income in standard calibrations of this class of models, the resulting distributional patterns are very similar to those in Figure 6.

Figure 6: WELFARE DECOMPOSITION BY ASSET QUINTILE

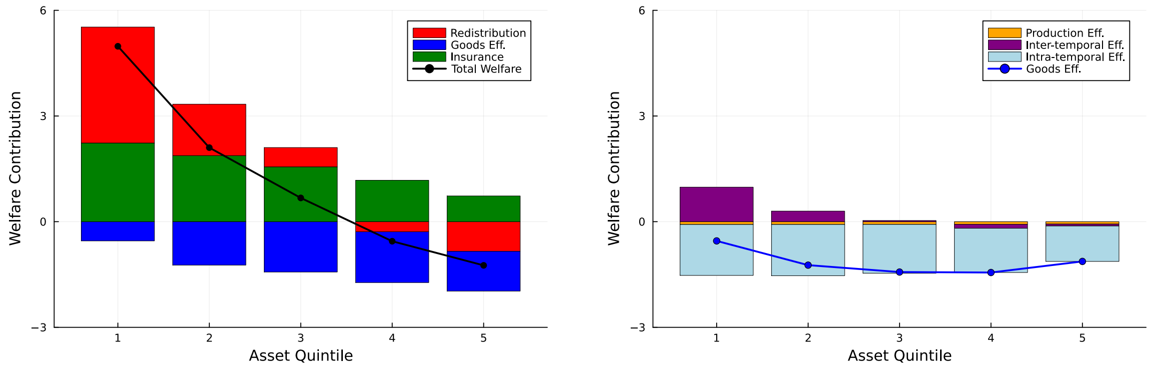

Notes: Welfare contributions by asset quintile. The left panel shows total welfare and its decomposition into redistribution, goods-efficiency, and insurance components. The right panel decomposes goods efficiency further into production, intertemporal, and intratemporal components of efficiency by asset quintile.

6.2 Public debt reform

In this example, we apply our decomposition to a fiscal reform described by Aguiar et al. (2024). Their reform is constructed to generate a Pareto improvement in a standard incomplete-markets economy by adjusting public debt and fiscal policy. In a setting like the one we study in this section, they show that with low interest rates, increasing public debt and appropriately adjusting taxes and transfers can expand the budget sets of all households.[25]

Although these reforms raise welfare for everyone by construction, it is not obvious whether gains come mainly from improvements in efficiency or insurance that benefit households broadly, or whether implicit redistribution allows some households to gain more than others. Our aim is to use our decomposition to quantify contributions of efficiency, insurance, and redistribution to the overall welfare gains, and to study how these components vary across households.

We specialize our environment to obtain what Aguiar et al. (2024) call a constant-K reform in which aggregate resources are held fixed. To achieve this, we set $\Psi = 0$ and assume that households supply one unit of labor inelastically. The status-quo economy is the steady state associated with $B_t = T_t = \tau_t = \varsigma_t = 0$. We denote the associated capital stock and return on savings by K^{ss} and r^{ss}, respectively.

The reform consists of (i) a sequence $\{B_t\}_t$ with $B_t \to \overline{B}$, and (ii) a sequence of capi-

[25]Aguiar et al. (2024) provide general conditions for a robust Pareto improvement in richer settings that allow for markup distortions.

tal wedges $\{\varsigma_t\}_t$ that ensure $K_t = K^{ss}$ for all $t \geq 0$. The associated equilibrium includes sequences $\{r_t\}_t$ and $\{T_t\}_t$ that clear asset markets and satisfy the government budget constraint. The reform constitutes a Pareto improvement if $r_t > r^{ss}$ and $T_t > 0$ for all t, thereby expanding the budget sets of all households.

Welfare decomposition. The steps for constructing our decomposition are like those for the tax reform and are detailed in Appendix B.2. Because the reform keeps aggregate output, consumption, and factor utilization fixed, it affects welfare only through its impacts on time paths of individual consumption. Redistribution is captured by the percentage change in the present value of the expected paths $\{\mathbb{E}_0 c_{i,t}\}_{i,t}$; goods-related efficiency captures deviations from intertemporal smoothing for those expected paths over the transition; and insurance captures deviations of realized consumption from its expected paths. Other forms of efficiency such as production efficiency and distortions in intratemporal labor-leisure allocation are irrelevant in this setting because aggregate quantities of capital and labor are held constant.

In the absence of ex-ante heterogeneity, expected consumption paths converge to the common steady-state level, $\mathbb{E}_0 c_{i,t} \to C^{ss}$. Thus, gains from goods-related efficiency are present only along the transition, while insurance gains are present both along the transition and in the new steady state. Given this structure, the relative contribution of the components depends on the size of the discount factor β.

Calibration. We adopt an illustrative calibration for the status-quo economy and construct a reform that generates a robust Pareto improvement. We use the same preference curvature, idiosyncratic skill process, and aggregate production technology as in the previous section. The status-quo economy is calibrated so that $1 + r^{ss} = 0.99$. A negative real interest rate is necessary for a robust Pareto improvement in the absence of additional features such as markups. The implied discount factor is high, $\beta = 0.997$. We verify that the paths satisfy $T_t > 0$ and $r_t > r^{ss}$ for all dates, confirming that the reform achieves a robust Pareto improvement. The reform raises public debt to 50% of output in the new steady state and the path of debt adjusts gradually over time. A complete list of parameters and transition paths for selected aggregates are provided in Appendix B.

Results. The reform raises welfare by 1.7% in consumption-equivalent units under the utilitarian criterion. Under utilitarian aggregation, nearly all of the gains—about 82.8%—

Figure 7: WELFARE DECOMPOSITION BY ASSET QUINTILE

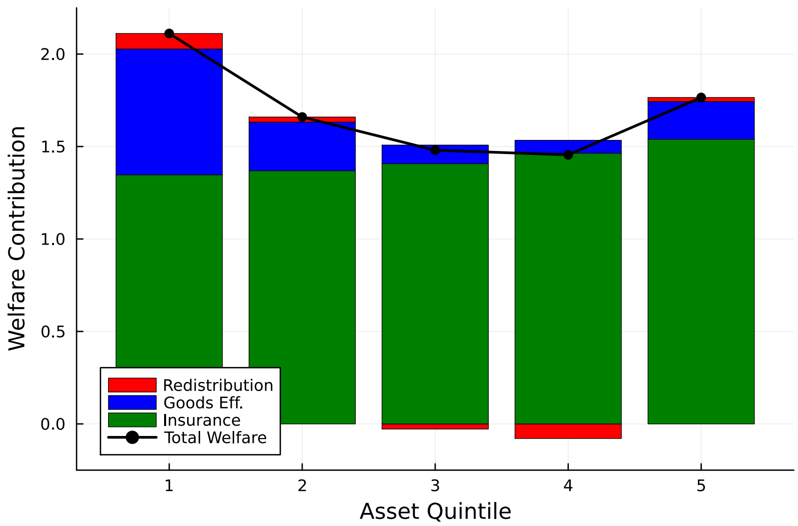

Notes: This figure presents the welfare decomposition by asset quintile for the alternative calibration that generates a robust Pareto improvement as in Aguiar et al. (2024). The contributions of redistribution, goods-related efficiency (intertemporal smoothing), and insurance across the wealth distribution are shown, together with the total welfare effect under the reform.

come from insurance, with a small 16.8% contribution from intertemporal smoothing, and a negligible 0.4% from redistribution. The dominant role of insurance reflects the permanent gains in risk sharing, amplified by the high value of β.

Figure 7 reports the individual-level decomposition across households grouped by initial wealth, as in the previous section. Consistent with the definition of a robust Pareto improvement, all households experience positive welfare gains (black line). Decomposing these gains shows that insurance and efficiency components dominate across the distribution. Redistribution is small but positive for two groups: households near the borrowing constraint (who value transfers more highly) and high-wealth households (who benefit from the higher returns).

7 Conclusion

Our welfare decomposition for heterogeneous-agent economies separates the effects of policy reforms into redistribution and various components of efficiency. By associating each feasible allocation with a vector of Pareto–Negishi weights, we capture implicit inequality and deduce a collection of wedges that capture distortions in production, allocation, and risk sharing. Our decomposition respects Pareto-consistency, in the sense that it detects efficiency losses only when the reform moves inside the Pareto frontier. It is also invariant to the choice of numeraire. It features a notion of redistribution that is connected to changes in real disposable income. Widely used alternative decompositions generally fail these Pareto-based requirements, depend on arbitrary normalizations, and consequently can blur the line between redistribution and efficiency. Two quantitative applications illustrate how our framework can be applied.

References

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L. Violante. 2019. "Education Policy and Intergenerational Transfers in Equilibrium." *Journal of Political Economy*, 127(6): 2569–2624.

Acikgöz, Ömer, Marcus Hagedorn, Hans Holter, and Yikai Wang. 2018. "The optimum quantity of capital and debt."

Aguiar, Mark, Manuel Amador, and Cristina Arellano. 2024. "Micro Risks and (Robust) Pareto Improving Policies." *American Economic Review*, 114(11): 3669–3713, DOI: `http://dx.doi.org/10.1257/aer.20210194`.

Aiyagari, S Rao. 1994. "Uninsured idiosyncratic risk and aggregate saving." *Quarterly Journal of Economics*, 109(3): 659–684.

Aiyagari, S. Rao, and Ellen R. McGrattan. 1998. "The optimum quantity of debt." *Journal of Monetary Economics*, 42(May 1997): 447–469.

Arrow, Kenneth J. 1953. "The Role of Securities in the Optimal Allocation of Risk." *Econometrica*, 31(1): 91–96.

Auerbach, Alan J. 1985. "The Theory of Excess Burden and Optimal Taxation." In *Handbook of Public Economics*. eds. by Alan J. Auerbach, and Martin Feldstein, 1, Amsterdam: North-Holland, , Chap. 2 61–127.

Auerbach, Alan J., and Jr. Hines, James R. 2002. "Taxation and Economic Efficiency." In *Handbook of Public Economics*. eds. by Alan J. Auerbach, and Martin Feldstein, 3, Amsterdam: Elsevier, , Chap. 21 1347–1421.

Baqaee, David, and Ariel Burstein. 2025. "Aggregate efficiency with heterogeneous agents."Technical report, National Bureau of Economic Research.

Baqaee, David Rezza, and Emmanuel Farhi. 2019. "Productivity and Misallocation in General Equilibrium*." *The Quarterly Journal of Economics*, 135(1): 105–163.

Benabou, Roland. 2002. "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?" *Econometrica*, 70(2): 481–517.

Bhandari, Anmol, David Evans, Mikhail Golosov, and Thomas J. Sargent. 2021. "Inequality, Business Cycles, and Monetary-Fiscal Policy." *Econometrica*, 89(6): 2559–2599.

Boar, Corina, and Virgiliu Midrigan. 2019. "Markups and inequality."Technical report, National Bureau of Economic Research.

Boar, Corina, and Virgiliu Midrigan. 2020. "Efficient redistribution."Technical report, National Bureau of Economic Research.

Bruggemann, Bettina. 2021. "Higher Taxes at the Top: The Role of Entrepreneurs." *American Economic Journal: Macroeconomics*, 13(3): .

Carroll, Christopher D. 2006. "The method of endogenous gridpoints for solving dynamic stochastic optimization problems." *Economics letters*, 91(3): 312–320.

Chien, YiLi, and Yi Wen. 2021. "Time-inconsistent optimal quantity of debt." *European Economic Review*, 140, p. 103913.

Chipman, John S., and James C. Moore. 1978. "The New Welfare Economics, 1939–1974." *International Economic Review*, 19(3): 547–584.

Cho, Jang-Ok, Thomas F. Cooley, and Hyung Seok E. Kim. 2015. "Business cycle uncertainty and economic welfare." *Review of Economic Dynamics*, 18(2): 185–200.

Conesa, Juan Carlos, Sagiri Kitao, and Dirk Krueger. 2009. "Taxing Capital? Not a Bad Idea after All!." *American Economic Review*, 99(1): 25–48.

Conesa, Juan Carlos, and Dirk Krueger. 1999. "Social Security Reform with Heterogeneous Agents." *Review of Economic Dynamics*, 2(4): 757–795.

Dávila, Eduardo, and Andreas Schaab. 2022. "Welfare Assessments with Heterogeneous Individuals." Working Paper 30571, National Bureau of Economic Research.

Debreu, Gerard. 1951. "The Coefficient of Resource Utilization." *Econometrica*, 19(3): 273–292.

Debreu, Gerard. 1959. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium.*: Yale University Press.

Diamond, Peter A. 1975. "A Many-Person Ramsey Tax Rule." *Journal of Public Economics*, 4(4): 335–342.

Diamond, Peter A., and James A. Mirrlees. 1971a. "Optimal Taxation and Public Production I: Production Efficiency." *American Economic Review*, 61(1): 8–27.

Diamond, Peter A., and James A. Mirrlees. 1971b. "Optimal taxation and public production II: Tax rules." *American Economic Review*, 61(3): 261–278.

Domeij, David, and Jonathan Heathcote. 2004. "On the Distributional Effects of Reducing Capital Taxes." *International Economic Review*, 45(2): 523–554.

Dyrda, Sebastian, and Marcelo Pedroni. 2021. "Optimal Fiscal Policy in a Model with Uninsurable Idiosyncratic Shocks." *Review of Economic Studies, Forthcoming.*

Erosa, Andres, and Tatyana Koreshkova. 2007. "Progressive taxation in a dynastic model of human capital." *Journal of Monetary Economics*, 54(3): 667–685.

Ferriere, Axelle, Philipp Grübener, Gaston Navarro, and Oliko Vardishvili. 2022. "On the optimal design of transfers and income-tax progressivity."

Floden, Martin. 2001. "The Effectiveness of Government Debt and Transfers as Insurance." *Journal of Monetary Economics*, 48(1): 81–108.

Gottardi, Piero, Atsushi Kajii, and Tomoyuki Nakajima. 2015. "Optimal Taxation and Debt with Uninsurable Risks to Human Capital Accumulation." *American Economic Review*, 105(11): 3443–70.

Guvenen, Fatih, Gueorgui Kambourov, Burhanettin Kuruscu, Sergio Ocampo-Diaz, and Daphne Chen. 2019. "Use It or Lose It: Efficiency Gains from Wealth Taxation." Working Paper 26284, National Bureau of Economic Research.

Harberger, Arnold C. 1971. "Three Basic Postulates for Applied Welfare Economics: An Interpretive Essay." *Journal of Economic Literature*, 9(3): 785–797.

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni Violante. 2017. "Optimal Tax Progressivity: An Analytical Framework." *Quarterly Journal of Economics.*

Hicks, John R. 1939. "The Foundations of Welfare Economics." *The Economic Journal*, 49(196): 696–712.

Hosseini, Roozbeh, and Ali Shourideh. 2019. "Retirement Financing: An Optimal Reform Approach." *Econometrica*, 87(4): 1205–1265.

Huang, He, Selahttin Imrohoroglu, and Thomas J. Sargent. 1997. "Two Computations to Fund Social Security." *Macroeconomic Dynamics*, 1(1): 7 – 44.

Kaldor, Nicholas. 1939. "Welfare Propositions of Economics and Interpersonal Comparisons of Utility." *The Economic Journal*, 49(195): 549–552.

Koehne, Sebastian, and Moritz Kuhn. 2015. "Should unemployment insurance be asset tested?" *Review of Economic Dynamics*, 18(3): 575–592.

Krueger, D., K. Mitman, and F. Perri. 2016. "Chapter 11 - Macroeconomics and Household Heterogeneity.": Elsevier, 843–921.

Krueger, Dirk, and Alexander Ludwig. 2013. "Optimal Progressive Labor Income Taxation and Education Subsidies When Education Decisions and Intergenerational Transfers are Endogenous." *American Economic Review*, 103(3): 496–501.

Krueger, Dirk, and Alexander Ludwig. 2016. "On the optimal provision of social insurance: Progressive taxation versus education subsidies in general equilibrium." *Journal of Monetary Economics*, 77 72–98.

Kuznets, Simon. 1948. "On the Valuation of Social Income: Reflections on Professor Hicks' Article." *The Economic Journal*, 58(230): 1–16.

Liu, Ernest. 2019. "Industrial Policies in Production Networks*." *The Quarterly Journal of Economics*, 134(4): 1883–1948, URL: `https://doi.org/10.1093/qje/qjz024`, DOI: `http://dx.doi.org/10.1093/qje/qjz024`.

Lundberg, Scott M, and Su-In Lee. 2017. "A unified approach to interpreting model predictions." *Advances in neural information processing systems*, 30.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green. 1995. *Microeconomic Theory.* New York: Oxford University Press.

McGrattan, Ellen R., and Edward C. Prescott. 2017. "On financing retirement with an aging population." *Quantitative Economics*, 8(1): 75–115.

McKenzie, Lionel W. 1959. "On the Existence of General Equilibrium for a Competitive Economy." *Econometrica*, 27(1): 54–71.

Meh, Césaire A. 2005. "Entrepreneurship, wealth inequality, and taxation." *Review of Economic Dynamics*, 8(3): 688–719.

Nakajima, Tomoyuki, and Shuhei Takahashi. 2020. "The Effectiveness of Consumption Taxes and Transfers as Insurance Against Idiosyncratic Risk." *Journal of Money, Credit and Banking*, 52(2-3): 505–530.

Negishi, Takashi. 1960. "Welfare Economics and Existence of an Equilibrium for a Competitive Economy." *Metroeconomica*, 12(2-3): 92–97.

Rohrs, Sigrid, and Christoph Winter. 2017. "Reducing Government Debt in the Presence of Inequality." *Journal of Economic Dynamics and Control*, 82 1–20.

Scitovsky, Tibor. 1941. "A Note on Welfare Propositions in Economics." *The Review of Economic Studies*, 9(1): 77–88.

Sen, Amartya. 1970. *Collective Choice and Social Welfare*. San Francisco: Holden-Day.

Seshadri, Ananth, and Kazuhiro Yuki. 2004. "Equity and efficiency effects of redistributive policies." *Journal of Monetary Economics*, 51(7): 1415–1447.

Shorrocks, Anthony F. 2013. "Decomposition procedures for distributional analysis: A unified framework based on the Shapley value." *The Journal of Economic Inequality*, 11(1): 99–126.

Uhlig, Harald. 1996. "A law of large numbers for large economies." *Economic Theory*, 8(1): 41–50.

Young, Eric R. 2010. "Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm and non-stochastic simulations." *Journal of Economic Dynamics and Control*, 34(1): 36–41.

Appendix

A Proofs for Section 3, 4, and 5

In this section, we provide the proofs for the propositions and theorems in the main text.

A.1 Proofs for Section 3

We begin by noting some useful properties of $V_i(p, y)$ as defined in (4), when $p_k > 0$ for all k. Strict concavity and strict monotonicity of the utility function implies that $V_i(p, y)$ is strictly concave in y. As u_i is twice continuously differentiable, $V_i(p, y)$ is also twice continuously differentiable in y for all $y \geq 0$ with derivative equal to

$$V_{i,y}(p, y) = \max_k \left\{ \frac{u_{i,k}(\mathsf{x}_i(p, y))}{p_k} \right\} \tag{29}$$

where $\mathsf{x}_i(p, y)$ is the Marshallian demand function for consumer i at prices p and income y.

Next, the following lemma documents the properties of the Lagrange multipliers P associated with Pareto optimal allocations.

Lemma 5. *Let $\alpha \in \mathbb{R}^I_+$ be a vector of Pareto weights. There exists a unique Pareto optimal allocation x that solves* (1) *for Pareto weights α. If $X_k > 0$ for all k then there exists a unique set of Lagrange multipliers $\mathsf{P} \in \mathbb{R}^K_+$ on the resource constraint $\sum_i x_{i,k} \leq X_k$, and mapping $\alpha \mapsto \mathsf{P}(\alpha)$ is continuous at α.*

Proof. For a given α, strict concavity of u_i and convexity of the production set $\mathcal{Y}$ imply that the objective function is strictly concave and the feasible set is convex.[26] Hence, the maximization problem (1) features a unique solution (x, X).

Let P be a vector of Lagrange multipliers on the resource constraint $\sum_i x_{i,k} \leq X_k$, and $\mu_{i,k}$ be the Lagrange multipliers on the non-negativity constraints $x_{i,k} \geq 0$. The Karush–Kuhn–Tucker conditions for the Pareto optimization problem with respect to consumption $x_{i,k}$ are

$$\alpha_i u_{i,k}(x_i) = \mathsf{P}_k - \mu_{i,k} \quad (\forall i, k)$$

where $\mu_{i,k} \geq 0$ and $\mu_{i,k} x_{i,k} = 0$.

[26] If $\alpha_i = 0$ for some i, the objective function is not strictly concave in x_i. However, since u_i does not enter the objective, the planner is indifferent over x_i and optimally sets $x_{i,k} = 0$ for all k to relax the resource constraints. The problem then reduces to maximizing over $\{x_j\}_{j \neq i}$, and the objective function is strictly concave in these remaining variables.

From these conditions, we have $\mathsf{P}_k = \alpha_i u_{i,k}(x_i) + \mu_{i,k} \geq \alpha_i u_{i,k}(x_i)$ for all i, with equality if $x_{i,k} > 0$. As $X_k > 0$ there must exist at least one i such that the equality holds. Thus, the vector of Lagrange multipliers P must satisfy

$$\mathsf{P}_k = \max_i \{\alpha_i u_{i,k}(x_i)\}, \tag{30}$$

which implies that P is uniquely determined. That $u_{i,k}(x_i) > 0$ for all i, k implies that $\mathsf{P}_k > 0$ for all k.

To establish continuity, note that the Theorem of the Maximum implies that the set of optimizers is upper-hemicontinuous in α. Since the solution $(x(\alpha), X(\alpha))$ is unique, the mapping $\alpha \mapsto x(\alpha)$ is continuous. The function $\alpha \mapsto \alpha_i u_{i,k}(x_i(\alpha))$ is continuous because u_i is continuously differentiable. Since $\mathsf{P}_k(\alpha)$ is the maximum of a finite number of continuous functions, $\mathsf{P}(\alpha)$ is continuous. □

When $X_k = 0$ for some k, the Lagrange multipliers P_k are not uniquely determined. In this case, we set P_k as smallest possible value, which is given by (30). This guarantees a mapping $\alpha \mapsto \mathsf{P}(\alpha)$ that is continuous for all $\alpha \in \mathbb{R}^I_+$.

Next, we prove Lemma 1.

A.1.1 Proof of Lemma 1

Proof. Let x be a Pareto optimal allocation that solves (1) for Pareto weights α. Let $\mathsf{P} \in \mathbb{R}^K$ be the vector of Lagrange multipliers on the resource constraints $\sum_i \tilde{x}_{i,k} = \tilde{X}_k$. We will prove the Lemma for $p = \mathsf{P}$, but the proof follows for any $p \propto \mathsf{P}$. Define individual and aggregate incomes by

$$y_i := \sum_k p_k x_{i,k}, \qquad Y := \sum_i y_i.$$

We prove parts (i)–(iii) directly.

(i) $Y = Y^{max}$ as in equation (5). Concavity of the utility functions and convexity of the production set implies that the maximization problem (1) features strong duality[27], so Problem (1) is equivalent to solving (2) which implies that X solves

$$\max_{\tilde{X} \in \mathcal{Y}} \sum_k p_k \tilde{X}_k.$$

Thus $\sum_k p_k X_k = \max_{\tilde{X} \in \mathcal{Y}} \sum_k p_k \tilde{X}_k =: Y^{max}$, and since $Y = \sum_k p_k X_k$, we obtain $Y =$

[27] Slater's condition is satisfied since $0 \in \mathcal{Y}$.

Y^{max}.

(iii) x_i solves (4) *with* $V_{i,y}(p, y_i) = 1/\alpha_i$ *if* $y_i > 0$ *and* $V_{i,y}(p, y_i) \leq 1/\alpha_i$ *if* $y_i = 0$. Consider consumer i's problem (4) at prices p and income y_i. Let λ_i be the Lagrange multiplier on the budget constraint. The consumer's first-order-conditions are

$$u_{i,k}(\tilde{x}_i) = \lambda_i p_k - \eta_{i,k} \quad (\forall i, k), \text{ and } \sum_k p_k \tilde{x}_{i,k} = y_i \text{ if } \lambda_i > 0, \text{ and } x_{i,k} = 0 \text{ if } \eta_{i,k} > 0. \quad (31)$$

The planner's first-order-conditions for (1) give $\alpha_i u_{i,k}(x_i) = p_k - \mu_{i,k}$ for all i, k. This implies that the equation (31) is satisfied by $\tilde{x}_i = x_i$ and $\lambda_i = 1/\alpha_i$ and $\eta_{i,k} = \mu_{i,k}/\alpha_i$ as $\sum_k p_k x_{i,k} = y_i$. Therefore x_i solves the consumer problem (4) at prices p and income y_i. For interior solutions, i.e. $x_{i,k} > 0$ for all k we have that $\alpha_i u_{i,k}(x_i) = p_k$ which guarantees (7).

The envelope theorem implies that

$$V_{i,y}(p, y_i) = \lambda_i = \frac{1}{\alpha_i} \quad (\forall i),$$

if $y_i > 0$. When $y_i = 0$ we have $x_i = 0$. Equation (29) implies

$$V_{i,y}(p, 0) = \max_k \{u_{i,k}(0)/p_k\}.$$

The planner's first order conditions guarantee that $u_{i,k}(0)/p_k = 1/\alpha_i - \frac{\mu_{i,k}}{\alpha_i p_k}$ which implies that $V_{i,y}(p, 0) \leq 1/\alpha_i$.

(ii) y solves (6). The Karush–Kuhn–Tucker conditions for the income-allocation problem (6) read $\alpha_i V_{i,y}(p, \tilde{y}_i) = \Lambda - \delta_i$ for some multiplier Λ on $\sum_i \tilde{y}_i \leq Y$ and $\delta_i \geq 0$ on $\tilde{y}_i \geq 0$. This is satisfied by $\tilde{y}_i = y_i$ with $\Lambda = 1$ since if $y_i > 0$ $V_{i,y}(p, y_i) = 1/\alpha_i$ and $\delta_i = 0$ and if $y_i = 0$ then $\alpha_i V_{i,y}(p, y_i) \leq 1$ so $\delta_i \geq 0$. Therefore, y satisfies the Karush–Kuhn–Tucker system and hence solves (6). □

A.1.2 Proof of Lemma 2

Proof. Let x be an interior Pareto optimal allocation and define $\alpha_i \propto u_{i,1}(x_i)^{-1}$ as the associated Pareto–Negishi weights, normalized so that $\sum_i \alpha_i = 1$. The proof of Lemma 1 establishes that α satisfies (8).

To establish uniqueness, suppose by contradiction that there exists $\alpha' \neq \alpha$ with $\sum_i \alpha'_i =$

1 also satisfying

$$\alpha_i' = \frac{V_{i,y}\big(\mathsf{P}(\alpha'), \sum_k \mathsf{P}_k(\alpha')x_{i,k}\big)^{-1}}{\sum_j V_{j,y}\big(\mathsf{P}(\alpha'), \sum_k \mathsf{P}_k(\alpha')x_{j,k}\big)^{-1}}.$$

Let x' denote the Pareto optimal allocation associated with weights α', and let $p' := \mathsf{P}(\alpha')$. Since $\alpha' \neq \alpha$ and preferences are strictly concave and the solutions are interior we have $x' \neq x$.

Because both x and x' are Pareto optimal, neither strictly dominates the other. Therefore, there exist consumer j such that $u_j(x_j) > u_j(x_j')$. As x_j' solves consumer j's utility maximization problem with prices p' and income $y_j' := \sum_k p_k' x_{j,k}'$, the fact that $u_j(x_j) > u_j(x_j')$ implies that x_j violates consumer j's budget constraint:

$$\sum_k p_k' x_{j,k} > \sum_k p_k' x_{j,k}'.$$

Meanwhile, aggregate output $X' := \sum_i x_i'$ maximizes aggregate income $\sum_k p_k' \tilde{X}_k$ over $\mathcal{Y}$ which implies that

$$\sum_{i,k} p_k' x_{i,k} \leq \sum_{i,k} p_k' x_{i,k}'.$$

Combining these two inequalities, there must exist some consumer j' for whom

$$\sum_k p_k' x_{j',k} < \sum_k p_k' x_{j',k}'.$$

Define incomes $y_i := \sum_k p_k' x_{i,k}$ and $y_i' := \sum_k p_k' x_{i,k}'$. By strict concavity of utility, the marginal utility of income $V_{i,y}(p, y)$ is strictly decreasing in y. As $y_j > y_j'$ and $y_{j'} < y_{j'}'$, we have

$$V_{j,y}(p', y_j)^{-1} > V_{j,y}(p', y_j')^{-1} \quad \text{and} \quad V_{j',y}(p', y_{j'})^{-1} < V_{j',y}(p', y_{j'}')^{-1}.$$

Taking ratios yields

$$\frac{V_{j,y}(p', y_j)^{-1}}{V_{j',y}(p', y_{j'})^{-1}} > \frac{V_{j,y}(p', y_j')^{-1}}{V_{j',y}(p', y_{j'}')^{-1}}.$$

By the fixed point equation, the left-hand side equals $\alpha_j'/\alpha_{j'}'$. Since α' are the Pareto–Negishi weights for allocation x', we have $\alpha_i' \propto V_{i,y}(p', y_i')^{-1}$, so the right-hand side also equals $\alpha_j'/\alpha_{j'}'$. This is a contradiction. □

For non-interior allocation, there will be multiple solutions to the fixed point equation (8) only if there are multiple Pareto–Negishi weights associated with the same Pareto optimal allocation. We provide a selection criterion when we discuss the uniqueness of the coordinate system below.

A.1.3 Proof of Lemma 3

Proof. Let x be any feasible interior allocation, i.e., $x_{i,k} > 0$ for all i, k. Our first claim is that there exists a fixed point to the non-linear equation (8). Let $\mathsf{B} = \{\tilde{\alpha} \in [0,1]^I : \sum_i \alpha_i = 1\}$ be the set of PN weights normalized to sum to 1. Define the function $f : \mathsf{B} \to \mathsf{B}$ by its i-th component:

$$f_i(\tilde{\alpha}) = \frac{V_{i,y}(\mathsf{P}(\tilde{\alpha}), \sum_k \mathsf{P}_k(\tilde{\alpha})x_{i,k})^{-1}}{\sum_j V_{j,y}(\mathsf{P}(\tilde{\alpha}), \sum_k \mathsf{P}_k(\tilde{\alpha})x_{j,k})^{-1}}.$$

Our assumptions on u guarantee that $V_{i,y} > 0$ and Lemma 5 implies $\mathsf{P}_k(\alpha) > 0$, which implies that $f(\tilde{\alpha})$ is a well-defined mapping from B to itself. Furthermore, as $\mathsf{P}(\tilde{\alpha})$ is continuous from Lemma 5, $f(\tilde{\alpha})$ is continuous. The Brouwer fixed point theorem implies that there exists a fixed point α of f which satisfies (8). Given α, both ξ and τ are uniquely determined by

$$\xi = 1 - \frac{\sum_k \mathsf{P}_k(\alpha)X_k}{\max_{\tilde{X}\in\mathcal{Y}} \sum_k \mathsf{P}_k(\alpha)\tilde{X}_k},$$

and

$$\tau_{i,k} = \frac{\mathsf{P}_1(\alpha)}{\mathsf{P}_k(\alpha)} \frac{u_{i,k}(x_i)}{u_{i,1}(x_i)} - 1.$$

Let $p = \gamma\mathsf{P}(\alpha)$ for some $\gamma > 0$ be any price vector such that $p \propto \mathsf{P}(\alpha)$. Let $y_i = \sum_k p_k x_{i,k}$ be the income of consumer i and $Y = \sum_i y_i$ be the aggregate income. Then

$$\frac{p_1}{p_k}\frac{u_{i,k}(x_i)}{u_{i,1}(x_i)} = \frac{\mathsf{P}_1(\alpha)}{\mathsf{P}_k(\alpha)}\frac{u_{i,k}(x_i)}{u_{i,1}(x_i)} = 1 + \tau_{i,k}$$

and

$$\frac{Y}{Y^{max}(p)} = \frac{\sum_k p_k X_k}{\max_{\tilde{X}\in\mathcal{Y}} \sum_k p_k \tilde{X}_k} = \frac{\sum_k \mathsf{P}_k(\alpha)X_k}{\max_{\tilde{X}\in\mathcal{Y}} \sum_k \mathsf{P}_k(\alpha)\tilde{X}_k} = 1 - \xi,$$

so ξ and τ satisfy (10) and (12).

The first-order conditions of the income-allocation problem (11), assuming the non-negativity constraints don't bind, are

$$\alpha_i V_{i,y}(p, \tilde{y}_i) = \Lambda. \tag{32}$$

The properties of the indirect utility function imply that $V_i(p, \tilde{y}_i) = V_i(\mathsf{P}(\alpha), \tilde{y}_i/\gamma)$ which implies that

$$V_{i,y}(p,\tilde{y}_i) = \frac{1}{\gamma} V_{i,y}(\mathsf{P}(\alpha), \tilde{y}_i/\gamma),$$

so (32) is satisfied by $\tilde{y}_i = y_i = \gamma \sum_k \mathsf{P}_k(\alpha) x_{i,k}$ with

$$\Lambda = \frac{1}{\gamma \sum_j V_{j,y}(\mathsf{P}(\alpha), \sum_k \mathsf{P}_k(\alpha) x_{j,k})^{-1}}.$$

Therefore, $\{y_i\}$ satisfies the Karush–Kuhn–Tucker system and hence solves (11). □

The same steps can be used when x is not interior, i.e. $x_{i,k} = 0$ for some i, k. At which point the allocation distortions are defined up to a set because of the non-negativity constraints:

$$\Upsilon_{i,k} = \left\{ \frac{\mathsf{P}_1(\alpha)}{\mathsf{P}_k(\alpha)} \frac{u_{i,k}(x_i) + \mu_{i,k}}{u_{i,1}(x_i) + \mu_{i,1}} - 1 \middle| \mu_{i,k} \geq 0, \mu_{i,k} x_{i,k} = 0 \right\}.$$

We select a unique $\tau_{i,k}$ from $\Upsilon_{i,k}$ by choosing the smallest distortion $\tau_{i,k} = \arg\min_{\tilde{\tau} \in \Upsilon_{i,k}} |\tilde{\tau}|$.

A.1.4 Coordinate system uniqueness

In this section we provide sufficient conditions such that for each (α, t) there is a unique x as well as discussing the selection criterion when there exist multiple solutions to the fixed point equation (8). We begin with conditions that guarantee that for each (α, t) there is a unique x.

Lemma 6. *Let $\mathsf{x}_i(q, m)$ be the Marshallian demand function for consumer i. Suppose that for all $p, q \in \mathbb{R}^K_{++}$ and $m \in \mathbb{R}_{++}$ individual demand satisfies*

$$\sum_k p_k \frac{\partial \mathsf{x}_{i,k}}{\partial m}(q, m) > 0 \tag{33}$$

then every (α, t) has a unique x.

Proof. Strict concavity of the utility function implies that $V_i(p, y)$ is strictly concave in y and the Marshallian demand correspondence is single-valued and continuous. For an arbitrary (α, ξ, τ) define $Y = (1 - \xi) Y^{max}(p)$ and $p = \mathsf{P}(\alpha)$. Strict concavity of $V_i(p, y)$ in y implies that the income-allocation problem (6) has a unique solution y_i.

We are searching for an allocation x that solves the following system of equations:

$$\sum_k p_k x_{i,k} = y_i$$
$$1 + \tau_{i,k} = \frac{p_1}{p_k} \frac{u_{i,k}(x_i) + \mu_{i,k}}{u_{i,1}(x_i) + \mu_{i,1}}, \quad \forall i, k.$$

where $\mu_{i,k} \geq 0$ and $\mu_{i,k} x_{i,k} = 0$. This is equivalent to finding a vector of incomes m_i such that

$$\sum_k p_k \mathsf{x}_i(q_i, m_i) = y_i.$$

where $q_{i,k} = p_k(1 + \tau_{i,k})$ and setting $x_i = \mathsf{x}_i(q_i, m_i)$. Holding fixed p, q_i let $\mathsf{y}_i(m) := \sum_k p_k \mathsf{x}_i(q_i, m)$. Continuity of x_i and (33) imply that y_i is continuous and strictly increasing in m. By construction $\lim_{m \to 0} \mathsf{y}_i(m) = 0$ and $\lim_{m \to \infty} \mathsf{y}_i(m) = \infty$. By the intermediate value theorem, there exists a unique m_i such that $\mathsf{y}_i(m_i) = y_i$. Setting $x_i = \mathsf{x}_i(q_i, m_i)$ yields a unique allocation x that solves the system of equations. □

Condition (33) is sufficient to guarantee that for each (α, t) there is a unique x. We can guarantee that this condition holds if all goods are normal since then $\frac{\partial \mathsf{x}_i}{\partial m} > 0$.

When there are multiple solutions to the fixed point equation (8), we apply the following selection criterion. For the status quo allocation x^*, we choose the solution α^* that minimizes the efficiency loss $\mathcal{W}(\alpha^*, 0) - \mathcal{W}(\alpha^*, t^*)$. When each Pareto-efficient allocation corresponds to a unique vector of Pareto–Negishi weights (up to normalization), this criterion uniquely determines α^*. For non-interior allocations, multiple Pareto–Negishi weights may correspond to the same point on the Pareto frontier; in this case, we select the weights closest to the welfare weights $\bar{\alpha}$ in Euclidean norm. For the reform allocation x^{**}, we select the solution α^{**} that is closest to α^* in Euclidean norm.

A.1.5 Proof of Lemma 4

Proof. Fix PN weights α and let $t = (\xi, \tau)$ denote distortions that satisfy $\xi \in [0, 1]$. Without loss of generality let $p = \mathsf{P}(\alpha)$, the case with $p \propto \mathsf{P}(\alpha)$ is identical. Define

$$x^* := \mathcal{X}(\alpha, 0), \quad x^{**} := \mathcal{X}(\alpha, t), \quad y_i^* := \sum_k p_k x_{i,k}^*, \quad y_i^{**} := \sum_k p_k x_{i,k}^{**}.$$

Let $Y^* := \sum_i y_i^*$ and $Y^{**} := \sum_i y_i^{**}$. By construction of the coordinates, x^* maximizes output at prices p so $Y^* = Y^{max}$. For x^{**}, the production wedge equals $\xi \geq 0$, so $Y^{**} =$

$(1-\xi)Y^{max} \leq Y^{max} = Y^*$.

Given prices p and PN weights α, the income distributions y^* and y^{**} solve, respectively,

$$y^* \in \arg\max_{\tilde{y}\geq 0} \sum_i \alpha_i V_i(p, \tilde{y}_i) \quad \text{s.t.} \quad \sum_i \tilde{y}_i = Y^*,$$
$$y^{**} \in \arg\max_{\tilde{y}\geq 0} \sum_i \alpha_i V_i(p, \tilde{y}_i) \quad \text{s.t.} \quad \sum_i \tilde{y}_i = Y^{**},$$

where V_i is the indirect utility in (4). The first-order-conditions imply $\alpha_i V_{i,y}(p, \tilde{y}_i^*) = \Lambda^*$ for some multiplier Λ, and similarly for y^{**}. As V_i is strictly concave there is a unique solution for each Y. Increasing income Y relaxes the constraint and therefore $Y^* > Y^{**}$ implies $\Lambda^* \leq \Lambda^{**}$, and thus

$$y_i^{**} \leq y_i^* \quad \text{for all } i.$$

That $\tau = 0$ implies that x^* solves the consumer problem (4) at prices p and income y_i^*. This implies that $u_i(x_i^*) = V_i(p, y_i^*)$. As $\sum_k p_k x_{i,k}^{**} = y_i^{**}$, x_i^{**} is feasible for (p, y^{**}), and we must have that $u_i(x_i^{**}) \leq V_i(p, y_i^{**})$. Because V_i is strictly increasing in income and $y_i^{**} \leq y_i^*$, we have for each i,

$$u_i(x_i^*) = V_i(p, y_i^*) \;\geq\; V_i(p, y_i^{**}) \;\geq\; u_i(x_i^{**}).$$

Multiplying by positive welfare weights $\bar{\alpha}_i$ and summing over i yields

$$\mathcal{W}(\alpha, 0) = \sum_i \bar{\alpha}_i u_i(x_i^*) \;\geq\; \sum_i \bar{\alpha}_i u_i(x_i^{**}) = \mathcal{W}(\alpha, t).$$

□

A.1.6 Proof of Proposition 1

Proof. We prove each property independently. We'll state the proof assuming there are unique solutions to the fixed point equation (8). The case where there are multiple solutions follows analogously using the selection criterion in footnote 5.

(i) Let x^* and x^{**} lie on the Pareto frontier and let $\alpha^* = \alpha(x^*)$ and $\alpha^{**} = \alpha(x^{**})$. Lemma 2 guarantees that $x^* = \mathcal{X}(\alpha^*, 0)$ and $x^{**} = \mathcal{X}(\alpha^{**}, 0)$. By definition of the Shapley-based efficiency component,

$$\mathsf{E}(x^*, x^{**}) = \tfrac{1}{2}\big[\mathcal{W}(\alpha^*, 0) - \mathcal{W}(\alpha^*, 0)\big] + \tfrac{1}{2}\big[\mathcal{W}(\alpha^{**}, 0) - \mathcal{W}(\alpha^{**}, 0)\big] = 0.$$

Hence movements along the Pareto frontier are attributed purely to redistribution.

(ii) We prove the case when x^* is Pareto efficient and x^{**} is not; the other case is analogous. Let $\alpha^* = \alpha(x^*)$, $\alpha^{**} = \alpha(x^{**})$ and $t^{**} = t(x^{**})$, then by construction we have that

$$x^* = \mathcal{X}(\alpha^*, 0), \qquad x^{**} = \mathcal{X}(\alpha^{**}, t^{**}).$$

By the Shapley definition,

$$\mathsf{E}(x^*, x^{**}) = \tfrac{1}{2}\big[\mathcal{W}(\alpha^*, t^{**}) - \mathcal{W}(\alpha^*, 0)\big] + \tfrac{1}{2}\big[\mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^{**}, 0)\big]. \tag{34}$$

Lemma 4 implies that both terms are weakly negative, so $\mathsf{E}(x^*, x^{**}) \leq 0$.

To show that the inequality is strict when $\bar{\alpha}_i > 0$ for all i, we need to show that the second bracket is strictly negative. Let

$$x^{**,PO} := \mathcal{X}(\alpha^{**}, 0), \quad p^{**} = \mathsf{P}(\alpha^{**}), \quad y_i^{**} = \sum_k p_k^{**} x_{i,k}^{**}, \quad y_i^{**,PO} = \sum_k p_k^{**} x_{i,k}^{**,PO}.$$

By construction, we have that $\sum_i y_i^{**} = (1 - \xi^{**}) Y^{max} \leq Y^{max} = \sum_i y_i^{**,PO}$. Strict concavity of $V_i(p, y)$ in y implies that a weakly tighter resource constraint in the income-allocation problem yields $y_i^{**} \leq y_i^{**,PO}$ for all i. Optimality of V_i implies that $u_i(x_i^{**}) \leq V_i(p^{**}, y_i^{**})$ since $\sum_k p_k^* x_{i,k}^{**} = \sum_k p_k^{**} x_{i,k}^{**} = y_i^{**}$. Chaining these inequalities yields

$$u_i(x_i^{**}) \leq V_i(p^{**}, y_i^{**}) \leq V_i(p^{**}, y_i^{**,PO}) = u_i(x_i^{**,PO}),$$

for all i. Since x^{**} is not Pareto optimal, there exists at least one j with $u_j(x_j^{**}) < u_j(x_j^{**,PO})$ and thus

$$\sum_i \bar{\alpha}_i u_i(x_i^{**}) \; < \; \sum_i \bar{\alpha}_i u_i(x_i^{**,PO}),$$

hence $\mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^{**}, 0) < 0$.

(iii) Invariance to price normalization follows directly from the proof of Lemma 3. The construction of $\alpha(x), \xi(x), \tau(x)$ only depends on $\mathsf{P}(\alpha(x))$, and thus are invariant to any rescaling of prices. Changing the numeraire requires choosing a different good l to set $\tau_{i,l} = 0$. Suppose instead we set $\tau_{i,l} = 0$ for some l then the distortions $\tilde{\tau}_{i,k}$ would be given by

$$1 + \tilde{\tau}_{i,k} = \frac{p_l}{p_k} \frac{u_{i,k}(x_i)}{u_{i,l}(x_i)}, \quad \forall i, k.$$

Let $\tilde{\mathcal{X}}(\alpha, \xi, \tilde{\tau})$ be the allocation implied by the new numeraire good l and distortions

$\tilde{\tau}$. By construction,

$$1+\tilde{\tau}_{i,k} = \frac{1+\tau_{i,k}}{1+\tau_{i,l}} \text{ and } 1+\tau_{i,k} = \frac{1+\tilde{\tau}_{i,k}}{1+\tilde{\tau}_{i,1}}.$$

which implies that $\tilde{\mathcal{X}}(\alpha,\xi,\tilde{\tau}) = \mathcal{X}(\alpha,\xi,\tau)$ if $1+\tilde{\tau}_{i,k} = \frac{1+\tau_{i,k}}{1+\tau_{i,l}} \forall i,k$. Thus the decomposition is invariant to the numeraire choice.

(iv) Reflexivity. By definition of the Shapley terms,

$$\begin{aligned}\mathsf{E}(x^{**},x^{*}) &= \tfrac{1}{2}\big[\mathcal{W}(\alpha^{**},t^{*}) - \mathcal{W}(\alpha^{**},t^{**})\big] + \tfrac{1}{2}\big[\mathcal{W}(\alpha^{*},t^{*}) - \mathcal{W}(\alpha^{*},t^{**})\big] \\ &= -\,\mathsf{E}(x^{*},x^{**}).\end{aligned}$$

The same antisymmetry holds for R. Since $W(x^{**}) - W(x^{*}) = \mathsf{R}(x^{*},x^{**}) + \mathsf{E}(x^{*},x^{**})$, dividing by $W(x^{**}) - W(x^{*})$ yields

$$\frac{\mathsf{R}(x^{*},x^{**})}{W(x^{**}) - W(x^{*})} = \frac{\mathsf{R}(x^{**},x^{*})}{W(x^{*}) - W(x^{**})}, \qquad \frac{\mathsf{E}(x^{*},x^{**})}{W(x^{**}) - W(x^{*})} = \frac{\mathsf{E}(x^{**},x^{*})}{W(x^{*}) - W(x^{**})}.$$

(v) Resource-based redistribution: α^{*} is defined by the fixed point

$$\alpha_i^{*} = \frac{V_{i,y}(\mathsf{P}(\alpha^{*}), \sum_k \mathsf{P}_k(\alpha^{*}) x_{i,k}^{*})^{-1}}{\sum_j V_{j,y}(\mathsf{P}(\alpha^{*}), \sum_k \mathsf{P}_k(\alpha^{*}) x_{j,k}^{*})^{-1}}.$$

As $p^{*} \propto \mathsf{P}(\alpha^{*})$, we have that $\sum_k p_k^{*} x_{i,k}^{*} = \sum_k p_k^{*} x_{i,k}^{**}$ implies

$$\sum_k \mathsf{P}_k(\alpha^{*}) x_{i,k}^{*} = \sum_k \mathsf{P}_k(\alpha^{*}) x_{i,k}^{**}.$$

Hence, α^{*} solves

$$\alpha_i^{*} = \frac{V_{i,y}(\mathsf{P}(\alpha^{*}), \sum_k \mathsf{P}_k(\alpha^{*}) x_{i,k}^{**})^{-1}}{\sum_j V_{j,y}(\mathsf{P}(\alpha^{*}), \sum_k \mathsf{P}_k(\alpha^{*}) x_{j,k}^{**})^{-1}},$$

which, by the selection procedure in footnote 5, implies that $\alpha^{**} = \alpha^{*}$ and thus $\mathsf{R}(x^{*},x^{**}) = 0$.

□

A.1.7 Relating ξ to Debreu's coefficient of resource utilization

In this section we establish the relationship between the index ξ and Debreu's coefficient of resource utilization ρ. We then construct an alternative welfare decomposition based on

Debreu's coefficient of resource utilization.

We now establish that, under appropriate conditions, the index $1-\xi$ coincides with Debreu's coefficient of resource utilization (Debreu, 1951).

Consider an economy with production possibilities described by $F(X-E) \leq 0$, where $X = \sum_i x_i$ is aggregate consumption and E is the endowment vector. Suppose that:

(a) The production function F is constant returns to scale (CRS)

(b) The relative Lagrange multipliers $\mathsf{P}_k(\alpha; E)/\mathsf{P}_1(\alpha; E)$ are independent of both α and E. Let p denote the common prices when $p_1 = 1$.

(c) The allocation wedges vanish: $\tau = 0$.

Proposition 3. *Under conditions (a)–(c), the index $1-\xi$ equals Debreu's coefficient of resource utilization ρ.*

Proof. Let x be a feasible allocation with aggregate consumption $X = \sum_i x_i$ satisfying $F(X-E^0) \leq 0$, where E^0 is the initial endowment. Define individual incomes $y_i = \sum_k p_k x_{i,k}$ and aggregate income $Y = \sum_i y_i = \sum_k p_k X_k$.

By CRS, the maximum aggregate income satisfies

$$Y^{\max}(E^0) = \max_{\tilde{X}: F(\tilde{X}-E^0)\leq 0} \sum_k p_k \tilde{X}_k = \sum_k p_k E^0_k,$$

where the second equality follows from the CRS assumption. Thus, by definition of ξ:

$$1-\xi = \frac{Y}{Y^{\max}(E^0)} = \frac{\sum_k p_k X_k}{\sum_k p_k E^0_k}.$$

Since $\tau = 0$, Lemma 3 implies that each agent's consumption bundle x_i solves the utility maximization problem (4) at prices p and income y_i. Therefore,

$$u_i(x_i) = V_i(p, y_i).$$

Following Debreu (1951), define the set of minimal endowments as

$$\mathcal{E}^{\min} = \left\{ E : \exists\, \tilde{x} \text{ with } u_i(\tilde{x}_i) \geq u_i(x_i)\ \forall i,\ F\left(\sum_i \tilde{x}_i - E\right) \leq 0,\ \text{and } E \text{ is minimal} \right\}.$$

An endowment E is minimal if there is no $E' \leq E$ with $E' \neq E$ that also supports such an allocation.

Consider any $E \in \mathcal{E}^{\min}$. There exists a Pareto-efficient allocation x^{PO} for the economy with endowment E such that $u_i(x_i^{PO}) \geq u_i(x_i)$ for all i. Since the relative Lagrange multipliers $\mathsf{P}_k(\alpha; E)/\mathsf{P}_1(\alpha; E)$ are independent of α and E, Lemma 1 implies

$$u_i(x_i^{PO}) = V_i(p, y_i^{PO}), \quad \text{where } y_i^{PO} = \sum_k p_k x_{i,k}^{PO}.$$

The constraint $u_i(x_i^{PO}) \geq u_i(x_i) = V_i(p, y_i)$ and the monotonicity of V_i in income imply $y_i^{PO} \geq y_i$ for all i. Hence,

$$\sum_i y_i^{PO} \geq \sum_i y_i = Y.$$

By CRS, the Pareto-efficient allocation satisfies $\sum_k p_k X_k^{PO} = \sum_k p_k E_k$, and since $\sum_k p_k X_k^{PO} = \sum_i y_i^{PO}$, we have

$$\sum_k p_k E_k = \sum_i y_i^{PO} \geq Y = \sum_k p_k X_k.$$

Minimality of E requires this inequality to bind, so $\sum_k p_k E_k = \sum_k p_k X_k$. Conversely, any E with $\sum_k p_k E_k = \sum_k p_k X_k$ and $F(X - E) \leq 0$ belongs to $\mathcal{E}^{\min}$ (the original allocation x itself is feasible and delivers the required utilities). Therefore,

$$\mathcal{E}^{\min} = \left\{ E : \sum_k p_k E_k = \sum_k p_k X_k \right\}.$$

Debreu's coefficient of resource utilization is

$$\rho = \max_{E \in \mathcal{E}^{\min}} \frac{\sum_k p_k E_k}{\sum_k p_k E_k^0} = \frac{\sum_k p_k X_k}{\sum_k p_k E_k^0} = 1 - \xi,$$

where the second equality uses the characterization of $\mathcal{E}^{\min}$ above. □

A.2 Alternative decompositions

In this section, we consider alternative global decompositions of the welfare change based on the (α, t) coordinates. First, we present an additive version of the decomposition. Next, we present a decomposition based on integrating marginal contributions. Third, we present an alternative decomposition based on Debreu's coefficient of resource utilization. Finally, in Appendix B.7 we compare the decompositions in our calibrated examples.

A.2.1 Additive decomposition

In this section we present an additive version of the decomposition. For a given status quo allocation x^* and a reform allocation x^{**}, we can write the welfare change as

$$\begin{aligned} W(x^{**}) - W(x^*) &= \mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^*, t^*) \\ &= (\mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^{**}, 0) + \mathcal{W}(\alpha^{**}, 0)) - (\mathcal{W}(\alpha^*, t^*) - \mathcal{W}(\alpha^*, 0) + \mathcal{W}(\alpha^*, 0)). \\ &= (\mathcal{W}(\alpha^*, 0) - \mathcal{W}(\alpha^*, t^*)) + (\mathcal{W}(\alpha^{**}, 0) - \mathcal{W}(\alpha^*, 0)) + (\mathcal{W}(\alpha^{**}, t^{**}) - \mathcal{W}(\alpha^{**}, 0)). \end{aligned}$$

The transition from x^* to x^{**} can be thought of as arising from 3 steps. First, a transition from (α^*, t^*) to $(\alpha^*, 0)$ which confers a welfare gain from improving efficiency. We let

$$\mathcal{L}(x^*) = \mathcal{W}(\alpha^*, 0) - \mathcal{W}(\alpha^*, t^*)$$

denote the efficiency gain in utils of this transition, which can be interpreted as the losses from distortions associated with the point x^*. Second, a transition along the Pareto frontier from $(\alpha^*, 0)$ to $(\alpha^{**}, 0)$ which confers a welfare change from redistribution. Finally, a transition from $(\alpha^{**}, 0)$ to (α^{**}, t^{**}) which confers a welfare loss $\mathcal{L}(x^{**})$ from worsening efficiency. Grouping these terms together, we get the following decomposition:

$$W(x^{**}) - W(x^*) = \underbrace{\mathcal{W}(\alpha^{**}, 0) - \mathcal{W}(\alpha^*, 0)}_{\mathsf{R}^{add}(x^*, x^{**})} + \underbrace{\mathcal{L}(x^*) - \mathcal{L}(x^{**})}_{\mathsf{E}^{add}(x^*, x^{**})}. \tag{35}$$

In addition to the properties of Proposition 1, the additive decomposition satisfies the following additive property

Lemma 7. *Let x^*, x^{**}, and x^{***} be three feasible allocations. Then*

$$\begin{aligned} \mathsf{R}^{add}(x^*, x^{**}) + \mathsf{R}^{add}(x^{**}, x^{***}) &= \mathsf{R}^{add}(x^*, x^{***}), \\ \mathsf{E}^{add}(x^*, x^{**}) + \mathsf{E}^{add}(x^{**}, x^{***}) &= \mathsf{E}^{add}(x^*, x^{***}). \end{aligned}$$

Proof. The proof follows directly from the construction of the decomposition

$$\begin{aligned} \mathsf{R}^{add}(x^*, x^{**}) + \mathsf{R}^{add}(x^{**}, x^{***}) &= (\mathcal{W}(\alpha^{**}, 0) - \mathcal{W}(\alpha^*, 0)) + (\mathcal{W}(\alpha^{***}, 0) - \mathcal{W}(\alpha^{**}, 0)) \\ &= (\mathcal{W}(\alpha^{***}, 0) - \mathcal{W}(\alpha^*, 0)) \\ &= \mathsf{R}^{add}(x^*, x^{***}). \end{aligned}$$

Similarly,

$$\begin{aligned}\mathsf{E}^{add}(x^*, x^{**}) + \mathsf{E}^{add}(x^{**}, x^{***}) &= \mathcal{L}(x^*) - \mathcal{L}(x^{**}) + \mathcal{L}(x^{**}) - \mathcal{L}(x^{***}) \\ &= \mathcal{L}(x^*) - \mathcal{L}(x^{***}) \\ &= \mathsf{E}^{add}(x^*, x^{***}).\end{aligned}$$

□

While this decomposition has the added feature of being additive, in practice it is not as easy to implement as the Shapley-value-based decomposition. We therefore do not recommend it for practical applications unless additivity is a necessary property.

A.2.2 Decomposition based on integrating the marginal distribution

In this section we present an alternative global decomposition based on integrating marginal contributions. For a given status quo allocation x^* and a reform allocation x^{**}, we assume that this reform arises from a series of reforms, $x(\theta)$, indexed by a perturbation parameter $\theta \in [0, 1]$. We assume that $x(0) = x^*$ and $x(1) = x^{**}$. Associated with each point $x(\theta)$ is a set of coordinates $(\alpha(\theta), t(\theta))$. We can then construct a global decomposition by integrating the marginal decomposition (14) along the path $x(\theta)$:

$$\mathsf{R}^{marg}(x^*, x^{**}) = \int_0^1 W_\alpha(\theta) \cdot \frac{d\alpha}{d\theta}(\theta) d\theta \quad \text{and} \quad \mathsf{E}^{marg}(x^*, x^{**}) = \int_0^1 W_t(\theta) \cdot \frac{dt}{d\theta}(\theta) d\theta. \quad (36)$$

By construction we have that $\mathcal{W}(x^{**}) - \mathcal{W}(x^*) = \mathsf{R}^{marg}(x^*, x^{**}) + \mathsf{E}^{marg}(x^*, x^{**})$. The marginal decomposition is also additive as long as all three points x^*, x^{**}, and x^{***} are along the same path $x(\theta)$. This follows directly from the additive property of integration. In cases where there is only one policy parameter, the marginal decomposition (36) may be a natural choice. However, when there are multiple policy parameters, there will be multiple possible paths $x(\theta)$ between x^* and x^{**}, and the choice of path will affect the decomposition.

A.2.3 Decomposition based on Debreu's coefficient of resource utilization

Here we present an alternative welfare decomposition based on Debreu's coefficient of resource utilization. We refer to this as the *BEGS-Debreu* decomposition as it is based on the notion of "coefficient of resource utilization" introduced by Debreu (1959).[28]

[28] Recently, Baqaee and Burstein (2025) extend this idea to arbitrary consumption sets that can reflect market incompleteness and other frictions.

To capture a notion of efficiency, we follow Debreu (1959) and define the scaled production set

$$\mathcal{Y}(\rho) = \{\rho X : X \in \mathcal{Y}\}$$

for $\rho \in [0, 1]$. The set $\mathcal{Y}(\rho)$ represents the production possibilities when a fraction $(1 - \rho)$ of resources is wasted.

For a given ρ, the *utility possibility set* $\mathcal{U}(\rho)$ is the set of utility vectors $\vec{u}$ achievable with some allocation x satisfying $\sum_i x_i \in \mathcal{Y}(\rho)$. The boundary of $\mathcal{U}(\rho)$ is the Pareto frontier for technology $\mathcal{Y}(\rho)$. Our assumptions guarantee that $\mathcal{U}(\rho)$ is convex, compact, and continuous in the Hausdorff metric. Moreover, $\mathcal{U}(\rho_1) \subset \mathcal{U}(\rho_2)$ whenever $\rho_1 < \rho_2$.

For any feasible allocation x with utility vector $\vec{u}$ (where $\vec{u}_i = u_i(x_i)$), we define the *efficiency coefficient* as

$$\rho(x) = \inf\{\rho : \vec{u} \in \mathcal{U}(\rho)\}.$$

By construction, $\rho(x)$ is well-defined and $\vec{u}$ lies on the boundary of $\mathcal{U}(\rho(x))$. Thus $\vec{u}$ is Pareto efficient for the scaled technology $\mathcal{Y}(\rho(x))$. We let $\omega(x)$ denote the Pareto–Negishi weights that rationalize $\vec{u}$ on this frontier, normalized so that $\sum_i \omega_i(x) = 1$.

The pair (ω, ρ) provides a coordinate system for utility vectors. Given an allocation x, we map it to $(\omega(x), \rho(x))$. Conversely, given coordinates (ω, ρ), we recover the utility vector by solving

$$\max_x \sum_i \omega_i u_i(x_i) \quad \text{subject to} \quad \sum_i x_i \in \mathcal{Y}(\rho).$$

Let $\tilde{u}(\omega, \rho)$ denote the utility vector that solves this problem. The associated social welfare function is

$$\mathcal{W}(\omega, \rho) = \sum_i \bar{\alpha}_i \tilde{u}_i(\omega, \rho).$$

This coordinate system has the following useful property for a special class of preferences.

Definition 1 (Scale-affine preferences)**.** Utilities $\{u_i\}_{i=1}^I$ are *scale-affine* if there exist functions $A : (0, \infty) \to (0, \infty)$ and $B_i : (0, \infty) \to \mathbb{R}$ such that, for all $\rho > 0$ and all $x \in \mathbb{R}_+^K$,

$$u_i(\rho x) = A(\rho)\, u_i(x) + B_i(\rho),$$

with the multiplicative term $A(\rho)$ common across agents.

One can show that CRRA and CARA utilities over homogeneous aggregators are scale-affine.

Lemma 8 (Monotonicity)**.** *When u_i are scale-affine, if x^* Pareto dominates x^{**}, then $\rho(x^*) > \rho(x^{**})$. Moreover, welfare is increasing in ρ:*

$$\mathcal{W}(\omega, \rho) \geq \mathcal{W}(\omega, \rho') \quad \textit{for all } \rho \geq \rho'.$$

Proof. We start with the first claim. Let $\vec{u}^* = (u_i(x_i^*))_{i=1}^I$ and $\vec{u}^{**} = (u_i(x_i^{**}))_{i=1}^I$ denote the utility vectors associated with allocations x^* and x^{**}, and let $\rho^* = \rho(x^*)$ and $\rho^{**} = \rho(x^{**})$. By construction, $\vec{u}^{**}$ lies on the boundary of $\mathcal{U}(\rho^{**})$. Since x^* Pareto dominates x^{**}, we have $u_i(x_i^*) \geq u_i(x_i^{**})$ for all i, with at least one strict inequality. Therefore $\vec{u}^* \notin \mathcal{U}(\rho^{**})$, which implies $\rho^* > \rho^{**}$.

Next, we prove the second claim. The welfare function $\mathcal{W}(\omega, \rho)$ is defined by the maximization problem

$$\mathcal{W}(\omega, \rho) = \max_{\tilde{x}_i} \sum_i \omega_i u_i(\tilde{x}_i) \quad \text{subject to } \sum_i \tilde{x}_i \in \mathcal{Y}(\rho).$$

Substituting $\tilde{x}_i = \rho \hat{x}_i$ and using the homogeneity of $\mathcal{Y}$, this is equivalent to

$$\max_{\hat{x}_i} \sum_i \omega_i u_i(\rho \hat{x}_i) \quad \text{subject to } \sum_i \hat{x}_i \in \mathcal{Y}(1).$$

When utilities are scale-affine, $u_i(\rho \hat{x}_i) = A(\rho) u_i(\hat{x}_i) + B_i(\rho)$. Since $A(\rho) > 0$ is common across agents, the maximizer $\hat{x}(\omega)$ is independent of ρ. Therefore,

$$\mathcal{W}(\omega, \rho) = A(\rho) \sum_i \omega_i u_i(\hat{x}_i(\omega)) + \sum_i \omega_i B_i(\rho).$$

Since $A(\rho)$ is increasing in ρ, we conclude that $\mathcal{W}(\omega, \rho)$ is increasing in ρ. $\square$

The coordinate ω captures redistribution through Pareto–Negishi weights, while ρ measures efficiency. We decompose the welfare change from x^* to x^{**} using Shapley values. The redistribution component is

$$\mathsf{R}^{BEGS-Deb}(x^*, x^{**}) = \tfrac{1}{2}\big[\mathcal{W}(\omega^{**}, \rho^*) - \mathcal{W}(\omega^*, \rho^*)\big] + \tfrac{1}{2}\big[\mathcal{W}(\omega^{**}, \rho^{**}) - \mathcal{W}(\omega^*, \rho^{**})\big], \quad (37)$$

and the efficiency component is

$$\mathsf{E}^{BEGS-Deb}(x^*, x^{**}) = \tfrac{1}{2}\big[\mathcal{W}(\omega^*, \rho^{**}) - \mathcal{W}(\omega^*, \rho^*)\big] + \tfrac{1}{2}\big[\mathcal{W}(\omega^{**}, \rho^{**}) - \mathcal{W}(\omega^{**}, \rho^*)\big]. \quad (38)$$

This yields the decomposition for the scale-affine class of preferences:

$$W(x^{**}) - W(x^*) = \mathsf{R}^{BEGS-Deb}(x^*, x^{**}) + \mathsf{E}^{BEGS-Deb}(x^*, x^{**}). \tag{39}$$

Proposition 4. *The decomposition* (39) *satisfies:*

(i) If x^* and x^{**} are both on the Pareto frontier given by $\mathcal{Y}$, then $\mathsf{E}^{BEGS-Deb}(x^*, x^{**}) = 0$.

(ii) If x^* is on the Pareto frontier but x^{**} is not, then $\mathsf{E}^{BEGS-Deb}(x^*, x^{**}) \leq 0$. If x^{**} is on the Pareto frontier but x^* is not, then $\mathsf{E}^{BEGS-Deb}(x^*, x^{**}) \geq 0$. When all social welfare weights $\bar{\alpha}_i$ are strictly positive, these inequalities are strict.

(iii) The decomposition is symmetric: the share attributed to redistribution (or efficiency) is unchanged when reversing the direction of the reform,

$$\frac{\mathsf{R}^{BEGS-Deb}(x^*, x^{**})}{W(x^{**}) - W(x^*)} = \frac{\mathsf{R}^{BEGS-Deb}(x^{**}, x^*)}{W(x^*) - W(x^{**})}, \qquad \frac{\mathsf{E}^{BEGS-Deb}(x^*, x^{**})}{W(x^{**}) - W(x^*)} = \frac{\mathsf{E}^{BEGS-Deb}(x^{**}, x^*)}{W(x^*) - W(x^{**})}.$$

Proof. We prove each property in turn:

(i) If x^* and x^{**} are both on the Pareto frontier of $\mathcal{Y}$, then by definition $\rho(x^*) = \rho(x^{**}) = 1$. Substituting $\rho^* = \rho^{**}$ into the expression for $\mathsf{E}^{BEGS-Deb}(x^*, x^{**})$ in equation (38) yields $\mathsf{E}^{BEGS-Deb}(x^*, x^{**}) = 0$.

(ii) If x^* is on the Pareto frontier, $\rho^* = 1$. If x^{**} is not, then x^{**} lies in the interior of the feasible set (or can be improved upon), implying $\rho^{**} < 1$. By Lemma 8, $\mathcal{W}(\omega, \rho)$ is increasing in ρ. Therefore, $\mathcal{W}(\omega^*, \rho^{**}) \leq \mathcal{W}(\omega^*, \rho^*)$ and $\mathcal{W}(\omega^{**}, \rho^{**}) \leq \mathcal{W}(\omega^{**}, \rho^*)$. Substituting these terms into equation (38) implies $\mathsf{E}^{BEGS-Deb}(x^*, x^{**}) \leq 0$.

If all welfare weights $\bar{\alpha}_i$ are strictly positive, then $\mathcal{W}(\omega, \rho)$ is strictly increasing in ρ, so the inequality becomes strict. The reverse case where x^{**} is on the frontier follows symmetrically.

(iii) From the definition,

$$\mathsf{E}^{BEGS-Deb}(x^{**}, x^*) = \tfrac{1}{2}\big[\mathcal{W}(\omega^{**}, \rho^*) - \mathcal{W}(\omega^{**}, \rho^{**})\big] + \tfrac{1}{2}\big[\mathcal{W}(\omega^*, \rho^*) - \mathcal{W}(\omega^*, \rho^{**})\big].$$

Rearranging terms shows $\mathsf{E}^{BEGS-Deb}(x^{**}, x^*) = -\mathsf{E}^{BEGS-Deb}(x^*, x^{**})$. Similarly, $\mathsf{R}^{BEGS-Deb}(x^{**}, x^*) = -\mathsf{R}^{BEGS-Deb}(x^*, x^{**})$ and $W(x^*) - W(x^{**}) = -(W(x^{**}) - W(x^*))$. Taking ratios establishes the result.

$\square$

One property this decomposition does not satisfy is the *income-based redistribution* property because the coordinate system is based on utilities rather than income. To illustrate, consider again the economy in Example 6 with two agents and idiosyncratic risk. As we noted in that discussion, the *income-based redistribution* property requires all welfare gains to be attributed to efficiency, so that our decomposition finds $\mathsf{R}(x^*, x^{**}) = 0$.

However, the decomposition (39) assigns some welfare gains to redistribution. The reform increases agent 2's utility while leaving agent 1's utility unchanged, causing the Pareto–Negishi weights to shift: $\omega_2(x^{**}) > \omega_2(x^*)$. In fact, the redistribution component can range from arbitrarily small to all of the welfare gain.

To see this concretely, assume log utility and let $c_1 = K$ be agent 1's allocation and $c_2(s)$ be agent 2's state-contingent allocation, normalized so that $\sum_s \Pr(s) c_2(s) = 1$ and total consumption is $C = K + 1$. Let c_2^{ce} denote agent 2's certainty equivalent, defined by $\ln(c_2^{ce}) = \sum_s \Pr(s) \ln(c_2(s))$. The Debreu coordinates are

$$\rho = \frac{K + c_2^{ce}}{K+1}, \quad \omega_1 = \frac{K}{K + c_2^{ce}}, \quad \omega_2 = \frac{c_2^{ce}}{K + c_2^{ce}},$$

and utilities become $\tilde{u}_i(\omega, \rho) = \ln(\omega_i) + \ln(\rho) + \ln(C)$. A reform that reduces agent 2's risk implies $\hat{c}_2^{ce} > 0$, with utility changes $\hat{u}_1 = 0$ and $\hat{u}_2 = \hat{c}_2^{ce}/c_2^{ce}$. For Pareto weights $\bar{\alpha}_i$, the welfare change decomposes as

$$\hat{W} = \bar{\alpha}_2 \frac{\hat{c}_2^{ce}}{c_2^{ce}} = \underbrace{\frac{\hat{c}_2^{ce}}{K + c_2^{ce}}}_{\hat{\mathsf{E}}} + \underbrace{\bar{\alpha}_1 \widehat{\ln \omega_1} + \bar{\alpha}_2 \widehat{\ln \omega_2}}_{\hat{\mathsf{R}}}. \tag{40}$$

As $K \to \infty$, $\hat{\mathsf{E}} \to 0$ and redistribution accounts for 100% of the welfare gain. As $\bar{\alpha}_2 \to 0$, $\hat{W} \to 0$ while $\hat{\mathsf{E}} > 0$, so efficiency accounts for $+\infty$% and redistribution for $-\infty$%.

The same result holds for Example 1. For ease of exposition, assume log-log preferences $u_i = \ln(c_i) + \ln(\ell_i)$. Agent 1 has K units of effective leisure; agent 2 has one unit and faces labor tax τ with lump-sum rebate. The aggregate resource constraint is $C + L \leq K + 1$. Agent 1 optimally sets $c_1 = \ell_1 = K/2$; agent 2's equilibrium is $c_2 = 1/(2-\tau)$ and $\ell_2 = (1-\tau)/(2-\tau)$.

To map this to Debreu coordinates, observe that on the Pareto frontier with efficiency ρ allocations satisfy $c_i = \ell_i = \frac{1}{2} x_i$ where $x_1 + x_2 \leq \rho(K+1)$ with corresponding Pareto weights $\omega_i \propto x_i$. The associated utilities are $u_i = 2\ln(\frac{1}{2} x_i)$ which implies that the distorted allocations have the same utilities as an allocation on the Pareto frontier with $x_1 = K$ and $x_2 = 2\sqrt{1-\tau}/(2-\tau)$. Therefore, the Debreu coordinates are

$$\rho = \frac{K + x_2}{K+1}, \quad \omega_1 = \frac{K}{K+x_2}, \quad \omega_2 = \frac{x_2}{K+x_2}.$$

A tax reduction ($\hat{\tau} < 0$) starting from $\tau > 0$ yields $\hat{x}_2 > 0$, with utility changes $\hat{u}_1 = 0$ and $\hat{u}_2 = 2\hat{x}_2/x_2$. The welfare decomposition is

$$\hat{W} = 2\bar{\alpha}_2 \frac{\hat{x}_2}{x_2} = \underbrace{\frac{2\hat{x}_2}{K+x_2}}_{\hat{\mathsf{E}}} + \underbrace{2\bar{\alpha}_1 \widehat{\ln \omega_1} + 2\bar{\alpha}_2 \widehat{\ln \omega_2}}_{\hat{\mathsf{R}}}. \tag{41}$$

As before, $K \to \infty$ implies redistribution accounts for 100%; $\bar{\alpha}_2 \to 0$ implies efficiency accounts for $+\infty$% and redistribution for $-\infty$%.

A.3 Separating insurance from aggregate and idiosyncratic risk

In this section we show how to decompose the insurance component into subcomponents capturing insurance against aggregate and idiosyncratic risk. We assume that there are N physical goods and $S = S^{agg} \times S^{idio}$ states of nature so that $K = N \times S^{agg} \times S^{idio}$. We let $s_i = (z, \epsilon_i)$ represent the state of nature with z representing the aggregate state common to all agents and ϵ_i representing the idiosyncratic state specific to agent i. Following Section 3.5, we let $x_{i,n}(s_i)$ be the allocation of good n to agent i in state $s_i = (z, \epsilon_i)$, let $p_n(s_i)$ be the price of good n in state s_i, and let $\tau_{i,n}(s_i)$ be the tax on good n for agent i in state s_i. We let $\Pr(s_i)$ be the probability that state s_i is realized and $u_i(x_i) = \sum_{s_i} \Pr(s_i)\, v_i(x_i(s_i), s_i)$ be the expected utility of agent i. Without loss of generality we can write $\Pr(s_i) = \Pr(z)\Pr(\epsilon_i|z)$.

We construct x_i^{ins} in the same way as in Section 3.5 by solving the optimization problem

$$\max_{\tilde{x}_i} u_i(\tilde{x}_i) \quad \text{s.t.} \quad \sum_{s_i} p_n(s_i)\tilde{x}_i(s_i) \leq \sum_{s_i} p_n(s_i) x_i(s_i) \quad \text{for all } n.$$

This allocation achieves the highest utility that agent i can attain if they were able to freely reallocate physical goods across *all* states of nature using the Arrow–Debreu state prices.

To separate insurance against aggregate and idiosyncratic risk, we construct the intermediate allocation $x_i^{ins,z}$ as the solution to

$$\max_{\tilde{x}_i} u_i(\tilde{x}_i) \quad \text{s.t.} \quad \sum_{\epsilon_i} p_n(z,\epsilon_i)\tilde{x}_i(z,\epsilon_i) \leq \sum_{\epsilon_i} p_n(z,\epsilon_i) x_i(z,\epsilon_i) \quad \text{for all } n \text{ and } z.$$

This allocation represents the best insurance that agent i can obtain if they are only allowed to reallocate their income across idiosyncratic states of nature conditional on the aggregate state z.

With the wedges $\tau^g_{i,n}$ and $\tau^{ins}_{i,n}(s_i)$ defined as in Section 3.5, we can decompose the insurance wedge into aggregate and idiosyncratic components:

$$\begin{aligned}1+\tau^{ins}_{i,n}(s_i) &= \frac{v_{i,n}(x_i(s_i),s_i)/v_{i,1}(x_i(1),1)}{v_{i,n}(x^{ins}_i(s_i),s_i)/v_{i,1}(x^{ins}_i(1),1)} \\ &= \underbrace{\frac{v_{i,n}(x^{ins,z}_i(s_i),s_i)/v_{i,1}(x^{ins,z}_i(1),1)}{v_{i,n}(x^{ins}_i(s_i),s_i)/v_{i,1}(x^{ins}_i(1),1)}}_{1+\tau^{ins,z}_{i,n}(z)} \times \underbrace{\frac{v_{i,n}(x_i(s_i),s_i)/v_{i,1}(x_i(1),1)}{v_{i,n}(x^{ins,z}_i(s_i),s_i)/v_{i,1}(x^{ins,z}_i(1),1)}}_{1+\tau^{ins,\epsilon_i}_{i,n}(z,\epsilon_i)}.\end{aligned}$$

In this construction, $\tau^{ins,z}$ captures insurance imperfections against aggregate risk, measured by deviations in the marginal rates of substitution of the partial insurance allocation $x^{ins,z}_i$ relative to the numeraire across states from full insurance. The term τ^{ins,ϵ_i} captures insurance imperfections against idiosyncratic risk, measured by deviations in the marginal rates of substitution of the actual allocation x_i relative to the partial insurance allocation $x^{ins,z}_i$.

This approach expresses welfare $W(x)$ in coordinates $(\alpha, \xi, \tau^g, \tau^{ins})$ where τ^{ins} is summarized by $(\tau^{ins,z}, \tau^{ins,\epsilon_i})$. By computing the Shapley value contributions of $\tau^{ins,z}$ and τ^{ins,ϵ_i} to the insurance component E^{ins}, we obtain the term $\mathsf{E}^{ins,z}$ that captures inefficiencies in insurance against aggregate risk and a term $\mathsf{E}^{ins,\epsilon_i}$ that captures inefficiencies in providing insurance against idiosyncratic risk. This extends (19) to

$$W(x^{**}) - W(x^*) = \mathsf{R}(x^*, x^{**}) + \mathsf{E}^{pr}(x^*, x^{**}) + \mathsf{E}^{g}(x^*, x^{**}) + \mathsf{E}^{ins,z}(x^*, x^{**}) + \mathsf{E}^{ins,\epsilon_i}(x^*, x^{**}). \tag{42}$$

A.4 Proofs for Section 4

A.4.1 Proof of Proposition 2

We begin by stating a more general result. When $\mathsf{P}(\alpha)$ is non-degenerate, an arbitrary perturbation of the policy $(\hat{G}, \hat{T}, \hat{\tau}, \hat{\varsigma})$ will be associated with a change in the implied PN weights $\hat{\alpha}$ and thus a change in relative prices $\hat{p}$. To interpret our decomposition, we distinguish between real and nominal changes in policy. Since price changes result from the implied change in PN weights $\hat{\alpha}$, nominal effects are isolated to the redistribution component, while real effects are captured by the productive and allocative efficiency components.

A given change in policy $(\hat{G}, \hat{T}, \hat{\tau}, \hat{\varsigma})$ will be associated with a change in net surplus $\hat{S}$. We define the surplus-to-GDP ratio as $\mathcal{S} = \frac{S}{\sum_k p_k A_k}$ and the real change in net surplus

$$\hat{S}^r = \widehat{\mathcal{S}} \sum_k p_k^* A_k \tag{43}$$

as the change in net surplus that would result if the ratio of surplus to GDP changed while prices were held constant.

We let $x_i^{net} = x_i - a_i$ be the net consumption of agent i given their initial endowments a_i. Net consumption is funded by net transfers T_i^{net} from the government: $T_i^{net} = \sum_k p_k x_i^{net}$. We define the real change in net transfers as the real component of nominal change in net transfers:

$$\hat{T}_i^{net,r} = \widehat{T_i^{net}} - \sum_k \hat{p}_k x_{i,k}^{net,*} = \sum_k p_k^* \hat{x}_{i,k}, \tag{44}$$

where the last equality follows from the fact that $\hat{a}_{i,k} = 0$ for all i, k since endowments are fixed.

As we are focused on the local decomposition around x^* we will assume that all coordinates $\alpha(x), \xi(x), \tau(x)$ are locally differentiable at x^*, as well as the pricing kernel $P(\alpha)$. Furthermore, we will assume that the Marshallian demand functions are locally differentiable at (q_i^*, m_i^*) for all i where $q_{i,k}^* = p_{i,k}^*(1 + \tau_{i,k}^*)$ and $m_i^* = \sum_k q_{i,k}^* x_{i,k}^*$. Finally, we will assume that all goods are normal goods so as to satisfy the assumptions of Lemma 6. Finally, we assume that the income distribution is non-degenerate so that $x_{i,k}^* > 0$ for some k for all i. With these assumptions, we can then express our decomposition locally in terms of the change in these real quantities.

Proposition 5. *Let η_i, $\vartheta_{i,k}^c$, ω_i be defined as $\eta_i = \sum_k p_k^* \tau_{i,k}^* \frac{\partial \mathsf{x}_{i,k}}{\partial m}$, $\vartheta_{i,k}^c = \sum_l \tau_{i,l}^* p_l^* x_{i,l}^* \zeta_{i,lk}^c$, and $\omega_i = \frac{V_{i,y}/V_{i,yy}}{\sum_{i'} V_{i',y}/V_{i',yy}}$, where derivatives of V_i are evaluated at (p^*, y_i^*). Then $\hat{\mathsf{R}}$, $\hat{\mathsf{E}}^{pr}$, $\hat{\mathsf{E}}^{al}$ in decomposition (17) satisfy*

$$\hat{\mathsf{R}} = \sum_i \bar{\alpha}_i \frac{u_{i,1}}{1 - \eta_i} \Big(\hat{T}_i^{net,r} - \omega_i \hat{S}^r + \sum_k \vartheta_{i,k}^c \widehat{\ln p_k} \Big),$$

$$\hat{\mathsf{E}}^{pr} = \sum_i \bar{\alpha}_i \frac{u_{i,1}}{1 - \eta_i} \omega_i \hat{S}^r, \qquad \hat{\mathsf{E}}^{al} = \sum_i \bar{\alpha}_i \frac{u_{i,1}}{1 - \eta_i} \sum_k \vartheta_{i,k}^c \ln(\widehat{1 + \tau_{i,k}}),$$

where $\hat{S}^r$ is defined in (43) *and $\hat{T}_i^{net,r}$ is defined in* (44).

Proof. For a given allocation x with coordinates (α, ξ, τ), let $p(\alpha) \in \mathbb{R}^K +$ be such that $p(\alpha) \propto \mathsf{P}(\alpha)$ with the normalization $p_1(\alpha) = 1$. Define after-tax prices and individual incomes by

$$q_{i,k}(\alpha,\tau) := (1+\tau_{i,k})\, p_k(\alpha), \qquad m_i(\alpha,\xi,\tau) := \sum_k q_{i,k}(\alpha,\tau)\, \mathcal{X}_{i,k}(\alpha,\xi,\tau).$$

Let $\mathsf{x}_i(q,m)$ be Marshallian demand functions. By construction, $\mathsf{x}_i(q_i(\alpha,\tau), m_i(\alpha,\xi,\tau)) = \mathcal{X}(\alpha,\xi,\tau)$ for all i, which implies that individual incomes $m_i(\alpha,\xi,\tau)$ are characterized by

$$\sum_k p_k(\alpha)\, \mathsf{x}_{i,k}\big(q_i(\alpha,\tau), m_i(\alpha,\xi,\tau)\big) \;=\; y_i(\alpha,\xi). \tag{45}$$

Here, $y(\alpha,\xi)$ is the distribution of income net of taxes chosen by the planner with PN weights α and aggregate resources $(1-\xi)Y^{max}(p(\alpha))$:

$$y(\alpha,\xi) \in \arg\max_{\tilde{y}\geq 0} \sum_i \alpha_i V_i\big(p(\alpha), \tilde{y}_i\big) \quad \text{s.t.} \;\; \sum_i \tilde{y}_i = (1-\xi)Y^{max}(p(\alpha)). \tag{46}$$

Social welfare can be written as

$$\mathcal{W}(\alpha,\xi,\tau) \;=\; \sum_i \bar{\alpha}_i\, V_i\big(q_i(\alpha,\tau),\, m_i(\alpha,\xi,\tau)\big). \tag{47}$$

All derivatives below are evaluated at the baseline (α^*,ξ^*,τ^*); we use $p^* := p(\alpha^*)$, and write $u_{i,1}$ for $u_{i,1}(x_i^*)$.

We start with the aggregate efficiency term (derivative w.r.t. ξ). By construction $\xi = -\mathcal{S}$. Differentiating (47) and using the envelope property of V_i in m yields

$$\frac{\partial \mathcal{W}}{\partial \xi}\,\hat{\xi} \;=\; \sum_i \bar{\alpha}_i\, V_{i,m}\, \frac{\partial m_i}{\partial \xi}\,\hat{\xi} \;=\; \sum_i \bar{\alpha}_i\, u_{i,1}\, \frac{\partial m_i}{\partial \xi}\,\hat{\xi}.$$

To obtain $\partial m_i/\partial\xi$, differentiate (45) (noting that q_i does not change with ξ):

$$\sum_k p_k^*\, \frac{\partial \mathsf{x}_{i,k}}{\partial m_i}\, \frac{\partial m_i}{\partial \xi}\,\hat{\xi} \;=\; \frac{\partial y_i}{\partial \xi}\,\hat{\xi}.$$

Let $1-\eta_i := \sum_k p_k^* \frac{\partial \mathsf{x}_{i,k}}{\partial m_i}$ (equivalently, $\eta_i = \sum_k p_k^* \tau_{i,k}^* \frac{\partial \mathsf{x}_{i,k}}{\partial m_i}$ using $\sum_k q_{i,k}^*\, \partial \mathsf{x}_{i,k}/\partial m_i = 1$). From (46), $\partial y_i/\partial\xi\,\hat{\xi} = -\omega_i Y^{max,*}\,\hat{\xi}$, where $\omega_i = \big(V_{i,y}/V_{i,yy}\big) \big/ \sum_j \big(V_{j,y}/V_{j,yy}\big)$. Hence

$$\frac{\partial m_i}{\partial \xi}\,\hat{\xi} = \frac{-\omega_i Y^{max,*}}{1-\eta_i}\,\hat{\xi},$$

which implies

$$\frac{\partial \mathcal{W}}{\partial \xi}\,\hat{\xi} = \sum_i \bar{\alpha}_i\,\frac{u_{i,1}}{1-\eta_i}\,\big(-\omega_i Y^{max,*}\,\hat{\xi}\big) \;=\; \sum_i \bar{\alpha}_i\,\frac{u_{i,1}}{1-\eta_i}\,\omega_i \hat{S}^r,$$

This yields the productive-efficiency term $\hat{\mathsf{E}}^{pr}$ in the statement.

Next, we turn to the allocative efficiency term (derivative w.r.t. τ). Differentiating (47) and using Roy's identity, $\frac{\partial V_i}{\partial q_{i,k}} = -V_{i,m} x_{i,k}$, gives

$$\begin{aligned}\frac{\partial \mathcal{W}}{\partial \tau}\cdot\hat{\tau} &= \sum_i \bar{\alpha}_i\Big(\sum_k V_{i,q_{i,k}}\,p_k^*\,\hat{\tau}_{i,k} + V_{i,m}\,\frac{\partial m_i}{\partial \tau}\cdot\hat{\tau}_i\Big)\\ &= \sum_i \bar{\alpha}_i\,V_{i,m}\Big(\frac{\partial m_i}{\partial \tau}\cdot\hat{\tau}_i - \sum_k p_k^* x_{i,k}^*\,\hat{\tau}_{i,k}\Big).\end{aligned}$$

From differentiating (45) w.r.t. τ and using (i) $\sum_k q_{i,k}^*\,\partial \mathsf{x}_{i,k}/\partial q_{i,k'} + x_{i,k'}^* = 0$, and (ii) the Slutsky decomposition $\partial \mathsf{x}_{i,l}/\partial q_{i,k} = \partial \mathsf{x}_{i,l}^c/\partial q_{i,k} - (\partial \mathsf{x}_{i,l}/\partial m_i)\,x_{i,k}^*$, one obtains

$$(1-\eta_i)\,\frac{\partial m_i}{\partial \tau}\cdot\hat{\tau}_i = \sum_{l,k} p_l^*\,\tau_{i,l}^*\,\frac{\partial \mathsf{x}_{i,l}^c}{\partial q_{i,k}}\,p_k^*\,\hat{\tau}_{i,k} + (1-\eta_i)\sum_k p_k^* x_{i,k}^*\,\hat{\tau}_{i,k}.$$

Therefore,

$$\frac{\partial m_i}{\partial \tau}\cdot\hat{\tau}_i - \sum_k p_k^* x_{i,k}^*\,\hat{\tau}_{i,k} = \frac{1}{1-\eta_i}\sum_k \vartheta_{i,k}^c\,\frac{\hat{\tau}_{i,k}}{1+\tau_{i,k}^*}, \qquad \vartheta_{i,k}^c := \sum_l \tau_{i,l}^* p_l^* x_{i,l}^*\,\zeta_{i,lk}^c,$$

where $\zeta_{i,lk}^c := \big(\partial \mathsf{x}_{i,l}^c/\partial q_{i,k}\big)\,(q_{i,k}^*/x_{i,l}^*)$ is the compensated elasticity. Using $V_{i,m} = u_{i,1}$ and $\ln(\widehat{1+\tau_{i,k}}) = \hat{\tau}_{i,k}/(1+\tau_{i,k}^*)$ gives

$$\frac{\partial \mathcal{W}}{\partial \tau}\cdot\hat{\tau} = \sum_i \bar{\alpha}_i\,\frac{u_{i,1}}{1-\eta_i}\sum_k \vartheta_{i,k}^c\,\ln(\widehat{1+\tau_{i,k}}),$$

which is $\hat{\mathsf{E}}^{al}$.

Finally, we turn to the Redistribution term (derivative w.r.t. α). Differentiating (47) w.r.t. α (prices and incomes both move) and using Roy's identity yields

$$\begin{aligned}\frac{\partial \mathcal{W}}{\partial \alpha}\cdot\hat{\alpha} &= \sum_i \bar{\alpha}_i\Big(\sum_k V_{i,q_{i,k}}\,(1+\tau_{i,k}^*)\,\hat{p}_k + V_{i,m}\,\frac{\partial m_i}{\partial \alpha}\cdot\hat{\alpha}\Big)\\ &= \sum_i \bar{\alpha}_i\,V_{i,m}\Big(\frac{\partial m_i}{\partial \alpha}\cdot\hat{\alpha} - \sum_k (1+\tau_{i,k}^*) x_{i,k}^*\,\hat{p}_k\Big),\end{aligned}$$

where $\hat{p}_k := \frac{\partial P_k}{\partial \alpha} \cdot \hat{\alpha}$. Differentiating (45) w.r.t. α and writing $\hat{y}_i = \sum_k (p_k^* \hat{x}_{i,k} + x_{i,k}^* \hat{p}_k)$ together with $\partial y_i / \partial \xi \, \hat{\xi} = \omega_i \hat{Y}^r$ gives

$$\sum_{l,k} p_l^* \frac{\partial \mathsf{x}_{i,l}}{\partial q_{i,k}} (1 + \tau_{i,k}^*) \, \hat{p}_k + \sum_l p_l^* \frac{\partial \mathsf{x}_{i,l}}{\partial m_i} \frac{\partial m_i}{\partial \alpha} \cdot \hat{\alpha} = \hat{T}_i^{net,r} - \omega_i \hat{S}^r, \qquad \hat{T}_i^{net,r} := \sum_k p_k^* \hat{x}_{i,k}.$$

Proceeding as in the τ-derivative case, we obtain

$$(1 - \eta_i) \frac{\partial m_i}{\partial \alpha} \cdot \hat{\alpha} = \hat{T}_i^{net,r} - \omega_i \hat{S}^r + \sum_{l,k} p_l^* \, \tau_{i,l}^* \frac{\partial \mathsf{x}_{i,l}^c}{\partial q_{i,k}} (1 + \tau_{i,k}^*) \, \hat{p}_k + (1 - \eta_i) \sum_k (1 + \tau_{i,k}^*) x_{i,k}^* \, \hat{p}_k.$$

Hence,

$$\frac{\partial m_i}{\partial \alpha} \cdot \hat{\alpha} - \sum_k (1 + \tau_{i,k}^*) x_{i,k}^* \, \hat{p}_k = \frac{1}{1 - \eta_i} \left(\hat{T}_i^{net,r} - \omega_i \hat{S}^r + \sum_k \vartheta_{i,k}^c \frac{\hat{p}_k}{p_k^*} \right)$$

Using $V_{i,m} = u_{i,1}$ then yields

$$\frac{\partial \mathcal{W}}{\partial \alpha} \cdot \hat{\alpha} = \sum_i \bar{\alpha}_i \frac{u_{i,1}}{1 - \eta_i} \Big(\hat{T}_i^{net,r} - \omega_i \hat{S}^r + \sum_k \vartheta_{i,k}^c \, \widehat{\ln p_k} \Big),$$

which is $\hat{\mathsf{R}}$ as stated. Collecting the three derivatives gives the proposition. □

The terms in Proposition 5 have the same interpretation as the terms in Proposition 2, except that they now refer to real quantities. The production efficiency term $\hat{\mathsf{E}}^{pr}$ arises from the change in real surplus, $\hat{S}^r$, induced by the change in the surplus to GDP ratio $\hat{\mathcal{S}}$. Consumer i gets an additional $\omega_i \hat{S}^r$ in real income, but the net amount needs to be adjusted by the fiscal externality and converted into utils to compute welfare. The redistribution term arises from two terms. The first term captures welfare gains from changes in the real net-transfers that depart from this rule. The second term captures the effect of changes in the deadweight loss induced by changes in relative prices which are a result of changes in inequality (as captured by the PN weights α). It follows the same structure as allocative efficiency term in Proposition 2.

When $\mathsf{P}(\alpha)$ is single-valued this last term vanishes and the real changes collapse to the nominal changes giving us Proposition 2.

A.4.2 Derivations for equations in Section 4.2

Here we provide additional details for the derivation of equation (25) and (27) in Section 4.2. We also derive the counterparts for the CARA economy later in this section.

Proof. Consider a single physical good ($n = 1$) indexed by idiosyncratic states s with probabilities $\Pr(s)$. We normalize $(n, s) = (1, 1)$ as the numeraire. Let consumption be $x_{i,1}(s)$. Preferences are CRRA with coefficient $\sigma > 0$ so that

$$u_i(x_i) = \sum_s \Pr(s) \frac{x_{i,1}(s)^{1-\sigma}}{1-\sigma} \equiv \sum_s \Pr(s) g(x_{i,1}(s)).$$

The technology is an endowment with feasibility $\sum_i \mathbb{E}[x_{i,1}] \leq A_1$. In this environment the efficient Arrow–Debreu prices equal probabilities, $p_1(s) = \Pr(s)$, independently of α.

Represent any feasible allocation via the triplet $(w, \bar{x}_1, \varepsilon)$ defined by

$$\bar{x}_1 := \frac{1}{I} \sum_i \mathbb{E}[x_{i,1}], \qquad w_i := \frac{\mathbb{E}[x_{i,1}]}{\bar{x}_1}, \qquad \varepsilon_i(s) := \frac{x_{i,1}(s)}{\mathbb{E}[x_{i,1}]}, \quad \mathbb{E}[\varepsilon_i] = 1.$$

Equivalently, $x_{i,1}(s) = w_i \, \bar{x}_1 \, \varepsilon_i(s)$. The mapping between (α, ξ, τ) and $(w, \bar{x}_1, \varepsilon)$ is one-to-one for each variable: (a) $\xi = 1 - \frac{I\bar{x}_1}{A_1}$ from feasibility; (b) with CRRA, $V_i(p, y) = \kappa(p) \frac{y^{1-\sigma}}{1-\sigma}$ so $V_{i,y} \propto y^{-\sigma}$ and hence $\alpha_i \propto y_i^\sigma = (\mathbb{E}[x_{i,1}])^\sigma \propto w_i^\sigma$; (c) wedges satisfy

$$1 + \tau_i(s) = \frac{p_1(1)}{p_1(s)} \frac{\Pr(s) x_{i,1}(s)^{-\sigma}}{\Pr(1) x_{i,1}(1)^{-\sigma}} = \left(\frac{\varepsilon_i(s)}{\varepsilon_i(1)} \right)^{-\sigma},$$

so wedges are equivalent to $\varepsilon_i(s)$. As a result, linearizing welfare with respect to (α, ξ, τ) is equivalent to linearizing with respect to $(w, \bar{x}_1, \varepsilon)$.

To obtain (25), write social welfare as $W = \sum_i \bar{\alpha}_i \, \mathbb{E}[g(x_{i,1})]$. Taking the first-order (Gateaux) variation around a baseline allocation $x_{i,1}^*(s)$ yields

$$\widehat{W} = \sum_i \bar{\alpha}_i \, \mathbb{E}\left[g'(x_{i,1}^*) \, x_{i,1}^* \, \widehat{\ln x_{i,1}} \right]. \tag{48}$$

With $x_{i,1}(s) = w_i \bar{x}_1 \varepsilon_i(s)$ we have the identity

$$\widehat{\ln x_{i,1}(s)} = \widehat{\ln \bar{x}_1} + \widehat{\ln w_i} + \widehat{\ln \varepsilon_i(s)}.$$

Substituting into (48) gives

$$\widehat{W} = \underbrace{\sum_i \bar{\alpha}_i \, \mathbb{E}[g'(x^*_{i,1})x^*_{i,1}] \, \widehat{\ln \bar{x}_1}}_{=: \hat{\mathsf{E}}^{pr}} + \underbrace{\sum_i \bar{\alpha}_i \, \mathbb{E}[g'(x^*_{i,1})x^*_{i,1}] \, \widehat{\ln w_i}}_{=: \hat{\mathsf{R}}}$$
$$+ \underbrace{\sum_i \bar{\alpha}_i \, \mathbb{E}[g'(x^*_{i,1})x^*_{i,1} \, \widehat{\ln \varepsilon_i(s)}]}_{=: \hat{\mathsf{E}}^{ins}} .$$

Since $w_i = \mathbb{E}[x_{i,1}]/\bar{x}_1$, $\widehat{\ln w_i} = \ln(\widehat{\mathbb{E}x_{i,1}/}\bar{x}_1)$. When feasibility binds ($I\bar{x}_1 = A_1$), $\widehat{\ln \bar{x}_1} = 0$, which yields (25).

Aggregate shocks. Now let the state s index *aggregate* shocks with physical probabilities $\Pr(s)$. Feasibility holds state-by-state:

$$\sum_i x_{i,1}(s) \leq A_1(s) \qquad \forall s,$$

and we write aggregate consumption in state s as $X_1(s) := \sum_i x_{i,1}(s)$ (so in a pure endowment economy, $X_1(s) = A_1(s)$).

Given homothetic preferences, the (efficient) Arrow–Debreu price vector is independent of α up to normalization and can be taken as

$$p_1(s) \propto \Pr(s) \, X_1(s)^{-\sigma}.$$

Normalize $\sum_s p_1(s) = 1$ and define the associated risk-neutral expectation $\widetilde{\mathbb{E}}[Z] := \sum_s p_1(s) \, Z(s)$. The (date-0) income of household i is its value at state prices:

$$y_i = \sum_s p_1(s) \, x_{i,1}(s) = \widetilde{\mathbb{E}} \, x_{i,1}.$$

Represent any allocation by the triplet $(w, \widetilde{\mathbb{E}}\bar{x}_1, \varepsilon)$ defined by

$$\widetilde{\mathbb{E}}\bar{x}_1 := \frac{1}{I}\sum_i \widetilde{\mathbb{E}}[x_{i,1}], \qquad w_i := \frac{\widetilde{\mathbb{E}}[x_{i,1}]}{\widetilde{\mathbb{E}}\bar{x}_1}, \qquad \varepsilon_i(s) := \frac{x_{i,1}(s)}{\widetilde{\mathbb{E}}[x_{i,1}]}, \quad \widetilde{\mathbb{E}}[\varepsilon_i] = 1.$$

Equivalently, $x_{i,1}(s) = w_i \, (\widetilde{\mathbb{E}}\bar{x}_1) \, \varepsilon_i(s)$.

As in the idiosyncratic case, the Pareto–Negishi weights depend only on the (value) incomes. With CRRA, $V_i(p, y) = \kappa(p)\frac{y^{1-\sigma}}{1-\sigma}$ so $V_{i,y} \propto y^{-\sigma}$ and hence $\alpha_i \propto y_i^{\sigma} = (\widetilde{\mathbb{E}}x_{i,1})^{\sigma}$, i.e.

inequality is governed by dispersion in $\widetilde{\mathbb{E}}x_{i,1}$.

Social welfare is $W = \sum_i \bar{\alpha}_i \,\mathbb{E}[g(x_{i,1})]$. The first-order (Gateaux) variation around a baseline allocation x^* is

$$\widehat{W} = \sum_i \bar{\alpha}_i \,\mathbb{E}\Big[g'(x^*_{i,1})\, x^*_{i,1}\, \widehat{\ln x_{i,1}}\Big].$$

Using $x_{i,1}(s) = w_i(\widetilde{\mathbb{E}}\bar{x}_1)\varepsilon_i(s)$ gives the identity

$$\widehat{\ln x_{i,1}(s)} = \widehat{\ln(\widetilde{\mathbb{E}}\bar{x}_1)} + \widehat{\ln w_i} + \widehat{\ln \varepsilon_i(s)}.$$

Substituting into the linearization and grouping terms yields

$$\widehat{W} = \underbrace{\sum_i \bar{\alpha}_i \,\mathbb{E}[g'(x^*_{i,1})x^*_{i,1}]\, \widehat{\ln(\widetilde{\mathbb{E}}\bar{x}_1)}}_{=:\,\hat{\mathsf{E}}^{pr}} + \underbrace{\sum_i \bar{\alpha}_i \,\mathbb{E}[g'(x^*_{i,1})x^*_{i,1}]\, \widehat{\ln w_i}}_{=:\,\hat{\mathsf{R}}} + \underbrace{\sum_i \bar{\alpha}_i \,\mathbb{E}[g'(x^*_{i,1})x^*_{i,1}\, \widehat{\ln \varepsilon_i(s)}]}_{=:\,\hat{\mathsf{E}}^{ins}}.$$

Finally, note that $\widetilde{\mathbb{E}}\bar{x}_1 = \widetilde{\mathbb{E}}X_1/I$, so

$$\widehat{\ln w_i} = \ln\Big(\widehat{\frac{\widetilde{\mathbb{E}}x_{i,1}}{\widetilde{\mathbb{E}}X_1}}\Big), \qquad \widehat{\ln \varepsilon_i(s)} = \ln\Big(\widehat{\frac{x_{i,1}(s)}{\widetilde{\mathbb{E}}x_{i,1}}}\Big).$$

For a purely redistributive feasible marginal reform under aggregate risk, state-by-state feasibility implies $\sum_i \hat{x}_{i,1}(s) = 0$ for all s, hence $\hat{X}_1(s) = 0$ for all s. Since prices $p_1(s) \propto \Pr(s)X_1(s)^{-\sigma}$ are then unchanged, $\widehat{\ln(\widetilde{\mathbb{E}}X_1)} = 0$ and therefore $\widehat{\ln(\widetilde{\mathbb{E}}\bar{x}_1)} = 0$, so $\hat{\mathsf{E}}^{pr} = 0$.

Replacing $g'(x) = v'(x)$ delivers (27):

$$\widehat{W} = \sum_i \bar{\alpha}_i \,\mathbb{E}\left[v'(x^*_{i,1})\, x^*_{i,1} \ln\Big(\widehat{\frac{\widetilde{\mathbb{E}}x_{i,1}}{\widetilde{\mathbb{E}}X_1}}\Big)\right] + \sum_i \bar{\alpha}_i \,\mathbb{E}\left[v'(x^*_{i,1})\, x^*_{i,1} \ln\Big(\widehat{\frac{x_{i,1}}{\widetilde{\mathbb{E}}x_{i,1}}}\Big)\right].$$

Derivations of counterparts of equations (25) and (27) for CARA preferences. Consider a single physical good ($n = 1$) indexed by idiosyncratic states s with probabilities $\Pr(s)$. We normalize $(n, s) = (1, 1)$ as the numeraire. Let consumption be $x_{i,1}(s)$. Preferences are CARA with coefficient $\gamma > 0$ so that

$$u_i(x_i) = \sum_s \Pr(s)\left(-\tfrac{1}{\gamma}e^{-\gamma x_{i,1}(s)}\right) \equiv \sum_s \Pr(s) h(x_{i,1}(s)).$$

The technology is an endowment with feasibility $\sum_i \mathbb{E}[x_{i,1}] \le A_1$. In this environment the efficient Arrow–Debreu prices equal probabilities, $p_1(s) = \Pr(s)$, independently of α.

Represent any feasible allocation via the triplet $(\bar{x}_1, \Delta, \delta)$ defined by

$$\bar{x}_1 := \frac{1}{I}\sum_i \mathbb{E}[x_{i,1}], \quad \Delta_i := \mathbb{E}[x_{i,1}] - \bar{x}_1, \quad \delta_i(s) := x_{i,1}(s) - \mathbb{E}[x_{i,1}], \quad \mathbb{E}[\delta_i] = 0.$$

Equivalently, $x_{i,1}(s) = \bar{x}_1 + \Delta_i + \delta_i(s)$. The mapping between (α, ξ, τ) and $(\bar{x}_1, \Delta, \delta)$ is one-to-one for each variable: (a) $\xi = 1 - \frac{I\bar{x}_1}{A_1}$ from feasibility; (b) with CARA $V_i(p, y_i) = -\frac{\kappa(p_1)}{\gamma}e^{-\gamma y_i}$ so $V_{i,y} = \kappa(p_1)e^{-\gamma y_i}$ and hence $\alpha_i \propto 1/V_{i,y} \propto e^{\gamma(\bar{x}_1+\Delta_i)} \propto e^{\gamma(\Delta_i)}$; (c) wedges satisfy

$$1 + \tau_i(s) = \frac{p_1(1)}{p_1(s)}\,\frac{\Pr(s)h'(x_{i,1}(s))}{\Pr(1)h'(x_{i,1}(1))} = \exp\big(-\gamma(\delta_i(s) - \delta_i(1))\big),$$

so, fixing the numeraire state to $s = 1$, wedges are equivalent to the profile $\delta_i(s)$. As a result linearizing welfare with respect to (α, ξ, τ) is equivalent to linearizing with respect to $(\bar{x}_1, \Delta, \delta)$.

To obtain the CARA analogue of (25), write social welfare as $W = \sum_i \bar{\alpha}_i\,\mathbb{E}[h(x_{i,1})]$ and take a first-order (Gateaux) variation around a baseline allocation $x_i^*(s)$. Using

$$\widehat{x_{i,1}(s)} = \widehat{\bar{x}_1} + \widehat{\Delta_i} + \widehat{\delta_i(s)}$$

we obtain

$$\widehat{W} = \sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x_{i,1}^*)\big]\widehat{\bar{x}_1} + \sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x_{i,1}^*)\big]\widehat{\Delta_i} + \sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x_{i,1}^*)\,\widehat{\delta_i}\big]. \tag{49}$$

Grouping terms gives the three components in exact parallel to the CRRA case (with levels in place of logs):

$$\widehat{W} = \underbrace{\sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x_{i,1}^*)\big]\widehat{\bar{x}_1}}_{=:\,\hat{\mathsf{E}}^{pr}} + \underbrace{\sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x_{i,1}^*)\big]\widehat{\Delta_i}}_{=:\,\hat{\mathsf{R}}} + \underbrace{\sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x_{i,1}^*)\,\widehat{\delta_i}\big]}_{=:\,\hat{\mathsf{E}}^{ins}}.$$

Thus the CARA decomposition mirrors the CRRA representation with: (i) aggregate efficiency acting through $\bar{x}_1$ (level); (ii) redistribution acting through expected levels via Δ_i; and (iii) insurance acting through changes in the dispersion of state deviations $\delta_i(s)$ around means.

To obtain the CARA analogue of (27), we follow the same steps as above. With aggregate risk feasibility now requires $\sum_i x_{i,1}(s) \leq A_1(s)$ for all s. In this environment the supporting price vector satisfies $p_1(s) \propto \Pr(s)e^{-\gamma A_1(s)}$, independently of α, which we normalize so that

$\sum_s p_1(s) = 1$. We let $\widetilde{\mathbb{E}}$ denote expectations computed with respect to this measure.

Any feasible allocation can be represented by the triplet $(\widetilde{\mathbb{E}}\bar{x}_1, \Delta, \delta)$ defined by

$$\widetilde{\mathbb{E}}\bar{x}_1 := \frac{1}{I}\sum_i \widetilde{\mathbb{E}}[x_{i,1}], \quad \Delta_i := \widetilde{\mathbb{E}}[x_{i,1}] - \widetilde{\mathbb{E}}\bar{x}_1, \quad \delta_i(s) := x_{i,1}(s) - \widetilde{\mathbb{E}}[x_{i,1}], \quad \widetilde{\mathbb{E}}[\delta_i] = 0.$$

Equivalently, $x_{i,1}(s) = \widetilde{\mathbb{E}}\bar{x}_1 + \Delta_i + \delta_i(s)$. The mapping between (α, ξ, τ) and $(\widetilde{\mathbb{E}}\bar{x}_1, \Delta, \delta)$ is one-to-one for each variable: (a) $\xi = 1 - \frac{I\widetilde{\mathbb{E}}\bar{x}_1}{\widetilde{\mathbb{E}}A_1}$ from feasibility; (b) with CARA $V_{i,y} \propto e^{-\gamma y}$ so $\alpha_i \propto e^{\gamma\widetilde{\mathbb{E}}x_{i,1}} \propto e^{\gamma(\Delta_i)}$; (c) wedges satisfy

$$1 + \tau_i(s) = \frac{p_1(1)}{p_1(s)}\frac{\Pr(s)h'(x_{i,1}(s))}{\Pr(1)h'(x_{i,1}(1))} = \exp\big(\gamma(A_1(s) - A_1(1))\big)\exp\big(-\gamma(\delta_i(s) - \delta_i(1))\big),$$

so wedges are equivalent to the profile $\delta_i(s)$. As a result linearizing welfare with respect to (α, ξ, τ) is equivalent to linearizing with respect to $(\widetilde{\mathbb{E}}\bar{x}_1, \Delta, \delta)$.

The same linearization steps yield

$$\widehat{W} = \underbrace{\sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x^*_{i,1})\big]\widehat{\widetilde{\mathbb{E}}\bar{x}_1}}_{=:\,\hat{\mathsf{E}}^{pr}} + \underbrace{\sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x^*_{i,1})\big]\widehat{\Delta_i}}_{=:\,\hat{\mathsf{R}}} + \underbrace{\sum_i \bar{\alpha}_i\,\mathbb{E}\big[h'(x^*_{i,1})\,\widehat{\delta_i}\big]}_{=:\,\hat{\mathsf{E}}^{ins}}.$$

Thus the CARA decomposition mirrors the CRRA representation with: (i) aggregate efficiency acting through $\widetilde{\mathbb{E}}\bar{x}_1$ (level); (ii) redistribution acting through expected levels via Δ_i; and (iii) insurance acting through changes in the dispersion of state deviations $\delta_i(s)$ around risk-neutral means. □

A.5 Additional details for Section 5

This section provides details for Example 5. The worker (agent 2) has utility $u_2(c_2, n_2) = \log(c_2) - \frac{\chi}{1+\gamma}n_2^{1+\gamma}$, where $n_2 = \bar{l} - l_2$ is labor supply and l_2 is leisure. Under Policy II, the worker faces budget constraint $c_2 = (1+\tau)n_2 + a_2$, where τ is a proportional subsidy funded by the tax T on the capitalist.

Lemma 9. $\operatorname{sign}\left(\frac{dl_2}{dT}\right) = -\operatorname{sign}(a_2)$.

Proof. The first-order condition is $(1+\tau)/c_2 = \chi n_2^\gamma$. Log-differentiating with respect to τ:

$$\frac{1}{1+\tau} - \frac{1}{c_2}\frac{\partial c_2}{\partial \tau} = \frac{\gamma}{n_2}\frac{\partial n_2}{\partial \tau}.$$

Differentiating the budget constraint: $\frac{\partial c_2}{\partial \tau} = n_2 + (1+\tau)\frac{\partial n_2}{\partial \tau}$. Substituting and rearranging:

$$\underbrace{(c_2 - (1+\tau)n_2)}_{=a_2} \cdot \frac{1}{1+\tau} = (\gamma c_2 + (1+\tau)n_2)\,\frac{1}{n_2}\frac{\partial n_2}{\partial \tau}.$$

Since $\gamma c_2 + (1+\tau)n_2 > 0$, we have $\text{sign}\left(\frac{\partial n_2}{\partial \tau}\right) = \text{sign}(a_2)$. Because $l_2 = 1 - n_2$ and T and τ move together, $\text{sign}\left(\frac{dl_2}{dT}\right) = -\text{sign}(a_2)$. □

Next we can prove the following two corollaries about the Dávila and Schaab (2022) and the Benabou (2002) and Floden (2001) decompositions.

Corollary 1. *If $a_2 < 0$ (> 0) and a_1 is sufficiently large, then the Benabou (2002) and Floden (2001) decomposition finds that $\frac{d\omega_{\mathsf{E}}}{dT} > 0$ (< 0) at $T = 0$.*

Proof. Benabou (2002) and Floden (2001) define the efficiency component ω_{E} such that $u((1+\omega_{\mathsf{E}})\bar{c}^*, \bar{\ell}^*) = u(\bar{c}^{**}, \bar{\ell}^{**})$, where $\bar{c} = \frac{1}{I}\sum_i c_i$ and $\bar{\ell} = \frac{1}{I}\sum_i l_i$. Thus, $\text{sgn}(\frac{d\omega_{\mathsf{E}}}{dT}) = \text{sgn}(\frac{du(\bar{c},\bar{\ell})}{dT})$. Differentiating utility with respect to T and using the resource constraint $d\bar{c} = -\frac{1}{2}d\ell_2$ (since $l_1 = 1$ is constant and $I = 2$):

$$\frac{du(\bar{c},\bar{\ell})}{dT} = u_c(\bar{c},\bar{\ell})\frac{d\bar{c}}{dT} + u_\ell(\bar{c},\bar{\ell})\frac{d\bar{\ell}}{dT} = \left(-\frac{1}{2}u_c(\bar{c},\bar{\ell}) + \frac{1}{2}u_\ell(\bar{c},\bar{\ell})\right)\frac{dl_2}{dT}.$$

Substituting the functional forms $u_c = 1/\bar{c}$ and $u_\ell = \chi(1-\bar{\ell})^\gamma$:

$$\frac{du(\bar{c},\bar{\ell})}{dT} = \frac{1}{2}\left(\chi(1-\bar{\ell})^\gamma - \frac{1}{\bar{c}}\right)\frac{dl_2}{dT}.$$

At $T = 0$, agent 2's optimality implies $1/c_2 = \chi(1-l_2)^\gamma$. Since $l_1 = 1$ and $l_2 < 1$, we have $\bar{\ell} = (1+l_2)/2 > l_2$, so $(1-\bar{\ell})^\gamma < (1-l_2)^\gamma$. Also $\bar{c} = (c_1+c_2)/2$. If a_1 is sufficiently large, then c_1 is large, making $\bar{c}$ large and $1/\bar{c}$ small. Specifically, if a_1 is large enough, $1/\bar{c} < \chi(1-\bar{\ell})^\gamma$, making the term in parentheses positive. If $a_2 < 0$, the Lemma implies $\frac{dl_2}{dT} > 0$. Thus, the product is positive, so $\frac{d\omega_{\mathsf{E}}}{dT} > 0$. □

Corollary 2. *If $a_2 < 0$ and $T > 0$ (< 0) then the Dávila and Schaab (2022) decomposition finds that the efficiency component $dC > 0$ for a marginal change $\hat{T} > 0$ (< 0).*

Proof. The marginal version of the Dávila and Schaab (2022) decomposition constructs the marginal change in welfare in consumption units as $dC = \sum_i dc_i$, where $dc_i = \hat{c}_i + \frac{u_{i,\ell}}{u_{i,c}}\hat{l}_i$. For agent 1, $dc_1 = \hat{c}_1$. For agent 2, the first-order condition $u_{2,\ell}/u_{2,c} = 1 + \tau$ implies

$dc_2 = \hat{c}_2 + (1+\tau)\hat{l}_2$. Summing these and using the resource constraint $\hat{c}_1 + \hat{c}_2 = -\hat{l}_2$, we find

$$dC = \hat{c}_1 + \hat{c}_2 + (1+\tau)\hat{l}_2 = -\hat{l}_2 + (1+\tau)\hat{l}_2 = \tau\hat{l}_2.$$

From the lemma, if $a_2 < 0$, then $dl_2/dT > 0$, so $\hat{l}_2$ has the same sign as $\hat{T}$. Thus, if T and $\hat{T}$ have the same sign (moving away from laissez-faire), then τ and $\hat{l}_2$ have the same sign, implying $dC = \tau\hat{l}_2 > 0$. □

B Additional details for Section 6

In this appendix, we provide additional details for the quantitative application in Section 6 of the main text. First, we review the numerical methods used to compute the baseline and post-reform allocations. Next, we provide details on how to implement our decomposition. Some routine but technical derivations are provided in the final subsection.

B.1 Computing the baseline and post-reform allocations

The numerical methods for computing the equilibrium path associated with a path of fiscal policy are standard, and we provide only a brief overview here.

We first solve for the steady-state equilibrium and then compute the transition path using a shooting algorithm. The key ingredients for the steady-state solution are the stationary aggregate capital and labor supply functions, $\mathcal{A}(\bar{r}, \bar{w}, \bar{Tr})$ and $\mathcal{L}(\bar{r}, \bar{w}, \bar{Tr})$. Similarly, the transition dynamics require the time-varying aggregate capital and labor supply functions, $\mathcal{A}_t(\{r_s, w_s, Tr_s\}_{s=0}^{\bar{T}})$ and $\mathcal{L}_t(\{r_s, w_s, Tr_s\}_{s=0}^{\bar{T}})$. Since the computation is similar for both applications, we detail the general procedure below.

Stationary aggregate supply. The stationary functions $\mathcal{A}(\bar{r}, \bar{w}, \bar{Tr})$ and $\mathcal{L}(\bar{r}, \bar{w}, \bar{Tr})$ are computed as follows:

1. **Household Problem.** Given prices $\{\bar{r}, \bar{w}\}$ and transfers $\bar{Tr}$, we solve for the household's optimal consumption $c(a, z)$ and labor supply $l(a, z)$ policy rules using the Endogenous Grid Method (EGM) of Carroll (2006) adapted to include labor supply.

2. **Stationary Distribution.** We solve for the stationary distribution vector μ using the histogram method of Young (2010). We discretize the state space into a grid of asset points $\{a_j\}_{j=1}^{J}$ and productivity points $\{z_k\}_{k=1}^{K}$. We index the combined state space by $i = (j, k)$, so that $a(i) = a_j$ and $z(i) = z_k$. Using the optimal policy rules,

we construct the transition matrix H over this discretized state space, where element $H_{i',i}$ represents the probability of transitioning from state i to state i'. The stationary distribution μ is the eigenvector associated with the unit eigenvalue of H.

3. **Aggregation.** We compute aggregate capital and effective labor by summing over the state space:

$$\mathcal{A}(\bar{r}, \bar{w}, \bar{Tr}) = \sum_i \mu_i a(i), \qquad \mathcal{L}(\bar{r}, \bar{w}, \bar{Tr}) = \sum_i \mu_i z(i) l(a(i), z(i)).$$

Transition aggregate supply. The transition functions $\mathcal{A}_t(\{r_s, w_s, Tr_s\}_{s=0}^{\bar{T}})$ and $\mathcal{L}_t(\{r_s, w_s, Tr_s\}_{s=0}^{\bar{T}})$ take sequences of prices and transfers as inputs. This procedure assumes that for $t > \bar{T}$, the sequence of prices and transfers equals the new steady-state values. The computation proceeds in three steps:

1. **Household Problem (Backward Iteration).** We solve for the sequence of optimal policy rules $\{c_t(a, z), l_t(a, z)\}_{t=0}^{\bar{T}}$ by backward induction. The terminal policy rules at $\bar{T}$ are given by the steady-state rules associated with the terminal prices. For $t < \bar{T}$, we iterate backward using the time-dependent Euler equation.

2. **Distribution (Forward Iteration).** Starting from the initial steady-state distribution μ_0, we simulate the distribution forward using the computed policy rules. At each date t, we construct the transition matrix H_t implied by $\{c_t, l_t\}$ and update the distribution: $\mu_{t+1} = H_t \mu_t$.

3. **Aggregation.** We compute the path of aggregates by summing over the distributions at each date:

$$\mathcal{A}_t = \sum_i \mu_{t,i} a(i), \qquad \mathcal{L}_t = \sum_i \mu_{t,i} z(i) l_t(a(i), z(i)).$$

With these functions defined, the steady-state equilibrium is found by iterating on prices until markets clear and the government budget is balanced. Similarly, the transition equilibrium is found by iterating on the sequences of prices and transfers.

The procedure above yields the full distribution of consumption and labor choices $x_{i,t}(s^t) = (c_{i,t}(s^t), l_{i,t}(s^t))$ for an agent starting with wealth $a(i)$ and productivity $z(i)$ at date 0. Specifically, we have a sequence of transition matrices $\{H_t\}$ and a sequence of grid point choices

$x_t(j) = (c_t(j), l_t(j))$,[29] which represent the optimal choices of an agent with state $(a(j), z(j))$ at date t. We can use these objects to compute conditional expectations of endogenous variables. Define $\mathbb{E}_0 x_t(i)$ as the vector of conditional expectations of x_t for an agent with assets $a(i)$ and productivity $z(i)$ at date 0. We can compute this vector by multiplying the transition matrices in reverse order:

$$\mathbb{E}_0 x_t = H_0' \times H_1' \times \cdots \times H_{t-1}' \times x_t.$$

These conditional expectations are a necessary ingredient for computing the welfare decomposition.

Calibration for Section 6.1. The full list of parameters used in the quantitative application is provided in Table 1. In Figure 8 we plot the transition of key macro aggregates for the optimal tax reform. The higher tax rate leads to declines in aggregate quantities and wages, while government revenues rise and transfers jump on impact before converging to their long-run level.

Numerical details for Section 6.2. To solve for the equilibrium in Section 6.2, we fix the path of the capital stock at its initial steady-state level, $K_t = K^{ss}$, and take a sequence of debt $\{B_t\}_{t=0}^{\bar{T}}$ as exogenous. This implies a sequence of aggregate assets $A_t = K^{ss} + B_t$. We search for a sequence of after-tax returns on savings $\{\tilde{r}_t\}_{t=0}^{\bar{T}}$ and lump-sum transfers $\{Tr_t\}_{t=0}^{\bar{T}}$ that satisfy the asset market clearing condition and the government budget constraint. Note that in this application, labor is inelastic ($L_t = 1$), so we need only evaluate the aggregate asset function $\mathcal{A}_t(\{\tilde{r}_s, w_s, Tr_s\}_{s=0}^{\bar{T}})$. The wage w_t is determined by the marginal product of labor at K^{ss}. Given a guess for $\{\tilde{r}_t, Tr_t\}$, we compute the implied aggregate assets from $\mathcal{A}_t$. We iterate on $\tilde{r}_t$ to ensure $\mathcal{A}_t = A_t$. The lump-sum transfers Tr_t are then determined residually from the government budget constraint, given the tax revenue from labor and capital and the path of debt issuance. The capital wedge ς_t is given by $1 + \varsigma_t = \tilde{r}_t/(F_K(K^{ss}, 1) - \delta)$. We iterate until convergence.

Figure 9 plots the transition paths for aggregate assets A, interest rates R, public debt B, and transfers Tr when the tax reform increases government debt. The reform increases public liquidity, raises interest rates while keeping them negative, and pushes up transfers. These changes expand household budget sets and deliver a Pareto improvement.

[29] For $t > \bar{T}$, we have $H_t = \bar{H}'$ and $x_t(j) = \bar{x}'(j)$, where $\bar{H}'$ and $\bar{x}'$ represent the transition matrices and grid point choices in the steady state of the economy following the tax reform.

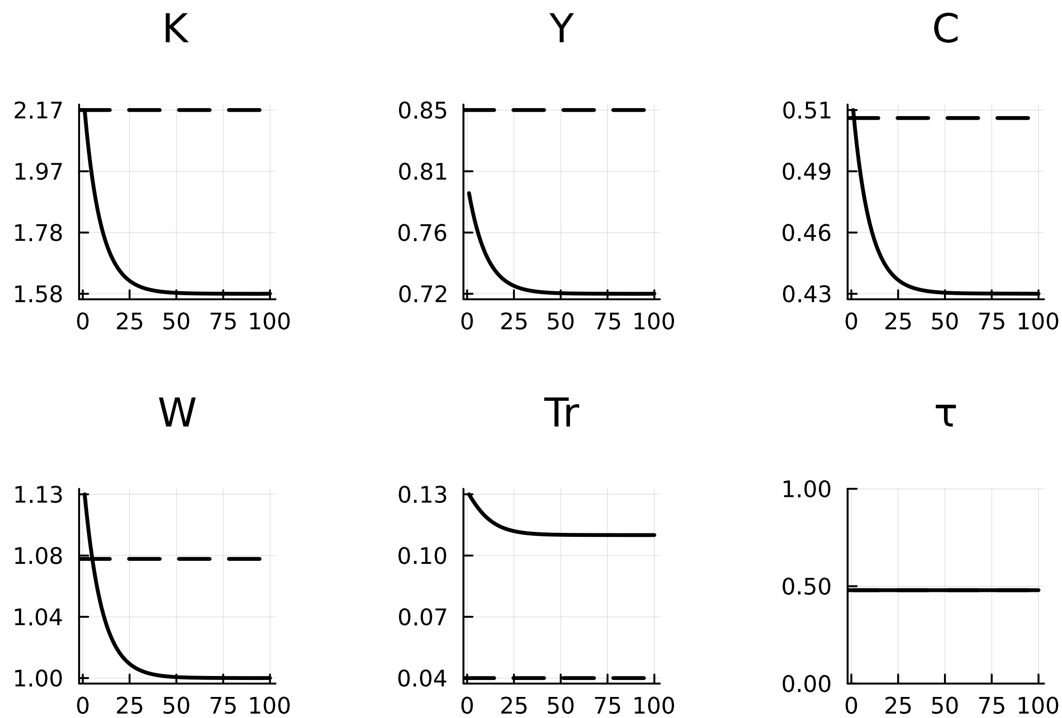

Figure 8: Transition paths for capital, output, consumption, wages, and transfers when the tax rate increases from 30% to the optimal rate of 48%. The dashed horizontal lines denote the initial steady-state values.

TABLE 1: Calibration and targets

Parameter	Value	Target
Preferences		
ψ	1.0	Log utility in consumption
γ	2.0	Frisch elasticity $1/\gamma = 0.5$
Ψ	17.89	Average hours $\bar{l} = 1/3$
β	0.96	After-tax net return $\approx 4\%$
Technology		
θ	0.36	Capital share
δ	0.10	Implies $I/K = 10\%$
Fiscal policy (baseline)		
τ	0.30	Marginal income tax $\approx 30\%$
G/Y	0.15	Government purchases share (excl. transfers)
B/Y	1.00	Government debt-to-output ratio
Idiosyncratic earnings process		
ρ_p	0.9695	Persistence of permanent component
σ_p	$\sqrt{0.0384}$	Std. dev. of permanent shock
σ_e	$\sqrt{0.053}$	Std. dev. of transitory shock

Notes: All parameters correspond to an annual calibration. Fiscal ratios are expressed relative to output. Earnings-process parameters follow Krueger et al. (2016).

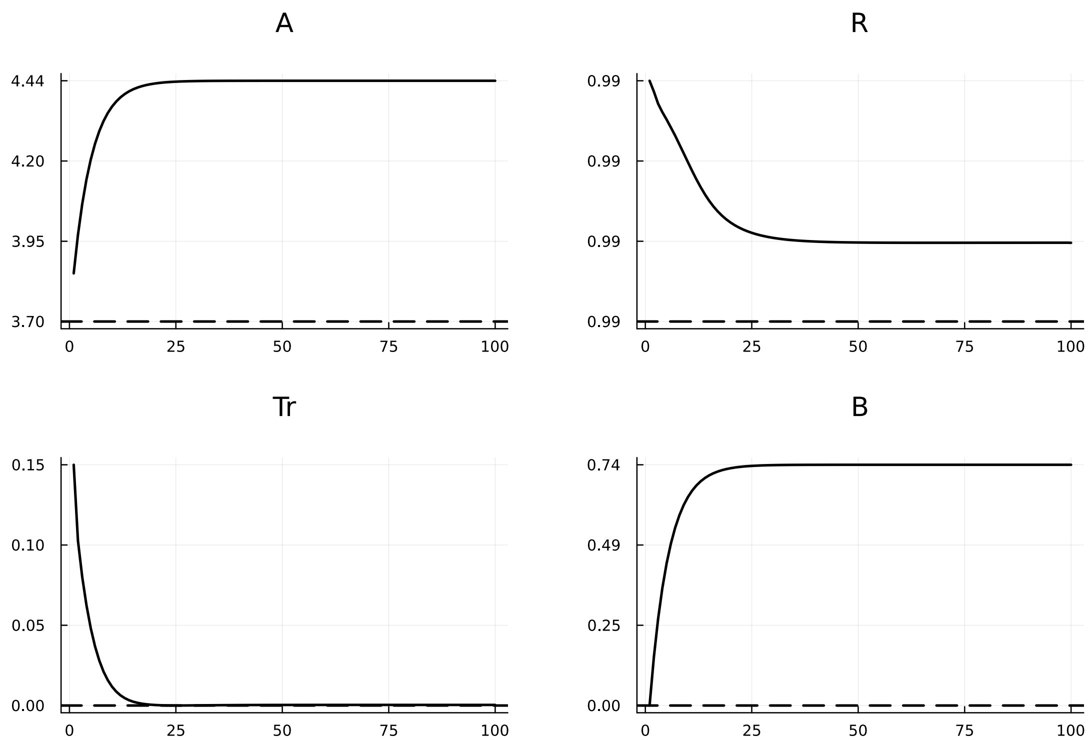

Figure 9: Transition paths for A, R, B, and Tr when debt increases. The reform raises debt from 0 to 50% of GDP. The dashed horizontal lines denote the initial steady-state values.

B.2 Welfare decomposition

Here we detail how to apply our decomposition to Section 6 settings. We assume that labor is supplied elastically as in the first application in Section 6.1. The case with inelastic labor supply in the application in Section 6.2 is immediate as a special case.

Goods, states, and expectations. Histories s^t record labor productivity $\epsilon_{i,t}(s^t)$, which is individual i's labor productivity following a history of shocks s^t. Physical goods are indexed by time t, and idiosyncratic states by histories s^t, so an allocation is $x = \{c_{i,t}(s^t),\, l_{i,t}(s^t)\}_{i,t,s^t}$.[30]

For any history-dependent random variable $z_{i,t}(s^t)$, define the date-0 expectation operator

$$\mathbb{E}_0[z_{i,t}] \;\equiv\; \sum_{s^t} \pi_t(s^t)\, z_{i,t}(s^t).$$

Utility. Define lifetime utility from an allocation x_i as

$$u(x_i) = \sum_{t=0}^{\infty} \beta^t\, \mathbb{E}_0\left[\frac{c_{i,t}^{1-\sigma}}{1-\sigma} - \Psi\, \frac{l_{i,t}^{1+\gamma}}{1+\gamma}\right].$$

Feasibility. Aggregates X are given by sequences $\{C_t, L_t\}_{t\geq 0}$ where

$$C_t = \int \mathbb{E}_0[c_{i,t}]\, di,$$
$$L_t = \int \mathbb{E}_0[\epsilon_{i,t} l_{i,t}]\, di.$$

The set of feasible aggregate allocations is denoted by $\mathcal{Y}(K_0)$ and defined by $X \in \mathcal{Y}(K_0)$ if there exist a strictly positive sequence of capital stocks $\{K_t\}_{t\geq 0}$ such that

$$C_t + G + K_{t+1} \leq A K_t^{\theta} L_t^{1-\theta} + (1-\delta)K_t \qquad \text{for all } t \geq 0,$$

with K_0 given.

B.2.1 Coordinate mapping $x \mapsto (\alpha, \tau, \xi)$

Given an allocation $x = \{c_{i,t}(s^t),\, l_{i,t}(s^t)\}_{i,t,s^t}$, we construct its welfare coordinates $(\alpha(x), \xi(x), \tau(x))$ as follows.

[30] Agents are endowed with $\bar{\ell} = 1$ unit of time, so leisure is $\bar{\ell} - l_{i,t}(s^t)$. We assume Ψ is calibrated such that the non-negativity constraint on leisure never binds. Hence defining an allocation in terms of labor supply $l_{i,t}(s^t)$ is without loss.

Pareto–Negishi planner. For any vector of weights $\alpha = \{\alpha_i\}_i$, the Pareto–Negishi (PN) planner solves

$$\max_{\{c_{i,t}^{PF}(s^t), l_{i,t}^{PF}(s^t), K_{t+1}^{PF}, C_t^{PF}, L_t^{PF}\}_{t,s^t}} \int \alpha_i \, u(x_i^{PF}) \, di$$

$$\text{s.t.} \quad C_t^{PF} + G + K_{t+1}^{PF} \leq A(K_t^{PF})^{\theta} (L_t^{PF})^{1-\theta} + (1-\delta)K_t^{PF}, \qquad \forall t \geq 0,$$

$$C_t^{PF} = \int \mathbb{E}_0[c_{i,t}^{PF}] \, di, \qquad L_t^{PF} = \int \mathbb{E}_0[\epsilon_{i,t} l_{i,t}^{PF}] \, di, \qquad \forall t \geq 0,$$

$$K_0^{PF} = K_0.$$

This problem determines a path of Pareto-frontier (PF) aggregates $\{K_t^{PF}(\alpha), L_t^{PF}(\alpha), C_t^{PF}(\alpha)\}_{t\geq 0}$, supporting prices $\{q_t^{PF}(\alpha), w_t^{PF}(\alpha)\}_{t\geq 0}$ for output and effective labor at date t in units of consumption at date 0, and individual PF allocations $\{c_{i,t}^{PF}(s^t;\alpha), l_{i,t}^{PF}(s^t;\alpha)\}_{t,s^t}$, where $x_i^{PF} \equiv \{c_{i,t}^{PF}(s^t), l_{i,t}^{PF}(s^t)\}_{t,s^t}$.

Indirect utility and value of resources. Given any price sequence (q^{PF}, w^{PF}), we can construct the corresponding state-contingent prices as

$$q_t^{PF}(s^t) = q_t^{PF} \, \pi_t(s^t), \qquad w_t^{PF}(s^t) = w_t^{PF} \, \pi_t(s^t) \, \epsilon_{i,t}(s^t).$$

With these prices, define the *value of resources* available to agent i under allocation x as the present value of consumption and leisure at PF prices:

$$y_i(q^{PF}, w^{PF}; x) := \sum_{t=0}^{\infty} \mathbb{E}_0\Big[q_t^{PF} \, c_{i,t} + w_t^{PF} \, \epsilon_{i,t} \, (\bar{\ell} - l_{i,t})\Big]. \tag{50}$$

The indirect utility function $V_i(y; q^{PF}, w^{PF})$ is defined as

$$V_i(y; q^{PF}, w^{PF}) = \max_{\tilde{c}_i, \tilde{l}_i} u(\tilde{x}_i) \quad \text{s.t.} \quad \sum_{t=0}^{\infty} \mathbb{E}_0\Big[q_t^{PF} \, \tilde{c}_{i,t} + w_t^{PF} \, \epsilon_{i,t} \, (\bar{\ell} - \tilde{l}_{i,t})\Big] \leq y,$$

where $\tilde{x}_i \equiv \{\tilde{c}_{i,t}(s^t), \tilde{l}_{i,t}(s^t)\}_{t,s^t}$.

Coordinate $\alpha(x)$. The PN weights $\alpha(x)$ solve the fixed point problem

$$\alpha_i \propto V_{i,y}^{-1}\Big(y_i(q^{PF}(\alpha), w^{PF}(\alpha); x); q^{PF}(\alpha), w^{PF}(\alpha)\Big), \qquad \int \alpha_i \, di = 1,$$

where $(q^{PF}(\alpha), w^{PF}(\alpha))$ are the supporting prices of the PN problem with weights α.

Coordinate $\xi(x)$. Define the PF aggregate value of resources as

$$Y^{\max}(\alpha) = \int y_i\big(q^{PF}(\alpha), w^{PF}(\alpha); x^{PF}(\alpha)\big)\, di.$$

The productive-efficiency wedge is then

$$\xi(x) = 1 - \frac{\int y_i\big(q^{PF}(\alpha(x)), w^{PF}(\alpha(x)); x\big)\, di}{Y^{\max}(\alpha(x))}. \tag{51}$$

Coordinate $\tau(x)$. The wedges $\tau(x)$ consist of goods wedges $\tau^g(x) := \{\tau^g_{i,c,t},\, \tau^g_{i,l,t}\}_{i,t}$ and insurance wedges $\tau^{ins}(x) := \{\tau^{ins}_{i,c,t}(s^t),\, \tau^{ins}_{i,l,t}(s^t)\}_{i,t,s^t}$, which are constructed from the full-insurance allocation $x_i^{ins} = \{c^{ins}_{i,t}(s^t), l^{ins}_{i,t}(s^t)\}_{t,s^t}$. These are obtained by maximizing $u(x_i^{ins})$ subject to holding fixed the value of consumption expenditures and effective labor earnings every period:

$$\sum_{s^t} q_t^{PF} \Pr(s^t)\, c^{ins}_{i,t}(s^t) = \sum_{s^t} q_t^{PF} \Pr(s^t) c_{i,t}(s^t),$$
$$\sum_{s^t} w_t^{PF} \Pr(s^t) \epsilon_{i,t}(s^t) l^{ins}_{i,t}(s^t) = \sum_{s^t} w_t^{PF} \Pr(s^t) \epsilon_{i,t}(s^t) l_{i,t}(s^t).$$

In our Section 6 settings, the full-insurance allocation is

$$c^{ins}_{i,t}(s^t) = \mathbb{E}_0[c_{i,t}], \qquad l^{ins}_{i,t}(s^t) = \epsilon_{i,t}(s^t)^{1/\gamma}\, \bar{l}^{ins}_{i,t}, \qquad \bar{l}^{ins}_{i,t} \equiv \frac{\mathbb{E}_0[\epsilon_{i,t} l_{i,t}]}{\Xi_{i,t}}.$$

where $\Xi_{i,t} \equiv \mathbb{E}_0\big[\epsilon_{i,t}^{(1+\gamma)/\gamma}\big]$.

The goods wedges as a function of the allocation x are

$$1 + \tau^g_{i,c,t} = \frac{\beta^t}{q_t^{PF}} \left(\frac{\mathbb{E}_0[c_{i,t}]}{c_{i,0}}\right)^{-\sigma}, \tag{52a}$$

$$1 + \tau^g_{i,l,t} = \frac{\beta^t}{w_t^{PF}} \left(\frac{\Psi\left(\bar{l}^{ins}_{i,t}\right)^{\gamma}}{c_{i,0}^{-\sigma}}\right). \tag{52b}$$

Insurance wedges as a function of the allocation x are

$$1 + \tau^{ins}_{i,c,t}(s^t) = \frac{1 + \tau^c_{i,t}(s^t)}{1 + \tau^g_{i,c,t}} = \left(\frac{c_{i,t}(s^t)}{c^{ins}_{i,t}}\right)^{-\sigma}, \tag{53a}$$

$$1 + \tau^{ins}_{i,l,t}(s^t) = \frac{1 + \tau^l_{i,t}(s^t)}{1 + \tau^g_{i,l,t}} = \left(\frac{l_{i,t}(s^t)}{\epsilon_{i,t}(s^t)^{1/\gamma} \bar{l}^{ins}_{i,t}}\right)^{\gamma}. \tag{53b}$$

Algorithm to compute coordinates. We outline a computational algorithm to compute $(\alpha(x), \tau(x), \xi(x))$ from a given allocation x. This algorithm uses the structure of the PN planner's problem and the CRRA/isoelastic preference structure to replace the fixed point over α with a fixed point over the PF paths $(K_t^{PF}, L_t^{PF})_{t\geq 0}$.

1. Guess a path $\{K_t^{PF}, L_t^{PF}\}_{t\geq 0}$ and compute the implied $\{C_t^{PF}\}_{t\geq 0}$ from the PF resource constraint. Compute marginal products $F_{K,t}^{PF}$ and $F_{L,t}^{PF}$, and supporting prices (q_t^{PF}, w_t^{PF}) using the planner's first-order conditions. Normalize $q_0^{PF} = 1$.

2. Given the allocation x, compute $y_i(q^{PF}, w^{PF}; x)$ for each agent using (50). Solve for the indirect-utility derivative $V_{i,y}\big(y_i; q^{PF}, w^{PF}\big)$. In Subsection B.5.1 we show that this can be done efficiently by solving a single nonlinear equation for each agent.

3. Compute the residual function $\text{Res}\big(\{K_t^{PF}, L_t^{PF}\}_{t\geq 0}\big)$ using the following two conditions (for all $t \geq 0$):
$$q_{t+1}^{PF} = \beta\, q_t^{PF} \left(\frac{C_{t+1}^{PF}}{C_t^{PF}}\right)^{-\sigma},$$
$$\int \Xi_{i,t} \left(\frac{F_{L,t}^{PF}}{\Psi}\right)^{1/\gamma} \left(\frac{(V_{i,y})^{1/\sigma}}{\int (V_{j,y})^{1/\sigma}\, dj}\right)^{-\sigma/\gamma} (C_t^{PF})^{-\sigma/\gamma}\, di = L_t^{PF}.$$

4. Adjust $\{K_t^{PF}, L_t^{PF}\}_{t\geq 0}$ using a root-finding algorithm (e.g., Newton–Raphson or Broyden's method) to drive $\text{Res}\big(\{K_t^{PF}, L_t^{PF}\}_{t\geq 0}\big)$ to zero.

5. Once converged, set
$$\alpha_i(x) = \frac{V_{i,y}^{-1}}{\int V_{j,y}^{-1}\, dj}.$$

6. Using the converged PF objects (q^{PF}, w^{PF}) and $\alpha(x)$, compute $\xi(x)$ from (51). Then compute the goods wedges $\tau^g(x)$ from (52) and the insurance wedges $\tau^{ins}(x)$ from (53).

B.3 Mapping coordinates $(\alpha, \tau, \xi) \mapsto$ welfare $\mathcal{W}$

To implement our decomposition, we need to evaluate welfare $\mathcal{W}$ at allocations corresponding to an arbitrary set of coordinates (α, τ, ξ). We detail this mapping below.

Utility decomposition and sufficient statistics. For each (i,t) define full-insurance benchmarks

$$c^{ins}_{i,t} \equiv \mathbb{E}_0[c_{i,t}], \qquad \Xi_{i,t} \equiv \mathbb{E}_0\big[\epsilon_{i,t}^{(1+\gamma)/\gamma}\big], \qquad \bar{l}^{ins}_{i,t} \equiv \frac{\mathbb{E}_0[\epsilon_{i,t} l_{i,t}]}{\Xi_{i,t}}, \qquad l^{ins}_{i,t}(s^t) \equiv \epsilon_{i,t}(s^t)^{1/\gamma}\, \bar{l}^{ins}_{i,t}.$$

Then the two terms that appear in $u(x_i)$ can be expressed as

$$\mathbb{E}_0\left[\frac{c_{i,t}^{1-\sigma}}{1-\sigma}\right] = \frac{(c^{ins}_{i,t})^{1-\sigma}}{1-\sigma}\, u^{ins}_{i,c,t},$$

and

$$\mathbb{E}_0\left[\frac{l_{i,t}^{1+\gamma}}{1+\gamma}\right] = \frac{(\bar{l}^{ins}_{i,t})^{1+\gamma}}{1+\gamma}\, u^{ins}_{i,l,t},$$

where the risk corrections $u^{ins}_{i,c,t}$ and $u^{ins}_{i,l,t}$ only depend on the insurance wedges and are given by a convenient set of sufficient statistics:

$$u^{ins}_{i,c,t} = \mathbb{E}_0\Big[\big(1+\tau^{ins}_{i,c,t}(s^t)\big)^{-(1-\sigma)/\sigma}\Big] = \mathbb{E}_0\left[\left(\frac{c_{i,t}}{c^{ins}_{i,t}}\right)^{1-\sigma}\right], \tag{54a}$$

$$u^{ins}_{i,l,t} = \mathbb{E}_0\Big[\epsilon_{i,t}(s^t)^{(1+\gamma)/\gamma}\,\big(1+\tau^{ins}_{i,l,t}(s^t)\big)^{(1+\gamma)/\gamma}\Big] = \mathbb{E}_0\left[\epsilon_{i,t}^{(1+\gamma)/\gamma}\left(\frac{l_{i,t}}{l^{ins}_{i,t}}\right)^{1+\gamma}\right]. \tag{54b}$$

Given this structure, we do the mapping from coordinates to welfare in two steps. First, we use (α, ξ, τ^g) to recover the expected allocations $\mathbb{E}_0[c_{i,t}]$ and $\mathbb{E}_0[\epsilon_{i,t} l_{i,t}]$. Second, we use the insurance wedges τ^{ins} to compute the risk corrections $u^{ins}_{i,c,t}$ and $u^{ins}_{i,l,t}$ using the sufficient statistics above. Combining these objects yields individual welfare $u(\chi_i(\alpha,\xi,\tau))$ and aggregate welfare $\mathcal{W}(\alpha,\xi,\tau)$.

Algorithm to compute welfare from coordinates.

1. Given α, compute the PF prices (w^{PF}, q^{PF}) and $Y^{\max}$. Use coordinate ξ to obtain aggregate value of resources $Y = (1-\xi)\, Y^{\max}$.

2. Use prices (w^{PF}, q^{PF}) and Y to allocate resource values $\{y_i\}_i$ across households. This can be done by solving the following equations, which are the first-order conditions for problem (6) in the main text:

$$\alpha_i\, V_{i,y}(y_i; q^{PF}, w^{PF}) = \Lambda, \qquad \int y_i\, di = Y.$$

An efficient algorithm to solve these equations is provided in Subsection B.5.2.

3. From the goods wedges τ^g, recover expected allocations as follows. For each i, guess $c_{i,0}$ and compute

$$\mathbb{E}_0[c_{i,t}] = \left(\frac{q_t^{PF}\left(1+\tau_{i,c,t}^g\right)}{\beta^t} \right)^{-1/\sigma} c_{i,0}, \qquad \mathbb{E}_0[\epsilon_{i,t} l_{i,t}] = \Xi_{i,t} \left(\frac{w_t^{PF}\left(1+\tau_{i,l,t}^g\right)}{\beta^t \Psi} \right)^{1/\gamma} c_{i,0}^{-\sigma/\gamma}.$$

Solve for $c_{i,0}$ using the single equation

$$\sum_{t=0}^{\infty} q_t^{PF}\, \mathbb{E}_0[c_{i,t}] + \sum_{t=0}^{\infty} w_t^{PF}\, \mathbb{E}_0[\epsilon_{i,t}(\bar{\ell} - l_{i,t})] = y_i,$$

where $\mathbb{E}_0[\epsilon_{i,t}(\bar{\ell} - l_{i,t})] = \bar{\ell}\, \mathbb{E}_0[\epsilon_{i,t}] - \mathbb{E}_0[\epsilon_{i,t} l_{i,t}]$ and $\mathbb{E}_0[\epsilon_{i,t}]$ is exogenous.

4. Compute the risk corrections using (54a) and (54b). Importantly, the utility corrections from risk depend on the sufficient statistics, which are of much lower dimension than the full history-dependent paths of wedges.

5. Using expected allocations and the risk corrections, compute individual welfare

$$u(\chi_i(\alpha,\xi,\tau)) = \sum_{t=0}^{\infty} \beta^t \frac{(\mathbb{E}_0[c_{i,t}])^{1-\sigma}}{1-\sigma} u_{i,c,t}^{ins} - \Psi \sum_{t=0}^{\infty} \beta^t \frac{(\bar{l}_{i,t}^{ins})^{1+\gamma}}{1+\gamma} u_{i,l,t}^{ins}, \qquad \bar{l}_{i,t}^{ins} = \frac{\mathbb{E}_0[\epsilon_{i,t} l_{i,t}]}{\Xi_{i,t}},$$

and aggregate welfare $\mathcal{W}(\alpha,\xi,\tau)$ using social welfare weights $\bar{\alpha}$ as

$$\mathcal{W}(\alpha,\xi,\tau) = \int \bar{\alpha}_i\, u(\chi_i(\alpha,\xi,\tau))\, di.$$

B.4 Shapley-value welfare decomposition

The Shapley-value decomposition satisfies

$$\Delta \mathcal{W} = \mathsf{R} + \mathsf{E}^{pr} + \mathsf{E}^{g,inter} + \mathsf{E}^{g,intra} + \mathsf{E}^{ins},$$

with each term obtained by applying a two-player Shapley value to an appropriately defined welfare-difference function normalized at the baseline. We detail the construction of each term below.

Two-player Shapley operator. For any function $F(x,y)$ with baseline (x^*,y^*) and counterfactual (x^{**},y^{**}), define

$$\mathrm{Sh}_x[F] \equiv \tfrac{1}{2}\big[F(x^{**},y^*) - F(x^*,y^*)\big] + \tfrac{1}{2}\big[F(x^{**},y^{**}) - F(x^*,y^{**})\big],$$

and define $\mathrm{Sh}_y[F]$ analogously.

Redistribution versus efficiency. Define the total welfare-difference function

$$\Delta W^{\mathrm{tot}}(\alpha,t) \equiv \mathcal{W}(\alpha,t) - \mathcal{W}(\alpha^*,t^*), \qquad t \equiv (\xi,\tau).$$

By construction, $\Delta W^{\mathrm{tot}}(\alpha^*,t^*) = 0$ and $\Delta W^{\mathrm{tot}}(\alpha^{**},t^{**}) = \Delta\mathcal{W}$. Applying the two-player Shapley value yields

$$\Delta\mathcal{W} = \mathsf{R} + \mathsf{E}, \qquad \mathsf{R} \equiv \mathrm{Sh}_\alpha[\Delta W^{\mathrm{tot}}], \quad \mathsf{E} \equiv \mathrm{Sh}_t[\Delta W^{\mathrm{tot}}].$$

The term R is the redistribution component and the term E is the efficiency component.

Productive efficiency versus allocative efficiency. Define the efficiency component function as

$$\Delta W^{\mathrm{eff}}(\xi,\tau) \equiv \tfrac{1}{2}\Big[\mathcal{W}(\alpha^*,\xi,\tau) - \mathcal{W}(\alpha^*,\xi^*,\tau^*)\Big] + \tfrac{1}{2}\Big[\mathcal{W}(\alpha^{**},\xi,\tau) - \mathcal{W}(\alpha^{**},\xi^*,\tau^*)\Big].$$

Then $\Delta W^{\mathrm{eff}}(\xi^*,\tau^*) = 0$ and $\Delta W^{\mathrm{eff}}(\xi^{**},\tau^{**}) = \mathsf{E}$. Applying the two-player Shapley value gives

$$\mathsf{E} = \mathsf{E}^{pr} + \mathsf{E}^{al}, \qquad \mathsf{E}^{pr} \equiv \mathrm{Sh}_\xi[\Delta W^{\mathrm{eff}}], \quad \mathsf{E}^{al} \equiv \mathrm{Sh}_\tau[\Delta W^{\mathrm{eff}}].$$

The term E^{pr} is the productive-efficiency component and the term E^{al} is the allocative-efficiency component.

Allocative efficiency: goods versus insurance wedges. Define the allocative-efficiency component function as

$$\Delta W^{\mathrm{al}}(\tau^g,\tau^{ins}) \equiv \tfrac{1}{2}\Big[\Delta W^{\mathrm{eff}}(\xi^*,\tau^g,\tau^{ins}) - \Delta W^{\mathrm{eff}}(\xi^*,\tau^{G*},\tau^{I*})\Big] + \tfrac{1}{2}\Big[\Delta W^{\mathrm{eff}}(\xi^{**},\tau^g,\tau^{ins}) - \Delta W^{\mathrm{eff}}(\xi^{**},\tau^{G*},\tau^{I*})\Big].$$

Then $\Delta W^{\text{al}}(\tau^{G*}, \tau^{I*}) = 0$ and $\Delta W^{\text{alloc}}(\tau^{G**}, \tau^{I**}) = \mathsf{E}^{alloc}$. Applying the two-player Shapley value yields

$$\mathsf{E}^{al} = \mathsf{E}^{g} + \mathsf{E}^{ins}, \qquad \mathsf{E}^{g} \equiv \text{Sh}_{\tau^{g}}[\Delta W^{\text{alloc}}], \quad \mathsf{E}^{ins} \equiv \text{Sh}_{\tau^{ins}}[\Delta W^{\text{alloc}}].$$

The term E^{g} captures allocative inefficiencies from goods wedges and the term E^{ins} captures allocative inefficiencies from insurance wedges.

Intertemporal versus intratemporal goods wedges. For some of the calculations in Section 6 it is convenient to write the goods wedges $\tau^{g} \equiv (\tau^{g}_{c}, \tau^{g}_{l})$ in terms of wedges on the intertemporal allocation of consumption and the intratemporal allocation of labor. For each agent i and date t, define

$$\tau^{\text{inter}}_{i,t} \equiv \frac{1+\tau^{g}_{i,c,t}}{1+\tau^{g}_{i,c,t+1}} - 1, \qquad\qquad \tau^{\text{intra}}_{i,t} \equiv \frac{1+\tau^{g}_{i,l,t}}{1+\tau^{g}_{i,c,t}} - 1.$$

The wedge $\tau^{\text{inter}}_{i,t}$ captures distortions in the intertemporal Euler equation for consumption, while $\tau^{\text{intra}}_{i,t}$ captures distortions in the intratemporal labor–consumption margin. There is a one to one correspondence between the goods wedges and the intertemporal and intratemporal wedges.

Define the goods-wedge component as

$$\begin{aligned}\Delta W^{g}(\tau^{\text{inter}}, \tau^{\text{intra}}) \equiv \tfrac{1}{2}\Big[\Delta W^{\text{al}}(\tau^{\text{inter}}, \tau^{\text{intra}}, \tau^{I*}) - \Delta W^{\text{al}}(\tau^{\text{inter}}*, \tau^{\text{intra}}*, \tau^{I*})\Big] \\ + \tfrac{1}{2}\Big[\Delta W^{\text{al}}(\tau^{\text{inter}}, \tau^{\text{intra}}, \tau^{I**}) - \Delta W^{\text{al}}(\tau^{\text{inter}}*, \tau^{\text{intra}}*, \tau^{I**})\Big].\end{aligned}$$

Then $\Delta W^{g}(\tau^{\text{inter}}*, \tau^{\text{intra}}*) = 0$ and $\Delta W^{g}(\tau^{\text{inter}}**, \tau^{\text{intra}}**) = \mathsf{E}^{g}$. Applying the two-player Shapley value yields

$$\mathsf{E}^{g} = \mathsf{E}^{g,inter} + \mathsf{E}^{g,intra}, \qquad \mathsf{E}^{g,inter} \equiv \text{Sh}_{\tau^{\text{inter}}}[\Delta W^{g}], \quad \mathsf{E}^{g,intra} \equiv \text{Sh}_{\tau^{\text{intra}}}[\Delta W^{g}].$$

The terms $\mathsf{E}^{g,inter}$ and $\mathsf{E}^{g,intra}$ capture intertemporal and intratemporal distortions arising from goods wedges.

B.5 Technical derivations

B.5.1 Finding $V_{i,y}$

To find $V_{i,y}$, let $\lambda_i^{-\sigma} = V_{i,y}$, where λ_i is proportional to the marginal utility of resources in the indirect utility problem. The optimal consumption and labor choices satisfy

$$\tilde{c}_{i,t}(s^t) = \left(\beta^{-t} q_t^{PF}\right)^{-1/\sigma} \lambda_i,$$
$$\Psi\, \tilde{l}_{i,t}(s^t)^{\gamma} = \beta^{-t} w_t^{PF} \epsilon_{i,t}(s^t)\, \lambda_i^{-\sigma}.$$

The second equation implies

$$\tilde{l}_{i,t}(s^t) = \left(\frac{\beta^{-t} w_t^{PF}}{\Psi}\right)^{1/\gamma} \epsilon_{i,t}(s^t)^{1/\gamma} \lambda_i^{-\sigma/\gamma}.$$

Substitute these into the budget constraint

$$\sum_{t=0}^{\infty} \mathbb{E}_0\Big[q_t^{PF} \tilde{c}_{i,t} + w_t^{PF} \epsilon_{i,t}(\bar{\ell} - \tilde{l}_{i,t})\Big] = y_i.$$

Noting that

$$\mathbb{E}_0[\epsilon_{i,t}\tilde{l}_{i,t}] = \Xi_{i,t} \left(\frac{\beta^{-t} w_t^{PF}}{\Psi}\right)^{1/\gamma} \lambda_i^{-\sigma/\gamma}, \qquad \Xi_{i,t} \equiv \mathbb{E}_0\big[\epsilon_{i,t}^{(1+\gamma)/\gamma}\big],$$

the budget constraint becomes

$$\mathcal{A}^V \lambda_i - \mathcal{B}_i^V \lambda_i^{-\sigma/\gamma} = y_i - \bar{\ell} \sum_{t=0}^{\infty} w_t^{PF} \mathbb{E}_0[\epsilon_{i,t}],$$

where

$$\mathcal{A}^V := \sum_{t=0}^{\infty} q_t^{PF} \left(\beta^{-t} q_t^{PF}\right)^{-1/\sigma}, \qquad \mathcal{B}_i^V := \sum_{t=0}^{\infty} w_t^{PF} \left(\frac{\beta^{-t} w_t^{PF}}{\Psi}\right)^{1/\gamma} \Xi_{i,t}.$$

This is a single equation in λ_i that can be solved numerically, yielding $V_{i,y} = \lambda_i^{-\sigma}$.

B.5.2 Finding y_i

Given the PN weights α, aggregate value of resources Y, and prices q_t^{PF}, w_t^{PF}, we can compute the distribution of resource values $\{y_i\}_i$ as follows. The individual resource values

Figure 10: WELFARE DECOMPOSITION BY INCOME QUINTILE

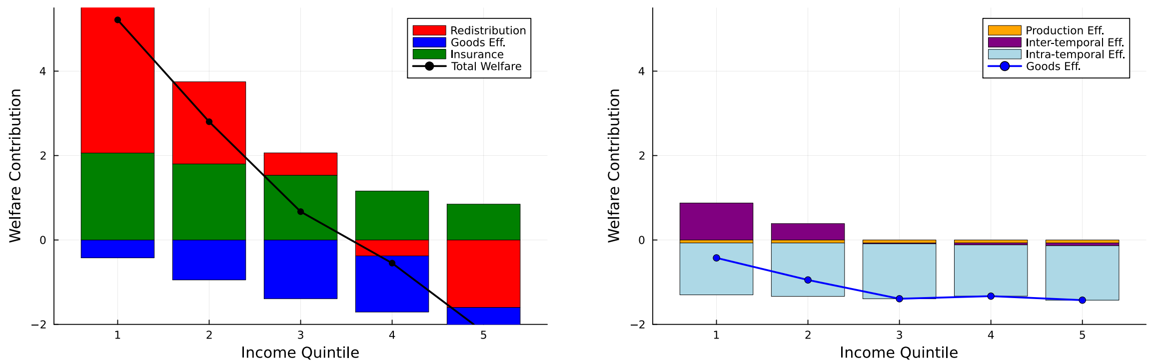

Notes: Welfare contributions by income quintile. The left panel shows total welfare and its decomposition into redistribution, goods-efficiency, and insurance components. The right panel decomposes goods efficiency further into production, intertemporal, and intratemporal components of efficiency by income quintile.

satisfy the constrained planner's optimality conditions

$$\alpha_i V_{i,y}(y_i; q^{PF}, w^{PF}) = \Lambda, \qquad \int y_i \, di = Y,$$

where Λ is the Lagrange multiplier on the aggregate resource constraint. From Subsection B.5.1, we have $V_{i,y} = \lambda_i^{-\sigma}$, so the first condition implies

$$\lambda_i = (\Lambda/\alpha_i)^{-1/\sigma}.$$

The indirect-utility budget constraint gives an expression for y_i as a function of λ_i:

$$y_i(\lambda_i) = \left(\sum_{t=0}^{\infty} q_t^{PF} \left(\beta^{-t} q_t^{PF}\right)^{-1/\sigma}\right) \lambda_i - \left(\sum_{t=0}^{\infty} w_t^{PF} \left(\frac{\beta^{-t} w_t^{PF}}{\Psi}\right)^{1/\gamma} \Xi_{i,t}\right) \lambda_i^{-\sigma/\gamma} + \bar{\ell} \sum_{t=0}^{\infty} w_t^{PF} \mathbb{E}_0[\epsilon_{i,t}].$$

Substituting $\lambda_i(\Lambda)$ into this equation expresses y_i as a function of the single unknown Λ. The aggregate constraint $\int y_i(\Lambda) \, di = Y$ can then be solved for Λ. Once Λ is found, we can compute each λ_i and consequently each y_i.

B.6 Additional distributional analysis

In the main text, we presented the distributional analysis of the welfare decomposition by wealth quintile. Here, we repeat the analysis grouping households by income rather than wealth. Given the high correlation between wealth and income in standard calibrations of this class of models, the resulting distributional patterns are very similar. See Figure 10.

B.7 Comparison of decompositions

Here we will present a comparison of the additive and marginal decompositions in the capitalist worker example of Section 5 and the tax reform example of Section 6. In both cases the marginal and additive decompositions are nearly identical to baseline Shapley-value-based decomposition.

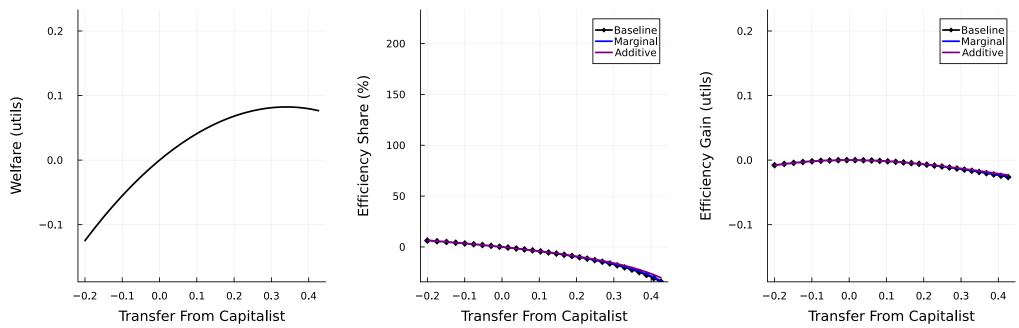

Figure 11: Comparison of Additive and Marginal Decompositions for the Worker Capitalist Example present in Section 5

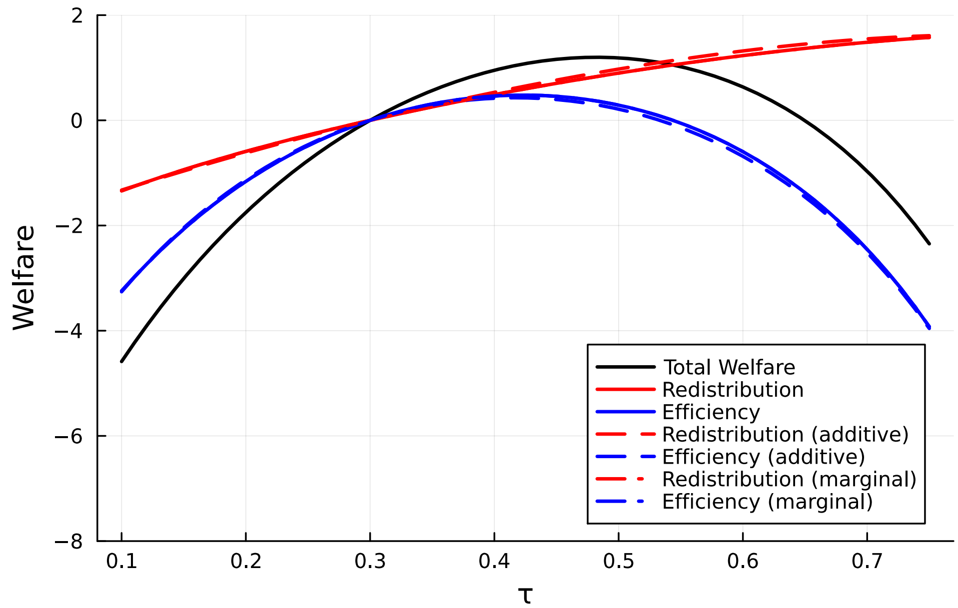

Figure 12: Comparison of Additive and Marginal Decompositions for the Tax Reform Example present in Section 6