Aggregate Efficiency with Heterogeneous Agents

David R. Baqaee
UCLA

Ariel Burstein
UCLA*

December 2025

Abstract

We study aggregate efficiency when households have heterogeneous preferences and outcomes. Our measure extends cost-benefit analysis and aggregate TFP to general equilibrium settings with heterogeneous households and potentially restricted redistribution tools. We ask what fraction of resources can be saved while keeping every agent at least as well off as in their status-quo allocation. We show how to convert this problem into an equivalent utility-maximization problem, enabling the use of tools and results normally applicable only in representative agent settings. We characterize changes in aggregate efficiency in terms of observables, like expenditures and price elasticities, and apply our results to study, among other things, the effects of productivity shocks, the costs of misallocation, and the effects of trade shocks, both with and without costly redistribution.

*We thank Rodrigo Adao, Manuel Amador, Andrew Atkeson, Natalie Bau, Pablo Fajgelbaum, Joao Guerreiro, Valentin Haddad, Keith Head, Oleg Itskhoki, Dirk Krueger, Yuhei Miyauchi, Elisa Rubbo, Aleh Tsyvinski, Gianluca Violante, Mike Waugh, and Pierre-Olivier Weill for their comments. This paper subsumes *Aggregate Welfare and Output with Heterogeneous Agents*, first circulated in 2022. We thank Lautaro Adamovsky for exceptional research assistance. We acknowledge financial support from NSF grant No. 2343691.

1 Introduction

Two central themes in economics are efficiency—how large the economic "pie" is and what determines its size—and equity—how that pie is, or ought to be, shared. The normative question of how resources should be divided among people is controversial and subjective. For this reason, it is useful to study efficiency in isolation from equity. With a single agent, this separation is trivial and automatic because distributional issues never arise. With heterogeneous consumers, however — especially when they differ in tastes or face different prices — the two questions are harder to separate.

In academic and theoretically-oriented settings, outcomes across heterogeneous agents are often summarized using a social welfare function. This approach, pioneered by Bergson (1938) and Samuelson (1947), mixes efficiency and equity considerations. Social welfare functions automatically take a stance not just on the size of the economic pie, but also on its optimal division. This means that social welfare functions depend on subjective Pareto-weights and are not invariant to monotone transformations of utility functions. This makes them unattractive for use in practical policy design, especially if heterogeneity in tastes is important (i.e. in cases where we cannot use "the same utility function" to represent every agent's preferences).

In more policy-oriented circles, changes in aggregate efficiency are typically measured by willingness-to-pay and cost-benefit analysis: "can the winners compensate the losers and still be better off?" This may be done explicitly, by estimating and summing compensating variations, or implicitly, by using market-based quantity indices like real GDP (which, under appropriate assumptions, are equivalent to summing compensating variations). These metrics measure the change in aggregate efficiency by the amount of money (at constant prices) left-over after winners compensate losers. Such measures are popular because they are objective, taking the view that a dollar is a dollar, regardless of who earns it, and because they are easy to communicate (Harberger, 1971).

However, it is well-known that standard cost-benefit analysis has theoretical problems. For one, the transfers implied by summing up compensating variations may be infeasible. This is either because there is no mechanism by which winners can compensate the losers (e.g. Antràs et al., 2012 and Schulz et al., 2023) or because, even if transfers can be made, relative prices may change in response to the transfers making compensations impossible (e.g. Boadway, 1974). Second, and at a more abstract level, in some settings, for example when decentralized markets are incomplete or absent, the notion of a compensating variation, or even prices themselves, may be undefined.

In this paper, we extend the cost-benefit approach to deal with these shortcomings,

building on ideas from Allais (1979), Debreu (1951), and Luenberger (1996). Our measure generalizes the popular Lucas (1987) consumption-equivalent beyond the one-agent setting, without taking a stance on how resources should be divided across people. To measure changes in aggregate efficiency we ask: *by how much can the consumption-possibility set shrink while keeping every agent at least indifferent to their status-quo allocation?* We call the resulting contraction factor the *aggregate consumption-equivalent* variation. It measures the resources left over after everyone has been compensated; efficiency rises whenever the same welfare as in the status-quo can be achieved with less resources. We do not refer to it as a measure of aggregate welfare to distinguish it from a social welfare function that embeds normative judgements about interpersonal utility comparisons.

Our measure answers a counterfactual question in terms of observables. Unlike social welfare functions, it does not introduce free parameters such as Pareto weights and is invariant to monotone transformations of utility (i.e. it depends only on ordinal properties of preference relations). Moreover, the measure does not take a normative stance on the optimal distribution of resources across individuals beyond compensating every agent relative to the status-quo. For example, if aggregate efficiency increases and there are extra resources left over after everyone has been compensated, our measure takes no stance on who should get those resources.

The numerical value of our efficiency measure is interpretable and, in the special case of a single agent, collapses to the familiar consumption-equivalent measure of Lucas. Because its definition neither presupposes the existence of markets nor depends on prices, it avoids the limitations and paradoxes of traditional cost-benefit analysis. The measure accommodates a wide range of decentralization mechanisms — such as competitive markets, search-and-matching, bargaining, and imperfect competition — as well as constraints on redistribution — including limited taxation or costly transfers. Moreover, our measure remains tractable across a broad class of models.

We characterize our measure of aggregate efficiency in terms of observables and generalize well-known representative-agent results to settings with heterogeneous-agents. This paper has two stand-alone companions —Baqaee and Burstein (2025a) and Baqaee and Burstein (2025b) — where we apply our framework to answer two different questions: the cost of misallocation due to financial market incompleteness (in both closed and open-economies) and changes in aggregate efficiency in random utility models with discrete choice. In both applications, household heterogeneity is a central feature of the problem.

The structure of the paper is as follows. In Section 2, we define our measure in abstract terms and present a key result: Theorem 1 converts the problem of calculating aggregate

consumption-equivalents into an equivalent fictional utility-maximization problem for an agent with homothetic preferences. This forms the basis for all other results in the paper because it allows us to port tools used to study the welfare of representative-agents with homothetic preferences, like Hulten (1978), Harberger (1964), Arkolakis et al. (2012), Petrin and Levinsohn (2012), and Baqaee and Farhi (2019b, 2020) to economies with heterogeneous agents.

The abstract environment in Section 2 does not impose much structure on how consumption possibility sets come about (for example, general equilibrium is a special case). In Section 3, we specialize the environment to a general equilibrium setting with potentially distorting wedges and lump-sum transfers. In this case, our measure can also be thought of as the maximum reduction in total factor productivity, scaling all factor endowments by a constant, that allows agents to be at least indifferent to the status-quo. Hence, our measure is a generalization of total factor productivity to economies with heterogeneous agents.

We also define some popular alternative measures of efficiency from the literature to facilitate comparison with our approach: real output (using a Divisia or chain-weighted index), Kaldor-Hicks efficiency (which compares total income to the sum of compensating incomes), and the welfare of a positive representative agent (if such an agent exists). We establish an important equivalence result: if all households have identical homothetic preferences and face the same relative prices, our measure of aggregate efficiency (with lump-sum transfers) coincides with those alternative measures. Outside of these common but restrictive assumptions, however, the measures generally differ.

A key result in Section 3 is that our aggregate efficiency measure can be calculated using a *compensated* equilibrium with an as-if representative agent. The compensated equilibrium is the general equilibrium of an economy with the same technologies and distortions as the real economy, but populated by a fictional representative agent with homothetic preferences. We show that the utility of this fictional representative agent measures the change in aggregate efficiency. This result makes it straightforward to use tools and methods from representative-agent economies to analyze aggregate efficiency with heterogeneous agents.

In Section 4 we restrict attention to perfectly competitive economies where both the first and second welfare theorem hold. We show that in such settings and to a first-order approximation, Hulten (1978) applies to our measure of aggregate efficiency unaltered. That is, to a first-order, our measure of changes in aggregate efficiency coincide with total factor productivity as measured by the Solow (1957) residual, and to Kaldor-Hicks efficiency. We then derive a nonlinear version of Hulten (1978) that applies to aggregate

consumption-equivalents, extending the nonlinear characterizations in Baqaee and Farhi (2019b) to economies with heterogeneous agents. We show that changes in aggregate efficiency depend only on expenditure shares and price elasticities. We also generalize the sufficient-statistics of Arkolakis et al. (2012), developed for single-agent economies, to quantify the gains from trade in economies with heterogeneous agents.

In Section 5 we consider distorted economies, and derive versions of Hsieh and Klenow (2009), Petrin and Levinsohn (2012), Harberger (1954, 1964), and Baqaee and Farhi (2020) that apply to economies with heterogeneous agents. We show that, once prices are not equal to marginal costs, then real output and Kaldor-Hicks efficiency are not reliable measures of aggregate efficiency even to a first-order. In particular, they can rise or fall in response to pure redistributions even as we move from one Pareto-efficient allocation to another, if prices are not equal to marginal costs. We also derive a version of the famous Harberger triangles formula that can be used to quantify misallocation with heterogeneous agents, and show that there is a sense in which misallocation losses in the heterogeneous agent model are lower than in representative agent models. This is because our Harberger triangles formula discards dispersion in wedges due to difference in average wedges paid by each household. Differences in the average wedge by household are equivalent to lump-sum transfers and do not imply (Pareto) inefficiency.

In Section 6 we consider economies with costly redistribution (i.e. without lump-sum transfers). We discuss how Theorem 1 can be used to analyze changes in aggregate efficiency when redistributive instruments are limited (e.g. revenues used for redistribution are raised via distortionary taxes). We show that, starting in perfect competition, the change in aggregate efficiency due to a change in primitives is, to a first-order, the same as Hulten (1978). To a second-order, the change in efficiency is equal to what would have happened with lump-sum transfers (characterized in Sections 4 and 5) minus the additional Harberger triangles caused by inefficient redistribution (which are zero if lump-sum transfers are available). We end the section with a quantitative example: how the rise of China affected the United States, as measured by the aggregate consumption equivalent. We show that the answer depends on the ease with which workers can move across sectors, and the range of redistributive tools available. Whereas the aggregate consumption equivalent change is positive when workers can move across sectors or if lump-sum transfers are available, it is negative if workers are restricted to working within narrow industries and redistribution is impossible or difficult.

Related literature. Our approach to measuring aggregate efficiency is related to willingness-to-pay based measures, which have a very long history dating all the way back to at least

Dupuit (1844). For example, the compensating variation, and sum of compensating variations, in Hicks (1939) and Kaldor (1939), are special cases. Furthermore, the notion of social surplus in Allais (1979), the coefficient of resource utilization in Debreu (1951, 1954), the measure of efficiency in Farrell (1957), and the benefit function in Luenberger (1996) are all related to our measure. Our contribution relative to these works is to provide a characterization without assuming either Pareto efficiency or markets, explicitly allowing for limited or costly redistribution, and applying our measure to modern models.

Our paper is also related to cost-benefit analysis, typically performed by using the sum of compensating variations, as in Harberger (1971), and related ideas like the marginal value of public funds (Hendren and Sprung-Keyser, 2020). The idea behind these measures is to ask: "after the winners compensate the losers using lump-sum transfers, is there still money left on the table?" Our measure of efficiency coincides with these measures when both welfare theorems hold and the consumption-possibility set is linear. However, outside of these cases, the two measures are different. First, if prices are not equal to marginal costs, then pure transfers can raise or lower real GDP and the sum of compensating variations, even when allocations are Pareto-efficient. Second, if the consumption-possibility set is nonlinear, then as shown by Boadway (1974), a pure transfer between agents can cause the sum of compensating variations to exceed zero. Our measure, which can be defined even when prices do not exist, does not have this property. Third, unlike the sum of compensating variations, our measure does not presuppose that lump-sum transfers are feasible. In this sense, our approach has similarities to Schulz et al. (2023), who generalize the sum of compensating variations to allow for limited redistribution.[1]

Our paper complements and differs from Schulz et al. (2023) in many ways, the most important being a difference in focus. They consider economies with a single consumption good, focusing their attention on a mechanism design problem where lump-sum taxes are unavailable because of asymmetric information. Although our formalism and definitions can be applied to such economies, we do not focus on these issues. Instead, we focus on allowing for multiple goods and heterogeneity in preferences and relative prices faced by consumers. This means that even with perfect information and lump-sum transfers, there are interesting questions about how to aggregate across consumers that consume and value different goods.[2]

[1] In response to a shock, they consider a tax reform that makes households indifferent to the status-quo and then measure the monetary value of aggregate welfare gains or losses by the fiscal surplus from this reform.

[2] A related approach, Auerbach and Kotlikoff (1987), quantifies aggregate efficiency in overlapping-generations economies by using lump-sum transfers to keep generations before a specified date at their

As mentioned above, a different approach to aggregation is to use a social welfare function to evaluate outcomes. A prominent example is the behind the veil-of-ignorance measure of Harsanyi (1955). Social welfare functions are by far the most common approach in the modern literature to aggregating across heterogeneous agents.[3] Our paper, which instead looks for and quantifies the potential for Pareto improvements (i.e. compensating everyone and looking to see if resources are left over), studies an alternative question that the one analyzed by this methodology.

Following in the social-welfare-function tradition, a recent set of papers, including Bhandari et al. (2021), Dávila and Schaab (2022, 2023), and Donald et al. (2023) provide approximate decompositions of changes in social welfare functions. Our goal in this paper is different: we do not provide decompositions of social welfare functions, but instead, define and characterize aggregate efficiency directly as an answer to a counterfactual question. The decompositions in the papers mentioned above contain components the authors refer to as capturing efficiency. However, since our objective is different, our notion of efficiency is also generically different to the efficiency components in these papers (even to a first-order in economies with distortions). Defining efficiency directly, instead of as part of an approximate decomposition, is useful because it means that we can also study large changes.[4]

In terms of the tools and methods, our paper is closely related to the literature that studies the macroeconomic consequences of microeconomic productivity changes and wedges. For productivity changes, this includes Gabaix (2011), Acemoglu et al. (2012), Baqaee and Farhi (2019b) and others. For wedges, this includes Harberger (1954), and more recently, Restuccia and Rogerson (2008), Hsieh and Klenow (2009), Bigio and La'O

status-quo utility level and increase the utility of all cohorts after that date by a common amount. Our measure is different because it keeps every agent indifferent to the status-quo, allows for limited redistribution, and measures efficiency in terms of resource-savings instead. It would be interesting to apply our measure to an overlapping-generations economy.

[3]There is a branch of the literature that assumes observed allocations can be rationalized by maximizing some social welfare function within some parametric class, estimates this function, and uses it to conduct policy analysis (see Heathcote and Tsujiyama, 2021 and the references therein). This is equivalent to assuming there exists a normative representative agent: a hypothetical single decision-maker whose utility function is maximized by observed allocations (Chapter 4 Mas-Colell et al., 1995). Our approach is different since we do not need to assume the existence of either a positive nor normative representative agent. Furthermore, even if a normative representative agent exists, there is nothing to say that its preferences should be privileged over any other social welfare function (see Example 7 below).

[4]Whereas infinitesimal changes in our measure of efficiency can be integrated to study large changes, integrals of components in a decomposition of social welfare are path-dependent. To see this point, suppose we approximately decompose changes in some function $y = f(x_1, x_2)$ into $dy \approx (\partial f/\partial x_1)dx_1 + (\partial f/\partial x_2)dx_2$. Then we can write non-infinitesimal changes as $\Delta y = \int(\partial f/\partial x_1)\Delta x_1 + \int(\partial f/\partial x_2)\Delta x_2$ but, unless $f(x_1, x_2)$ is linear in x_1 and x_2, the size of each component of this nonlinear decomposition depends on the arbitrary path of integration.

(2016), Liu (2017), Baqaee and Farhi (2020), among others. We relax the assumption typically maintained in both of these literatures that households have common preferences and face common prices.[5]

Finally, because we use the gains from trade as one of our examples, our paper is also related to the gains from trade with heterogeneous agents. Much of the work on international trade with heterogeneous agents focuses on the distributional effects of trade. Some examples of papers that also calculate aggregate welfare are Antras et al. (2017) and Galle et al. (2023) (using an Atkinson (1970)-style social welfare function with inequality aversion), Kim and Vogel (2020) (using the sum of compensating variations), and Rodríguez-Clare et al. (2022) (using a population-weighted average of welfare gains across regions), all of which differ from our measure of aggregate efficiency for reasons already discussed.

2 Abstract Definition and Characterization

Consider an economy populated by agents indexed by $h \in \{1, ..., H\}$. Agent h has ordinal preferences $\succeq_h$ over commodity vectors $\boldsymbol{c}_h \in \mathbb{R}^N$, where N is the number of goods. Assume preferences are represented by utility functions $u_h(\boldsymbol{c}_h)$.[6] A *consumption allocation* is a matrix $\boldsymbol{c} \in \mathbb{R}^{H \times N}$ whose hth row, denoted by $\boldsymbol{c}_h$, equals the consumption vector of agent h.

Fix some consumption allocation, denoted by $\boldsymbol{c}^0$, as the *status-quo*. We think of $\boldsymbol{c}^0$ as the data generated by some initial equilibrium outcome. Define $\mathcal{C} \subset \mathbb{R}^{H \times N}$ to be some set of feasible consumption allocations, determined by deeper primitives like technologies, policies, and so on. We think of $\mathcal{C}$ as the counterfactual we are interested in studying.

Definition 1 (Aggregate Consumption Equivalent Variation)**.** The *Aggregate Consumption-Equivalent Variation* of the set $\mathcal{C}$ relative to the status-quo $\boldsymbol{c}^0$ is the maximum proportional contraction of $\mathcal{C}$ such that every agent can be kept at least indifferent to the status-quo allocation. Formally,

$$A \equiv \max \left\{ Z \in \mathbb{R} : \text{there is } \boldsymbol{c} \in Z^{-1}\mathcal{C} \text{ and } \boldsymbol{c}_h \succeq_h \boldsymbol{c}_h^0 \text{ for every } h \right\}. \tag{1}$$

If $A > 1$, then the consumption possibility set $\mathcal{C}$ contains a strict Pareto-improvement

[5]One exception is Bornstein and Peter (2024), who study misallocation with differences in tastes and markups across households. In their setting, symmetry and the law-of-large numbers implies that every households' problem is identical despite the fact that households have different preferences.

[6]That is, for each household $u_h(\boldsymbol{c}_h) \geq u_h(\boldsymbol{c}_h')$ if, and only if, $\boldsymbol{c}_h \succeq_h \boldsymbol{c}_h'$. We assume that preferences are continuous and locally nonsatiated.

relative to the status-quo, and if $A < 1$, then at every point in $\mathcal{C}$, at least one agent is worse off than in the status-quo. The cardinal value of A is interpretable. For concreteness, say, $A = 1.01$, then this means that it is possible to make everyone at least as well off as in the status-quo and discard 1% of every good (or more precisely, $1 - 1/A$ percent). Agents may not be consuming the same bundle as in the status-quo after they are compensated — we only require that they be indifferent to the status-quo. If there is a single household and $\mathcal{C}$ is a consumption allocation, then A is the same as the consumption-equivalent variation of Lucas (1987). As in both Lucas (1987) and Debreu (1951), the definition of A treats all commodities symmetrically by shifting the consumption possibility set, $\mathcal{C}$, proportionately in every dimension. Figure 1 illustrates this for a simple economy with two households, indexed by h and h', and two consumption goods, one consumed only by h and the other only by h'.[7]

The measure in Definition 1 has some desirable properties: (1) it answers a counterfactual question about observable phenomena with interpretable units. (How much of every good is left-over after everyone is compensated?) The answer to this question is invariant to monotone transformations of utility functions, and only relies on ordinal properties of preference relations. (2) Our measure does not take a stance on how social surplus or losses should be divided among agents. That is, while we can assign a numerical efficiency value to every feasible set of consumption allocations, $\mathcal{C}$, we do not attempt to pick a specific allocation among the possibilities as being socially "optimal." (3) Our definition places no restrictions on the set of redistributive tools that are available or the mechanism by which allocations are decentralized (e.g. spot markets, search, bargaining, etc.). The instruments of redistribution are implicitly embedded into the shape of the feasible set $\mathcal{C}$.[8]

[7]Note that the choice of how to scale the set $\mathcal{C}$ is present even with a single agent. For example, do we shrink or expand every element in a consumption vector to reach indifference, as in the consumption-equivalent variation of Lucas (1987), or only some components. Similarly, do we scale the whole budget set proportionally, as in the compensating variation of Hicks (1939), or do we change individual prices of certain commodities to reach indifference. In this sense, this choice is not special to having heterogeneous agents or not. One reason to focus on a radial contraction is because, for the class of general equilibrium models studied in Section 3, expanding and contracting the feasible consumption set proportionally is isomorphic to scaling productivity of all factors (including quasi-fixed factors capturing decreasing returns to scale). This is an intuitive way to measure efficiency — efficiency is measured by the reduction in total factor producticity necessary to reach indifference to status-quo. The reason our measure of aggregate efficiency differs from standard approaches like using social welfare functions or chained real consumption in multi-agent settings is not due to the fact that we scale the consumption possibility set proportionally. For example, our measure differs even if there is only a single consumption good (so there is no choice of how to expand or shrink the consumption set because there is only one consumption good). In some settings, it may be interesting to expand or shrink $\mathcal{C}$ in a particular direction, rather than uniformly. This requires generalizing our definition to allow for non-radial expansions, as in the transferable surplus notion in Allais (1979). Proposition 11 in the appendix shows that the aggregate consumption variation based on uniform and non-uniform expansions are always on the same side of one.

[8]For example, for convex competitive economies, if lump-sum transfers are available, then the second

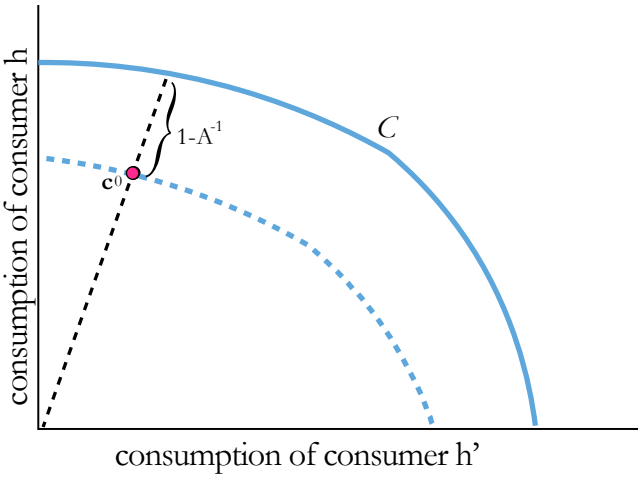


Figure 1: Aggregate efficiency is measured by the maximal radial contraction of the feasible set necessary to achieve indifference.

(4) This abstract definition is applicable to a wide variety of models and circumstances, allowing for, among other things, non-homothetic preferences, discrete choice, risk, and dynamics.[9] (5) As shown by Lemma 1, it coincides with traditional measures like real output, Kaldor-Hicks efficiency, and Lucas' consumption-equivalent variation in cases where those measures are well-behaved.

To characterize A, we prove a useful theorem, which we repeatedly use in the rest of the paper. This theorem proves that calculating A is equivalent to solving a utility-maximization problem for some fictitious agent. To state this result, we first define *individual consumption-equivalent* transformations of individual preferences.[10]

Definition 2. Let $u_h(\boldsymbol{c}_h)$ denote a utility representation for agent h. The *individual consumption-equivalent* function $\tilde{u}_h(\boldsymbol{c}_h)$ is implicitly defined by

$$u_h(\frac{\boldsymbol{c}_h}{\tilde{u}_h}) = u_h(\boldsymbol{c}_h^0).$$

The individual consumption-equivalent function, $\tilde{u}_h$, is homogenous of degree one in consumption by construction. If the preference relation $\succeq_h$ is homothetic, then $\tilde{u}$ is a cardinalization of $\succeq_h$ — in this case, $\tilde{u}_h$ ranks consumption bundles in the same order as

welfare theorem implies that $\mathcal{C}$ is the set of all Pareto-efficient allocations. However, the definition can also be applied to economies where such transfers are not available.

[9]Our measure of aggregate efficiency can be used to order feasible sets under a given status-quo. The ordering of two feasible sets may flip for different status-quos, similar to Scitovszky (1941). For a fixed status-quo, our measure of aggregate efficiency gives a unique ordering of feasible sets. For example, if $\mathcal{C}' \subseteq \mathcal{C}$, then $A(\boldsymbol{c}^0, \mathcal{C}') \leq A(\boldsymbol{c}^0, \mathcal{C})$.

[10]The individual consumption-equivalent function is also called the *distance* function in the duality literature on optimization (see, for example, Cornes, 1992).

$\succeq_h$. By construction, $\tilde{u}_h$ is homogenous of degree one and normalized to equal to 1 at $\boldsymbol{c}_h^0$. The magnitude of $\tilde{u}_h(\boldsymbol{c}_h)$ is interpretable — it measures the amount the consumption bundle $\boldsymbol{c}_h$ has to be scaled to make the household exactly indifferent to the status-quo. In this sense, $\tilde{u}_h(\boldsymbol{c}_h)$ is the household's consumption-equivalent variation relative to the status-quo.

Example 1 (Single good). Suppose there is a single consumption good, so $u_h(c_h)$ is some increasing function. In this case,

$$\tilde{u}_h(\boldsymbol{c}_h) = \frac{c_h}{c_h^0},$$

regardless of the functional form of u_h. See Appendix D for a more involved example using non-homothetic CES preferences.

If $\succeq_h$ is non-homothetic, then $\tilde{u}_h$ is *not* a cardinalization of $\succeq_h$ (i.e. $\tilde{u}_h$ does not rank consumption allocations according to $\succeq_h$). Figure 2 graphically depicts indifference curves of $\tilde{u}_h$ — they are radial expansions of the status-quo indifference curve defined by $u_h(\boldsymbol{c}_h) = u_h(\boldsymbol{c}_h^0)$. When $\succeq_h$ is homothetic, all indifference curves are radial expansions, so that the ranking produced by $\tilde{u}_h$ coincides with the one produced by $\succeq_h$.

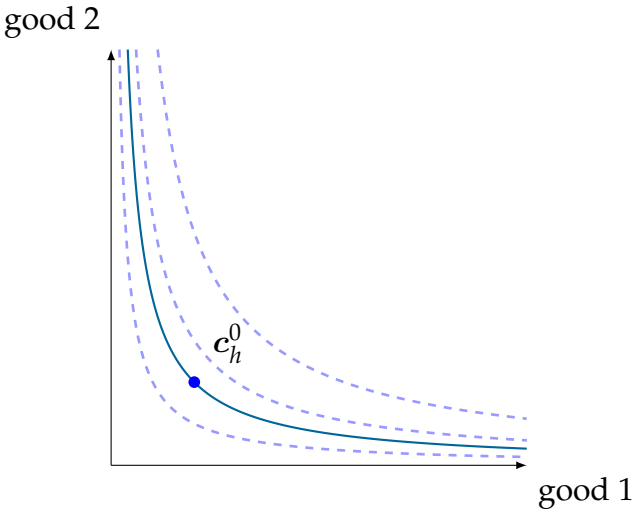


Figure 2: The solid blue line is the indifference curve $u_h(\boldsymbol{c}_h) = u_h(\boldsymbol{c}_h^0)$ and the dashed lines are the indifference curves of $\tilde{u}_h$.

Define a fictitious compensated agent as follows.

Definition 3. The *compensated agent* is an agent whose preferences are represented by

$$U(\boldsymbol{c}) = \min_h \{\tilde{u}_h(\boldsymbol{c}_h)\},$$

where $\tilde{u}_h$ are individual consumption-equivalent functions.

The utility function $U(\boldsymbol{c})$ is homogeneous of degree one by construction. We call this agent compensated because, as we show in Appendix A, the budget shares of this agent are the average of the compensated budget shares of all the agents weighted by each agent's compensating income.

Note that, the function $U(\boldsymbol{c})$ is not a social welfare function. There are two reasons for this. First, U does not depend on households' true utility functions, instead it depends on the individual consumption-equivalent functions. Second, U depends on the minimum growth in individual consumption-equivalents relative to status-quo, rather than the level of utility (i.e. $U(\boldsymbol{c})$ is not a Rawlsian social welfare function).

We can now state the main result of this section.

Theorem 1 (Aggregate Efficiency by Utility Maximization)**.** *The aggregate consumption equivalent is equal to the value of $\mathcal{C}$ to the compensated agent:*

$$A = \max_{\boldsymbol{c} \in \mathcal{C}} U(\boldsymbol{c}). \tag{2}$$

Figure 3 graphically illustrates the content of Theorem 1. Rather than proportionally shifting $\mathcal{C}$ to reach the indifference point, Theorem 1 states that we can instead maximize $U(\boldsymbol{c})$ — by shifting out the indifference curves of the compensated agent — until we reach the boundary of $\mathcal{C}$. The utility of this fictional agent is numerically identical to the maximal reduction in $\mathcal{C}$ needed to reach indifference.

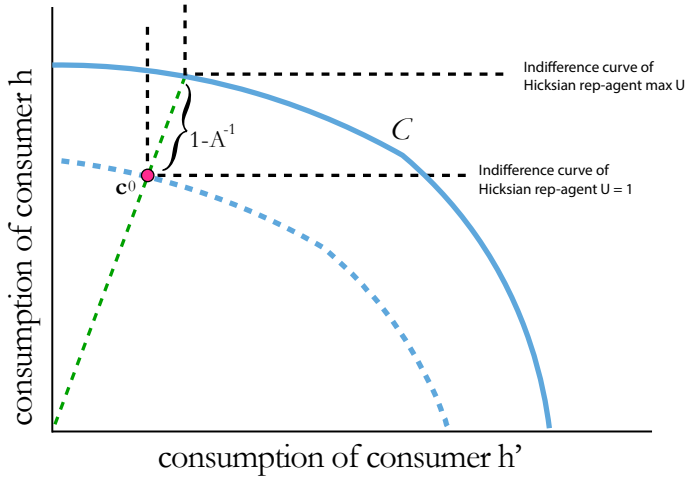


Figure 3: The increase in the utility of the compensated agent also measures the amount by which $\mathcal{C}$ needs to shrink to ensure indifference.

Theorem 1 is crucial because it converts the problem of calculating aggregate efficiency

in (1) into an equivalent utility-maximization problem. Since utility-maximization problems are common in economics, this means that Theorem 1 allows us to easily convert results about representative agent problems into ones about aggregate efficiency with heterogeneous agents.

If the maximizer in (2), $\boldsymbol{c}^*$, is such that $\tilde{u}_h(\boldsymbol{c}_h^*) = \tilde{u}_{h'}(\boldsymbol{c}_{h'}^*)$ for every h and h', then we say that *all households can be exactly compensated*. This is because there exists a consumption allocation, $\boldsymbol{c}^*/A$ that makes every agent *exactly* indifferent to the status-quo and $(1 - 1/A)\boldsymbol{c}^*$ of each consumption good is left-over.

Theorem 1 guarantees that maximizing $U(\boldsymbol{c})$ yields the same number as the maximization problem that defines A. However, unlike a social welfare maximization problem, there is no sense in which the allocation $\boldsymbol{c}^*$ that maximizes $U(\boldsymbol{c})$ is the optimal allocation because primitive problem defining A is stated in terms of shrinking a possibility frontier, not choosing an allocation inside it. In other words, if $A > 1$, then there is a range of allocations that can engineer a Pareto improvement and we do not take a stance on which one society should choose.

The rest of the paper uses Theorem 1 in different contexts to characterize changes in aggregate efficiency in terms of observables. Although Theorem 1 applies very generally, we focus on cases where the consumption possibility set is determined by general equilibrium (possibly with distortions and limited transfers).

3 Decentralized Equilibrium with Lump-Sum Transfers

In this section we set up a general-equilibrium framework that admits distortive wedges. For simplicity of exposition, throughout the rest of the body of the paper, we assume that while preferences can be heterogeneous, they are homothetic. In Appendix E, we show how all results can be extended to cover non-homothetic preferences as well. We begin by assuming that the consumption possibility set is the set of equilibrium consumption allocations achievable by lump-sum transfers, and consider the case where transfers are restricted in Section 6.

We show, via Theorem 1, how to compute the aggregate consumption-equivalent variation in these settings. For comparison, we also define three alternative metrics that are popular in the literature: (1) chain-weighted real output, (2) the sum of compensating variations (Kaldor-Hicks/Cost-Benefit), and (3) the welfare of a positive representative agent. We establish the restrictive conditions under which all four measures coincide. The remainder of the paper then examines the more general cases in which these conditions fail, and where our measure no longer aligns with the conventional ones.

3.1 Environment and Equilibrium

Each household maximizes utility subject to a budget constraint

$$\max\ u_h(\boldsymbol{c}_h) \text{ such that } \sum_i p_i c_{hi} \leq \sum_f \omega_{hf} w_f L_f + T_h, \tag{3}$$

where the left-hand side is total expenditures and the right-hand side is total income. As in Arrow-Debreu, commodities could be indexed by time and state of nature. On the left-hand side, p_i is the price of i and c_{hi} is the quantity of good i purchased by household h. On the right-hand side, households derive income from factors and lump-sum transfers. Households h owns a share ω_{hf} of factor f, where w_f is the wage and L_f is the total quantity of factor f. Lump-sum transfers are T_h.

Producer i chooses its inputs to minimizes costs

$$\min \sum_j p_j y_{ij} + \sum_f w_f l_{if}, \tag{4}$$

subject to production technology

$$y_i = z_i F_i \left(\{y_{ij}\}, \{l_{if}\} \right), \tag{5}$$

where y_i is the quantity of output, F_i is a constant-returns production function, y_{ij} are intermediate inputs used by i produced by j, and l_{if} are primary factors used by i. The assumption that F_i has constant-returns is without loss of generality, since we can capture decreasing returns using producer-specific factors. The parameter z_i is a Hicks neutral productivity shifter. The price of i is equal to an exogenous markup or tax, $\mu_i > 0$, times i's marginal cost of production

$$p_i = \mu_i mc_i. \tag{6}$$

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

The resource constraint for goods and factors is

$$\sum_j y_{ji} + \sum_h c_{hi} \leq y_i, \qquad \text{and} \qquad \sum_i l_{if} \leq z_f L_f, \tag{7}$$

where z_f, when f indexes a factor, controls the endowment of efficiency units of factor f. Finally, net transfers across households are equal to the revenues generated by the

wedges:

$$\sum_h T_h = \sum_i p_i y_i \left(1 - \frac{1}{\mu_i}\right). \tag{8}$$

Buyer-seller-specific productivity and wedges. Although we assume that z_i is Hicks neutral and wedges are on gross output only, both of these assumptions are made without loss of generality. This is because we can recreate buyer-seller productivity changes and wedges by relabeling. Specifically, we can treat firm or household i's purchases of an input from j as a distinct good (made linearly using j's output). A productivity shock or a wedge on this good is then isomorphic to a buyer-seller specific productivity shock or wedge. We make the assumption that z_i is Hicks neutral and assume all wedges take the form of taxes on gross output to simplify the notation.

We now define a general equilibrium with wedges.[11]

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

3.2 Different Aggregate Measures in General Equilibrium

We specialize the definition of aggregate-consumption equivalents to this environment. We also define some other popular measures of aggregate economic activity used in the literature. To do so, we index all exogenous parameters (productivities, wedges, transfers) by a scalar t and let $t = 0$ denote the status-quo allocation. For any equilibrium price or quantity X, we write $X(t)$ to denote its dependence on the exogenous parameters.[12]

We begin by defining our measure of efficiency.

Aggregate Consumption-Equivalent Variation. To define $A(t)$, we need to specify the status-quo and the consumption possibility set. The status-quo $\boldsymbol{c}^0$ is the equilibrium allocation at $t = 0$, which we take to be the data generating process for consumption. The

[11] This notion of general equilibrium is the same one used by Baqaee and Farhi (2020), extended to allow for multiple households.

[12] In the case of multiple equilibria, we assume there is an equilibrium selection mechanism. The nature of this equilibrium selection mechanism is not relevant for $A(t)$, because $A(t)$ is unique given t and the status-quo.

consumption possibility set $\mathcal{C}(t)$ is the set of consumption allocations that can be supported as part of an equilibrium given technologies $\boldsymbol{z}(t)$, wedges $\boldsymbol{\mu}(t)$, and lump-sum transfers:

$$\mathcal{C}(t) = \left\{ \boldsymbol{c} \in \mathbb{R}^{H \times N} : \boldsymbol{p}, \boldsymbol{w}, \boldsymbol{y}, \boldsymbol{c}, \boldsymbol{T}, \boldsymbol{x}, \boldsymbol{l} \text{ satisfying (3), (4), (5), (6), (7), (8)} \right\}.$$

The set $\mathcal{C}(t)$ depends on t because technologies $\boldsymbol{z}(t)$ and wedges $\boldsymbol{\mu}(t)$ are indexed by t. The aggregate consumption-equivalent is the maximum uniform contraction of $\mathcal{C}(t)$ that allows every agent to be at least indifferent to the status-quo. Given the additional structure provided by general equilibrium, we can also think of $A(t)$ as the maximum contraction of total factor productivity, multiplying all z_f for factors by $1/A(t)$, such that every agent can be kept at least indifferent to the status-quo. This means $A(t)$ converts the shock at t into an equivalent change in total factor productivity that brings everyone back to the status-quo.

If the economy at t is efficient, $\boldsymbol{\mu}(t) = \mathbf{1}$, then our results go through, and the value of $A(t)$ is the same, if instead $\mathcal{C}(t)$ is the set of all technologically feasible consumption allocations. This is a consequence of the second welfare theorem (i.e. every Pareto-efficient allocation can be reached via lump-sum transfers). However, if the economy is inefficient, then swapping $\mathcal{C}(t)$ for the technologically feasible set of allocations changes the value of $A(t)$ — in particular, using the former measures the distance to the set of distorted equilibrium allocations, whereas the latter measures the distance to the Pareto-efficient frontier.

For comparison, we define some other common measures of aggregate efficiency.

Real Output. In national income accounting, real output is measured using approximations to the Divisia (1925) index. Aggregate real output, defined using the Divisia index, is

$$\log Y(t) = \int_0^t \sum_i \frac{p_i(s)c_i(s)}{\sum_{i'} p_{i'}(s)c_{i'}(s)} \frac{d \log c_i(s)}{ds} ds,$$

where $c_i(s) = \sum_{h \in H} c_{hi}(s)$ denotes aggregate consumption of good i at $s \in [0, t]$.

Kaldor-Hicks/Cost-Benefit Efficiency. Another popular aggregate measure is the Kaldor-Hicks efficiency measure. This measure compares the sum of compensating incomes to aggregate income at t. If the sum of compensating incomes is less than aggregate income, then the winners can hypothetically compensate the losers and there can still be money left-over. The amount of money left over is a measure of the increase in efficiency. This method is the foundation of most of cost-benefit style analyses in applied welfare

economics and policy evaluation in public finance and industrial organization.

Let $e_h(\boldsymbol{p}, u_h)$ be an expenditure function representing preferences $\succeq_h$. The Kaldor-Hicks measure of efficiency at t is

$$A^{KH}(t) = \frac{\sum_h e_h(\boldsymbol{p}(t), u_h(t))}{\sum_h e_h(\boldsymbol{p}(t), u_h(0))}. \tag{9}$$

Note that, by construction, Kaldor-Hicks efficiency at the status-quo is equal to one: $A(0) = 1$.

Consumption-equivalent of Representative Agent. Another well-known aggregate measure, when a representative agent exists, is the consumption-equivalent variation used by Lucas (1987). A *representative agent* is a hypothetical single consumer such that the demand of the representative agent for each good, given prices and total income, coincides with equilibrium quantity of that good, given the same prices and aggregate income.[13]

If a representative agent exists, define the consumption-equivalent for the representative agent, $A^{RA}(t)$, to be

$$u^{RA}\left(\boldsymbol{c}^{RA}(t)/A^{RA}(t)\right) = u^{RA}\left(\boldsymbol{c}^{RA}(0)\right),$$

where u^{RA} is the utility function of the representative agent. In words, $A^{RA}(t)$ is the amount by which the aggregate consumption bundle in t must be contracted to make the positive representative agent exactly indifferent to the status-quo. As with all the other measures, $A^{RA}(0) = 1$ by construction.

3.3 Characterizing $A(t)$ Using Theorem 1

We characterize aggregate efficiency via Theorem 1. To do so, we define a compensated equilibrium, which is a useful fictional construct for proving results and constructing sufficient statistics formulas.[14]

Definition 5 (Compensated Equilibrium)**.** A *compensated equilibrium with an as-if representative agent* is the general equilibrium of an economy with the same technologies, resource constraints, and wedges as the original economy but where there is a representative agent

[13]For a formal definition, see Appendix B.

[14]This notion of the compensated equilibrium has many antecedents in prior work, for example it nests the concept in Jones (2002) and Johansson et al. (2022), the Hicksian equilibrium in Baqaee and Burstein (2023), the synthetic equilibrium in Debreu (1951), and is closely related to the adjusted price function in Luenberger (1996). A major difference relative to these notions is that our compensated equilibrium need not be efficient.

with preferences as in Definition 3. For any equilibrium variable $X(t)$, denote the same variable in the compensated equilibrium by $X^{\text{comp}}(t)$.

It is important to note that the compensated equilibrium with an as-if representative agent is not of direct interest, but is instead a useful stepping-stone to calculating changes in aggregate efficiency. We refer to this as the compensated equilibrium for short.

The following result, which is a consequence of Theorem 1, shows that aggregate efficiency is as simple as calculating welfare in an economy with a representative agent.

Theorem 2 (Aggregate Efficiency Using Compensated Equilibrium)**.** *The aggregate consumption equivalent variation can be calculated using the compensated equilibrium:*[15]

$$A(t) = U(\boldsymbol{c}^{comp}(t)) = Y^{comp}(t) = A^{KH,comp}(t) = A^{RA,comp}(t),$$

if either households can be exactly compensated or there are no wedges $\boldsymbol{\mu} = \mathbf{1}$*.*

Theorem 2 follows from the fact that the planner's transfers in (2) coincide with the expenditure choices of a compensated representative agent (who takes prices as given) whenever either of the following sufficient conditions holds: (i) households are exactly compensated, or (ii) the equilibrium is efficient $\boldsymbol{\mu} = \mathbf{1}$.[16] Moreover, $\boldsymbol{c}^{\text{comp}}(t)/A(t)$ is the allocation that compensates agents and leaves $1 - 1/A(t)$ of every consumption good left-over. If neither condition holds, then Equation (2) still applies, but the solution may not be representable as a compensated equilibrium with an "as-if" representative agent. We maintain the assumptions required for Theorem 2 throughout, except in Section 6, where we abandon lump-sum transfers and work directly with Theorem 1.

In words, Theorem 2 implies that aggregate efficiency, $A(t)$, can be computed by solving for changes in the utility of the compensated representative agent in a Walrasian equilibrium where the representative agent takes prices as given. The first equality, $A(t) = U(\boldsymbol{c}^{\text{comp}}(t))$, is the key; subsequent identities follow mechanically from the fact that the compensated economy has a single representative agent. The importance of Theorem 2 lies in the fact that it allows every tool and result used to calculate welfare in homothetic representative agent economies to be used to calculate aggregate efficiency with heterogeneous preferences. (We show in Appendix E that the result also extends when households' preferences are non-homothetic).

[15] If there are multiple compensated equilibria, pick the one with the highest $U(\boldsymbol{c}^{\text{comp}})$.

[16] Either condition is sufficient, and both may hold simultaneously. When $\boldsymbol{\mu} = 1$, the result follows from the welfare theorems. If $\boldsymbol{\mu} \neq 1$ but households are exactly compensated, the prices associated with the planner's solution in (2) still induce the representative agent to choose the same allocation.

Theorem 2 is expressed in terms of endogenous variables in the compensated equilibrium. Solving that equilibrium is essentially the same as solving a representative–agent model, which is well-understood problem. For this reason, we present the full characterization of variables in the compensated equilibrium in Appendix Section F.[17] However, we do note the following useful fact about the compensated equilibrium in this section.

Lemma 1 (Compensated equilibrium at the status-quo)**.** *At the status-quo $t = 0$, prices and quantities in the compensated equilibrium coincide with those in the decentralized equilibrium.*

This lemma, which guarantees that compensated and decentralized variables coincide at the status-quo, is important for calibrating the model when solving for the compensated equilibrium — expenditure shares at the compensated equilibrium must coincide with observations in the absence of any shock (i.e. at the status-quo).

Before deploying Theorem 2 to construct heterogeneous-agent generalizations of well-known results, we first point out an important, but highly restrictive, special case, where our measure of aggregate efficiency coincides with the other popular alternatives.

3.4 A Miraculous Consensus

To set aside household heterogeneity, a standard benchmark is to assume every household has identical homothetic preferences and faces the same relative prices (i.e. there are no household-specific wedges). Under this condition we obtain the following.

Proposition 1 (Miraculous Consensus)**.** *If households have identical homothetic preferences, and face the same relative prices, then a positive representative agent exists and*

$$A(t) = Y(t) = A^{KH}(t) = A^{RA}(t).$$

The first equality can break down if lump-sum transfers are not available.

In words, the change in aggregate efficiency, measured by the consumption-equivalent variation, matches the change in chain-weighted index of real output, Kaldor-Hicks (cost-benefit) efficiency, and the consumption-equivalent of the positive representative agent all in the decentralized equilibrium. That is, under these assumptions, one can compute $A(t)$ without relying on the compensated equilibrium. Each of the underlying assumptions is individually essential: relaxing any one breaks the equivalence. The remainder of the paper explores those more general settings and illustrates, through examples, why

[17] In Appendix A we also provide the expenditure function of the compensated agent in the compensated equilibrium, since the expenditure function is a useful way to solve general equilibrium models.

our efficiency measure avoids the paradoxes of the alternatives when consensus fails. We consider non-identical preferences in Section 4, allow for different households to pay different relative prices for the same goods in Section 5, and consider limits to redistribution in Section 6.[18]

Proposition 1 is a consequence of Theorem 2. The reason is that, under the stated assumptions, the price and quantity of each good in the compensated equilibrium coincides with those in the decentralized equilibrium. This implies that $A(t) = Y^{\text{comp}}(t) = Y(t)$, since real output in the compensated and decentralized equilibrium only depend on the prices and quantities of goods. The remaining two equalities are then standard.[19]

4 Competitive Economies with Lump-Sum Transfers

In this section, we characterize how aggregate efficiency responds to changes in productivity when both welfare theorems hold. This means that, for the remainder of this section, we assume that all wedges, $\boldsymbol{\mu}(t)$, are all equal to one (first welfare theorem holds) and lump-sum transfers are available (second welfare theorem holds). We consider distorted economies in Section 5 and economies without lump-sum transfers in Section 6.

We begin this section by providing some general comparative static results. We then apply these results to some analytical examples to build intuition.

4.1 Comparative Statics for Changes in Technologies

Denote the *Domar* weight of each producer or factor i by

$$\lambda_i(t) = \frac{p_i(t)y_i(t)}{\sum_{i'} p_i(t)c_i(t)}\mathbf{1}\{i \text{ is a producer}\} + \frac{w_i(t)z_i(t)}{\sum_{i'} p_i(t)c_i(t)}\mathbf{1}\{i \text{ is a factor}\}.$$

[18] In this paper, we are focused on household heterogeneity, but even with a single household, the miraculous consensus breaks down when preferences are non-homothetic. This point is discussed in detail by Baqaee and Burstein (2023). Intuitively, when preferences are non-homothetic, even for a single agent, scaling the production possibility set ($A(t)$), the budget constraint ($A^{KH}(t)$), and the equilibrium consumption allocation ($A^{RA}(t)$) do not coincide with one another since, as we shrink resources, the household would want to change the bundle of goods they consume. See also Appendix E.

[19] Under the stated assumptions, there is a positive representative agent with homothetic preferences, and it follows from Shephard's lemma that $Y(t) = A^{RA}(t)$ (see, e.g., Baqaee and Burstein, 2023). The final equality in Proposition 1 follows from the fact that the indirect utility function of each agent, $v(\boldsymbol{p}, I_h)$ can be written as $I_h/P(\boldsymbol{p})$, where $P(\boldsymbol{p})$ is an ideal price index. It then follows that, $v^{RA}(\boldsymbol{p}, \sum_h I_h)$, can be written as $(\sum_h I_h)/P^{RA}(\boldsymbol{p})$. Hence, $A^{RA}(t) = v^{RA}(\boldsymbol{p}(t), \sum_h I_h(t))/v^{RA}(\boldsymbol{p}(0), \sum_h I_h(0)) = \sum_h I_h(t)/(\sum_h I_h(0)) \times P(\boldsymbol{p}(0))/P(\boldsymbol{p}(t)) = A^{KH}(t)$.

This the sales of i divided by total final expenditures. Recall that for factor i, the quantity of the factor is z_i. The following is a well-known result characterizing changes in real output in the decentralized equilibrium.

Proposition 2 (Hulten's Theorem)**.** *The change in chain-weighted real output is*

$$\log Y(t) = \int_0^t \sum_i \lambda_i(s) \frac{d \log z_i}{ds} ds. \tag{10}$$

The most well-known consequence of this result, given by differentiating with respect to t, is that $d \log Y / dt = \sum_i \lambda_i d \log z_i / dt$. This formula, which generalizes Solow (1957), shows that the elasticity of real output to the productivity of producer i or the quantity of factor i is just the Domar weight of i. Using Theorem 2, we can easily state a version of Hulten's theorem that applies to our measure of aggregate efficiency instead.

Proposition 3 (Compensated Hulten's Theorem)**.** *The change in aggregate efficiency at t is*

$$\log A(t) = \int_0^t \sum_i \lambda_i^{comp}(s) \frac{d \log z_i}{ds} ds. \tag{11}$$

In Appendix F, we characterize $\lambda_i^{comp}(s)$ explicitly as a function of the productivity changes $\Delta \log z$, elasticities of substitution, and expenditure shares.

Differentiating (11) with respect to t and evaluating at $t = 0$ shows that, to a first-order approximation, the change in aggregate efficiency, $\Delta \log A$, coincides with the change in real output in the competitive equilibrium $\Delta \log Y$.

Corollary 1 (First Order Changes in Aggregate Efficiency)**.** *To a first-order approximation, the change in aggregate efficiency is*

$$\Delta \log A \approx \sum_i \lambda_i^{comp}(0) \Delta \log z_i = \sum_i \lambda_i(0) \Delta \log z_i \approx \Delta \log Y \approx \Delta \log A^{KH}.$$

The second equality follows from Lemma 1, which states that prices and quantities in the compensated equilibrium $t = 0$ are equal to those in decentralized economy. In other words, the first-order version of Hulten's theorem applies unaltered to aggregate efficiency. The final equality, which is standard, shows that real output is also equal to Kaldor-Hicks efficiency to a first-order approximation.

Baqaee and Farhi (2019b) show that, to a second-order approximation, changes in real

output are given by

$$\Delta \log Y \approx \sum_i \lambda_i \Delta \log z_i + \frac{1}{2} \sum_i \Delta\lambda_i \Delta \log z_i.$$

Differentiating (11) twice with respect to t and evaluating at $t = 0$, gives the following extension of Baqaee and Farhi (2019b) to multi-agent settings.

Proposition 4 (Second Order Changes in Aggregate Efficiency)**.** *To a second-order approximation, the change in aggregate efficiency due to changes in primitives is*

$$\Delta \log A \approx \sum_i \lambda_i \Delta \log z_i + \frac{1}{2} \sum_i \Delta\lambda_i^{comp} \Delta \log z_i,$$

where λ_i and $\Delta\lambda_i^{comp}$ are evaluated at status-quo. In Appendix F, we write $\Delta\lambda_i^{comp}$ explicitly as a function of the productivity changes $\Delta \log z$, microeconomic elasticities of substitution, and expenditure shares in the status-quo.

Proposition 4 shows that, if the economy is efficient, then discrepancies between aggregate efficiency $\Delta \log A$ and real output $\Delta \log Y$ start at the second-order, since, generically $\Delta\lambda \neq \Delta\lambda^{\text{comp}}$ (explicit formulas are in Appendix F).[20] The gap between $\Delta\lambda_i$ and $\Delta\lambda_i^{\text{comp}}$ arises because the compensated agent's demand responds differently to shocks than aggregate demand in the decentralized economy. The compensated agent's demand is constructed to preserve indifference with status-quo for all households, whereas the decentralized aggregate demand comes from utility maximization by households whose incomes evolve according to equilibrium changes in relative factor prices.

4.2 Some Paradoxes of Real Output and Kaldor-Hicks

Having characterized $A(t)$ in economies where the first and second welfare theorem hold, we now contrast $A(t)$ with real output, $Y(t)$, and Kaldor-Hicks efficiency $A^{KH}(t)$.

Real Output. It is well-known that Divisia-based indices, like real output, suffer from paradoxes unless households have identical and homothetic preferences Hulten (1973). Specifically, if households do not have identical and homothetic preferences, then generically, the value of real output $Y(t)$ can be any positive number, regardless of the technol-

[20] To derive an expression for $\Delta\lambda^{\text{comp}}$ in terms of microeconomic primitives, we use the fact that $\Delta\lambda^{\text{comp}}$ is the change in Domar weights in a special case of the environment considered by Baqaee and Farhi (2019b) where the consumption growth of each agent is treated as-if it is a final good, and there is a Leontief final demand aggregator over final goods.

ogy parameters $\boldsymbol{z}(t)$ depending on the path of integration.[21] Hence, although $A(t)$ and $Y(t)$ coincide up to a first-order approximation at $t = 0$, they are not the same nonlinearly.

Kaldor-Hicks. We now turn our attention to Kaldor-Hicks efficiency. Of course, if lump-sum transfers are not available, then $A(t)$ and $A^{KH}(t)$ are different, and we explore the case with limited redistribution in Section 6. However, even when lump-sum transfers are available, our measure need not coincide with $A^{KH}(t)$. The following proposition illustrates this fact.

Proposition 5 (Paradox for Kaldor-Hicks Efficiency)**.** *For any change in technologies (movements of the Pareto efficient frontier), the change in aggregate efficiency, measured using aggregate consumption-equivalent variation, is weakly less than Kaldor-Hicks efficiency:*

$$\Delta \log A \leq \Delta \log A^{KH}.$$

For pure redistributions (movements along the Pareto efficient frontier), the change in aggregate efficiency, measured using aggregate consumption-equivalent variation, is zero, whereas the change in Kaldor-Hicks efficiency can be positive:

$$\Delta \log A = 0 \leq \Delta \log A^{KH}.$$

These inequalities are strict outside of knife-edge cases. The final inequality is a restatement of the so-called Boadway (1974) paradox. It states that $A^{KH}(t)$ assigns strictly positive value to pure redistributions when relative prices respond to transfers.[22]

Figure 4 illustrates the Boadway paradox using a two-good, two-consumer economy. Intuitively, redistributions lower the relative price of those goods that are more valued by the losers. Hence, in the new equilibrium, it is relatively cheap to compensate these households. Of course, such compensations are, in practice, infeasible because if they were to occur, then relative prices would rise for those households that need compensation.[23] Graphically, A^{KH} evaluates efficiency by shifting the line tangent to the frontier at $\boldsymbol{c}_1$ (the dashed line) until it reaches the status-quo $\boldsymbol{c}_0$. The increase in A^{KH} depends on how much the tangent line must be shifted inwards to reach the status-quo. In contrast,

[21] The technical reason is that $\log Y(t)$ is a line integral, and unless preferences are identical and homothetic, the vector field defined by Domar weights is not conservative, making $\log Y(t)$ dependent on the path of integration. See Baqaee and Burstein (2023) for related discussions.

[22] See also Blackorby and Donaldson (1990) for a related critique of the sum of compensating variations as a measure of efficiency.

[23] See Jones (2002) for a detailed discussion.

A uses the frontier itself. Since both $\boldsymbol{c}_0$ and $\boldsymbol{c}_1$ are on the same frontier, $\Delta \log A = 0$ in this case.

In fact, it is possible to construct examples where the production possibility set of the economy shrinks, $\mathcal{C}(t) \subset \mathcal{C}(0)$, so that $\Delta \log A < 0$, but $\Delta \log A^{KH} > 0$. This happens if the decentralized equilibrium associated with $\mathcal{C}(t)$ has very different relative prices to $\mathcal{C}(0)$. In this case, it is relatively cheap for the winners to compensate the losers under the new prevailing prices.

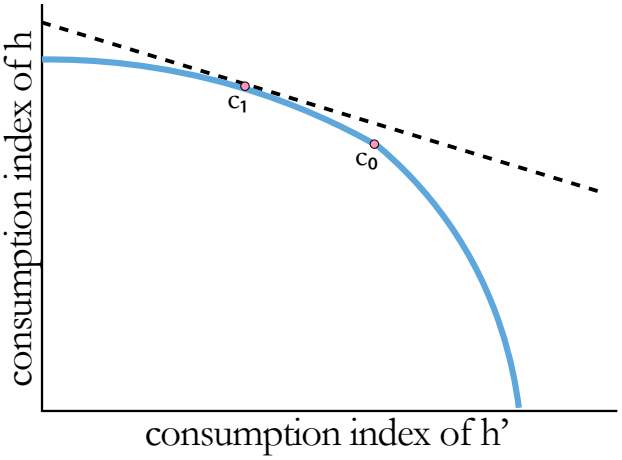


Figure 4: The sum of compensating variations is less than aggregate income since $\boldsymbol{c}^0$ is below the dashed straight line.

There are some special cases where $A^{KH}(t)$ and $A(t)$ coincide in perfectly competitive economies. First, $A(t)$ and $A^{KH}(t)$ coincide to a first-order approximation at $t = 0$ (assuming lump-sum transfers are available). Second, $A^{KH}(t)$ and $A(t)$ coincide nonlinearly if relative prices do not depend on final demand. This happens when the economy's production possibility frontier is a hyperplane, so the marginal rate of transformation, and hence relative prices, are determined by technology only.[24] Under these conditions, the change in aggregate efficiency, $\Delta \log A$, coincides with the change in Kaldor-Hicks efficiency as defined in Equation (9). Given our assumption that production functions are all constant-returns to scale, the production possibility frontier becomes a hyperplane whenever there is only one primary factor of production.

Of course, there is another important reason (besides endogeneity of prices) why our measure of efficiency can differ from the Kaldor-Hicks measure. The Kaldor-Hicks measure, by summing up compensating variations, implicitly assumes that lump-sum trans-

[24] See Proposition 10 in Appendix C for a formal statement.

fers are available, so that winners can costlessly compensate the losers (assuming relative prices are constant). By contrast, our definition of aggregate efficiency naturally extends to allow for limited redistribution, as we discuss further in Section 6.

4.3 Analytical Examples

To build some intuition, we work through some analytical examples for how $\Delta \log A$ responds to changes in technologies. Appendix D provides more detailed derivations.

Example 2 (Regional Productivity Shocks). Consider households in different regions, indexed by h, with preferences over tradeable goods and locally produced nontradeable services:

$$u_h(\boldsymbol{c}_h) = c_{hg}^{\alpha} c_{hs}^{1-\alpha}, \qquad \sum_h c_{hg} = z_g, \qquad c_{hs} = z_{hs}.$$

The first equation shows that utility in each region depends on goods and services with the expenditure share on goods equal to α. The second equation is the resource constraint for goods, which clear at the aggregate level, since goods are traded. The third equation is the resource constraint for services, which clear region-by-region, since services are not traded. The parameters z_g and z_{hs} control the endowments of goods and services.

To calculate the status-quo (as well real output), we must make assumptions on factor ownership. Suppose that households in region h own the local endowment of services and own a share χ_h of the aggregate endowment of the traded good. This implies that, in equilibrium, χ_h is the expenditures of each household as a share of total consumption expenditures. The Domar weight on goods is $\lambda_g = \sum_h \chi_h \alpha = \alpha$ and the Domar weight on services in region h is $\lambda_{hs} = \chi_h(1-\alpha)$.

When productivities change, according to Proposition 2, the change in real output is

$$\Delta \log Y = \alpha \Delta \log z_g + (1-\alpha)\mathbb{E}_\chi \left[\Delta \log z_s\right],$$

where $\mathbb{E}_\chi \left[\Delta \log z_s\right]$ is the average productivity shock to services weighted by the vector $\boldsymbol{\chi}$. This expression is exact because in the decentralized equilibrium Domar weights do not change ($\Delta\lambda = 0$). Furthermore, since the Domar weights are constant in the competitive equilibrium, there is a positive representative agent with Cobb-Douglas preferences over goods and services in all regions:

$$u^{RA}(\boldsymbol{c}) = c_g^{\alpha} \prod_h c_{hs}^{\chi_h(1-\alpha)}.$$

Since the positive representative agent has homothetic preferences, we also have that $\Delta \log Y = \Delta \log A^{RA}$.

According to Proposition 4, the change in aggregate efficiency, to a second-order, is

$$\Delta \log A \approx \alpha \Delta \log z_g + (1-\alpha)\mathbb{E}_\chi \left[\Delta \log z_s\right] - \frac{1}{2}\frac{(1-\alpha)^2}{\alpha} Var_\chi \left[\Delta \log z_s\right] \leq \Delta \log Y = \Delta \log A^{RA}.$$

The miraculous consensus of Proposition 1 fails because the agents do not have the same preferences. As predicted by Corollary 1, $\Delta \log Y$ and $\Delta \log A$ do coincide to a first-order, since the discrepancy scales in the square of $\Delta \log z_s$. The second-order approximation shows that $\Delta \log A$ is a concave envelope of $\Delta \log Y$ around the status-quo $\Delta \log \boldsymbol{z} = 0$ — amplifying negative shocks and mitigating positive shocks to services relative to real output. Intuitively, negative shocks to services are "more costly" since the losses cannot be shared across regions, whereas the positive representative agent is "willing" to substitute between regions with a unit elasticity.

Our next example uses Theorem 1 to apply a version of the popular Arkolakis et al. (2012) (ACR) formula to economies with heterogeneous preferences.

Example 3 (Gains from Trade with Heterogeneous Preferences). Consider a country with different consumers, h, that value domestic and foreign goods differently:

$$u_h(\boldsymbol{c}_h) = \left[(\alpha_h)^{\frac{1}{\theta_h}} c_{hd}^{\frac{\theta_h - 1}{\theta_h}} + (1-\alpha_h)^{\frac{1}{\theta_h}} c_{hf}^{\frac{\theta_h - 1}{\theta_h}}\right]^{\frac{\theta_h}{\theta_h - 1}}.$$

The parameter α_h controls home bias for household h, $\theta_h > 1$ is the compensated Armington elasticity. The domestic good is produced linearly from a labor endowment and trade is balanced. We consider the gains from trade relative to autarky by raising iceberg trade costs to infinity.

The country trades with the rest of the world in the status-quo. We measure the efficiency loss from autarky via the increase in the autarky consumption possibility set needed to keep every consumer indifferent to the status-quo. With one domestic good, this is simply the increase in the aggregate quantity of the domestic good.

Let s_{hd} denote household h's budget share on the domestic good in the status-quo. Replicating the argument from ACR, for the compensated representative agent in the compensated equilibrium, losses from autarky are

$$\Delta \log A = -\log \mathbb{E}_\chi \left[(s_{hd})^{\frac{1}{1-\theta_h}}\right] \leq 0. \tag{12}$$

Note that $(s_{hd})^{\frac{1}{1-\theta_h}}$ is the ACR formula for the gains from trade for a single agent. The equation above shows that aggregate efficiency losses are the average of these individual losses weighted by expenditures in the status-quo (denoted by χ).[25, 26]

To get more intuition, consider a second-order approximation of $\Delta \log A$ around autarky:[27]

$$\Delta \log A \approx \mathbb{E}_\chi \left[\frac{\log s_{hd}}{\theta_h - 1} \right] - \frac{1}{2} Var_\chi \left[\frac{\log s_{hd}}{\theta_h - 1} \right]. \tag{13}$$

The first term is just an "average" version of the ACR formula in logs — the log ACR formula is applied household-by-household and then averaged using households' share of aggregate income χ_h. The second term is the Jensen's inequality term, and it lowers aggregate efficiency if there is any heterogeneity in households' exposure to traded goods either due to variance in expenditure shares, s_{hd}, or trade elasticities, θ_h. In this sense, heterogeneity raises the costs of autarky since some households are more exposed to trade than average.

5 Distorted Economies with Lump-Sum Transfers

We now relax the assumption in the previous section that there are no wedges, and we characterize how aggregate efficiency responds to changes in productivities and wedges. One important comparative static we focus on is the efficiency losses from misallocation — the increase in aggregate efficiency caused by the elimination of all wedges.

5.1 Comparative Statics for Changes in Technologies and Wedges

Theorem 2 means that we can convert results about real output into results about aggregate consumption-equivalent variation by applying them to variables in the compensated

[25] Compare (12) to the representative-agent ACR formula. Suppose that agents have common homothetic CES preferences with Armington elasticity θ. Then the losses from autarky are $\Delta \log A^{RA} = \log \left(\mathbb{E}_\chi \left[s_{hd} \right]^{\frac{1}{1-\theta}} \right)$. If the Armington elasticities are the same, $\theta_h = \theta$, then the losses are larger with heterogeneous agents due to Jensen's inequality. For intuition, consider the case where some household h consumes no home goods, i.e., $s_{hd} = 0$ for some h. In this case, $\Delta \log A = -\infty < \Delta \log A^{RA}$ because it is impossible to compensate h in autarky.

[26] Interestingly, the non-homotheticity plays no role in the sense that the formula is the same even if agents have non-homothetic CES utility functions (see Appendix D for a worked-out derivation). For example, if there is a single agent with non-homothetic CES preferences, the ACR formula holds without change as long as we use the compensated trade elasticity. If preferences are non-homothetic, then there is a distinction between the compensated and uncompensated trade elasticities. If we have estimates of the latter, one must use Slutsky's equation to first convert them into the compensated elasticities (see, e.g. Auer et al., 2024).

[27] This is an approximation in $\frac{\log s_{hd}}{\theta_h - 1}$ around $s_{hd} = 1$. To derive this, we follow the strategy in Theorem 3 of Baqaee and Farhi (2019a) who consider the gains from trade with a homothetic representative agent.

equilibrium. For example, consider the following generalization of Petrin and Levinsohn (2012).

Proposition 6 (Changes in Aggregate Efficiency with Wedges)**.** *In response to changes in wedges and productivities, the change in aggregate efficiency is*

$$\Delta \log A = \int_0^t \sum_i \lambda_i^{comp}(s) \left[\left(1 - \frac{1}{\mu_i(s)} \right) \frac{d \log y_i^{comp}}{ds} + \frac{1}{\mu_i(s)} \frac{d \log z_i}{ds} \right] ds.$$

In Appendix F, we characterize $\lambda_i^{comp}(s)$ and $d \log y_i^{comp} / ds$ explicitly as a function of the productivity changes $\Delta \log z$, elasticities of substitution, and expenditure shares using the results in Baqaee and Farhi (2020).

The compensated Domar weights, λ^{comp}, and quantities, $d \log y^{\text{comp}}$, can be computed using standard methods for inefficient economies with homothetic representative agents (see Appendix F).

We can contrast Proposition 6 with Harberger's social welfare formula. In his classic paper, Harberger (1971) argued that the welfare effect of a policy that changes wedges can be computed as

$$\Delta \log Y(t) = \int_0^t \sum_i \left[p_i(s) - mc_i(s) \right] \frac{dy_i}{ds} ds \; = \; \int_0^t \sum_i \lambda_i(s) \left(1 - \frac{1}{\mu_i(s)} \right) \frac{d \log y_i}{ds} ds, \quad (14)$$

where the equality uses the fact that final expenditure is the numeraire ($\sum_i p_i(s) c_i(s) = 1$ for every s). In words, he argued that whenever a good's marginal benefit, $p_i(s)$, exceeds its marginal cost, $mc_i(s)$, then expanding its quantity (holding others fixed) raises aggregate output.

We see from these equations that if $\Delta \log y_i \neq \Delta \log y_i^{\text{comp}}$, then real output and aggregate efficiency differ to a first-order in inefficient economies. (Note that, if agents have the same preferences and face the same relative prices, then $\Delta \log y_i = \Delta \log y_i^{\text{comp}}$ and real output coincides with A consistent with Lemma 1).

The following example illustrates the first-order difference between aggregate efficiency, as measured by $A(t)$, and real output/Kaldor-Hicks in a simple example. We show that if prices are not equal to marginal costs, then real output and Kaldor-Hicks efficiency can rise, to a first-order, in response to pure redistribution, even when the status-quo allocation is Pareto efficient.

Example 4 (First-Order Aggregate Efficiency vs. Real Output/Kaldor-Hicks). Consider an economy with two agents, 1 and 2, each buying a single consumption good, c_1 and c_2,

produced linearly from labor and sold at markup μ_1 and μ_2. A lump-sum transfer from 2 to 1 raises the spending on good 1 and lowers the spending on good 2. Holding constant technologies, aggregate efficiency, measured using $A(t)$, is unchanged. This is because wedges in this economy do not move the allocation off the Pareto-frontier — instead, they redistribute resources from agent 1 to agent 2 given their incomes.

Using Equation (14), it is easy to show that the effect on real output (and also Kaldor-Hicks efficiency) from a pure transfer is, to a first-order,

$$\Delta \log Y \approx \Delta \log A^{KH} \approx \bar{\mu} \left[\frac{\mu_1 - \mu_2}{\mu_1 \mu_2} \right] \Delta \lambda_1, \tag{15}$$

where $\bar{\mu}$ is the harmonic sales-weighted average of markups and $\Delta\lambda_1$ is the change in the share of income (including transfers) earned by household 1. Hence, a pure redistribution from agent 2 to agent 1 raises real output and Kaldor-Hicks efficiency if the markup paid by agent 1 is higher than the one paid by agent 2. However, there is no way to compensate agent 1 using lump-sum transfers without going back to the status-quo equilibrium (which means there is nothing left-over for agent 2). Since the status-quo allocation of this economy is Pareto-efficient, any gains to agent 1 come at the cost of agent 2. Accordingly, and in contrast to real output and Kaldor-Hicks, $\Delta \log A$ in this example is equal to zero.

We can push this example even farther: suppose again that the transfer occurs at the same time as a decline in the productivity of labor by $\Delta \log z < 0$. In this case, aggregate efficiency falls by the same amount: $\Delta \log A = \Delta \log z$. However, the change in real output and Kaldor-Hicks efficiency is given by (15) plus $\Delta \log z$, which can be either positive or negative. That is, it is possible that real output and Kaldor-Hicks efficiency assign a strictly higher number to a strictly smaller frontier to a first-order.

Having studied the first-order properties of $A(t)$ in models with distortions, we now turn our attention to measuring the distance to the Pareto efficient frontier.

5.2 Misallocation and the Distance to Pareto Frontier

We now focus on a particular counterfactual: we apply Proposition 6 to compute the economic waste caused by distortions. Let the status-quo equilibrium feature wedges $\boldsymbol{\mu}(0) \neq \mathbf{1}$. We compute the economic waste caused by distortions by eliminating wedges, $\boldsymbol{\mu}(t) = \mathbf{1}$, and computing $A(t)$. This quantifies efficiency losses in terms of the fraction of every consumption good that can be saved by eliminating wedges while keeping every agent indifferent to the status-quo. Equivalently, this is also the fraction of factors that can be saved when all distortions are eliminated.

With a complete structural model, $A(t)$ can be computed using Theorem 2. However, below we derive an approximation that is more intuitive and requires less information to be applied.

Proposition 7 (Harberger Triangles)**.** *To a second-order approximation in* $\log \mu$*, the change in aggregate efficiency is*

$$\Delta \log A \approx -\frac{1}{2} \sum_i \lambda_i d \log y_i^{comp} \log \mu_i, \tag{16}$$

where $d \log y_i^{comp} \equiv \sum_j \frac{\partial \log y_i^{comp}}{\partial \log \mu_j} \log \mu_j$ *is the change in the quantity of* i *caused by the wedges in the compensated equilibrium. The approximation error is order* $\log \mu^3$*. The derivatives and expenditure shares in* (16) *are evaluated at the status-quo.*[28] *We provide an explicit formula for* $d \log y_i^{comp}$ *in terms of microeconomic primitives in Appendix F.*

This proposition generalizes deadweight loss triangles to measure aggregate efficiency losses from wedges in general equilibrium economies with heterogeneous agents. The proof relies on translating results from Baqaee and Farhi (2024) using Theorem 2.

There are two advantages to using Proposition 7 over and above simply applying Proposition 6 using a fully-spelled out structural model. First, the Harberger triangles formula can be used to get analytical intuition for misallocation costs through the use of loglinearized expressions, as demonstrated below. Second, it is possible to populate the terms in (7) with considerably fewer assumptions about the primitives of the economy — e.g. the drivers of distortions, productivity processes, and so on.

The intuition for (16) is familiar — a wedge on i is more costly the higher is the Domar weight and the more elastic is the quantity of i relative to the wedge. However, compared to a representative agent model with homothetic preferences, the relevant notion of elasticity here is the one in the compensated equilibrium, not the decentralized one.

We provide some pen-and-paper examples to build intuition.

Example 5 (Misallocation when Markups Vary by Household). Consider the misallocation problem studied by Hsieh and Klenow (2009), but suppose there are multiple agents. Each agent h has CES preferences over consumption goods with elasticity of substitution θ_h. We consider a situation in which each agent pays potentially a different markup μ_{hi} on each good i.[29] Suppose that all consumption goods are ultimately produced linearly

[28] Usually, such quadratic expansions are evaluated at the undistorted point. However, since λ_i and $d \log y_i^{comp}$ are multiplied by one power of $\log \mu$, evaluating these terms at status-quo wedges changes the expression only at the third order. Hence, the stated approximation remains valid to a second-order.

[29] Formally, hi indexes the intermediary between good i and household h. We assume that this intermediary charges a markup of μ_{hi} on its marginal cost. The intermediary's marginal cost is just the price of good i.

from a single common primary factor called labor, which is inelastically supplied.

We can apply Proposition 7 to write the aggregate efficiency losses, up to a second order approximation as

$$\Delta \log A \approx \frac{1}{2}\mathbb{E}_{\chi}\left[\theta_h Var_{\boldsymbol{b}_h}\left[\log \boldsymbol{\mu}_h | h\right]\right], \tag{17}$$

where the expectation uses the vector of household income shares, $\boldsymbol{\chi}$, and the variance uses household budget shares over goods, $\boldsymbol{b}_h$, as weights.[30] Since $A > 1$ means the efficient consumption possibility set can be contracted (scaled by $1/A$) while holding everyone indifferent, the larger $\Delta \log A$, the greater the losses from markups. If all agents have the same preferences and face the same wedges, then the expectation in (17) disappears, and the equation collapses to the single agent case, equation (19), in Baqaee and Farhi (2020).

In words, the reduction in efficiency caused by the markups depends on the average variance in markups paid by each household multiplied by that household's elasticity of substitution. Intuitively, if θ_h is very high, then dispersion in markups faced by h causes a greater reduction in aggregate efficiency. Furthermore, aggregate efficiency falls by more if richer households (those with higher χ_h) are exposed to more markup dispersion. Importantly, this expression does not depend on the *level* of markups paid by each household. A proportional scaling of all markups paid by any household would leave this expression unchanged because increasing all markups on a single household is equivalent to a lump-sum tax on that household, and has no effect on aggregate efficiency.

The next example applies equation (17) to study the efficiency losses due to imperfect insurance.

Example 6 (Misallocation Due to Financial Market Incompleteness). Consider agents with expected utility

$$u_h(\boldsymbol{c}_h) = \sum_s \frac{c_h(s)^{1-1/\theta}}{1-1/\theta}.$$

States of nature, indexed by s, are all equally likely. The coefficient of relative risk aversion is $1/\theta$ (or equivalently, the elasticity of substitution across states is θ).

Each agent h has income $y_h(s) = a_h + \epsilon_h(s)$, where $\epsilon_h(s)$ is an idiosyncratic shock that sums to zero across agents, $\sum_h \epsilon_h(s) = 0$ for every s, with mean zero for each agent $\mathbb{E}[\epsilon_h(s)|h] = 0$. The status-quo allocation is financial autarky, so h's consumption in state

[30] Formally written out, (17) is $= \frac{1}{2}\sum_h \chi_h \theta_h \sum_n b_{hn}\left[\log \mu_{hn} - \sum_{n'} b_{hn'} \log \mu_{hn'}\right]^2$.

s is $c_h^0(s) = y_h(s)$. The aggregate resource constraint for the economy is $\sum_h c_h(s) = \sum_h a_h$, because, by assumption, $\sum_h \epsilon_h(s) = 0$ for every s.

To decentralize this allocation using wedges, suppose that there are complete state-contingent markets with household-by-state wedges $\mu_h(s)$. Household h's budget constraint can be written as

$$\sum \mu_h(s) c_h(s) = a_h.$$

The wedges that decentralize the status-quo allocation must satisfy

$$\frac{c_h^0(s)}{c_h^0(s')} = \left[\frac{\mu_h(s)}{\mu_h(s')}\right]^{-\theta}.$$

Substituting these wedges into equation (17) implies that the gains from completing financial markets are approximately given by:

$$\Delta \log A \approx \frac{1}{2}\theta \mathbb{E}_\chi \left[Var\left[\log \mu_h(s)|h\right]\right] = \frac{1}{2}\frac{1}{\theta}\mathbb{E}_\chi \left[Var\left[\log c_h(s)|h\right]\right],$$

where χ_h is household h's expected share of consumption in the status-quo, $a_h / \sum_{h'} a_{h'}$. This formula disregards inequality due to dispersion in the persistent component of income (dispersion in consumption caused by a_h) because the variance is conditional on household h. Instead, misallocation depends on the average conditional consumption variance, weighted by household income. Holding the consumption process fixed, the gains from completing financial markets are larger, the higher is the risk aversion $1/\theta$ parameter.

The example above is simple, but hints at a much more general idea. Baqaee and Burstein (2025b) build on this basic idea to analyze and quantify the losses from financial market imperfections in both open and closed economies. In that paper, we discuss how to allow for dynamics, labor-leisure choice, capital accumulation, borrowing constraints, and international trade.

The final example in this section shows that misallocation, as measured by $\Delta \log A$, need not equal to changes in real output, $\Delta \log Y$, or changes in the welfare of a representative agent, $\Delta \log A^{RA}$, even in cases where a representative agent exists.

Example 7 (Distance to Frontier: Real Output and Positive Representative Agent). Suppose each agent h's has CES preferences over consumption goods with elasticity of substitution θ_h. Suppose each agent consumes a different selection of goods but all goods are produced linearly from the labor endowment. The markup on the ith good consumed

by household h is denoted by μ_{hi}. Suppose that we eliminate markups, each households' share of income, χ_h, stays constant. Since the distribution of income is constant, there is a positive (and in this case, also normative) representative agent with Cobb-Douglas preferences across each households' consumption bundle (i.e. an agent whose utility is maximized by observed allocations). The change in the welfare of this representative agent, in consumption-equivalent terms, is equal to the change in chain-weighted real output, and both are equal to a second-order approximation to

$$\Delta \log Y = \Delta \log A^{RA} \approx \underbrace{\frac{1}{2}\mathbb{E}_{\chi}\left[\theta_h Var_{\boldsymbol{b}_h}\left[\log \boldsymbol{\mu}_h | h\right]\right]}_{\approx \Delta \log A} + \frac{1}{2} Var_{\boldsymbol{\chi}}\left[\mathbb{E}_{\boldsymbol{b}_h}\left[\log \boldsymbol{\mu}_h | h\right]\right],$$

where we use (17). The change in real output and the welfare of the representative agent are (weakly) larger than the change in aggregate efficiency. One limiting case of this is where all markup-variation is at the household level. In this case, $\Delta \log A = 0$ because the economy is already on the efficient frontier and eliminating markups is purely redistributive. However, in this example, $\Delta \log Y = \Delta \log A^{RA} > 0$.

---

6 Aggregate Efficiency with Limits to Redistribution

In this section, we extend the analyses in Sections 3 to 5 to allow for imperfect redistributive tools. This is another advantage of our approach relative to measures based on adding up willingness-to-pay across all households (e.g. as in Kaldor-Hicks). Intuitively, when we add up willingness-to-pay, we implicitly assume that winners can compensate losers. In Section 4, we illustrated one issue with this approach: monetary compensations can change relative prices so that, in practice, the necessary compensations are infeasible. In this section, we focus on a second issue — monetary compensations may be infeasible because lump-sum transfers are not available.

Theorem 1 applies regardless of what redistributive tools are available. In this section, we apply Theorem 1 to the case where redistribution can only be achieved via linear taxation in general equilibrium with wedges.

6.1 General Solution with Linear Taxes

Consider a decentralized equilibrium with technologies $\boldsymbol{z}$ and wedges $\boldsymbol{\mu}$. We allow a vector of linear taxes $\boldsymbol{\tau}$ on different goods, and let the vector $\boldsymbol{T}$ dictate the amount of tax revenues sent to each household. We require budget-balance, so that total tax revenues

must equal total transfers to households. The set of feasible tax-and-transfer schemes is $\mathcal{T}$. As before, $A(t)$ is the maximum contraction of the consumption possibility set that allows every agent to remain indifferent to the status-quo.

Recall that Theorem 1 implies that aggregate efficiency satisfies:

$$A(t) = \max_{\boldsymbol{c} \in \mathcal{C}(t)} U(\boldsymbol{c}), \tag{18}$$

where $\mathcal{C}(t)$ is the set of consumption allocations that can be supported as equilibria given technologies, $\boldsymbol{z}(t)$, wedges, $\boldsymbol{\mu}(t)$, and feasible tax-and-transfer schemes. As before, we index technologies and wedges by a scalar t and let $t = 0$ denote the status-quo.

Figure 5 illustrates $A(t)$ using a two household example. In the figure, the status-quo allocation, $\boldsymbol{c}(0)$, and the decentralized allocation without transfers, $\boldsymbol{c}(t)$, are denoted by red circles. The solid blue line shows the feasible consumption possibility frontier at t given distortionary taxation, and the dashed line indicates the frontier at t given unrestricted lump-sum taxes. The two frontiers touch at the decentralized point, since the decentralized point does not engender any distortionary redistributive taxation. However, the solid blue set is strictly smaller than the dashed line since distortionary taxation limits the set of feasible redistributions. The change in efficiency, $\Delta \log A$, is still the largest radial contraction of $\mathcal{C}(t)$ that allows every household to be made at least indifferent to the status-quo. Since the larger is the possibility set $\mathcal{C}(t)$ the more it must be contracted to reach indifference, aggregate efficiency gains are larger with better redistributive tools.

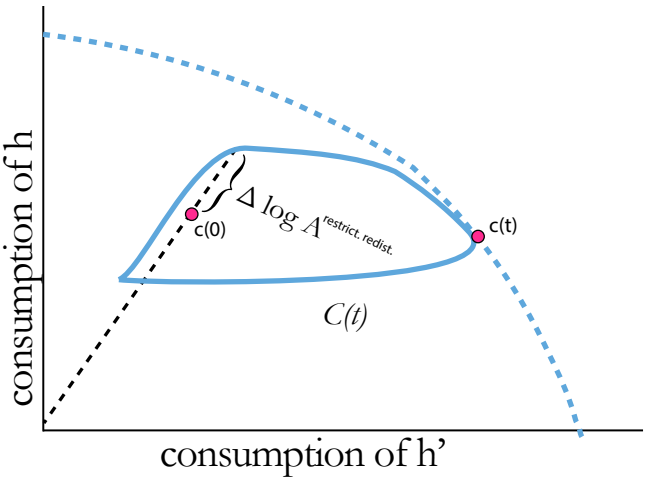


Figure 5: Aggregate efficiency is measured by the maximal radial expansion of the feasible set necessary to achieve indifference.

The following generalization of Proposition 4 makes this point.

Proposition 8 (Productivity Shocks with Limited Redistribution)**.** *Consider a perfectly competitive status-quo without linear taxes and wedges. If households can be exactly compensated, then response of aggregate productivity to a productivity shock,* $\Delta \log \boldsymbol{z}$*, to a second-order approximation, is given by*

$$\Delta \log A = \sum_i \left(\lambda_i + \frac{1}{2} \sum_j \frac{\partial \lambda_i^{comp}}{\partial \log z_j} \Delta \log z_j \right) \Delta \log z_i + \frac{1}{2} \sum_i \lambda_i \left(\sum_j \frac{\partial \log y_i^{comp}}{\partial \log \tau_j} \Delta \log \tau_j^* \right) \Delta \log \tau_i^*, \tag{19}$$

where $\boldsymbol{\tau}^*(t)$ *implement the solution in* (18)*.*

The first set of summands are exactly as in Proposition 4. The second set of summands, which are new and non-positive, capture the inefficiency caused by imperfect redistribution. These are the sum of Harberger triangles associated with the linear taxes in $\tau^*(t)$. If lump-sum taxation is not feasible, then linear taxes must be used, $\Delta \log \tau^* \neq 0$, and aggregate efficiency with limited redistribution is lower than with lump-sum taxation by exactly the sum of deadweight loss triangles. That is, the response of aggregate efficiency to productivity shocks is the same as it would be if lump-sum transfers were possible minus the deadweight loss triangles associated with distortionary taxes. The simplicity of Equation (19) follows from the fact that the status-quo is undistorted. This ensures that (1) there are no interactions of taxes with pre-existing distortions, (2) the cross-partials between $d \log \tau^*$ and $d \log z$ are all zero due to the envelope theorem.

The following simple corollary, obtained by ignoring the second order terms, shows that Hulten's theorem holds, without change, even with limited redistribution, as long as outcomes are interior.

Corollary 2 (Hulten's Theorem with Limited Redistribution)**.** *Consider a perfectly competitive status-quo without linear taxes and wedges. If households can be exactly compensated, the response of aggregate productivity to a productivity shock,* $\Delta \log \boldsymbol{z}$*, to a first-order approximation, is given by*

$$\Delta \log A = \sum_i \lambda_i \Delta \log z_i.$$

Intuitively, the losses from costly-redistribution are second-order, and hence to a first-order approximation, only the direct effects of the productivity shock matter (assuming we start at a competitive equilibrium).

6.2 Analytical Example: Losses from Autarky

We now study the efficiency gains from international trade, accounting for limited redistribution. We use the approximation in Proposition 8 to provide intuition. We check the

numerical performance of the second-order approximation by computing exact results using (18).

Example 8 (Gains from Trade with Limited Redistribution). We revisit Example 3, which studied the gains from trade, but this time we incorporate limits to redistribution. We add a labor-leisure margin, and assume redistribution can only be done by taxing consumption.

Suppose there are two households, and household h has nested-CES preferences over domestic consumption goods, c_{hd}, imported consumption goods, c_{hf}, and leisure l_h:

$$u_h(\boldsymbol{c}_h) = \left[(1-\gamma_h)^{\frac{1}{\rho}}\left[(1-\alpha_h)^{\frac{1}{\theta}} c_{hd}^{\frac{\theta-1}{\theta}} + \alpha_h{}^{\frac{1}{\theta}} c_{hf}^{\frac{\theta-1}{\theta}}\right]^{\frac{\theta}{\theta-1}\frac{\rho-1}{\rho}} + \gamma_h^{\frac{1}{\rho}} l_h^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}}.$$

The model in Example 3 did not feature the leisure good. The inner nest combines domestic and foreign consumption goods with Armington elasticity θ and home bias controlled by the parameter α_h. The outer nest combines the goods bundle with leisure with elasticity of substitution ρ and share parameter γ_h. The parameter ρ controls the Frisch elasticity of labor supply.

Household h is endowed with one unit of time and a_h efficiency units of labor and faces a budget constraint:

$$\tau p_d c_{hd} + \tau p_f c_{hf} = w a_h (1 - l_h) + T_h,$$

where p_d and p_f denote the price of each consumption good, w_h is the wage per efficiency unit, τ is the gross-tax rate on consumption, and T_h is a lump sum transfer. Budget balance requires $(\tau - 1)\sum(p_d c_{hd} + p_f c_{hf}) = \sum T_h$. The domestic consumption good is produced linearly with labor, so the resource constraint for domestic consumption is $\sum_h c_{hd} = \sum_h a_h(1 - l_h)$, with $p_d = w$. The resource constraint for leisure is $l_h \leq 1$.

The status-quo is a competitive equilibrium without taxes in which the country trades with the rest of the world. We consider the gains from trade relative to autarky by raising iceberg trade costs to infinity. The efficiency loss from autarky, $\Delta \log A$, measures the increase in the autarky consumption possibility set needed to keep every consumer indifferent to the status-quo. The consumption possibility set, which consists of allocations of consumption and leisure for each agent, encodes the distortionary impact of consumption taxes required to transfer income between households. We compare two cases: (1) lump-sum taxation is available and the second welfare theorem holds; (2) lump-sum taxation is not available, $T_h \geq 0$, and linear consumption taxes must be used.

Let

$$\Omega_{hd} = \frac{p_d^0 c_{hd}^0 + p_f^0 c_{hf}^0}{w^0 a_h}$$

denote household h's budget share on consumption in the status-quo as a share of the value of h's total time endowment (the remainder is implicit expenditures on leisure). For simplicity of exposition, and since it is fairly realistic, we assume that both households work the same number of hours in the status-quo, which implies that $\Omega_{hd} = \Omega_d$ does not vary by household. Let s_{hd} denote household h's share of expenditures on the domestic consumption good relative to all consumption goods:

$$s_{hd} = \frac{p_d^0 c_{hd}^0}{p_d^0 c_{hd}^0 + p_f^0 c_{hf}^0}.$$

Lump-Sum Taxation. With lump-sum taxation, using Proposition 4, we can write the losses from autarky to a second-order approximation as

$$\Delta \log A^{\text{lump-sum}} \approx \underbrace{\Omega_d \mathbb{E}_\chi \left[\frac{\log s_h}{\theta - 1}\right]}_{\text{1st order}} \underbrace{- \frac{1}{2}\Omega_d^2 Var_\chi \left(\frac{\log s_h}{\theta - 1}\right) + \frac{1}{2}(\rho - 1)\Omega_d(1 - \Omega_d)\mathbb{E}_\chi \left[\left[\frac{\log s_h}{\theta - 1}\right]^2\right]}_{\text{2nd order with lump-sum taxation}}.$$

This expression is identical to Equation (13) in Example 3 when there is no leisure, $\Omega_d = 1$. The first and second summands are the same as in (13) but are now scaled by Ω_d to account for the fact that households also consume leisure. The final summand, which is absent in (13), accounts for complementarities/substitutabilities between consumption and leisure. If consumption and leisure are complements, $\rho < 1$, then the reduction to consumption caused by autarky reduces the value of leisure through complementarity, raising losses from autarky further.

Linear Taxation. Now consider the case where lump-sum taxation is unavailable so that lump-sum transfers must be non-negative: $\boldsymbol{T} \geq 0$, financed by a uniform consumption tax. Proposition 8 now implies that, to a second-order approximation,

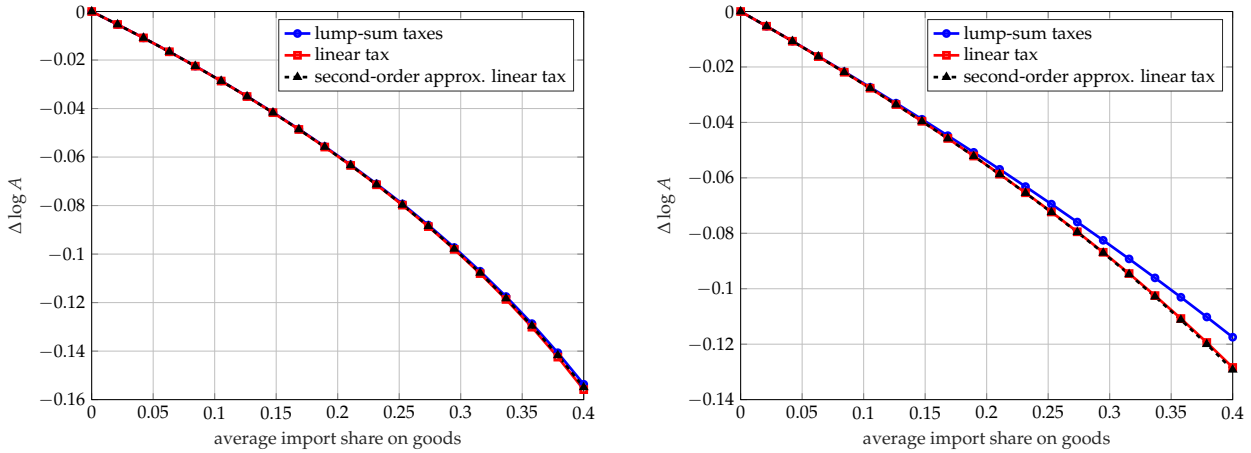
$$\Delta \log A^{\text{linear tax}} \approx \Delta \log A^{\text{lump-sum}} - \underbrace{\frac{1}{2}\rho\Omega_d(1 - \Omega_d)(d \log \tau^*)^2}_{\text{2nd order losses from distorting taxes}},$$

where τ^* is the optimal consumption tax in Equation (18).

Index the two households by h and h' and suppose that h is more exposed to foreign goods: $s_{hd} < s_{h'd}$. This means that, in the decentralized equilibrium, household h is more negatively affected by the trade shock than h'. In this case, the optimal feasible tax-and-transfer from (18) sends all collected tax revenues to h.

The tax required for the compensation is $d \log \tau^* = \frac{\chi_h}{\theta - 1} [\log s_{h'd} - \log s_{hd}] > 0$, to a first-order, where χ_h is h's share of aggregate income. The required consumption tax is larger the bigger is the heterogeneity in exposure to the trade shock, and the larger is household h's share of aggregate income. A given tax is more distorting the higher is ρ, which controls substitution between consumption and leisure (ρ can be interpreted as the Frisch elasticity of labor supply), and the closer is Ω_d to $1/2$. If Ω_d is equal to either one (households do not value leisure) or zero (households do not value consumption), then there is no distortion from the tax.

Figure 6 numerically illustrates the performance of the second-order approximation to the exact solution with distortionary linear taxes, and compares them both to the solution with lump-sum taxes. The second-order approximation performs well even for large shocks. Panel (a) uses $\rho = 0.5$, so consumption and leisure are complements and the Frisch elasticity of labor supply is a reasonable 0.5. Since ρ is low, distortionary taxes are able to achieve an outcome that is roughly as good as lump-sum taxes. Panel (b) uses a much higher $\rho = 3$. In this case, the gap between the lump-sum and linear taxation scenarios is larger since consumption taxes reduces labor and increase leisure, which causes efficiency to fall.



(a) $\rho = 0.5$ (b) $\rho = 3.0$

Figure 6: A numerical example of the losses from autarky with and without distortionary redistribution. The other parameter values are $\Omega_d = 0.5$, $\chi_h = 0.5$, $\theta_h = 3$, $s_{hd} = 3s_{h'd}$.

To summarize: losses from autarky are larger if some households are more badly affected than others, especially if efficient redistributive tools are not available to compensate the households that are more badly affected.

In Appendix D, we consider another example, where skill-biased technical change affects different workers differently, and show how the second order approximation in Proposition 8 performs.

6.3 Quantitative Example: the China Shock

Our final example uses an off-the-shelf quantitative trade model to quantify the aggregate efficiency effects of international shocks. We consider the effect of the rise of China on the United States using the Armington model in Baqaee and Farhi (2024). Each country has different factor endowments, and we treat owners of different factor endowments as different agents. The rise of China, which we model by raising Chinese productivity, changes relative wages amongst the different domestic factors, with different consequences for different agents. We study how aggregate efficiency for the U.S. changes as China becomes more productive, as measured by the aggregate consumption equivalent, and we study how this number changes depending on assumptions about factoral mobility (across sectors) and the availability of tax and transfer tools.

Summary of Calibration. The calibration follows Baqaee and Farhi (2024). The model has 40 countries and 30 industries in each country. Production by each industry is a nested CES aggregator combining four domestic primary factors (low-, medium-, high-skill labor, and capital) with intermediate inputs. The intermediate input bundle used by each industry is a nested CES aggregator cover all industries and origin countries. All households in each country consume the same domestic consumption bundle, which is a homothetic nested CES aggregator over all industries and origin countries. (We abstract from heterogeneity in preferences within countries). The initial expenditure shares are calibrated according to the World Input-Output Database in 2008. Since tariffs in 2008 were quite low, the model is calibrated assuming there are no import tariffs.[31]

[31]The elasticity of substitution between primary factors is set to one. The elasticity of substitution between value-added and intermediates is 0.5. Each country-industry pair has a unique bundle of intermediate inputs sourced from different industries with elasticity of substitution 0.2. Each country also has a unique consumption bundle, with elasticity of substitution across industries of 0.9. Every destination country-industry pair purchases a unique bundle of inputs from other industries sourced from different origin countries. The Armington trade elasticity is equal to 5.

Consumption Possibility Set. In 2008 (the calibration year), China's GDP is roughly 5% of the world's. By 2023, this number had risen to 18%. We model the rise of China through an increase in Chinese factor-augmenting productivity growth (roughly tripling the efficiency units of all Chinese factor endowments) to ensure that China's share of world GDP rises to 18%.

We consider how this shock affects the United States, measured using the aggregate consumption equivalent change, under four different scenarios. Each scenario defines a different consumption possibility set $\mathcal{C}$ for U.S. consumers, depending on the tools available to the government. These scenarios are:

i. *Tariffs with lump-sum transfers:* the government can raise a uniform tariff on imports, and also has access to unrestricted lump-sum transfers (i.e. lump-sum transfers can be positive or negative).

ii. *Tariffs with targeted rebates:* the government can raise a uniform tariff on imports and has full discretion on how to rebate any additional tariff revenues (i.e. if tariff revenues rise after the shock, in units of world GDP, then the government can choose who to rebate that additional revenue to).

iii. *Tariffs with non-targeted rebates:* the government can raise a uniform tariff on all imports but any revenues generated by tariffs are rebated back to domestic households in proportion to their pre-shock initial share of aggregate income.

iv. *Laissez-faire*: there are no redistributive tools available. The consumption possibility set is a single point corresponding to the decentralized equilibrium in the U.S..

In scenarios i., ii. and iii., the consumption possibility set is the set of equilibrium consumption allocations for U.S. consumers given different levels of uniform import tariffs and feasible transfers. It is this set that we expand or contract to ensure all U.S. households can be made at least indifferent to the status-quo.[32] The consumption possibility set in each scenario is smaller than the preceding, which implies the efficiency gains will be smaller as well.

We now turn to the choice of the status-quo. We assume that the 2008 data correspond to laissez-faire in which all taxes are zero. If we were to treat this as the status-quo, then

[32]That is, for this exercise, we use a partial equilibrium consumption possibility set — we do not consider how expanding or shrinking the U.S. consumption possibility set would, in general equilibrium, be achieved and what effects it would have on world markets. We do this for two reasons. First, it demonstrates the flexibility of our approach, which can nest partial equilibrium exercises as a special case. Second, for this exercise, since we are only interested in the welfare of U.S. households, rather than all consumers in the world, the partial equilibrium exercise of expanding or shrinking U.S. consumption set is more easily interpretable.

even without the China shock our efficiency measure would detect $A > 1$, reflecting the gains from optimal tariffs. Our focus is not on quantifying these tariff gains, but to illustrate how an imperfect redistributive tool — here, an import tariff —influences the aggregate efficiency implications of the China shock.

We therefore assume that, if an import tariff is available (scenarios ii, iii, iv), then that import tariff is already set to maximize the quantity of the U.S. consumption good prior to the China shock. This ensures that the status-quo allocation is not Pareto dominated by other allocations in the feasible status-quo consumption possibility set (i.e. in the absence of the China shock, aggregate efficiency is 1 by construction in every scenario). Nevertheless, to keep the status-quo allocation similar to the data in 2008 — which we calibrate the lessaiz faire model to — we assume that U.S. tariffs provoke symmetric retaliation from the rest of the world. As a result, the optimal U.S. tariff in the status-quo is small anyhow (2.3% in the case without factor mobility and 1.0% in the case with factor mobility), ensuring that expenditure shares in the status-quo are close to the 2008 data prior to the shock.

Solution Strategy. To solve for aggregate efficiency, note that if we substitute the definition of $U(\boldsymbol{c})$ into (18), we get

$$A(\boldsymbol{z}) = \max \left\{ \min_h \tilde{u}_h \left(\boldsymbol{c}_h(\boldsymbol{z}, \boldsymbol{\tau}, \boldsymbol{T}) \right) : (\boldsymbol{\tau}, \boldsymbol{T}) \in \mathcal{T} \right\},$$

where $\boldsymbol{c}_h(\boldsymbol{z}, \boldsymbol{\tau}, \boldsymbol{T})$ are equilibrium U.S. consumption vectors given taxes $\boldsymbol{\tau}$, transfers, $\boldsymbol{T}$, and technologies $\boldsymbol{z}$. (We omit dependence on wedges because all non-import-tariff wedges $\boldsymbol{\mu}$ are equal to one). We make this max-min problem more computationally tractable, by noting that it is equivalent to

$$A(\boldsymbol{z}) = \max \left\{ x : (\boldsymbol{\tau}, \boldsymbol{T}) \in \mathcal{T}, \tilde{u}_h \left(\boldsymbol{c}(\boldsymbol{z}, \boldsymbol{\tau}, \boldsymbol{T}) \right) \geq x \right\},$$

where we replace the inner minimum with a series of constraints instead. This reduces the max-min problem to a simpler linear-programming problem that can be solved easily on a computer.

Results. The change in aggregate efficiency due to the China shock is shown in Table 1. We consider two specifications of the model, labelled "immobile" and "mobile" factors. When factors are "immobile", primary factors (labor and capital) in each country-industry pair cannot move across industries. When factors are mobile, there is one na-

tional market for each factor type (low-, medium-, high-skill labor and capital) and all industries in a country hire from the national market. Comparing these two specifications reveals the importance of reallocation for determining aggregate efficiency.

Table 1: Effect of China Shock on the United States

Scenario	Immobile Factors		Mobile Factors	
	$\Delta \log A$	Δ Tariff (p.p.)	$\Delta \log A$	Δ Tariff (p.p.)
Tariffs & lump-sum transfers	0.008	+0.1	0.010	+0.3
Tariffs & targeted rebates	-0.010	+7.4	0.010	+0.4
Tariffs & non-targeted rebates	-0.179	+11.4	0.008	+1.6
Laissez-faire	-0.235	–	0.008	–

Tariff changes are expressed in percentage points.

The first row shows the change in aggregate efficiency for the U.S., assuming there are lump-sum transfers. In this case, Proposition 1 holds, and the response of $\Delta \log A$ coincides with the response of aggregate real consumption in the U.S.. This also coincides with real consumption by the representative U.S. household, as well as the change in the Kaldor-Hicks measure of efficiency. In this case, $\Delta \log A$ rises by around 1 log point in response to the China shock.

The second row considers the case where import tariffs are available, but redistribution can only draw on excess tariff revenues. In this case, $\Delta \log A$ falls by 1 log point if factors cannot move across sector. This is because the China shock causes real wages in some sectors to fall. To compensate these households, we must raise tariff revenues by around 7%. These tariffs, which trigger retaliatory tariffs from the rest of the world, cause overall efficiency to decline. In other words, in order to make every U.S. household at least as well off after the China shock requires expanding the consumption possibility set by 1 log point. The picture is very different if factors are mobile, since in this case, real wages do not decline, and hence large tariffs are not needed to offset real income losses suffered by some subset of households.

The third row considers the case where tariff revenues can only be distributed according to a fixed distribution (pre-shock initial income shares). In this case, since the redistributive tools are much more severely restricted, it is much harder to compensate losers (it may not be possible to equalize the change in real consumption across consumers). Accordingly, aggregate efficiency falls by 17.9 log points, since compensating workers in losing industries requires instituting very large tariffs (of around 11.4%), which then trigger a large trade war. Once again, these restrictions on redistributive policy are much

less important if factors are mobile across sectors, with the overall efficiency gains being barely affected.

The last and final row considers the Laissez-faire case, where $\Delta \log A$ simplifies to being simply the change in real consumption for the workers whose real wages decline the most in Lessaiz-faire. When factors are immobile, this is "Textile and Leather Products," where real wages decline by around 23 log points in response to the China shock. The contrast with the mobile factor case is very stark: when factors are mobile, heterogeneous effects among factors is substantially attenuated, and so the change in aggregate efficiency closely approximates that under full redistribution.

The results in Table 1 show that the change in aggregate efficiency depends strongly on (a) the extent to which shocks have asymmetric effects on households, and (b) the redistributive tools available. Note that although the redistributive tools are very important for quantifying the change in aggregate consumption equivalents in Table 1, we do not take a stance on how these redistributive tools should be used in practice. For example, if lump-sum transfers are available, then the consumption possibility set can be contracted by 1% and everyone can be kept at least indifferent. Hence, after the shock and compensations, there is a 1% surplus of the consumption good. We measure this surplus without taking a stance on how it should be distributed.

7 Conclusion

We generalize the cost-benefit approach to aggregate efficiency to environments with heterogeneous agents, general equilibrium, and limited redistribution. Our measure, which converts shocks into a welfare-equivalent change in total factor productivity, collapses to the Solow residual and Kaldor-Hicks efficiency when households have the same homothetic preferences and face the same relative prices.

We show how to compute this measure by maximizing the utility of some fictional as-if agent. This provides a method to translate theorems and tools about representative-agent economies to study aggregate efficiency in economies with heterogeneous agents, including, for example, Hulten (1978), Harberger (1964), Hsieh and Klenow (2009), Arkolakis et al. (2012), Baqaee and Farhi (2019b), and Baqaee and Farhi (2020). This also allows us to calculate aggregate efficiency using only information on observables like expenditure shares and price elasticities of supply and demand curves.

In two stand-alone companion papers, we apply the theoretical results of this paper to other contexts where household heterogeneity is central. Baqaee and Burstein (2025b) consider losses in aggregate efficiency from financial market incompleteness, within and

across borders. Baqaee and Burstein (2025a) characterize aggregate efficiency in random utility models with discrete choice, focusing on spatial economies, where households make different choices due to differences in their preferences. An interesting extension, which we do not pursue in this paper but pursue in ongoing work, is to study policy problems where maximizing aggregate efficiency is the objective of the policymaker.[33]

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica 80*(5), 1977–2016.

Allais, M. (1979). La théorie générale des surplus.

Antràs, P., D. Chor, T. Fally, and R. Hillberry (2012). Measuring the upstreamness of production and trade flows. *The American Economic Review 102*(3), 412–416.

Antras, P., A. de Gortari, and O. Itskhoki (2017). Globalization, inequality and welfare. *Journal of International Economics 108*(C), 387–412.

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012). New trade models, same old gains? *American Economic Review 102*(1), 94–130.

Atkinson, A. B. (1970). On the measurement of inequality. *Journal of economic theory 2*(3), 244–263.

Auer, R., A. Burstein, S. Lein, and J. Vogel (2024). Unequal expenditure switching: Evidence from switzerland. *Review of Economic Studies 91*(5), 2572–2603.

Auerbach, A. J. and L. J. Kotlikoff (1987). *Dynamic Fiscal Policy*. Cambridge University Press.

Baqaee, D. and A. Burstein (2025a). Aggregate efficiency with discrete choice. Technical report.

Baqaee, D. and E. Farhi (2019a, July). Networks, barriers, and trade. Working Paper 26108, National Bureau of Economic Research.

Baqaee, D. and E. Rubbo (2023). Micro propagation and macro aggregation. *Annual Review of Economics 15*(1), 91–123.

Baqaee, D. R. and A. Burstein (2023). Welfare and output with income effects and taste shocks. *The Quarterly Journal of Economics 138*(2), 769–834.

Baqaee, D. R. and A. Burstein (2025b, September). Efficiency costs of incomplete markets. Working Paper 34233, National Bureau of Economic Research.

Baqaee, D. R. and E. Farhi (2019b). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. *Econometrica 87*(4), 1155–1203.

Baqaee, D. R. and E. Farhi (2020). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics 135*(1), 105–163.

Baqaee, D. R. and E. Farhi (2024). Networks, barriers, and trade. *Econometrica 92*(2), 505–541.

Bergson, A. (1938, February). A reformulation of certain aspects of welfare economics. *The Quarterly Journal of Economics 52*(2), 310–334.

Bhandari, A., D. Evans, M. Golosov, and T. Sargent (2021). Efficiency, insurance, and redistribution effects of government policies. Technical report, Working paper.

Bigio, S. and J. La'O (2016). Financial frictions in production networks. Technical report.

[33]The maximization problem over tax instruments in Section 6 is not an optimal policy problem. Instead, here the maximization is only used to define aggregate efficiency.

Blackorby, C. and D. Donaldson (1990). A review article: The case against the use of the sum of compensating variations in cost-benefit analysis. *canadian Journal of Economics*, 471–494.

Boadway, R. W. (1974). The welfare foundations of cost-benefit analysis. *The Economic Journal 84*(336), 926–939.

Bornstein, G. and A. Peter (2024). Nonlinear pricing and misallocation. Technical report, National Bureau of Economic Research.

Comin, D., D. Lashkari, and M. Mestieri (2021). Structural change with long-run income and price effects. *Econometrica 89*(1), 311–374.

Cornes, R. (1992). *Duality and modern economics*. CUP Archive.

Dávila, E. and A. Schaab (2022). Welfare assessments with heterogeneous individuals. Technical report, National Bureau of Economic Research.

Dávila, E. and A. Schaab (2023). Welfare accounting. Technical report, National Bureau of Economic Research.

Debreu, G. (1951). The coefficient of resource utilization. *Econometrica: Journal of the Econometric Society*, 273–292.

Debreu, G. (1954). A classical tax-subsidy problem. *Econometrica: Journal of the Econometric Society*, 14–22.

Dekle, R., J. Eaton, and S. Kortum (2008). Global rebalancing with gravity: measuring the burden of adjustment. *IMF Staff Papers 55*(3), 511–540.

Divisia, F. (1925). L'indice monétaire et la théorie de la monnaie. *Revue d'économie politique 39*(4), 842–861.

Donald, E., M. Fukui, and Y. Miyauchi (2023). Unpacking aggregate welfare in a spatial economy. Technical report, Tech. rep.

Dupuit, A. J. É. J. (1844). De la mesure de lútilité des travaux publics. *Annales des ponts et chaussées 8*.

Farrell, M. J. (1957). The measurement of productive efficiency. *Journal of the royal statistical society series a: statistics in society 120*(3), 253–281.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*(3), 733–772.

Galle, S., A. Rodríguez-Clare, and M. Yi (2023). Slicing the pie: Quantifying the aggregate and distributional effects of trade. *The Review of Economic Studies 90*(1), 331–375.

Harberger, A. C. (1954). Monopoly and resource allocation. In *American Economic Association, Papers and Proceedings*, Volume 44, pp. 77–87.

Harberger, A. C. (1964). The measurement of waste. *The American Economic Review 54*(3), 58–76.

Harberger, A. C. (1971). Three basic postulates for applied welfare economics: an interpretive essay. *Journal of Economic literature 9*(3), 785–797.

Harsanyi, J. C. (1955). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of political economy 63*(4), 309–321.

Heathcote, J. and H. Tsujiyama (2021). Optimal income taxation: Mirrlees meets ramsey. *Journal of Political Economy 129*(11), 3141–3184.

Hendren, N. and B. Sprung-Keyser (2020). A unified welfare analysis of government policies. *The Quarterly journal of economics 135*(3), 1209–1318.

Hicks, J. R. (1939). The foundations of welfare economics. *The economic journal 49*(196), 696–712.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and manufacturing TFP in China and India. *The quarterly journal of economics 124*(4), 1403–1448.

Hulten, C. R. (1973). Divisia index numbers. *Econometrica: Journal of the Econometric Society*, 1017–1025.

Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, 511–518.

Johansson, P.-O., B. Kriström, et al. (2022). The boadway paradox revisited: The case of compensated equilibrium. *Center for Environmental and Resource Economics, CERE, The Boadway Paradox Revisited: The Case of Compensated Equilibrium (June 13, 2022)*.

Jones, C. (2002). The" boadway paradox" revisited.

Kaldor, N. (1939). Welfare propositions of economics and interpersonal comparisons of utility. *The economic journal 49*(195), 549–552.

Kim, R. and J. Vogel (2020). Trade and welfare (across local labor markets). Technical report, National Bureau of Economic Research.

Liu, E. (2017). Industrial policies and economic development. Technical report.

Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.

Luenberger, D. G. (1996). Welfare from a benefit viewpoint. *Economic Theory 7*(3), 445–462.

Mas-Colell, A., M. D. Whinston, J. R. Green, et al. (1995). *Microeconomic theory*, Volume 1. Oxford university press New York.

Petrin, A. and J. Levinsohn (2012). Measuring aggregate productivity growth using plant-level data. *The RAND Journal of Economics 43*(4), 705–725.

Restuccia, D. and R. Rogerson (2008). Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic dynamics 11*(4), 707–720.

Rodríguez-Clare, A., M. Ulate, and J. P. Vasquez (2022). *Trade with nominal rigidities: Understanding the unemployment and welfare effects of the China shock*. CESifo, Center for Economic Studies & Ifo Institute.

Samuelson, P. A. (1947). *Foundations of Economic Analysis*. Cambridge, MA: Harvard University Press.

Schulz, K., A. Tsyvinski, and N. Werquin (2023). Generalized compensation principle. *Theoretical Economics 18*(4), 1665–1710.

Scitovszky, T. (1941). A note on welfare propositions in economics. *The Review of Economic Studies 9*(1), 77–88.

Solow, R. M. (1957). Technical change and the aggregate production function. *The review of Economics and Statistics*, 312–320.

Appendix

A Expenditure Function of Compensated Agent **47**

B Definition of Positive & Normative Representative Agent **48**

C Proofs **48**

D Additional Examples and Derivations **55**

E Non-homothetic Preferences **59**

F Explicit Characterization of Compensated Equilibrium **61**
F.1 Input-Output Notation 61
F.2 Differential Hat-Algebra 63

Appendix A Expenditure Function of Compensated Agent

The following proposition characterizes the expenditure function of the compensated agent.

Proposition 9 (Dual Representation of Hicksian Representative Agent)**.** *The expenditure function associated with $U(\boldsymbol{c})$, in Definition 3, denoted by $E(p, U)$ is*

$$E(\boldsymbol{p}, U) = \left(\sum_{h \in H} e_h(\boldsymbol{p}; u_h(\boldsymbol{c}_h^0)) \right) U.$$

By Shephard's lemma, the budget share of the Hicksian representative agent on good i, denoted b_i^{comp}, is

$$b_i^{comp}(\boldsymbol{p}) = \frac{\partial \log E(\boldsymbol{p}, U)}{\partial \log p_i} = \sum_h \frac{e_h(\boldsymbol{p}, u_h^0)}{\sum_{h'} e_{h'}(\boldsymbol{p}, u_{h'}^0)} b_{hi}(\boldsymbol{p}, u_h^0),$$

where $b_{hi}(\boldsymbol{p}, u_h^0)$ is the compensated budget share of household i at the status-quo indifference curve $u_h^0 \equiv u_h(\boldsymbol{c}_h^0)$.

In words, the compensated agent's budget share on each good i is the average compensated budget share of all households, where each household is weighted according to its compensating income, $e_h(\boldsymbol{p}, u_h^0)$.

Given compensated aggregate budget shares, $b_i^{\text{comp}}(\boldsymbol{p})$, we can solve for equilibrium variables in the compensated equilibrium including prices $\boldsymbol{p}^{\text{comp}}$. Setting aggregate spending to be the numeraire in the compensated equilibrium, and using Theorem 2, we know that

$$A(t) = U(t) = \frac{1}{\sum_{h \in H} e_h(\boldsymbol{p}^{\text{comp}}(t); u_h(\boldsymbol{c}_h^0))} = A^{KH,comp}(t).$$

Appendix B Definition of Positive & Normative Representative Agent

We follow the definitions in Mas-Colell et al. (1995). We say that $u^{RA} : \mathbb{R}^N \to \mathbb{R}$ is a *positive representative agent* if the Marshallian demand curves generated by u^{RA}, given prices and total income, coincide with equilibrium allocations given the same prices and aggregate income:

$$\arg\max_{\boldsymbol{c}} \{u^{RA}(\boldsymbol{c}) : \sum_i p_i(t) c_i \leq I(t)\} = \sum_h \arg\max_{\boldsymbol{c}_h} \{u_h(\boldsymbol{c}_h) : \sum_i p_i(t) c_{hi} \leq I_h(t)\}.$$

The positive representative agent, $u^{RA} : \mathbb{R}^N \to \mathbb{R}$, is a *normative representative agent* relative to the social welfare function W if for every $(\boldsymbol{p}(t), I(t))$, the distribution of wealth across households, denoted by $\{I_h(t)\}$, also maximizes

$$W(v_1(\boldsymbol{p}(t), I_1(t)), \ldots, v_H(\boldsymbol{p}(t), I_H(t)))$$

subject to $\sum_{h=1}^H I_h(t) = I(t)$, where v_h is the indirect utility function of agent h.

Appendix C Proofs

Proof of Theorem 1. Denote the solution to (1) by Z^*, and an allocation that attains this solution (it does not need to be unique) by $\boldsymbol{c}^* \in (Z^*)^{-1}\mathcal{C}$. By the definition of Z^* and $U(\cdot)$, and given local nonsatiaton, $U(\boldsymbol{c}^*) = U(\boldsymbol{c}^0) = 1$. Define the value to the compensated agent of the set $\mathcal{C}$ by $V(\mathcal{C}) = \max_{\boldsymbol{c} \in \mathcal{C}} U(\boldsymbol{c})$. Denote the solution to $V(\mathcal{C})$ by $\boldsymbol{c}^{**}$, with $\boldsymbol{c}^{**} \in \mathcal{C}$. Since $U(\cdot)$ is homogeneous of degree 1, it follows that $V((Z^*)^{-1}\mathcal{C}) = (Z^*)^{-1}V(\mathcal{C})$, with solution $(Z^*)^{-1}\boldsymbol{c}^{**}$. Note that $V\left((Z^*)^{-1}\mathcal{C}\right) \geq 1$ because $\boldsymbol{c}^* \subseteq (Z^*)^{-1}\mathcal{C}$ and $U(\boldsymbol{c}^*) = 1$. Moreover, $V\left((Z^*)^{-1}\mathcal{C}\right) = 1$ because, if it were strictly higher than 1, the solution to (1)

would be higher than Z^*. It thus follows that:

$$V(\mathcal{C}) = \frac{V(\mathcal{C})}{V\left((Z^*)^{-1}\mathcal{C}\right)} = Z^*.$$

□

Proof of Theorem 2. Recall that $C(t)$ is the set of consumption allocations that can be attained in general equilibrium given some transfers and compensated demands. Define

$$\boldsymbol{c}^{**}(t) \in \operatorname{argmax}\left\{U(\boldsymbol{c}) : \boldsymbol{c} \in C(t)\right\}.$$

From Theorem 1, we know that $U(\boldsymbol{c}^{**}(t)) = A(t)$. We suppress t to simplify notation.

First, suppose that $\boldsymbol{\mu} = \mathbf{1}$. Define $\tilde{\mathcal{C}}$ to be the set of technologically feasible consumption allocations and note that $\mathcal{C} \subset \tilde{\mathcal{C}}$. We know that

$$c^{**} \in \arg\max\{U(\boldsymbol{c}) : \boldsymbol{c} \in \tilde{\mathcal{C}}\},$$

because the solution to this problem is Pareto-efficient, and by the second welfare theorem, Pareto efficient allocations can be decentralized using some transfers. Furthermore, we know that the aggregate budget set in the compensated equilibrium, denoted by $B(\boldsymbol{p}^{\text{comp}})$ contains $\tilde{\mathcal{C}}$. Hence $U(\boldsymbol{c}^{\text{comp}}) \geq U(\boldsymbol{c}^{**})$. To derive a contradiction, suppose that $U(\boldsymbol{c}^{\text{comp}}) > U(\boldsymbol{c}^{**})$. This implies that $\tilde{u}_h(\boldsymbol{c}_h^{\text{comp}}) > \tilde{u}_h(\boldsymbol{c}_h^*)$ for every h and $\boldsymbol{c}^{\text{comp}}$ is technologically feasible. This violates Pareto efficiency of $\boldsymbol{c}^{**}$. Hence $U(\boldsymbol{c}^{\text{comp}}) = U(\boldsymbol{c}^*) = A$.

Now suppose that $\boldsymbol{\mu} \neq \mathbf{1}$ and that every household is exactly compensated. There exists some $\boldsymbol{p}^{**}(t)$ and $\boldsymbol{T}^{**}(t)$ supporting $\boldsymbol{c}^{**}(t)$ with aggregate income $I(t)$. Let $p_h^{CPI}(\boldsymbol{p}^{**}(t))$ be the price index associated with the homothetic aggregator $\tilde{u}_h(\boldsymbol{c}_h)$. We know that

$$\sum_h \tilde{u}_h(\boldsymbol{c}_h^{**}) p_h^{CPI}(\boldsymbol{p}^{**}) = I(t),$$

because preferences are locally nonsatiated. Exact compensation implies that

$$\tilde{u}_h\left(\boldsymbol{c}_h^{**}\right) = \tilde{u}_{h'}\left(\boldsymbol{c}_{h'}^{**}\right).$$

Substituting this into the previous equation implies that

$$\tilde{u}_h\left(\boldsymbol{c}_h^{**}\right) = \frac{I(t)}{\sum_h p_h^{CPI}(\boldsymbol{p}^{**}(t))}.$$

Consider the choices of the compensated agent given the same price vector $\boldsymbol{p}^{**}$ and total income $I(t)$. Since $\boldsymbol{p}^{**}$ is a competitive equilibrium with transfers, $p_h^{CPI}(p^{**}) > 0$, which implies that the compensated representative agent chooses:

$$\tilde{u}_h\left(\boldsymbol{c}_h^{\text{comp}}\right) = \tilde{u}_{h'}\left(\boldsymbol{c}_{h'}^{\text{comp}}\right).$$

Since the compensated representative agent exhausts all income, we have

$$\tilde{u}_h\left(\boldsymbol{c}_h^{\text{comp}}\right) = \frac{I(t)}{\sum_h p_h^{\text{CPI}}(\boldsymbol{p}^{**}(t))}.$$

Hence, the consumption allocations coincide and

$$A(t) = U(\boldsymbol{c}^{**}(t)) = U(\boldsymbol{c}^{comp}(t)).$$

This proves the first equality. Since the compensated representative agent has a homogeneous of degree 1 utility function, the remaining equalities follow as a consequence of standard results for representative agent economies. □

Proof of Lemma 1. Consider the decentralized equilibrium at the status-quo. We wish to show that the same prices and allocations are a compensated equilibrium. This simply requires showing that given decentralized price vector $\boldsymbol{p}(0)$ and aggregate income $I(0)$, the compensated representative agent makes the same consumption choices. This follows from the fact that (i) $\boldsymbol{c}(0)$ is affordable to the compensated representative agent; (ii) any other consumption choice that dominates $\boldsymbol{c}(0)$ is unaffordable to the compensated representative agent.

(ii) follows from the fact that if $U(\boldsymbol{c}') > U(\boldsymbol{c}(0))$, then it must be that $\boldsymbol{c}_h' \succ_h \boldsymbol{c}_h(0)$ for every h. The latter fact implies that $\boldsymbol{p}(0) \cdot \boldsymbol{c}_h' > \boldsymbol{p}(0) \cdot \boldsymbol{c}_h(0)$ for every h (otherwise $\boldsymbol{c}(0)$ would not be a decentralized equilibrium). Summing this across all h guarantees that $\boldsymbol{c}'$ is unaffordable because $\boldsymbol{p}(0) \cdot \sum_h \boldsymbol{c}_h' > \boldsymbol{p}(0) \cdot \sum_h \boldsymbol{c}_h = I(0)$. □

Proof of Proposition 1. By Theorem 2, we know that

$$A(t) = Y^{comp}(t).$$

If preferences are identical, homothetic, and all households face the same relative prices, then the distribution of spending across households has no effect on equilibrium relative prices. Hence, $\boldsymbol{p}^{comp}(t) = \boldsymbol{p}(t)$. Given these prices, and homotheticity of preferences, we

also know that

$$\sum_h \boldsymbol{c}^{comp}(t) = \sum_h \boldsymbol{c}(t).$$

From this, it follows that

$$A(t) = Y^{comp}(t) = Y(t).$$

Since there is a positive representative agent with homothetic preferences, it follows from standard results that

$$Y(t) = A^{RA}(t).$$

Finally, letting $u(\boldsymbol{c})$ be the homogeneous of degree one representation of the utility function of every agent (since all agents have the same preferences), we have that

$$\begin{aligned} A^{KH}(t) &= \frac{\sum_h e_h(\boldsymbol{p}(t), u_h(t))}{\sum_h e_h(\boldsymbol{p}(t), u_h(0))} = \frac{\sum_h u_h(t)}{\sum_h u_h(0)} = \frac{\sum_h u(\boldsymbol{c}_h(t))}{\sum_h u(\boldsymbol{c}_h(0))}, \\ &= \frac{\sum_h u(\boldsymbol{c}(t)\chi_h(t))}{\sum_h u(\boldsymbol{c}(0)\chi_h(0))}, \end{aligned}$$

where $\chi_h(t)$ is household h's share of aggregate expenditures at t,

$$\begin{aligned} &= \frac{\sum_h \chi_h(t) u(\boldsymbol{c}(t))}{\sum_h \chi_h(0) u(\boldsymbol{c}(0))}, \\ &= \frac{u(\boldsymbol{c}(t))}{u(\boldsymbol{c}(0))} = A^{RA}(t). \end{aligned}$$

□

Proof of Proposition 2. This is restatement of Hulten (1978). □

Proof of Proposition 3. This follows from combining Theorem 2 with Proposition 2. □

Proof of Proposition 4. This follows from combining Theorem 2 with Proposition 3 from Baqaee and Farhi (2019b). □

Proof of Proposition 5. Kaldor-Hicks efficiency can be defined using the aggregate consumption-equivalent variation where the set we shift is the aggregate budget set. Specifically,

$$A^{KH}(t) = \max \left\{ \phi \in \mathbb{R} : \text{there is } \boldsymbol{c} \in \phi^{-1}\mathcal{B}(\boldsymbol{p}(t), I(t)) \text{ and } u_h(\boldsymbol{c}_h) \geq u_h(\boldsymbol{c}_h^0) \text{ for every } h \right\},$$

where $\mathcal{B}(\boldsymbol{p}(t), I(t)) = \{\boldsymbol{c} : \boldsymbol{p}(t) \cdot \sum_h \boldsymbol{c_h} \leq I(t)\}$ and $I(t) = \sum_h I_h(t)$. In the absence of distortions ($\boldsymbol{\mu} = 1$), $\boldsymbol{p}(t)$ and $I(t)$ are prices and income in the competitive equilibrium.

We first show that $\mathcal{C}(\boldsymbol{z}(t))$ is contained in the aggregate budget set $\mathcal{B}(\boldsymbol{p}(t), I(t))$. Suppose that there is a feasible consumption allocation $\boldsymbol{c}' \in \mathcal{C}(\boldsymbol{z}(t))$ that violates the aggregate budget constraint at equilibrium prices. That is, $\boldsymbol{p}(t) \cdot \boldsymbol{c}' > I(t)$, where $I(t) = \boldsymbol{w}(t) \cdot \boldsymbol{L} + \Pi(t)$, and $\Pi(t)$ denotes aggregate profits which are equal to zero in equilibrium. Hence, $\Pi(t) < \boldsymbol{p}(t) \cdot \boldsymbol{c}' - \boldsymbol{w}(t) \cdot \boldsymbol{L}$. Aggregate profits Π' under any feasible allocation $\boldsymbol{c}'$ are given by $\Pi' = \boldsymbol{p}(t) \cdot \boldsymbol{c}' - \boldsymbol{w}(t) \cdot \boldsymbol{L}$. By the inequality above, $\Pi(t) < \Pi'$. This is a contradiction, since aggregate profits are maximized in a competitive equilibrium given prices (see Proposition 5.E.1 in Mas-Colell et al. (1995)).

If $\mathcal{C}' \subseteq \mathcal{C}$ then $A(\boldsymbol{c}^0, \mathcal{C}') \leq A(\boldsymbol{c}^0, \mathcal{C})$. To see this, let c^* and ϕ^* be the solution to $A(\boldsymbol{c}^0, \mathcal{C}')$. That is, $c^* \in \mathcal{C}'$ and $u_h(\boldsymbol{c}_h^*) \geq u_h(\boldsymbol{c}_h^0)$ for every h. Since $c^* \in \phi^* \mathcal{C}$, the statement above follows. Hence, because the Pareto frontier $\mathcal{C}(\boldsymbol{z}(t))$ is contained in the aggregate budget set $\mathcal{B}(\boldsymbol{p}(t), I(t))$, it follows that $A(t) \leq A^{KH}(t)$. Finally, pure redistributions leave the Pareto frontier unchanged, so $\mathcal{C}(t)$ is unchanged with t and $A(t) = 1$. □

Proof of Proposition 6. This is a consequence of Theorem 2 and Petrin and Levinsohn (2012). □

Proof of Proposition 7. Index wedges by t, and denote the status-quo with wedges by $\boldsymbol{c}^0(t)$. Let $t = 0$ denote the point where $\boldsymbol{\mu}(0) = \mathbf{1}$. The set $\mathcal{C}(0)$ is then the Pareto efficient frontier. From Proposition 6, we have that

$$\log A(\boldsymbol{c}_0(t), C(0)) = -\int_0^t \sum_i \lambda_i^{\text{comp}}(s) \left(1 - \frac{1}{\mu_i(s)}\right) \frac{d \log y_i^{\text{comp}}}{ds} ds.$$

Differentiate the expression above to get

$$\frac{d}{dt}\left[\log A\right] = -\sum_i \lambda_i^{\text{comp}}(t) \left(1 - \frac{1}{\mu_i(t)}\right) \frac{d \log y_i^{\text{comp}}}{dt}.$$

Differentiate a second time to get

$$\frac{d^2}{dt^2}\left[\log A\right] = -\sum_i d\lambda_i^{\text{comp}}(t) \left(1 - \frac{1}{\mu_i(t)}\right) \frac{d \log y_i^{\text{comp}}}{dt} - \sum_i \lambda_i^{\text{comp}}(t) \frac{1}{\mu_i(t)} \frac{d \log \mu_i}{dt} \frac{d \log y_i^{\text{comp}}}{dt} - \sum_i \lambda_i^{\text{comp}}(t) \left(1 - \frac{1}{\mu_i(t)}\right) \frac{d^2 \log y_i^{\text{comp}}}{dt^2}.$$

Evaluate these derivatives at $t = 0$ and write the second-order Taylor approximation:

$$\log A \approx 0 - \frac{1}{2}\sum_i \lambda_i^{\text{comp}}(0)\frac{1}{\mu_i(0)}\frac{d\log\mu_i}{dt}dt\frac{d\log y_i^{\text{comp}}(0)}{dt}dt.$$

To a second-order, this can also be written as

$$\Delta\log A \approx 0 - \frac{1}{2}\sum_i \lambda_i^{\text{comp}}(t)\frac{d\log\mu_i}{dt}dt\frac{d\log y_i^{\text{comp}}(t)}{dt}dt \approx -\frac{1}{2}\sum_i \lambda_i^{\text{comp}}(t)\Delta\log\mu_i\Delta\log y_i^{\text{comp}},$$

since the differences are higher-order. □

Proof of Proposition 8. Let $\boldsymbol{\tau}^*(t)$ and $\boldsymbol{T}^*(t)$ to be the maximizers of (18). Using $\boldsymbol{\tau}^*(t)$, we provide a slightly more general definition of the compensated equilibrium.

Definition 6 (Compensated Equilibrium)**.** A *compensated equilibrium* is the general equilibrium of an economy with the same technologies, resource constraints, wedges, and linear taxes $\boldsymbol{\tau}^*(t)$ but where there is a representative agent with preferences as in Definition 3. For any equilibrium variable $X(t)$, denote the same variable in the compensated equilibrium by $X^{\text{comp}}(t)$.

The following result, which is a consequence of Theorem 1, generalizes Theorem 2 to allow for limited redistribution.

Theorem 3 (Aggregate Efficiency Using Compensated Equilibrium)**.** *If households are exactly compensated, then aggregate efficiency can be calculated using the compensated equilibrium:*

$$A(t) = Y^{comp}(t) = A^{KH,comp}(t) = A^{RA,comp}(t).$$

In words, aggregate efficiency, $A(t)$, can be computed by solving for changes in real output, or welfare of the Hicksian representative agent, in the compensated equilibrium.[34] Once again, this means that tools and results used to calculate welfare in homothetic representative agent economies can be converted into results about aggregate efficiency with heterogeneous and non-homothetic preferences.

The main challenge lies in knowing the necessary taxes $\tau^*(t)$ which the proposition takes as given. However, given these taxes, then the change in every price and quantity in the compensated equilibrium can be calculated as a function of t by applying the results in Baqaee and Farhi (2020).

[34]The fact that Kaldor-Hicks efficiency, $A^{KH,comp}(t)$, in the compensated equilibrium with distortionary taxes τ^* coincides with the rest follows trivially from the fact that the compensated equilibrium has a single agent with homothetic preferences. It is important to note that $A^{KH,comp}(t)$ is not the same as $A^{KH}(t)$.

Proof of Theorem 3. Taking $\tau^*(t)$ as given, we can apply Theorem 2 where $\tau^*(t)$ are treated in the same way as the exogenous wedges $\boldsymbol{\mu}(t)$. □

Combine Theorem 3 with Proposition 4 from Baqaee and Rubbo (2023).

□

Proposition 10 (Equivalence of Kaldor-Hicks and Aggregate Efficiency)**.** *If there is one primary factor, so that relative prices are independent of demand, then* $A(t) = A^{KH}(t)$.

Proof. With one primary factor of production and constant-returns technologies, it is well-known that relative prices do not depend on final demand. Hence, the vector of equilibrium prices and aggregate income in the decentralized (multi-agent) economy $\boldsymbol{p}(t)$ and $I(t)$ are also compensated equilibrium prices and aggregate income. Theorem 2 implies that $A(t) = A^{KH,comp}(t)$. Since relative prices and aggregate income are the same, it follows that $A^{KH,comp}(t) = A^{KH}(t)$. □

Aggregate equivalent variation based on uniform versus non-uniform expansions Our baseline definition of aggregate equivalent variation shrinks or expands the consumption possibility set $\mathcal{C}$ proportionately in every dimension. We now consider define an alternative definition that shrinks or expands the set in a particular direction. We show that the two measures cannot lie on the opposite side of 1.

We begin by defining these two alternative measures. For $t > 0$ and a set $\mathcal{C} \subseteq \mathbb{R}_+^{H\times N}$ define the uniform scaling

$$t^{-1}\mathcal{C} \;\equiv\; \left\{\, t^{-1}\boldsymbol{c} : \boldsymbol{c} \in \mathcal{C} \,\right\}, \quad \text{where } (t^{-1}\boldsymbol{c})_h = t^{-1}\boldsymbol{c}_h \text{ for all } h.$$

Given a subset of goods $\mathcal{M} \subseteq \{1,\ldots,N\}$ and $t > 0$, define the partial scaling

$$t^{-1}_{\mathcal{M}}\mathcal{C} \;\equiv\; \left\{ \boldsymbol{c}' \in \mathbb{R}_+^{H\times N} \;:\; \exists\, \boldsymbol{c} \in \mathcal{C} \text{ such that } c'_{hn} = \begin{cases} t^{-1}c_{hn}, & n \in \mathcal{M}, \\ c_{hn}, & n \notin \mathcal{M} \end{cases} \text{ for all } h, n \right\}.$$

Let

$$\mathcal{A} \;\equiv\; \left\{ a > 0 :\; \exists\, \boldsymbol{c} \in a^{-1}\mathcal{C} \text{ with } \boldsymbol{c}_h \succeq_h \boldsymbol{c}_h^0 \;\forall h \right\}.$$

The uniform CEV measure is $A := \max \mathcal{A}$. Let

$$\mathcal{B} \;\equiv\; \left\{ b > 0 :\; \exists\, \boldsymbol{c} \in b^{-1}_{\mathcal{M}}\mathcal{C} \text{ with } \boldsymbol{c}_h \succeq_h \boldsymbol{c}_h^0 \;\forall h \right\}.$$

The non-uniform CEV measure is $B := \max \mathcal{B}$.

Proposition 11 (Uniform versus non-uniform expansions)**.** *A and B cannot lie on opposite sides of* 1*.*

Proof. If $A > 1$, take $\tilde{\boldsymbol{c}} \in A^{-1}\mathcal{C}$ with $\tilde{\boldsymbol{c}}_h \succeq_h \boldsymbol{c}_h^0$ for all h. Set $\boldsymbol{c} := A\tilde{\boldsymbol{c}} \in \mathcal{C}$. By weak monotonicity, $\boldsymbol{c}_h \succeq_h \boldsymbol{c}_h^0$ for all h, hence $1 \in \mathcal{B}$ and $B \geq 1$ (so $B < 1$ is impossible). Conversely, if $B > 1$, take $\tilde{\boldsymbol{c}} \in B_{\mathcal{M}}^{-1}\mathcal{C}$ with $\tilde{\boldsymbol{c}}_h \succeq_h \boldsymbol{c}_h^0$; undoing the partial scaling and using weak monotonicity yields some $\boldsymbol{c} \in \mathcal{C}$ with $\boldsymbol{c}_h \succeq_h \boldsymbol{c}_h^0$, hence $1 \in \mathcal{A}$ and $A \geq 1$ (so $A < 1$ is impossible). This also implies that if $A < 1$ then $B \leq 1$, and if $B < 1$ then $A \leq 1$. □ □

Appendix D Additional Examples and Derivations

Individual consumption-equivalent for non-homothetic CES. Consider a household with non-homothetic CES preferences, as in Comin et al. (2021),

$$u_h(\boldsymbol{c}_h) = \left(\sum_i (c_{hi})^{\frac{\eta-1}{\eta}} \left(u_h(\boldsymbol{c}_h)\right)^{\xi_i} \right)^{\frac{\eta}{\eta-1}}.$$

where η is the compensated elasticity of substitution and ξ_i controls income effects. Then $\tilde{u}_h(\boldsymbol{c}_h)$ is homothetic CES given by

$$\tilde{u}_h(\boldsymbol{c}_h) = \frac{1}{u_h^0} \left(\sum_i (c_{hi})^{\frac{\eta-1}{\eta}} \left(u_h^0\right)^{\xi_i} \right)^{\frac{\eta}{\eta-1}},$$

where $u_h^0 \equiv u_h(\boldsymbol{c}_h^0)$ is treated as a constant. If ξ_i are the same for every i, then $\tilde{u}_h$ and u_h are both cardinalizations of the same preference rankings.

Derivations in Example 2 The individual consumption-equivalent function associated to

$$u_h(\boldsymbol{c}_h) = c_{hg}^{\alpha} c_{hs}^{1-\alpha},$$

is

$$\tilde{u}_h(\boldsymbol{c}_h) = \frac{c_{hg}^{\alpha} c_{hs}^{1-\alpha}}{\left(c_{hg}^0\right)^{\alpha} \left(c_{hs}^0\right)^{1-\alpha}}.$$

The compensated representative agent assuming interior outcomes sets

$$A = \tilde{u}_h = \tilde{u}_{h'}$$

for all h' subject to

$$\sum_h c_{hg} = z_g, \qquad c_{hs} = z_{hs}.$$

Hence, $\{c_{hg}\}_{h=1}^H$ solves

$$A = \frac{c_{hg}^{\alpha} z_{hs}^{1-\alpha}}{\left(c_{hg}^0\right)^{\alpha}\left(c_{hs}^0\right)^{1-\alpha}} = \frac{c_{h'g}^{\alpha} z_{h's}^{1-\alpha}}{\left(c_{h'g}^0\right)^{\alpha}\left(c_{h's}^0\right)^{1-\alpha}}, \qquad \text{for all } h'$$

subject to $\sum_h c_{hg} = z_g$. The solution is

$$A = \left(\frac{z_g}{\sum_h z_{hs}^{\frac{\alpha-1}{\alpha}} c_{hg}^0 \left(z_{hs}^0\right)^{\frac{1-\alpha}{\alpha}}}\right)^{\alpha}.$$

In status-quo, $c_{hg}^0 = \chi_h^0 z_g^0$,[35] so

$$A = \left(\frac{z_g / z_g^0}{\sum_h \chi_h^0 \left(z_{hs}/z_{hs}^0\right)^{\frac{\alpha-1}{\alpha}}}\right)^{\alpha}.$$

A second-order approximation yields the expression in the text.

Derivations in Example 3 For completeness, we allow for non-homotheticities in this example. The homothetic special case is when $\zeta_h = 0$. The individual consumption-equivalent function associated to the utility function

$$u_h(\boldsymbol{c}_h) = \left[(\alpha_h)^{\frac{1}{\theta_h}} \left(u_h(\boldsymbol{c}_h)\right)^{\zeta_h} c_{hd}^{\frac{\theta_h-1}{\theta_h}} + (1-\alpha_h)^{\frac{1}{\theta_h}} c_{hf}^{\frac{\theta_h-1}{\theta_h}}\right]^{\frac{\theta_h}{\theta_h-1}},$$

is

$$\tilde{u}_h(\boldsymbol{c}_h) = \frac{1}{u_h^0}\left[(\alpha_h)^{\frac{1}{\theta_h}} \left(u_h^0\right)^{\zeta_h} c_{hd}^{\frac{\theta_h-1}{\theta_h}} + (1-\alpha_h)^{\frac{1}{\theta_h}} c_{hf}^{\frac{\theta_h-1}{\theta_h}}\right]^{\frac{\theta_h}{\theta_h-1}},$$

[35] To see that χ_h^0 is also region h's share in total income, note that (under the Cobb-Douglas specification of this example) the first-order condition for c_{hg} and c_{hs} is $p_{hs}^0 z_{hs}^0 = \frac{1-\alpha}{\alpha}\chi_h^0 z_g^0$, where we normalize the price of the tradable good to 1. Therefore, $\chi_h^0 z_g^0 + p_{hs}^0 z_{hs}^0 = \frac{1}{\alpha}\chi_h^0 z_g^0$ and $\frac{\chi_h^0 z_g^0 + p_{hs}^0 z_{hs}^0}{\sum_{h'} \chi_{h'}^0 z_g^0 + p_{h's}^0 z_{h's}^0} = \frac{\chi_h^0}{\sum \chi_{h'}^0} = \chi_h^0$.

In autarky, $c_{hf} = 0$, so

$$\tilde{u}_h\left([c_{hd}, 0]\right) = \frac{1}{u_h^0} c_{hd} \left[\left(\alpha_h\right)^{\frac{1}{\theta_h}} \left(u_h^0\right)^{\zeta_h} \right]^{\frac{\theta_h}{\theta_h - 1}}$$

The domestic expenditure share of household h in the status-quo is

$$s_{hd}^0 \equiv \frac{p_d^0 c_{hd}^0}{p_d^0 c_{hd}^0 + p_f^0 c_{fd}^0} = \frac{p_d^0 c_{hd}^0}{p_h^0 c_h^0} = \frac{\left(\alpha_h\right)^{\frac{1}{\theta_h}} \left(u_h^0\right)^{\zeta_h} \left(c_{hd}^0\right)^{\frac{\theta_h - 1}{\theta_h}}}{\left(u_h^0\right)^{\frac{\theta_h - 1}{\theta_h}}}$$

Solving for $\left[\left(\alpha_h\right)^{\frac{1}{\theta_h}} \left(u_h^0\right)^{\zeta_h} \right]$ and substituting into the individual consumption-equivalent function yields

$$\begin{aligned}
\tilde{u}_h\left([c_{hd}, 0]\right) &= \frac{c_{hd}}{c_{hd}^0} \left(s_{hd}^0\right)^{\frac{\theta_h}{\theta_h - 1}} \\
&= \frac{p_d^0}{p_h^0 c_h^0} \frac{p_h^0 c_h^0}{p_d^0 c_{hd}^0} c_{hd} \left(s_{hd}^0\right)^{\frac{\theta_h}{\theta_h - 1}} \\
&= \frac{p_d^0}{p_h^0 c_h^0} c_{hd} \left(s_{hd}^0\right)^{\frac{1}{\theta_h - 1}} \\
&= \frac{p_d^0 y_d^0}{p_h^0 c_h^0} \frac{c_{hd}}{y_d^0} \left(s_{hd}^0\right)^{\frac{1}{\theta_h - 1}} \\
&= \frac{1}{\chi_h^0} \frac{c_{hd}}{y_d^0} \left(s_{hd}^0\right)^{\frac{1}{\theta_h - 1}},
\end{aligned}$$

where $\chi_h^0 = p_h^0 c_h^0 / p_d^0 y_d^0$ is the share of $h's$ expenditures in total income (assuming balanced trade), and y_d^0 is the aggregate quantity of the home produced good in the status-quo (which is consumed and exported).

The compensated representative agent assuming interior outcomes sets

$$\tilde{u}_h = \tilde{u}_{h'}$$

for all h' subject to

$$\sum c_{hd} = y_d = y_d^0,$$

where y_d is the total output of domestic good in autarky (which is equal to that in status-

quo). Combining, we obtain

$$\Delta \log A = \log \tilde{u}_h = \log \frac{1}{\sum \chi_h^0 \left(s_{hd}^0\right)^{-\frac{1}{\theta_h - 1}}},$$

which is the expression in the text.

Example with skill-biased technical change and costly redistribution. We now consider a simple example with skill-biased technical change that raises the real wage of high-skill workers but lowers the real-wage for low-skill workers. We compare how the response of aggregate efficiency changes depending on the redistributive tools available. Suppose that output (and consumption) are a CES aggregate of the output of manufacturing and services:

$$c = y = \left[\gamma_1^{\frac{1}{\rho}} y_m^{\frac{\rho-1}{\rho}} + (1-\gamma_1)^{\frac{1}{\rho}} y_s^{\frac{\rho-1}{\rho}}\right]^{\frac{\rho}{\rho-1}},$$

where each sector's output is a CES aggregate of low- and high-skill labor

$$y_o = \left[\alpha_o^{\frac{1}{\sigma}} \left(z_{o1} l_{o1}\right)^{\frac{\sigma-1}{\sigma}} + (1-\alpha_o)^{\frac{1}{\sigma}} \left(z_{o2} l_{o2}\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}},$$

where l_{o1} is low- and l_{o2} is high-skill labor. The resource constraints are that

$$\sum_h c_h = c, \qquad \sum l_{o1} = l_1, \qquad \sum l_{o2} = l_2.$$

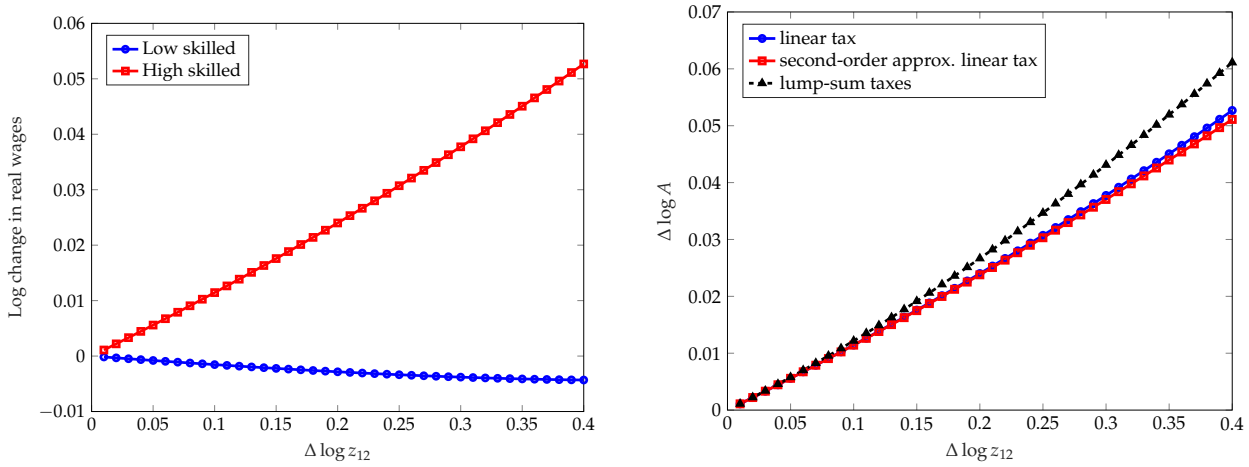
We assume that workers are much more substitutable than sectors: $\rho \ll \sigma$. We also assume that manufacturing is more intensive in low-skill labor use than services.

Consider an increase in automation or the productivity of capital, which we capture via an increase in the productivity of high-skill labor in manufacturing: $\Delta \log z_{m2} > 0$. This is a reduced-form representation for the idea that high-skill labor in manufacturing is equipped by capital, and hence an increase in the quality of capital makes high-skill more productive.[36]

Again, we contrast two scenarios: (1) lump-sum taxation is available, (2) lump-sum transfers must be non-negative and the government can only levy a linear tax on machine use in manufacturing, which we capture as a linear tax, τ, on manufacturing's use of high-skill labor.

[36]For example, high-skill labor and capital are combined in a Leontief nest together called equipped labor, and then equipped labor is substitutable with low-skill labor. We can then think of altering the productivity of equipped labor by varying the productivity of capital.

Figure 7 illustrates the results in a numerical example. Panel 7a shows that skill-biased technical change raises the real wage for high-skill workers and lowers them for low-skill workers in the decentralized equilibrium. The fact that low-skill wages decline means that they need to be compensated via transfers financed by either lump-sum or distortionary taxes. Panel 7b shows the increase in efficiency depending on which taxes are used. As expected, the increase in aggregate efficiency is lower if only distortionary redistributive tools are available. Panel 7b also shows that the second-order approximation is very accurate. In the absence of any redistributive tools whatsoever, aggregate efficiency in this example actually declines because the low-skill workers are worst off and there is no feasible way to compensate them.



(a) real wages

(b) aggregate efficiency

Figure 7: A numerical example of skill-biased technical change. The parameter values are $\rho = 1$, $\sigma = 8$, $\gamma = 0.5$, $\alpha_{m1} = 0.9$, and $\alpha_{s1} = 0.5$. We normalize steady-state quantities so that the CES share parameters are equal to expenditure shares in the status-quo.

Appendix E Non-homothetic Preferences

Our results apply to non-homothetic preferences without change, with one subtlety, which arises even with a single agent. In general equilibrium and when preferences are non-homothetic, scaling the productivity of factors may no longer be equivalent to scaling the set of equilibrium consumption allocations with transfers.

To see this, let $\mathcal{C}(t, Z)$ be the set of consumption allocations that can be supported as equilibria given transfers and a total factor-augmenting productivity shifter Z. If preferences are homothetic, scaling the consumption possibility set by Z coincides with scaling

factor efficiency by Z:

$$Z\mathcal{C}(t,1) = \mathcal{C}(t,Z).$$

However, with non-homothetic preferences, if we scale factor endowments, households may want to change the mix of goods they consume. This means that the above equality will not hold. Hence, in such cases, a factor-scaling definition of $A(t)$:

$$A(t) = \max\left\{Z : \boldsymbol{c} \in \mathcal{C}(t, Z^{-1}) \text{ with } \boldsymbol{c}_h \succeq_h \boldsymbol{c}_h^0 \text{ for every } h\right\}.$$

would not coincide with a consumption-set scaling definition:

$$A(t) = \max\left\{Z : \boldsymbol{c} \in Z^{-1}\mathcal{C}(t,1) \text{ with } \boldsymbol{c}_h \succeq_h \boldsymbol{c}_h^0 \text{ for every } h\right\},$$

because the former allows households to reoptimize their consumption bundles. This distinction arises even with a single agent — scaling the equilibrium consumption bundle in a single agent economy is not the same as scaling the factor productivities, if that agent has non-homothetic preferences. All of our theoretical results go through unchanged, if we use the factor-scaling definition, even when preferences are non-homothetic.

We make two further remarks.

i. To make the two definitions give the same answer with non-homothetic preferences, one could define $\tilde{\mathcal{C}}(t)$ to be the set of equilibrium consumption allocations given transfers and household preferences $\tilde{u}_h(\boldsymbol{c}_h)$ instead of $u_h(\boldsymbol{c}_h)$. In this case, scaling the consumption possibility set $\tilde{\mathcal{C}}(t)$ would give the same value of $A(t)$ as scaling the factor productivity.

ii. One special but important case is when $\mathcal{C}(t,Z)$ is defined to be the set of all technologically feasible allocations (e.g. relevant when studying Pareto-efficient counterfactuals, like the distance to the frontier or productivity shocks in efficient economies). So, instead of $\mathcal{C}(t,Z)$ being the set of equilibrium consumption allocations given some transfers and factor productivity Z, we instead define $\mathcal{C}(t,Z)$ to be the set of all technologically feasible consumption allocations. In this case, we recover

$$Z\mathcal{C}(t,1) = \mathcal{C}(t,Z)$$

even if preferences are non-homothetic, and so the definition of aggregate efficiency does not depend on whether we scale the consumption set or primary factor efficiency.

Appendix F Explicit Characterization of Compensated Equilibrium

Theorem 2 and Theorem 3 (in Appendix C) show that calculating changes in aggregate efficiency can be boiled down to solving for the compensated equilibrium. This section provides some formulas for calculating variables in the compensated equilibrium. To do so, we rely on the differential hat algebra approach in Baqaee and Farhi (2020), which characterizes equilibria of representative agent economies with wedges using differential equations. Alternatively, one could also use exact-hat algebra methods, as in Dekle et al. (2008).

For concreteness, assume that all production and utility functions are nested-CES. (Non-CES economies can be analyzed in a similar way following the non-CES extensions in Baqaee and Farhi (2019b)). To make the notation more compact, represent the economy in such a way that each producer, i, is associated with a single elasticity of substitution θ_i (by treating each sub-nest as a separate producer).

F.1 Input-Output Notation

Stack the expenditure shares of the representative household, all producers, and all factors into the $(H+N+F)\times(H+N+F)$ input-output matrix Ω. The first H rows correspond to the households consumption baskets. The next N rows correspond to the expenditure of each producer on every other producer and factor as a share of its sales (where the sales price is always inclusive of the wedge and tax). The last F rows correspond to the expenditure shares of the primary factors (which are all zeros, since primary factors do not require any inputs). With some abuse of notation, the heterogeneous agent input-output matrix can be written as

$$\Omega = \left[\begin{array}{ccc|ccc|ccc} 0 & \cdots & 0 & b_{11} & \cdots & b_{1N} & 0 & \cdots & 0 \\ \vdots & \cdots & \vdots & & \cdots & & & \cdots & \\ 0 & \cdots & 0 & b_{H1} & \cdots & b_{HN} & 0 & \cdots & 0 \\ \hline 0 & \cdots & 0 & \Omega_{11} & \cdots & \Omega_{1N} & \Omega_{1N+1} & \cdots & \Omega_{1N+F} \\ \vdots & \cdots & \vdots & & \ddots & & & & \\ 0 & \cdots & 0 & \Omega_{N1} & & \Omega_{NN} & \Omega_{NN+1} & \cdots & \Omega_{NN+F} \\ \hline 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \\ \vdots & \cdots & \vdots & \vdots & \cdots & \vdots & \vdots & \cdots & \vdots \\ 0 & \cdots & 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \end{array}\right].$$

Note that our convention is that rows (not columns) record costs relative to revenues inclusive of wedges and taxes. If wedges and taxes are greater than one, then the rows of this matrix will generally sum to a number less than one. The Leontief inverse matrix is the $(H+N+F)\times(H+N+F)$ matrix defined as

$$\Psi \equiv (I-\Omega)^{-1} = I + \Omega + \Omega^2 + \ldots,$$

where I is the identity matrix. The Leontief inverse matrix $\Psi \geq I$ records the *direct and indirect* exposures through the supply chains in the production network.

Denote the distribution of expenditures by each household by χ, which is an $(H+N+F)\times 1$ vector. The first H elements are equal to each household's share of aggregate consumption expenditures, and the remaining $N+F$ elements are all zeros. As a matter of accounting identities, the vector of Domar weights satisfies:

$$\lambda' = \chi'\Psi.$$

In this equation λ is a $(H+N+F)\times 1$ vector. The first H elements are equal the expenditures of each household relative to aggregate consumption expenditures, χ', the next $N+F$ elements are equal to the sales of each good and factor relative to aggregate consumption expenditures.

Let μ and τ denote the diagonal matrices whose iith element is equal to μ_i and τ_i respectively. Recall that μ are exogenous wedges, whereas τ are linear taxes that can be used for redistribution. Define the cost-based Leontief inverse to be

$$\tilde{\Psi} = (I-(\tau\mu)\Omega)^{-1}.$$

Note that the cost-based Leontief inverse coincides with Ψ in the absence of wedges. Intuitively, $\tilde{\Psi}$ is a version of the Leontief inverse that calculates exposures of i to j in terms of cost shares rather than revenue shares (revenues exceed costs if wedges and taxes are greater than one).

For any non-negative vector a, define

$$Cov_a(b,c) = \mathbb{E}_a[bc] - \mathbb{E}_a[b]\mathbb{E}_a[c] = \sum_i \frac{a_i}{\sum_{i'} a_{i'}} b_i c_i - \sum_i \frac{a_i}{\sum_{i'} a_{i'}} b_i \sum_i \frac{a_i}{\sum_{i'} a_{i'}} c_i,$$

where $\mathbb{E}_a[\cdot]$ denotes averages of vectors weighted by the elements of a. For any matrix X, denote its ith row and column by $X_{(i,:)}$ and $X_{(:,i)}$.

F.2 Differential Hat-Algebra

The next proposition characterizes compensated variables in terms of initial expenditure shares, wedges, and shocks.

Proposition 12 (Differential Equations for Compensated Equlibrium)**.** *Let aggregate spending be the numeraire. Then, assuming that households can be exactly compensated, the compensated equilibrium satisfies the following system of differential equations. For each $i \in H + N + F$, the compensated price satisfies*

$$d \log p_i^{comp} = \sum_j \tilde{\Psi}_{ij}^{comp} [d \log \mu_j \tau_j^* - d \log z_j] + \sum_{f \in F} \tilde{\Psi}_{if}^{comp} d \log \lambda_f^{comp}. \tag{20}$$

Compensated Domar weights for goods and factors satisfy

$$d\lambda_l^{comp} = \sum_j \lambda_j^{comp} (1 - \theta_j) \mu_j^{-1} Cov_{\Omega_{(j,:)}^{comp}} \left(d \log p^{comp}, \Psi_{(:,l)}^{comp} \right) + Cov_{\chi^{comp}} \left(d \log \chi^{comp}, \Psi_{(:,l)}^{comp} \right)$$
$$- \sum_j \lambda_j \left(\Psi_{jl} - \mathbf{1}[j = l] \right) d \log \mu_j \tau_j^*. \tag{21}$$

Changes in compensated expenditure shares for household h satisfy

$$d \log \chi_h^{comp} = d \log p_h^{comp}, \tag{22}$$

where $d \log p_h^{comp}$ is the price of the consumption bundle for household h. The compensated input-output matrix satisfies

$$d\Omega_{ij}^{comp} = (1 - \theta_i) \left(d \log p_j^{comp} - \mathbb{E}_{\Omega_{(i,:)}^{comp}} [d \log p^{comp}] \right) - d \log \mu_i. \tag{23}$$

Finally, $d \log y_i^{comp}$ is given by $d \log \lambda_i^{comp} - d \log p_i^{comp}$. The initial conditions are given by Lemma 1 that all prices and expenditures are equal to the ones in the competitive equilibrium for $t = 0$.

Equation (20), (21), and (23) are standard and identical to expressions in Baqaee and Farhi (2020). They are loglinearizations of marginal cost-functions, market clearing conditions, and demand curves respectively. The key equation, which distinguishes the compensated equilibrium from the decentralized equilibrium is (22). Whereas in the decentralized equilibrium changes in household expenditures are determined by changes in the income of each household, in the compensated equilibrium, they are determined by the choices of the compensated agent (who tries to equate individual consumption-

equivalents across agents). The term $d \log p_h^{\text{comp}}$, which is pinned down by (20), is the change in the compensated price index of household h.

The taxes $\boldsymbol{\tau}^*(t)$ are given by the maximizers of the problem in (18). If only lump-sum transfers are used for redistribution, as in Sections 4 and 5, then $\boldsymbol{\tau}^*(t) = 0$, and Proposition 12 fully characterizes the compensated equilibrium in terms of exogenous parameters: $\boldsymbol{z}(t)$ and $\boldsymbol{\mu}(T)$. If lump-sum transfers are unavailable, then solving for $\tau^*(t)$ requires specifying more details about the set of available tax instruments. Specifically, we would need to add the log-linearized first-order conditions for the tax instruments from (18) as additional equations in Proposition 12 to pin down how τ^* evolves.

There is one case where this optimization problem can be avoided. If there are only $H-1$ taxes available, and outcomes are interior, then (22) can pin down $\tau^*(t)$. For example, suppose that there are $H-1$ taxes, and the share of revenues from the ith tax sent to household h are given by α_{ih}:

$$T_h(t) = \sum_i \alpha_{ih} \left(1 - \frac{1}{\tau_i^*(t)}\right) \lambda_i(t).$$

Log-differentiating household h's budget constraint gives:

$$d \log \chi_h^{\text{comp}} = \sum_f \frac{\omega_{hf} \lambda_f^{\text{comp}}}{\chi_h^{\text{comp}}} d \log \lambda_f^{\text{comp}} + \frac{dT_h}{\chi_h^{\text{comp}}},$$

differentiating $T_h(t)$ above, and substituting it into the log-linearized budget constraint gives $H-1$ additional equations which, assuming regularity conditions, will pin down $d \log \boldsymbol{\tau}^*$.

Generally, solving the system of linear equations in Proposition 12 requires inverting a system of equations. When there is a single primary factor of production and we evaluate these derivatives at a perfectly competitive point, then the change in efficiency can be solved out easily up to a second-order.

Proposition 13 (Aggregate Efficiency with One Factor)**.** *Consider a competitive economy with a single primary factor of production. The change in aggregate efficiency in response to a vector of productivity shocks,* $\Delta \log \boldsymbol{z}$ *and changes in wedges* $\Delta \log \boldsymbol{\mu}$ *is*

$$\begin{aligned}\Delta \log A \approx &\sum_i \lambda_i \Delta \log z_i + \frac{1}{2} \sum_{i \in N+H} \lambda_i (\theta_i - 1) Var_{\Omega(i,:)} \left(\sum_k \Psi_{(:,k)} \log z_k \right) \\ &- \frac{1}{2} \sum_{i \in N+H} \lambda_i \theta_i Var_{\Omega(i,:)} \left(\sum_k \Psi_{(:,k)} \Delta \log(\mu_k \tau_k^*) \right).\end{aligned}$$

to a second-order approximation in $\Delta \log z$ *and* $\Delta \log \mu$*.*

There are three summands. The first one is just Hulten's theorem. The second summand is a nonlinear adjustment due to changes in Domar weights. The second summand is also equal to: $1/2 \sum_k \left[\sum_j \partial \lambda_k^{\text{comp}} / \partial \log z_j \Delta \log z_j \right] \Delta \log z_k$. If the compensated Domar weight for k rises due to productivity shocks, then the shock to k is more important. This happens if exposure to k is heterogeneous, captured by the variance term, and if elasticities of substitution, θ_i, are far from unity. The final summand are the Harbeger triangles caused by the taxes and wedges. The triangles are larger the higher are elasticities of substitution, θ_i, and the more heterogeneous are exposures to the taxes and wedges, captured by the variance terms.