NBER WORKING PAPER SERIES

WELFARE AND OUTPUT WITH INCOME EFFECTS AND DEMAND INSTABILITY

David Baqaee
Ariel Burstein

Working Paper 28754
http://www.nber.org/papers/w28754

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2021, Revised May 2021

We thank Sihwan Yang for superb research assistance. We thank Andy Atkeson, Natalie Bau, and Jon Vogel for helpful comments. We are grateful to Emmanuel Farhi and Seamus Hogan, both of whom passed away tragically before this paper was written, for their insights and earlier conversations on these topics. This paper received support from NSF grant No. 1947611. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Welfare and Output with Income Effects and Demand Instability
David Baqaee and Ariel Burstein
NBER Working Paper No. 28754
May 2021, Revised May 2021
JEL No. E0,E01,E1,E21,E23,E3

ABSTRACT

We provide a non-parametric characterization of how welfare responds to changes in budget and production possibility sets when preferences are non-homothetic or subject to shocks, in both partial and general equilibrium. We generalize Hulten's theorem, which is the basis for constructing aggregate quantities, to this context. We identify a new bias in measures of real consumption. This bias depends on the covariance of price changes and expenditure changes due to income effects or preference shocks. We apply our results to long-run and short-run phenomena. In the long-run, we show that structural transformation, if caused by income effects, is roughly twice as important for welfare than what is implied by standard measures of Baumol's cost disease. In the short-run, we show that when firms' demand shocks are correlated with their supply shocks, industry-level price and output indices are biased, and this bias does not disappear in the aggregate. Finally, we show that correlated supply and demand shifters make real GDP and aggregate TFP unreliable metrics for measuring production and productivity, and illustrate this using the Covid-19 crisis.

David Baqaee
Department of Economics
University of California at Los Angeles
Bunche Hall
Los Angeles, CA 90095
and CEPR
and also NBER
baqaee@econ.ucla.edu

Ariel Burstein
Department of Economics
Bunche Hall 8365
Box 951477
UCLA
Los Angeles, CA 90095-1477
and NBER
arielb@econ.ucla.edu

A data appendix is available at http://www.nber.org/data-appendix/w28754

1 Introduction

How does a change in the economic environment affect welfare? For example, how does the welfare of a consumer change when the budget constraint changes, or how does the welfare of a nation change when the production possibility frontier of the economy changes? At first blush, answering this question seems very difficult, perhaps requiring detailed information about the nonlinear functions describing preferences and technologies.

Under the strong assumptions of homotheticity and preference stability, standard theory offers a simple non-parametric formula to answer these questions. Consider a change in prices and income that occurred (or will occur) over some time horizon t_0 to t_1. Indexing individual goods by i, the change in welfare comparing the allocation in t_0 to t_1 is approximately equal to

$$\Delta \text{Welfare} \approx \Delta \log I - \sum_{t=t_0}^{t_1} \sum_i b_i(t) \left(\log p_i(t+1) - \log p_i(t) \right), \tag{1}$$

where $\Delta \log I$ is the log change in nominal income, and $b_i(t)$ and $p_i(t)$ are the budget share and price of good i at time t. In words, the change in nominal income deflated by the expenditure-share weighted change in prices approximates the change in welfare. The fact that the expenditure shares are updated at every period t between t_0 and t_1 reduces substitution bias, and eliminates it in the continuous-time limit. The general equilibrium version of this equation is Hulten's theorem, which states that, in a closed and efficient economy, changes in aggregate welfare are approximately equal to sum of microeconomic productivity shocks weighted by sales over GDP.

Variations on equation (1), called *chain-linked* indices, are foundational to macroeconomics. Such indices are used to calculate most types of real economic activity and price deflators, ranging from aggregates like output (real GDP), total factor productivity (TFP), private consumption and investment, to less aggregated objects like industry-level measures of production and inflation. The fact that these indices approximate changes in welfare and production under homotheticity and stability justifies their recommended use in the United Nations' System of National Accounts.[1]

While homotheticity and preference stability are highly convenient assumptions, they have counterfactual implications: homotheticity requires that income effects be uniform, that is, the income elasticity of demand must equal one for every good; stability requires

[1] See OECD et al. (2004) and references therein, in particular Chapters 15 and 17, for a comprehensive overview and discussion of price and quantity indices and their relation to welfare.

that consumers only change their spending decisions in response to changes in incomes and relative prices. In this paper, we provide sufficient statistics that adjust equation (1) when these assumptions are relaxed.[2] Our baseline welfare measure is the equivalent variation at fixed final preferences, which answers the question: *"holding fixed preferences, how much must the consumers' initial endowment change to make them indifferent between their choice sets at t_0 and t_1?"*

We study this problem in both partial equilibrium, where choice sets are defined in terms of budget sets (prices and income are exogenous), and in general equilibrium, where choice sets are defined in terms of production possibility frontiers (prices and income are endogenous).[3] We provide exact and approximate characterizations of the adjustment to equation (1) in partial equilibrium and a welfare-counterpart to Hulten's theorem in general equilibrium when preferences are non-homothetic and or unstable.

The generalized version of equation (1) and Hulten's theorem use budget shares and sales shares calculated under the final indifference curve, rather than actual shares observed along the transition path. Hence, any adjustment to these formulas is caused by the fact that they *undercount* the substitution caused by either income effects or preference instability. That is, with non-homotheticities or preference instability, chained indices, like real consumption, real GDP, or TFP, suffer from exactly the type of substitution bias that they were designed to eliminate. We show that this bias is larger if changes in expenditures caused by income effects or taste shocks are correlated with changes in prices. If, on the other hand, demand shifters are uncorrelated with price changes (or, in general equilibrium, if demand shifters are orthogonal to supply shifters), then no adjustment is required.

For example, suppose that over the post-war period, services grew relative to manufacturing due to aging or the fact that services are luxuries compared to manufactured goods. Then measured real consumption does not correctly account for changes in market shares over this period. Intuitively, if expenditure shares change due to changes in demand, then when we compare the past to the present, we must use demand curves that are relevant for choices today, not the ones that were relevant in the past. For example, if richer or older households prefer to spend their income on healthcare, and households in 2021 are richer and older than they were in 1950, then when we compare the economy's

[2] As we discuss in detail in Section 2, preference instability is driven by any factor that changes preference rankings over bundles of goods at fixed prices and income, e.g. age, health, advertising, fads. In the literature, preference instability and non-homotheticities are typically studied independently. We analyze them jointly in this paper because both generate the same type of biases in measures of real consumption. Our results are relevant when either of these forces is active.

[3] For the macro problem, we consider neoclassical economies with representative agents.

productive capacity in 1950 to 2021, we must account for the fact that the richer and older households of today demand a different set of goods than the younger and poorer households of 1950. The gap between real consumption and welfare is large if spending on healthcare increases due to aging or income effects and the relative price of healthcare changes relative to that of other sectors.

Similarly, during the Covid-19 crisis, household demand changed in favor of low-contact goods. At the same time, the market price of these goods rose relative to high-contact goods. In this case, chained real consumption and inflation fail to correctly account for changing expenditure shares. Intuitively, when we construct the basket of goods to measure how the cost of living has changed for a consumer who values low-contact goods, we must account for the fact that this consumer would have consumed more of these goods in the earlier periods. That is, compared to a chained index, we must place a higher weight on changes in expenditure shares that occurred due to the demand shocks.

Our partial equilibrium welfare measure answers a microeconomic question for an infinitesimal agent who cannot alter market-level prices through her choices. Our general equilibrium welfare measure answers a macroeconomic question for a collection of agents whose collective decisions alter market-level prices. When preferences are homothetic and stable, macroeconomic changes in welfare are equal to microeconomic changes in welfare. However, we show that these two measures are not equal when household preferences are non-homothetic or unstable, if production possibility sets are nonlinear. Intuitively, some points on a budget constraint, which may be feasible for an individual agent, are not feasible for society as a whole due to curvature in the production possibility frontier.

Our results for welfare and the gap between welfare and real consumption are expressed in terms of measurable sufficient statistics. In both partial and general equilibrium, we show that computing the change in welfare caused by changes in prices (in partial equilibrium) or technology (in general equilibrium) does not require direct knowledge of the taste shocks or the shape of non-homotheticities. Instead, what we must know are the expenditure shares and the elasticities of substitution at the final allocation. For the micro problem, these are the household expenditure shares and the elasticities of substitution in consumption.[4] For the macro problem, these are the input-output table and the elasticities of substitution in both production and consumption. These results can be used both for ex-post accounting and ex-ante counterfactuals.

[4]These elasticities of substitution can be derived from Hicksian demand curves. As is well known, see for example Hausman (1981), equivalent and compensating variations (given price and income changes) can be derived by integrating Hicksian demand curves. We allow for unstable preferences, endogenize prices, and provide a characterization of the gap between welfare and measured real consumption.

In general equilibrium, for very simple economies with one factor and no intermediates, the difference between welfare and real GDP is approximately half the covariance of supply and demand shocks. This formula can be generalized to more complex economies. We show how the details of the production structure, like input-output linkages, complementarities in production, and decreasing returns to scale, interact with non-homotheticities and preference shocks to magnify the gap between welfare and GDP. There is no reason to expect the discrepancies between welfare and GDP that we emphasize to get aggregated away. In fact, these discrepancies can become larger the more we disaggregate. In this sense, our results are related to the literature studying the macroeconomic implications of production networks and disaggregation (e.g. Gabaix, 2011; Acemoglu et al., 2012; Baqaee and Farhi, 2019c).

We illustrate the relevance of our results for understanding short-run and long-run phenomena by means of three applications. In our first application, we analyze the importance of non-homotheticity or preference instability for measures of long-run productivity growth. Since Baumol (1967), an enduring stylized fact about economic growth has been the observation that industries with slow productivity growth tend to become larger as a share of the economy over time. This phenomenon, known as Baumol's cost disease, implies that aggregate growth is increasingly determined by productivity growth in slow-growth industries since, over time, the industrial mix of the economy shifts to favor these industries. To be specific, from 1948 to 2014, aggregate TFP in the US grew by 60%. If the US economy had kept its original 1948 industrial structure, then aggregate TFP would have grown by 78% instead. We show that if structural transformation is caused solely by income effects and demand instability, then welfare-relevant TFP grew by only 47%. This is because measured aggregate TFP does not fully account for substitution caused by changes in demand, and hence the increase in the welfare-relevant measure of aggregate TFP is much lower than what is measured.

In our second application, we consider a firm-level specification of our model. We show that when firms' demand shocks are correlated with their supply shocks, there is a gap between welfare-relevant and measured changes in industry-level output and prices. We show that these biases, which can be sizable even at annual frequency, do not disappear as we aggregate up to the level of real GDP even if firms and industries are infinitesimal. At annual frequency, the gap between welfare and real GDP due only to firm-level supply and demand shocks could be as high as 1%, and this gap gets larger at lower frequencies if firm-level supply and demand shocks are persistent, becoming unboundedly large in the limit when the shocks are random walks. If we start with industry-level (rather than firm-level) data, we are ruling out the existence of these biases by construction.[5]

In our final application, we show that if changes in household spending patterns are in part driven by demand shifters, then real GDP and aggregate TFP can become unreliable metrics for measuring changes in production and productivity. What should be irrelevant details, like the order of supply and demand shocks, can cause real GDP to be different between the initial and final periods even if initial and final prices and quantities are the same. To illustrate these results, we consider the large changes in household spending patterns towards low-contact goods and services over the first few months of 2020 due to the Covid-19 pandemic. Since these changes in spending patterns were not entirely driven by market prices, real GDP is path dependent. For example, if the path of supply and demand shocks during the recovery does not look exactly like the initial collapse in reverse, measured real GDP can be as much as 6% higher or lower than what it was before the crisis, even if every price and quantity returns to its pre-Covid value.

Of course, there are other reasons, besides instability and non-homotheticity, why (1) can fail to accurately measure welfare. Many of the well-known reasons why the approximation in (1) fails can be thought of as being due to missing prices and quantities. For example, it is well-known that (1) fails to properly account for the creation and destruction of goods if we cannot measure the quantity of goods continuously as their price falls from or goes to their choke price (Hicks, 1940; Feenstra, 1994; Hausman, 1996; Aghion et al., 2019); equation (1) also fails to properly account for changes in the quality of goods (see Syverson, 2017); finally, (1) fails to properly account for changes in non-market components of welfare, like changes in leisure and mortality (see Jones and Klenow, 2016), or changes in the user cost of durables. In all of these cases, the problem is that some of the relevant prices or quantities in the consumption bundle are missing or mismeasured, and correcting the index involves imputing a value for these missing prices or quantities. In this paper, we abstract from these issues and assume that prices and quantities have been correctly measured. If prices and quantities are mismeasured or missing, then our results would apply to the quality-adjusted, corrected, version of prices instead of observed prices. That is, the corrections we derive are different to the ones that are equivalent to adjustments in prices.[6]

[5]These aggregation biases are not unique to firm-level data, a similar logic also applies to the use of sectoral aggregates in place of more disaggregated industry-level data, whereby sectoral measures of TFP are contaminated by substitution bias caused by demand instability and income effects.

[6]Our approach to calculate ex-post welfare changes requires well-measured price changes and elasticities of substitution in the final period. For ex-post welfare measurement, when information on prices is missing or mismeasured, if preferences are non-homothetic an alternative approach is to infer changes in welfare by relying on changes in prices, expenditures, price elasticities, and Engel curve slopes for only a subset of goods, given assumptions on separability and stability in preferences (see e.g. Hamilton, 2001

Relatedly, preference instability and mismeasured prices (i.e. unobserved quality change) are sometimes viewed as alternative means to the same end. This is because they can both be used to justify why demand curves shift over time, even holding prices and incomes fixed. That is, both can rationalize changes in behavior that are not triggered by changes in observed prices or incomes. However, while they have similar implications for changes in prices and quantities, they have very different implications for welfare. When there are unobserved changes in quality, the gap between welfare and real consumption is caused by a difference between measured and welfare-relevant *prices*. We show that in the case of non-homotheticities and taste shocks, the gap between welfare and real consumption is caused by a difference between measured and welfare-relevant *expenditure shares*.

Other related literature. This paper contributes to the literatures on growth and productivity accounting, multi-sectoral and disaggregated macroeconomics, as well as the literature on structural transformation. We discuss the way our paper complements and relates to these literatures in turn.

A key assumption in growth accounting is the existence of a stable and homothetic final aggregator. As shown by, for example Hulten (1973) among others, chain-linked indices are meaningful if, and only if, a homothetic and stable final aggregator exists. Therefore, this assumption is ubiquitous in growth accounting, and also appears in almost all papers that study aggregate outcomes using disaggregated input-output models.[7] We show that when preferences are non-homothetic or unstable, then the "ideal" price deflator is shock-dependent. This allows us to provide a generalization of Domar (1961) and Hulten (1978) that measures changes in welfare in situations when preferences are unstable or non-homothetic. Using this, we can construct exact and approximate characterizations of how welfare responds to shocks in general equilibrium, a question which is of central importance in the literature on disaggregated and production network models.[8]

Our approach contrasts with the one in Redding and Weinstein (2020). They show that variations in sales are difficult to explain via shifts in supply curves alone, and shifts in demand curves are an important source of variation in the data. They interpret changes

and, more recently, Atkin et al., 2020). In addition to ex-post measurement with non-homotheticities, in this paper we study the implications of instability of preferences (that generate shifts in expenditures correlated with price changes) and we also consider counterfactuals.

[7]See, for example, the review paper by Carvalho and Tahbaz-Salehi (2018) and the references therein.

[8]The biases we identify, and the failure of Hulten's theorem, are not caused by inefficiencies (e.g. markups, wedges, taxes). Baqaee and Farhi (2019b) analyze how growth accounting must be adjusted in inefficient economies. Whereas incorporating inefficiencies in production does not affect our micro welfare results, how they interact with demand instability and non-homotheticity in general equilibrium is beyond the scope of this paper.

in demand curves as being due to changes in tastes, but unlike us, they treat changes in tastes as being equivalent to changes in price. Operationally, this makes the taste shocks behave like quality shocks. They estimate changes in taste/quality necessary to explain variations in product-level data. However, this only determines changes in the *relative* size of demand shocks across goods, and it does not pin down changes in the overall level of these shocks. Redding and Weinstein (2020) pin down the overall level of the shocks by assuming they are mean zero (see Martin, 2020 for a discussion of this assumption). Our approach is different in that we do not compare utils before and after the taste shocks. Instead we compute changes in equivalent variation keeping preferences over goods constant for the variation, as advocated by Fisher and Shell (1968) and Samuelson and Swamy (1974). This approach does not require any assumptions about the overall level of the taste shocks in terms of utils.[9]

Our paper is also related to the literature on structural transformation and Baumol's cost disease. As explained by Buera and Kaboski (2009) and Herrendorf et al. (2013), this literature advances two microfoundations for structural transformation. The first explanation is all about relative prices differences: if demand curves are not unit-price-elastic, then changes in relative prices change expenditure shares (e.g. Ngai and Pissarides, 2007; Acemoglu and Guerrieri, 2008; Buera et al., 2015). The second explanation emphasizes non-homotheticities, or income effects, whereby households spend more of their income on some goods as they become richer (e.g. Kongsamut et al., 2001; Boppart, 2014; Comin et al., 2015; Alder et al., 2019). Our results suggest that settling this question has important implications for welfare. From a welfare perspective, structural transformation driven by relative price changes should be treated differently to structural transformation driven by non-homotheticity or demand instability. In particular, measures of real production or consumption must be adjusted for substitution bias in the latter case, but no adjustment is necessary in the former case.[10]

[9]Given CES preferences, Martin (2020) estimates using scanner level data large differences in annual price changes between price indices based on fixed initial tastes and final tastes. Other papers studying the relationship between conventional index numbers and welfare in the presence of preference instability include Feenstra and Reinsdorf (2007) who show that the Sato-Vartia index is equal to the CES price index evaluated at some intermediate level of taste shifters, and Caves et al. (1982) who show that when preferences are homothetic, translog, but unstable, measured Tornqvist-based indices correspond to a geometric average of welfare changes under initial and final preferences. A similar result to the latter holds locally in our context, as we show in Appendix B, but in the body of the paper our focus is different. We characterize welfare (in partial and general equilibrium) at either initial or final preferences and using either EV or CV, rather than averaging these different measures.

[10]For welfare analysis with non-homothetic preferences in other contexts such as cross-country real income comparisons and gains from trade, see Feenstra et al. (2009) and Fajgelbaum and Khandelwal (2016).

The structure of the paper is as follows. In Section 2, we set up the microeconomic problem and provide exact and approximate characterizations of the difference between welfare and measured real consumption changes. In Section 3, we set up the macroeconomic general equilibrium model and provide exact and approximate characterizations of the difference between welfare and measured real output changes. Whereas in section 3 we present our macro results in terms of endogenous sufficient statistics, in Section 4 we solve for these endogenous sufficient statistics in terms of microeconomic primitives and consider some simple but instructive analytical examples. Our applications are in Section 5. We discuss some extensions in Section 6 and conclude in Section 7. Proofs are in the appendix.

2 Microeconomic Changes in Welfare and Consumption

In this section, we consider changes in budget constraints in partial equilibrium. We ask how consumers value these changes, and compare these measures of welfare with measures of real consumption. We provide exact and approximate results. We model the equilibrium determination of prices in Section 3.

2.1 Definition of Welfare and Real Consumption

In this subsection we define welfare and real consumption. Measuring changes in welfare using equivalent variation is standard when preferences are stable. However, measuring welfare changes in the presence of unstable preferences is less common and therefore we discuss this issue in some detail.

Consider a set of preference relations, $\{\succeq_x\}$, over bundles of goods. These preferences are indexed by x, which represents anything that affects preference rankings over bundles of goods. For example, x could be calendar time, age, exposure to advertising, or state of nature. For every x, we represent the preference relation $\succeq_x$ by a utility function $u(c;x)$, where $c \in \mathbb{R}^N$ and N is the number of goods in the consumption bundle. Since the consumer makes no choices over x, we do not need to specify how $u(c;x)$ varies with x. Moreover, preferences over x, if they exist, are not revealed by choices.[11]

There are two properties of preferences that are analytically convenient benchmarks throughout the rest of the analysis.

[11] In Section 6, we discuss situations in which x is endogenously chosen and valued by the consumer, such as leisure, but its price and quantity are not being measured. We also discuss situations in which x is endogenously chosen by firms, such as advertising.

Definition 1 (Homotheticity)**.** Preferences over goods c are *homothetic* if, for every positive scalar $a > 0$ and every feasible c and x, we can write

$$u(ac; x) = au(c; x).$$

Definition 2 (Stability)**.** Preferences over goods c are *stable* if there exists a time-invariant function $\Phi(\cdot)$ such that the utility function can be written as $u(c; x) = U(\Phi(c); x)$ for every feasible c and x.

If preferences are stable, x can change over time (e.g. households get higher or lower utils from all goods) but, since x is separable from c, these changes do not impact preferences over bundles of goods c. If preferences are not stable, we say that they are unstable.

Given preferences encapsulated in u, the indirect utility function of the consumer, for any value of x, is

$$v(p, I; x) = \max_c \{u(c; x) : p \cdot c = I\}.$$

where p is a price vector over goods and I is expenditures (which we interchangeably refer to as income).

Consider shifts in the budget set as prices and income change from p_{t_0} and I_{t_0} to p_{t_1} and I_{t_1}. Here, t_0 and t_1 simply index the vector of prices and income being compared. Motivated by our applications, we refer to this index as time.[12] This change in the budget set is accompanied by changes in x from x_{t_0} to x_{t_1}.

Since utility is only defined up to monotone transformations, changes in utility do not have meaningful units. When prices are exogenous, we measure changes in utility using corresponding changes in income. Our baseline measure of microeconomic welfare is defined as follows.

Definition 3 (Micro Welfare)**.** The change in welfare measured using the *micro equivalent variation* with *final preferences* is $EV^m(p_{t_0}, I_{t_0}, p_{t_1}, I_{t_1}; x_{t_1}) = \phi$ where ϕ solves

$$v(p_{t_1}, I_{t_1}; x_{t_1}) = v(p_{t_0}, e^{\phi} I_{t_0}; x_{t_1}). \tag{2}$$

In words, EV^m is the change in income (in logs), under initial prices p_{t_0}, that a consumer with preferences $\succeq_{x_{t_1}}$ would need to be indifferent between the budget set defined by initial prices $(p_{t_0}, e^{\phi} I_{t_0})$ and the new budget set defined by new prices and income

[12] The vector p includes all relevant prices in the preference relation. If the preference relation is intertemporal, then p includes the path of current and future prices. In our applications, we abstract from intertemporal choices.

(p_{t_1}, I_{t_1}).[13] The new budget set is preferred to the initial one, if and only if, EV^m is positive. The superscript m represents the fact that this is the *micro* equivalent variation, since we take prices as given.

Discussion of our welfare criterion. As pointed out by Fisher and Shell (1968), the welfare criterion in Definition 3 is different to an alternative one that would attempt to measure the change in income that a consumer would need to be as well off in t_0 as in t_1. Whereas (2) refers to a choice that the consumer can in principle make, the latter criterion relies on an arbitrary intertemporal comparison of utils which do not correspond to any choice.[14]

To be concrete, suppose that x represents the age of the consumer. Clearly, we cannot meaningfully compare the amount of utils an individual derives from playing with toys during childhood to the amount of utils that same individual derives from drinking wine during adulthood. Since consumers never make choices about how old they are, their preferences across consumption goods consumed at different ages are not revealed by their choices. In the words of Heraclitus: "No man ever steps in the same river twice, for it's not the same river and he's not the same man." A comparison of utils during childhood to utils during adulthood is as meaningless as a comparison of utils between two different individuals. On the other hand, if we fix the consumer's age x, we can meaningfully compare the consumer's preferences about the bundles of goods they consumed at different points in their life or different consumption streams that they may face in the future.

This approach, of holding x constant, is different to the one taken when x represents some form of quality change. Intuitively, quality adjustments are more applicable to situations where the consumer can conceivably make choices between the good at differing

[13]In principle, we could also measure changes in welfare using compensating (instead of equivalent) variation, or by using initial (rather than final) preferences (see e.g. Balk, 1989 for a discussion of the various ways one can define welfare changes). Combining EV with final preferences (CV with initial preferences) is natural since this requires preserving the shape of the indifference curve at the final (initial) allocation. In the body of the paper, we focus on EV using final preferences since equivalent variation is more commonly used and final preferences are more relevant than initial preferences, but we characterize these other welfare measures in Appendix B. See also Remark 2.

[14]By an intertemporal comparisons of utils, we do not refer to dynamic decisions made at a point in time. To see this, imagine a dynamic decision maker with an intertemporal utility function. Suppose that after every period, we multiply the entire intertemporal utility function of the consumer by a large positive number. Doing this has no observable implications, since at every point in time, the consumer faces the exact same trade-offs as if we had not multiplied their utility. However, it dramatically alters the amount of utils received in one period versus another. Note that assumptions on time discounting (which discipline saving decisions) are about future versus present payoffs evaluated at a given point in time, so they do not pin down the level utils in one time period versus the level utils in another time period.

levels of quality. For example, if a box of chocolates undergoes quality change so that each box now contains twice as many chocolates, the consumer can conceivably make choices between the old and new boxes that reveal how much they value the quality change. Taste changes, on the other hand, do not involve meaningful choices from the consumer's perspective — if a consumer decides that she prefers dark chocolate to white chocolate, it does not make sense to ask how she would trade off consuming white chocolate in the past, when she preferred white chocolate, to consuming dark chocolate in the present, when she prefers dark chocolate. Instead, holding fixed her preferences, we can ask how she trades off white chocolate against dark chocolate. Therefore, the welfare implications of changes in unmeasured quality and changes in tastes are very different.

Of course, this is not to say that quality changes do not happen in reality. Rather, that we will abstract from quality change in our analysis. If there is quality change that can be represented as an unobservable price reduction, then all of our results will apply to the quality-adjusted "correctly-measured" prices instead of the market prices.

Real Consumption. Having defined changes in welfare, we now define changes in real consumption. The change in real consumption corresponds to what national income accountants and statistical agencies do when given data on the evolution of prices p and consumption bundles c. We assume that this data is *perfect* — completely accurate, comprehensive, adjusted for any necessary quality changes, and available in continuous time. This is because the biases associated with imperfections in the data, like the lack of quality adjustment, missing prices, or infrequent measurement, are different to the biases we study.

Definition 4 (Real consumption)**.** For some path of prices that unfold as a function of time t, the change in *real consumption* from t_0 to t_1 is defined to be

$$\Delta Y = \int_{t_0}^{t_1} \sum_{i \in N} p_i(t) \frac{dc_i}{dt} dt. \tag{3}$$

That is, changes in real consumption are cumulated changes in consumption goods measured at constant instantaneous prices.[15] Equation (3) is called a *Divisia* quantity index. In practice, since perfect data is not available in continuous time, statistical agencies approximate this integral via a (Riemann) sum using chained indices (e.g. Fisher or Tornqvist). We abstract from the imperfections of these approximations in this paper.

[15]For any variable z, we denote by dz its change over infinitesimal time intervals, so that $\Delta z = \int_{t_0}^{t_1} dz$.

In log terms, (3) can be rewritten as

$$\Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} b_i(t) \frac{d \log c_i}{dt} dt = \int_{t_0}^{t_1} \sum_{i \in N} b_i d \log c_i, \tag{4}$$

where $b_i(t)$ is the budget share of good i given prices, income, and preferences at time t. The last equation on the right-hand side simplifies notation by suppressing dependence on t in the integral. Using the budget constraint, we can express changes in real consumption in terms of changes in income deflated by price changes,

$$\Delta \log Y = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i d \log p_i = \Delta \log I - \Delta \log P^Y. \tag{5}$$

In other words, changes in real consumption are equal to changes in income minus changes in the consumption price deflator. Notice that changes in real consumption (or the consumption price deflator) potentially depend on the entire path of prices and quantities between t_0 and t_1 and not just the initial and final values. This is unlike welfare changes, EV^m, which depend only on initial and final prices and incomes and not on their entire path.

2.2 Relating Welfare and Consumption

We consider how real consumption and welfare change in response to changes in the budget set and the indifference curves of the consumer. We first consider globally exact results and then local approximations. The results are stated in terms of changes in prices and income, which we endogenize in Sections 3 and 4.

Global results. When preferences are unstable or non-homothetic, the following lemma shows that changes in welfare can be expressed as changes in income deflated by a shock-dependent price index. Changes in this price index are equal to budget-share weighted price changes, as in expression (5) for real consumption. However, whereas the price deflator for real consumption is based on observed budget shares (given prices, income, and preferences over time), the price deflator for welfare is based on hypothetical budget shares (at fixed utility level and fixed preferences).

To state this, define the expenditure function for any value of x by

$$e(p, u; x) = \min_c \{p \cdot c : u(c; x) = u\}.$$

The budget share of good i (given prices, preferences, and a level of utility) is

$$b_i(p, u; x) \equiv \frac{p_i c_i(p, u; x)}{e(p, u; x)} = \frac{\partial \log e(p, u; x)}{\partial \log p_i}, \tag{6}$$

where the second equality, Shephard's lemma, establishes a connection between budget shares and elasticities of the expenditure function. Note that when preferences are homothetic, then the expenditure function can be written as $e(p, u; x) = e(p; x)\, u$ and, hence, budget shares do not depend on u.

The following lemma characterizes changes in microeconomic welfare.

Lemma 1 (Micro Welfare). *Given any change in prices, income, and preferences, micro welfare changes are given by*

$$EV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i^{ev} d \log p_i = \Delta \log I - \Delta \log P^{EV}, \tag{7}$$

where $b_i^{ev}(p) \equiv b_i(p, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})$ *denotes budget shares at prices* p*, but fixing final preferences* x_{t_1} *and final utility* $v(p_{t_1}, I_{t_1}; x_{t_1})$.

Lemma 1 follows from the observation that EV^m can be re-expressed, using the expenditure function, as

$$EV^m = \log \frac{e(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})}{e(p_{t_0}, v(p_{t_0}, I_{t_0}; x_{t_1}); x_{t_1})} = \Delta \log I - \log \frac{e(p_{t_1}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})}{e(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})},$$

and recognizing that the second term can be written as the integral in (7).[16]

Compared to real consumption (5), which weights price changes by observed budget shares, EV^m weights price changes by hypothetical budget shares evaluated at current prices but for fixed final preferences, x_{t_1}, and final utility, $v(p_{t_1}, I_{t_1}; x_{t_1})$. Welfare depends on budget shares at x_{t_1} since only these preferences matter for EV^m. Welfare depends on budget shares evaluated at final utility, $v(p_{t_1}, I_{t_1}; x_{t_1})$, since EV^m adjusts the level of income in t_0 to make consumers as well off as they are in t_1. If welfare increases from t_0 to t_1, consumers must be given more income in t_0 to make them indifferent between t_0 and t_1. As we give consumers more income in t_0, the shape of their indifference curve changes until it mirrors the one in t_1. This means that the shape of the indifference curve relevant for the comparison is the one at t_1.

[16]By definition, EV^m only depends on initial and final prices and income, given t_1 preferences. This means that the integral in (7) is path independent and can be computed under any continuously differentiable path of prices that go from p_{t_0} to p_{t_1}. When comparing EV^m and real consumption, we consider the integral under the realized set of prices over time.

We can reinterpret the hypothetical budget shares b^{ev} as corresponding to those of a fictional consumer with homothetic and stable preferences with expenditure function $e^{ev}(p,u) = e(p, v_{t_1}; x_{t_1}) \frac{u}{v_{t_1}}$, where $v_{t_1} = v(p_{t_1}, I_{t_1}; x_{t_1})$. This implies that we can calculate changes in welfare given changes in prices based on budget shares $b^{ev}(p)$, *without* needing to know income elasticities or the nature of demand shocks. This is because the fictional consumer has homothetic and stable preferences, which means that all income elasticities are equal to one and there are no demand shocks. To compute $b^{ev}(p)$, we need to know the terminal budget shares and elasticities of substitution at the terminal equilibrium, as discussed in the following remark.

Remark 1 (Implementation)**.** To illustrate how Lemma 1 can be used to calculate EV^m with only knowledge of the elasticities of substitution, we consider a non-homothetic CES example as in Comin et al. (2015) or Fally (2020).[17] In this case, the following differential equation pins down changes in budget shares:

$$d\log b_i = (1-\theta_0)\left(d\log p_i - \mathbb{E}_b[d\log p]\right) + (\varepsilon_i - 1)\left(d\log I - \mathbb{E}_b[d\log p]\right) + d\log x_i, \quad (8)$$

where $\mathbb{E}_b(\cdot)$ is the expectation using budget shares as probability weights. The elasticity ε_i is the income elasticity of good i, and θ_0 is the elasticity of substitution across goods. The term $d\log x_i$ is a demand shifter, a residual that captures changes in expenditure shares not attributable to changes in income or prices. Note that when ε_i is equal to 1 for every i, final demand is homothetic, and when x_i is constant for all i, final demand is stable.[18]

For ex-post welfare questions, where $b(t_1)$ is observable, we can construct $b^{ev}(p)$ between t_0 and t_1 by iterating on

$$d\log b_i^{ev} = (1-\theta_0)\left(d\log p_i - \mathbb{E}_{b^{ev}}[d\log p]\right), \quad (9)$$

starting at t_1 and going back to t_0. These are changes in budget shares which are only due to substitution effects. Given b^{ev}, we can apply Lemma 1. Hence, demand shocks and income elasticities are not required to calculate EV^m.[19]

[17] Our results that only elasticities of substitution are necessary to calculate EV^m can be generalized to arbitrary non-CES functional forms, but since the intuition for the more general case is very similar to the CES case, we leave the more general non-parametric results in Appendix F.

[18] Since b_i are expenditure shares that always add up to one, it must necessarily be the case that $\mathbb{E}_b[d\log x] = 0$, and $\mathbb{E}_b[\varepsilon] = 1$. See Appendix C for a derivation of this log-linearization of Marshallian demand.

[19] In practice, estimating the elasticity of substitution θ_0 may require knowing the income elasticities. However, if the expenditure share of each good is sufficiently small, then θ_0 can be estimated without knowledge of income elasticities. Auer et al. (2021) estimate the relevant elasticities and apply Lemma 1 to measure the heterogeneous welfare effects of changes in foreign prices in the presence of demand non-

For ex-ante counterfactuals, where $b(t_1)$ is not known, we must first iterate on equation (8) from t_0 to t_1 to obtain $b(t_1)$. This first step requires knowledge of demand shocks and income elasticities. Once we calculate b_{t_1}, we repeat the procedure above and apply (9) to get the path of welfare-relevant budget shares b^{ev}.

Remark 2 (Compensating Variation under Initial Preferences)**.** Our baseline measure of welfare changes is equivalent variation under final preferences. An alternative would be to use compensating variation under initial preferences. Every result in the paper can be translated into compensating variation under initial preferences simply by reversing the flow of time. In particular, whereas Lemma 1 preserves the shape of the indifference curve at the final allocation, the compensating variation counterpart to Lemma 1 preserves the shape of the indifference curve at the initial allocation. Hence, calculating compensating variation requires knowledge of budget shares and elasticities of substitution at the initial allocation, whereas equivalent variation requires knowledge of budget shares and elasticities of substitution at the final allocation. See Appendix B for more details.[20]

We now contrast changes in real consumption and welfare.

Proposition 1 (Consumption vs. Welfare)**.** *Given any continuously differentiable change in prices, income, and preferences, the difference between welfare changes and real consumption is*

$$EV^m - \Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} \left(b_i - b_i^{ev}\right) d \log p_i = \left(t_1 - t_0\right) \mathbb{E}_t Cov\left(b - b^{ev}, d \log p\right),$$

where the covariance is calculated across goods at a point in time, and the expectation is calculated across time between t_0 and t_1.

An immediate consequence of Proposition 1 is the well-known result that real consumption is equal to changes in equivalent variation if, and only if, preferences are homothetic and stable. This is because when preferences are stable and homothetic, budget shares do not depend on x or changes in utility u over time. Hence, whenever preferences are homothetic and stable, $b_i^{ev}(t) = b_i(t)$ for every path of shocks and every t. In other words, we have the following corollary.

Corollary 1 (Homothetic and Stable Preferences)**.** *Welfare changes equal real consumption, $EV^m = \Delta \log Y$, if, and only if, preferences are homothetic and stable.*

homotheticities.

[20]In Appendix B we show that, up to a second-order approximation (but not globally), changes in real consumption equal a simple average of equivalent variation under final preferences and compensating variation under initial preferences.

Since changes in welfare depend only on initial and final values of primitives, Corollary 1 also implies that real consumption is path-independent whenever preferences are homothetic and stable.

When preferences are non-homothetic or unstable, observed budget shares not only reflect price changes but also non-price changes (that is, changes in x and changes in u). This generates discrepancies between observed and hypothetical budget shares, and hence between real consumption and welfare.

For a given path of price and income changes, welfare exceeds real consumption if, on average, $Cov\left(b - b^{ev}, d\log p\right) < 0$. That is, if changes in expenditure shares due to changes in u or x favor goods whose prices are falling more rapidly. Welfare changes are smaller than changes in real consumption if this pattern is reversed. If deviations between expenditure shares and relative price movements are orthogonal, then real consumption and welfare are equal. For example, with Cobb-Douglas preferences where expenditure shares are subject to exogenous shocks, real consumption and welfare are equal if, and only if, shocks to prices and shocks to consumer preferences are orthogonal.

To gain more intuition, we now characterize changes in real consumption and welfare using a second-order approximation around initial choices.

Local results. We consider local approximations of the objects of interest as the time period goes to zero, $t_1 - t_0 = \Delta t \to 0$. Throughout the rest of the paper, a second-order approximation means that the remainder term is of order Δt^3. We focus on second-order approximations to capture the interaction between price changes and expenditure-switching, which is the source of the gaps between real consumption and welfare changes.

We begin by stating the results in terms of Hicksian budget shares, and then we re-express them in terms of Marshallian (observable) budget shares. We start by characterizing the change in real consumption.

Lemma 2 (Approximate Consumption)**.** *Up a to a second order approximation, the change in real consumption is*

$$\Delta \log Y \approx \Delta \log I - b' \Delta \log p - \frac{1}{2} \sum_{i \in N} \left[\Delta \log p' \frac{\partial b_i}{\partial \log p} + \Delta \log x' \frac{\partial b_i}{\partial \log x} + \Delta \log v \frac{\partial b_i}{\partial \log u} \right] \Delta \log p.$$

This lemma, which is standard, shows that a second-order approximation accounts for the fact that budget shares change over time. The first term in the square brackets reflects changes in budget shares due to changes in relative prices and the next two terms correspond to changes in budget share due to non-price factors: preferences (under unstable preferences) and utility (under non-homothetic preferences).[21]

As discussed above, welfare is measured using changes in budget shares at fixed final preferences and utility. Thus, when comparing our welfare measures to real consumption starting at initial choices, we must add in changes in budget shares due to non-price factors, as indicated in the following proposition:

Proposition 2 (Approximate Welfare vs. Consumption)**.** *To a second-order approximation, the change in welfare is given by*

$$EV^m \approx \Delta \log Y - \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial b_i}{\partial \log x} + \Delta \log v \frac{\partial b_i}{\partial \log u} \right] \Delta \log p. \tag{10}$$

To a first order, changes in welfare equal real consumption. To understand the (second-order) gap between welfare and real consumption changes, consider first the case of homothetic but unstable preferences. Whereas changes in real consumption only take into consideration changes in budget shares in response to changes in utility parameters as the shock unfolds over time, changes in welfare must account for these changes from the start. Therefore, changes in budget shares due to non-price factors are multiplied by $1/2$ in real consumption, but they are multiplied by 1 in welfare. In other words, real consumption does not sufficiently account for substitution caused by preference instability. For example, the additional reduction in welfare (at new preferences) from a price increase in a good i with increasing demand ($d \log x \frac{\partial b_i}{\partial \log x} d \log p_i > 0$) is not fully reflected in real consumption, implying $EV^m < \Delta \log Y$.[22]

Similar reasoning applies in the case of stable but non-homothetic preferences, since changes in budget shares due to non-homotheticities should be incorporated in welfare immediately but are reflected in real consumption only gradually. For example, a reduction in the price of a good for which income effects are relatively weak ($d \log v \frac{\partial b_i}{\partial \log v} d \log p_i > 0$) implies a smaller increase in welfare than in real consumption ($EV^m < \Delta \log Y$).

Lemma 2 and Proposition 2 are both expressed in terms of Hicksian elasticities. We now re-express these results in terms of Marshallian elasticities using the non-homothetic CES aggregator introduced in Remark 1.

[21] The terms $\Delta \log x$ and $\Delta \log u$ need only be first-order approximations since they are multiplied by $\Delta \log p$ (and we only need to keep terms that are of order Δt^2). However, for the first term $-b \Delta \log p$, the primitive shock in prices must be approximated up to the second order, that is, $\Delta \log p \approx (\partial \log p / \partial t) \Delta t + 1/2 (\partial \log p^2 / dt^2) \Delta t^2$.

[22] A non-zero correlation between prices and demand shifters may emerge endogenously if firms have non-constant returns to scale or if firms invest in advertisement in response to productivity shocks. We consider the first possibility in Example 4 in Section 4 and discuss the second in Section 6.

Proposition 3 (Approximate Micro using Marshallian Demand)**.** *Consider some perturbation in demand $\Delta \log x$, prices $\Delta \log p$, and income $\Delta \log I$. Then, to a second-order approximation, the change in real consumption is*

$$\begin{aligned}\Delta \log Y \approx \Delta \log I - \mathbb{E}_b\left[\Delta \log p\right] - \frac{1}{2}(1-\theta_0) Var_b\left(\Delta \log p\right) \\ - \frac{1}{2} Cov_b\left(\Delta \log x, \Delta \log p\right) - \frac{1}{2}\left(\Delta \log I - \mathbb{E}_b\left[\Delta \log p\right]\right) Cov_b\left(\varepsilon, \Delta \log p\right),\end{aligned} \tag{11}$$

and the change in welfare is

$$EV^m \approx \Delta \log Y - \frac{1}{2} Cov_b\left(\Delta \log x, \Delta \log p\right) - \frac{1}{2}\left(\Delta \log I - \mathbb{E}_b\left[\Delta \log p\right]\right) Cov_b\left(\varepsilon, \Delta \log p\right), \tag{12}$$

where $Cov_b(\cdot)$ is the covariance using the initial budget shares as the probability weights.

We begin by considering the change in real consumption in (11). To a first order, the change in real consumption is just the change in income deflated by prices: $\Delta \log I - \mathbb{E}_b[\Delta \log p]$. The remaining terms capture nonlinearities associated with expenditure-switching. Since these are second-order, they are multiplied by $1/2$. We discuss these terms one-by-one. If goods are substitutes, $\theta_0 > 1$, then variance in relative prices boosts expenditure shares of cheaper goods and this increases measured real consumption. The second line captures the changes due to changes in demand. Intuitively, if the composition of demand shifts in favor of goods that happen to become relatively cheap, either due to non-homotheticity $Cov_b\left(\varepsilon, \Delta \log p\right)\left(\Delta \log I - \mathbb{E}_b\left[\Delta \log p\right]\right) < 0$ or demand shocks $Cov_b(\Delta \log x, \Delta \log p) < 0$, then real consumption increases.

Now consider changes in welfare in (12). As expected, the first-order terms are identical. The remaining terms capture the nonlinear response of welfare to price shocks. Note that if preferences are stable and homothetic, then welfare changes coincide with changes in real consumption. However, if preferences are unstable or non-homothetic, real consumption strays from welfare whenever price changes covary with non-price changes in demand. This happens because real consumption "undercounts" expenditure-switching due to the changes in demand.

3 Macroeconomic Changes in Welfare and Consumption

In the previous section we showed how changes in budget sets affect welfare when preferences are unstable and non-homothetic. For these problems, the frontier of the consumer's choice set is linear, since prices do not respond to the choices of the consumer. At the level of a whole society however, choice sets need not be linear. The production possibility set associated with an economy may have a nonlinear frontier. In this case, relative prices respond endogenously to choices made by consumers. In this section, we extend our analysis to allow for nonlinear production possibility frontiers (PPFs). The analysis in this section collapses to the one in Section 2 when the PPF of the economy is the same as the budget constraint (as happens in very simple general equilibrium models).

We first update our definitions of welfare, now at the macroeconomic level, and we introduce some basic structure and notation. We then present expressions for real GDP and welfare at the macroeconomic level, first globally and then locally in terms of endogenous sufficient statistics. In the next section, Section 4, we solve for these endogenous objects in terms of observable primitives.

3.1 Definition of Welfare and Real GDP

Consider a perfectly competitive neoclassical production economy with a representative agent. Each good $i \in N$ has a production function

$$y_i = A_i G_i \left(\{m_{ij}\}_{j \in N}, \{l_{if}\}_{f \in F} \right),$$

where G_i is a neoclassical production function, m_{ij} are intermediate inputs used by i and produced by j, and l_{if} denotes primary factor inputs used by i for each factor $f \in F$. The exogenous scalar A_i is a Hicks-neutral productivity shifter. Without loss of generality, we assume that G_i has constant returns to scale since decreasing returns to scale can be captured by adding producer-specific factors. Furthermore A_i is Hicks-neutral without loss of generality. This is because we can capture non-neutral (biased) productivity shocks to input j for producer i by introducing a fictitious producer that buys from j and sells to i with a linear technology. A Hicks-neutral shock to this fictitious producer is equivalent to a non-neutral technology shock to i.[23]

Let A be the $N \times 1$ vector of technology shifters and L be the $F \times 1$ vector of primary factor endowments. The production possibility set (and its associated frontier) is the set of feasible consumption bundles that can be attained given A and L.

For each A, L, and x, we denote equilibrium prices and aggregate income by $p(A, L, x)$

[23] Input-specific technology shocks shift input demand. In this sense, such shocks are similar to preference shocks, since they change sales holding fixed prices. However, such shocks are not a problem for measuring welfare or GDP since the effect of these shocks will be reflected in marginal costs, prices, and quantities. In other words, unlike utility, which does not have units that can be compared across time, physical production can be compared across time.

and $I(A, L, x)$. These equilibrium prices and incomes are unique up to the choice of a numeraire.

Define the *macro indirect utility* function as the maximum amount of utility the economy can deliver

$$V(A, L; x) = \max_{c}\{u(c; x) : c \text{ is feasible}\}.$$

Whereas the micro indirect utility takes prices as given and lets consumers pick any point in their budget set (even if such a point is not feasible at the economy-wide level), the macro indirect utility function takes the PPF as the primitive and lets society pick feasible points in the production possibility set. The first welfare theorem implies that the competitive equilibrium decentralizes the planning problem above with prices determined in equilibrium.[24]

We generalize our macroeconomic measure of welfare in the following way.

Definition 5 (Macro Welfare)**.** The change in welfare measured using the *macro equivalent variation* with *final preferences* is $EV^M(A_{t_0}, L_{t_0}, A_{t_1}, L_{t_1}; x_{t_1}) = \phi$ where ϕ solves

$$V(A_{t_1}, L_{t_1}; x_{t_1}) = V(A_{t_0}, e^{\phi}L_{t_0}; x_{t_1}).$$

In words, to compare the initial PPF, defined by (A_{t_0}, L_{t_0}), to the new PPF, defined by (A_{t_1}, L_{t_1}), we look for the proportional change in initial factor endowments required to make a planner with preferences $\succeq_{x_{t_1}}$ indifferent between the PPF defined by $(A_{t_0}, e^{\phi}L_{t_0})$ and the new PPF. Intuitively, EV^M expresses utility changes in terms of factor endowments. This is convenient in general equilibrium since it can be stated without reference to (endogenous) prices. In this sense, EV^M is similar to consumption-equivalents commonly used to measure welfare in macroeconomics.[25]

To see the difference between macro and micro notions of welfare, consider the example from the introduction. Suppose that in the final period households are both richer and older, and so they prefer to spend more of their income on healthcare services. Microeconomic welfare is measured by the endowment a single consumer, living in t_1, would have to be given to make her willing to go back to an economy with t_0 prices. This means that, in this hypothetical, the consumer would spend more of her budget on healthcare than she actually did in the initial period, because now she is both richer and older. If we did this for all consumers, the fact that in this hypothetical consumers collectively want to

[24]When the decentralized equilibrium is inefficient or preferences are non-aggregable, we can still rely on the micro welfare change defined in Section 2, which requires neither assumption. We discuss non-aggregable preferences in Section 6.

[25]We do not state our results in terms of consumption equivalents because when preferences are non-homothetic or unstable, households' desired consumption bundle is not stable.

spend more on healthcare would change the price of healthcare services in general equilibrium. Intuitively, since consumers were young in the initial period, healthcare services were relatively cheap due to low demand. For the older and wealthy consumers in the final period, relative prices in the initial period are therefore very attractive. However, if the older and wealthy consumers were transported to the initial economy, the fact that they demand more healthcare would raise healthcare prices, and this would mean that they would not be able to consume as much healthcare services. The issue is that using the initial budget set to represent the initial PPF is deceptive, since the initial budget ceases to be relevant when consumers are richer and older. Our macroeconomic notion of welfare accounts for the endogenous changes in prices by comparing the initial and final PPFs rather than the initial and final budget sets. To compare initial and final PPFs, we scale factor quantities instead of nominal income, since a proportional shift in factor quantities results in a proportional shift in the PPF and is interpretable without reference to prices.

Note that when relative prices do not respond to consumers' choices, then macro and micro welfare are always the same. Similarly, if preferences are homothetic and stable, then macro and micro welfare are the same. Intuitively, in this example, when preferences are homothetic and stable, households do not spend relatively more on healthcare in the hypothetical compared to what they actually spent in the initial period. Proposition 10 in Appendix A formalizes this result. Appendix E provides a quantitative example of the difference between micro and macro welfare in the context of a Covid-19 application.

As in Section 2, to study this problem we index the path of technologies, factor inputs, and preferences $A(t)$, $L(t)$ and $x(t)$, by time t. The definition of ΔY is the same as before: $\Delta Y = \int_{t_0}^{t_1} \sum_{i \in N} p_i dc_i$. In the general equilibrium model and (its applications), we refer to ΔY as real GDP in this economy since it coincides with the chain-weighted change in real GDP.[26]

[26]We abstract from international trade, government spending, and investment, all of which drive a gap between nominal GDP and nominal consumption. The most rigorous way to deal with investment is to cast the model in intertemporal terms and index goods by period of time in which they are consumed, along the lines of Basu et al. (2012). Theoretically, this presents no issues and our results apply to such economies unchanged. In practice, this changes the way we would map our model to the data, where nominal GDP would now be defined in net-present-value terms and prices would have to be multiplied by the (potentially stochastic) discount factor. A simpler way to justify our static application of these formulas is to either assume that consumption and investment bundles of goods are the same (so consumers value investment goods in the same way as they value consumption goods, and we are measuring the change in the output of the final good) or to treat investment as a static intermediate input (assuming full depreciation every period).

3.2 Relating Welfare and Real GDP

We now characterize changes in real GDP and welfare, first globally and then locally. The results in this subsection are the general equilibrium counterparts to those in Section 2. They are "reduced-form" in the sense that they are not expressed in terms of primitives. In Section 4, we explicitly solve for these sufficient statistics in terms of observable primitives.

Define the sales shares relative to GDP of each good or factor i to be

$$\lambda_i = \frac{p_i y_i}{I}\mathbf{1}(i \in N) + \frac{w_i L_i}{I}\mathbf{1}(i \in F),$$

where w_i and L_i are the price and quantity of factor i. The sales share λ_i is often referred to as a *Domar* weight. Note that referring λ_i as a "share" is an abuse of language since $\sum_{i \in N} \lambda_i > 1$ whenever there are intermediate inputs. On the other hand $\sum_{i \in F} \lambda_i = 1$.

Global Results. The following results show that changes in real GDP and welfare can both be represented as sales-weighted averages of technology changes. Real GDP uses actual sales shares over time, while welfare uses sales shares in an artificial economy in which budget shares only respond to price changes.

Lemma 3 (Real GDP)**.** *Given a change in technologies, factor quantities, and preferences, the change in real GDP is*

$$\Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} \lambda_i d \log A_i + \int_{t_0}^{t_1} \sum_{i \in F} \lambda_i d \log L_i, \tag{13}$$

where λ are sales shares which are functions of A, L, and x.

In (13), the first N summands are equal to measured TFP, and the last F summands are the growth in real GDP caused by changes in factor inputs. Lemma 3 shows that changes in real GDP are equal to sales-weighted changes in technology and factor inputs. This is a slight generalization of Hulten (1978) to environments with unstable and non-homothetic final demand.

Next, we show that a Hulten-style result also exists for changes in welfare. Define $\lambda^{ev}(A, L)$ to be sales shares in a fictional economy with the PPF (A, L) but where consumers have stable homothetic preferences represented by the expenditure function $e^{ev}(p, u) = e(p, v_{t_1}, x_{t_1})\frac{u}{v_{t_1}}$ where $v_{t_1} = v(p_{t_1}, I_{t_1}; x_{t_1})$, similar to Section 2.

Proposition 4 (Macro Welfare). *Changes in macro welfare are*

$$EV^M = \int_{t_0}^{t_1} \sum_{i \in N} \lambda_i^{ev} d \log A_i + \int_{t_0}^{t_1} \sum_{i \in F} \lambda_i^{ev} d \log L_i. \tag{14}$$

According to Proposition 4, growth accounting for welfare should be based on hypothetical sales shares evaluated at current technology but for fixed final preferences and final utility. This should be contrasted with real GDP in (13), which uses sales shares evaluated at current technology and current preferences. As with real GDP, the first N summands of (14) are changes in *welfare-relevant TFP* and the last F summands are changes in welfare due to changes in factor inputs. We discuss some salient implications of this proposition below.

The first implication is that for welfare questions, the only information we need about preferences are expenditure shares and elasticities of substitution at the final allocation, since the fictional consumer in Proposition 4 has stable preferences with income elasticities all equal to one.[27]

Second, Proposition 4 implies that if the path of technologies and factor quantities is continuously differentiable, then real GDP is equal to the change in welfare if, and only if, preferences are homothetic and stable (in which case $\lambda(A, L, x) = \lambda^{ev}(A, L)$ for every A, L, and x). That is, Corollary 1, introduced for our micro welfare notion, holds in general equilibrium.

Third, as stated in the following corollary, movements on the surface of a PPF driven by changes in preferences have no effect on macroeconomic welfare or real GDP.

Corollary 2 (Demand Shocks Only). *In response to changes in preferences, x, that keep the PPF unchanged, $A(t) = A(t_0)$ and $L(t) = L(t_0)$ for $t \in [t_0, t_1]$,*

$$\Delta \log Y = EV^M = 0.$$

However, micro welfare changes, EV^m, may be nonzero.

Since the production possibility set is not changing, macro welfare (defined for fixed preferences) does not change. Quantities and prices do, however, change between t_0 and t_1 in response to changes in preferences over these goods. Micro welfare changes are typically non-zero when prices change, as shown in Section 2. These results are not contradictory: the micro welfare metric assumes that consumers can choose any bundle

[27] Following the observation made in Remark 2, for compensating variation at initial preferences, we need to know elasticities of substitution at the initial allocation instead of the final one.

in their budget set at given prices (hence welfare changes as prices change). On the other hand, the macro welfare metric takes into account the fact that such choices may not be feasible for society as a whole. Finally, movements along the surface of a PPF have no effect on real GDP because demand-driven changes in output raise some quantities and reduce others, and these effects exactly cancel out.

While real GDP and macroeconomic welfare changes are the same so long as we stay on the surface of a given PPF, the two are not equal when the PPF shifts. This is because real GDP is based on a path of sales shares λ that take into consideration technology shocks as well as changes in preferences and non-homotheticities in final demand. However, changes in welfare are based on a path of sales shares λ^{ev} that only take into consideration technology shocks. Therefore, if productivity rises for goods for which sales shares fall due to non-technological factors, then $EV^M < \Delta \log Y$.

To get more intuition for Proposition 4, in the following section, we use a second-order approximation to characterize changes in real GDP and welfare.

Local Results. We characterize, up to a second order approximation (as $t_1 - t_0 = \Delta t \to 0$), the response of real GDP and welfare to technology and preference shocks, now taking into account the endogenous evolution of sales shares. To make the formulas more compact and without loss of generality, when we write local approximations we abstract from shocks to factor endowments.[28]

Lemma 4 (Approximate Real GDP)**.** *Up to to a second order approximation, the change in real GDP is*

$$\Delta \log Y \approx \lambda' \Delta \log A + \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} + \Delta \log A' \frac{\partial \lambda_i}{\partial \log A} \right] \Delta \log A_i. \tag{15}$$

Equation (15) resembles the one in Lemma 2, but it is based on sales shares and technology shocks rather than budget shares and price changes. The first term in (15) corresponds to the Hulten-Domar formula. The terms in square brackets reflect nonlinearities due to changes in sales shares. Intuitively, if sales shares decrease for those goods with higher productivity growth, then real GDP growth slows down due to substitution effects. This type of effect, known as Baumol's cost disease, is an important driver of the slow-down in aggregate productivity growth.

[28] Shocks to factor endowments are a special case of TFP shocks. To represent a factor endowment shock as a TFP shock, we add fictitious producers that buy the factor endowments on behalf of the other producers and shock their productivity.

The following proposition compares the change in welfare with the change in real GDP.

Proposition 5 (Approximate Macro Welfare vs. GDP)**.** *Up to a second order approximation,*

$$EV^M \approx \Delta \log Y + \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} + \Delta \log A' \frac{\partial \log v}{\partial \log A} \frac{\partial \lambda_i}{\partial \log v} \right] \Delta \log A_i. \tag{16}$$

The intuition underlying the gap between macro welfare and real GDP in Proposition 5 is similar to that in Proposition 2 for our micro results. Specifically, real GDP takes into consideration changes in sales shares *along* the equilibrium path. These changes in sales shares could be induced by technology shocks but they could also be due to changes in preferences and non-homotheticities. However, welfare measures treat changes in shares due to technology shocks differently than changes in shares due to demand shocks or non-homotheticities. In both cases, real GDP "undercorrects" for changes in shares caused by non-homotheticities or changes in preferences. In particular, welfare is lower than real GDP if technology growth is lower in goods where sales shares rise due to preference changes or non-homotheticities.

In contrast to Proposition 2, there can be a gap between real GDP and welfare even if all productivity shocks are the same. Specifically, suppose that productivity growth is common across all goods $\Delta \log A_i = \Delta \log A$ and denote the gross output to GDP ratio by $\lambda^{sum} = \sum_{i \in N} \lambda_i \geq 1$. Then Proposition 5 implies that the gap between real GDP and welfare is

$$EV^M - \Delta \log Y \approx \frac{1}{2} \Delta \log A \left[\Delta \lambda^{sum} - \frac{\partial \lambda^{sum}}{\partial \log A} \Delta \log A \right], \tag{17}$$

where the term in square brackets is the change in the gross output to GDP ratio due to demand-side forces only. In particular, if demand shifts towards sectors with higher value-added as a share of sales, then $EV^M < \Delta \log Y$ when technology shocks are positive. Intuitively, this happens because welfare is less reliant on intermediates than real GDP, and hence real GDP is more sensitive to productivity shocks. Of course, in the absence of intermediate inputs, this effect disappears because λ^{sum} will always equal one.

4 Structural Macro Results and Analytic Examples

The results in the previous section are reduced-form in the sense that they take changes in prices and sales shares as given and are written using compensated (Hicksian) demand. In this section, we solve for changes in these endogeneous objects in terms of observable sufficient statistics. We first consider economies with linear PPFs and then economies with nonlinear PPFs. We provide some analytical examples to provide more intuition. Although our results are local, one can use them to conduct global counterfactuals, as we do in Section 5, along the lines of Baqaee and Farhi (2019a).

We now spell out a macroeconomic model and solve for changes in prices and shares in general equilibrium. For clarity, we restrict attention to nested-CES economies. The general case is in Appendix F, and the intuition is very similar.

Nested-CES economies. Household preferences are represented by a non-homothetic CES aggregator, which imply that budget shares vary according to (8). Recall that θ_0 is the elasticity of substitution across consumption goods and ε is the vector of income-elasticities. Production also uses nested-CES aggregators. Nested-CES economies can be written in many different equivalent ways, since they may have arbitrary patterns of nests. We adopt the following representation. We assume that each good $i \in N$ is produced with the production function

$$y_i = A_i G_i\left(\left\{m_{ij}\right\}_{j\in N}, \left\{l_{if}\right\}_{f\in F}\right) = A_i\left(\sum_{j\in N}\omega_{ij} m_{ij}^{\frac{\theta_i-1}{\theta_i}} + \sum_{f\in F}\omega_{if} l_{if}^{\frac{\theta_i-1}{\theta_i}}\right)^{\frac{\theta_i}{\theta_i-1}},$$

where the parameters ω_{ij} and ω_{if} are constants. Any nested-CES production network can be represented in this way if we treat each CES aggregator as a separate producer (see Baqaee and Farhi, 2019c).

Input-output matrix. We stack the expenditure shares of the representative household, all producers, and all factors into the $(1 + N + F) \times (1 + N + F)$ input-output matrix Ω. The first row corresponds to the household. To highlight the special role played by the representative agent, we index the household by 0, which means that the first row of Ω is equal to the household's budget shares introduced above ($\Omega_0 = b'$, with $b_i = 0$ for $i \notin N$).[29] The next N rows correspond to the expenditure shares of each producer on every other producer and factor. 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).

Leontief inverse matrix. The Leontief inverse matrix is the $(1 + N + F) \times (1 + N + F)$ matrix defined as

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

[29] We expand the vector of demand-shifters $\Delta \log x$ and income elasticities ε to be $(1 + N + F) \times 1$, where $\Delta \log x_i = \varepsilon_i = 0$ if $i \notin N$.

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. We partition Ψ in the following way:

$$\Psi = \left[\begin{array}{c|ccc|ccc} 1 & \lambda_1 & \cdots & \lambda_N & \Lambda_1 & \cdots & \Lambda_F \\ \hline 0 & \Psi_{11} & \cdots & \Psi_{1N} & \Psi_{1N+1} & \cdots & \Psi_{1N+F} \\ 0 & & \ddots & & & & \\ 0 & \Psi_{N1} & & \Psi_{NN} & \Psi_{NN+1} & \cdots & \Psi_{NN+F} \\ \hline 0 & 0 & \cdots & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \cdots & \vdots & \vdots & 1 & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 1 \end{array}\right].$$

The first row and column correspond to final demand (good 0). The first row is equal to the vector of sales shares for goods and factors λ'. To highlight the special role played by factors, we interchangeably denote their sales share by the $F \times 1$ vector Λ. The next N rows and columns correspond to goods, and the last F rows and columns correspond to the factors. Define the $(1 + N + F) \times F$ matrix Ψ^F as the submatrix consisting of the right F columns of Ψ, representing the network-adjusted factor intensities of each good. The sum of network-adjusted factor intensities for every good i is equal to one, $\sum_{f \in F} \Psi_{if} = 1$ because the factor content of every good is equal to one. In our results below we will use the identities $\lambda' = b'\Psi$ and $\Lambda' = b'\Psi^F$.

4.1 One-Factor Models (Linear PPF)

For intuition, we start by focusing on general equilibrium economies with only one factor of production. When the economy has a single factor of production (or equivalently, a linear PPF), then macro and micro welfare changes are the same. Furthermore, in this case, we can characterize both the change in real GDP and the change in welfare up to a second order.

Proposition 6 (Approximate Macro Welfare vs GDP: Single Factor)**.** *Consider some perturbation in technology, $\Delta \log A$, and final demand, $\Delta \log x$. When the economy has one factor of* *production, the change in real GDP is*

$$\Delta \log Y \approx \sum_{i \in N} \lambda_i \Delta \log A_i + \frac{1}{2} \sum_{j \in N} \lambda_j (\theta_j - 1) Var_{\Omega^{(j)}} \left(\sum_{i \in N} \Psi_{(i)} \Delta \log A_i \right)$$
$$+ \frac{1}{2} Cov_{\Omega^{(0)}} \left(\Delta \log x + (\sum_{i \in N} \lambda_i \Delta \log A_i) \varepsilon, \sum_{i \in N} \Psi_{(i)} \Delta \log A_i \right), \tag{18}$$

where the summations are evaluated over all goods and factors, so that i and $j \in \{0\} + N + F$, and the $Cov_{\Omega^{(j)}}(\cdot)$ is the covariance using the jth row of Ω as the probability weights and $\Psi_{(i)}$ is the ith column of the Leontief inverse. The difference between welfare and GDP is

$$EV^M - \Delta \log Y \approx \frac{1}{2} Cov_{\Omega^{(0)}} \left(\Delta \log x + (\sum_{i \in N} \lambda_i \Delta \log A_i) \varepsilon, \sum_{i \in N} \Psi_{(i)} \Delta \log A_i \right). \tag{19}$$

Proposition 6 is a general equilibrium counterpart to Proposition 3. We discuss (18) and (19) in turn, starting with (18). The first term in Equation (18) is the Hulten-Domar term. The other terms are second-order terms resulting from the fact that sales shares change in response to shocks. The first one of these terms captures nonlinearities due to the fact that sales shares can respond to changes in relative prices caused by technology shocks (these effects were emphasized by Baqaee and Farhi, 2019c). The terms on the second line of (18), which are the ones we focus on in this paper, capture changes in sales shares due to changes in preferences or non-homotheticities.

Equation (19) shows that while real GDP correctly accounts for substitution due to supply shocks, it needs to be corrected for substitution due to changes in final demand due to demand shocks or non-homotheticities. Whereas in partial equilibrium, the gap between welfare and real GDP is proportional to the covariance of supply and demand shocks (see Proposition 3), equation (19) shows that in general equilibrium, the relevant statistic is the covariance of demand shocks with a network-adjusted notion of supply shocks not supply shocks per se. Furthermore, Proposition 6 shows that the elasticities of substitution are irrelevant for the gap between welfare and real GDP in one-factor models. This is because relative prices do not change as the equilibrium moves along a linear PPF in response to demand-driven forces. Therefore, demand shocks do not trigger expenditure switching due to the endogenous response of relative prices. When we relax the linearity of the PPF, we see that the elasticities of substitution in production do, in general, affect the gap between welfare and GDP.

We now work through some simple examples to illustrate the intuition in Proposition 6.

Example 1 (Correlated Supply and Demand Shocks)**.** We start with the simplest possible example, a one sector model without any intermediates. In this case, sales shares are just budget shares $\lambda_i = b_i = \Omega_{0i}$ and $\Psi_{(i)}$ is the ith column of the $1+N+F$ identity matrix $I_{(i)}$.Therefore, Proposition 6 implies

$$EV^M - \Delta \log Y \approx \frac{1}{2} \left(Cov_b \left(\Delta \log x, \Delta \log A \right) + Cov_b \left(\varepsilon, \Delta \log A \right) \mathbb{E}_b [\Delta \log A] \right).$$

Hence, welfare changes are greater than the change in real GDP if productivity and demand shocks are positively correlated. This could happen either because preferences exogenously change to favor high productivity goods, $Cov_b \left(\Delta \log x, \Delta \log A \right) > 0$, or preferences endogenously change to favor high productivity growth goods due to non-homotheticities, $Cov_b \left(\varepsilon, \Delta \log A \right) \Delta \log Y > 0$. When shifts in demand are orthogonal to shifts in supply, to a second-order approximation, real GDP measures welfare correctly.

Example 2 (Input-Output Connections)**.** For models with linear PPFs, input-output connections affect the gap between real GDP and welfare in two ways: (1) the impact of technology shocks is bigger when there are input-output linkages because $\Psi_{(i)} \geq I_{(i)}$ and $\lambda_i \geq b_i$; (2) the production network "mixes" the shocks, and this may reduce the correlation of supply and demand shocks by making the technology shocks more uniform. However, since it is the covariance (not the correlation) of the shocks that matters, this means the effects are, at least theoretically ambiguous.

To see these two forces, consider the three economies depicted in Figure 1. Each of these economies has a roundabout structure. Panel 1a depicts a situation where each producer uses only its own output as an input, Panel 1b a situation where all producers use the same basket of goods (denoted by M) as an intermediate input, and Panel 1c a situation where each producer uses the output of the other producer as an input. We compute the correction to GDP necessary to arrive at welfare for each of these cases using Proposition 6. For clarity, we focus on demand shocks caused by instability rather than non-homotheticity, though it should be clear that this does not affect any of the intuitions.

For Panel 1a, we get

$$EV^M - \Delta \log Y \approx \frac{1}{2} Cov_b(\Delta \log x_i, \Omega_{iL}^{-1} \Delta \log A_i),$$

where the covariance is computed across goods $i \in N$ and Ω_{iL} is the labor share for i. Hence, as intermediate inputs become more important, the necessary adjustment becomes larger. This is because, for a given vector of preference shocks, the movement in

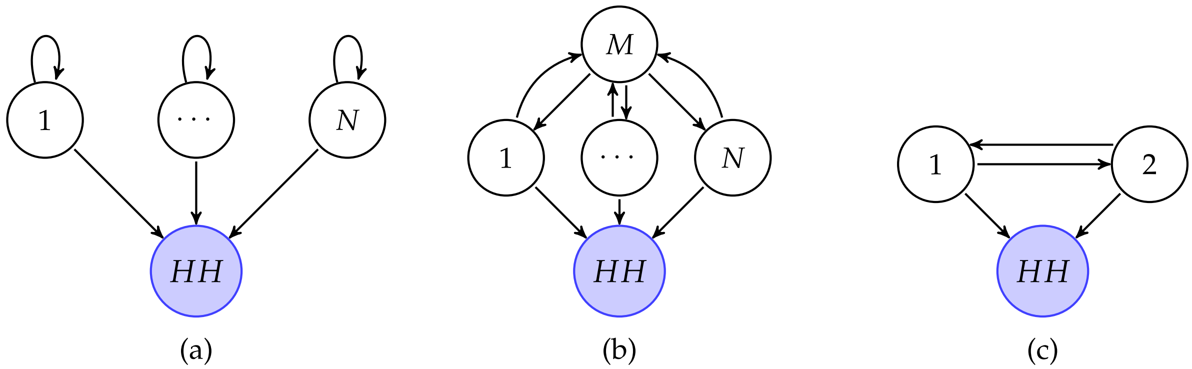

Figure 1: Three different kinds of round-about economy. The arrows represent the flow of goods. The only factor is labor which is not depicted in the diagram.

sales shares is now larger due to the roundabout nature of production.[30]

On the other hand, for Panel 1b, we get[31]

$$EV^M - \Delta \log Y \approx \frac{1}{2}\left(Cov_b(\Delta \log x_i, \Delta \log A_i) - Cov_b(\Delta \log x_i, \Omega_{iL}) \frac{\sum_{i \in N} \Delta \log A_i}{\sum_{i \in N} \Omega_{iL}} \right).$$

Hence, in this case, if the labor share Ω_{iL} is the same for all $i \in N$, then the intermediate input share is irrelevant. Intuitively, in this case, all producers buy the same share of materials, so a shock to the composition of household demand does not alter the sales of any producer through the supply chain, and hence only the first-round non-network component of the shocks matters.[32]

Finally, consider Panel 1c. For clarity, focus on the case where only producer 1 gets a productivity shock ($\Delta \log A_2 = 0$). In this case, the difference between real GDP and welfare is

$$EV^M - \Delta \log Y \approx \frac{1}{2} \frac{1}{1 - \Omega_{12}\Omega_{21}} Cov_b \left(\Delta \log x, \begin{bmatrix} 1 \\ \Omega_{21} \end{bmatrix} \right) \Delta \log A_1.$$

As the intermediate input share Ω_{21} approaches one, the adjustment goes to zero (since the covariance term goes to zero). Intuitively, as Ω_{21} goes to one, the increase in demand

[30] As discussed after Equation (17), if all productivity shocks are the same, there may still be an adjustment due to heterogeneity in labor shares. In particular, if demand shocks are higher for sectors with higher labor shares, then $EV^M < \Delta \log Y$ when technology shocks are positive.

[31] For this example, we assume that there are no productivity shocks to the intermediate bundle $\Delta \log A_M = 0$ and we assume that $\Omega_{iM} = 1/N$ for each $i \in N$.

[32] As in Footnote 30, if the labor share is heterogeneous across producers, there is an additional adjustment which depends on the covariance between demand shocks and labor shares. If the demand shocks reallocate expenditures towards sectors with high labor shares, then welfare becomes less sensitive to productivity shocks than real GDP.

for the first producer from a change in preferences is exactly offset by a reduction in demand from the second producer who buys inputs from the first producer. In this limiting case, changes in consumer preferences have no effect on the overall sales share of the first producer.

These three examples serve to illustrate that the effect of input-output networks on the adjustment are theoretically ambiguous but potent.

4.2 Multi-Factor Models (Nonlinear PPF)

To characterize the response of output and welfare to shocks when the PPF is nonlinear we rely on the following result.[33]

Proposition 7 (Sales and Prices in General Equilibrium). *Consider some perturbation in final demand $d\log x$ and technology $d\log A$. Then changes in prices of goods and factors are*

$$d\log p_i = -\sum_{j\in N}\Psi_{ij}d\log A_j + \sum_{f\in F}\Psi^F_{if}d\log\lambda_f. \tag{20}$$

Changes in sales shares for goods and factors are

$$\begin{aligned}\lambda_i d\log\lambda_i = &\sum_{j\in\{0\}+N}\lambda_j(\theta_j-1)Cov_{\Omega^{(j)}}\left(-d\log p, \Psi_{(i)}\right) \\ &+ Cov_{\Omega^{(0)}}\left(d\log x, \Psi_{(i)}\right) + Cov_{\Omega^{(0)}}(\varepsilon, \Psi_{(i)})\left(\sum_{k\in N}\lambda_k d\log A_k\right).\end{aligned} \tag{21}$$

Unlike the previous propositions, Proposition 7 pins down changes in prices and sales shares only up to a first-order approximation. First-order approximations of changes in prices and sales shares are all we need to plug into the reduced-form expressions in Proposition 2 and Proposition 5. We briefly describe the intuition for (20) and (21).

Equation (20) captures how productivity shocks and changes in factor prices travel downstream through supply chains to affect the marginal cost of downstream produces. Equation (21), on the other hand, captures how changes in relative prices and final demand travel upstream through supply chains to determine the sales of each producer. The first line in (21) corresponds to changes in final demand and intermediate good demand for good i due to changes in relative (goods and factor) prices. The second line of (21) captures instability of final demand (first term on the second line) and non-homotheticity (second term on the second line).

[33]Proposition 7 extends Baqaee and Farhi (2019c) to economies where final demand is non-homothetic and unstable.

We now work through some simple examples to illustrate how nonlinear PPFs affect our results from earlier.

Example 3 (Decreasing Returns to Scale)**.** Consider the one-sector model without intermediate inputs in Example 1 but now suppose that production functions are non-constant-returns-to-scale. Specifically, the production for good i is

$$y_i = A_i L_i^{\gamma},$$

where L_i is labor and γ need not equal 1. Furthermore, suppose that preferences are homothetic ($\varepsilon_i = 1$ for every i), but potentially unstable ($\Delta \log x \neq 0$). To apply our theorems to this economy, where producers have non-constant-returns production functions, we introduce a set of producer-specific factors in inelastic supply, and suppose that each producer has a Cobb-Douglas production function that combines a common factor with elasticity γ and a producer-specific factor with elasticity $1-\gamma$. This means that our economy has $1+N$ factors.

We apply Proposition 5 to compute the difference between welfare and real GDP. To do this, we first use Proposition 7 to compute changes in sales shares due to demand shocks:

$$\frac{\partial \lambda_i}{\partial \log x} \cdot d\log x = Cov_{\Omega^{(0)}}\left(d\log x, \Psi_{(i)}\right) + (\theta_0 - 1) Cov_{\Omega^{(0)}}\left(-\frac{d\log p}{d\mathrm{log}x} d\log x, \Psi_{(i)}\right).$$

The factor content of every good $i \in N$ is given by $\Psi^F_{if} = \gamma$ when f is the common factor and $\Psi^F_{if} = 1-\gamma$ when f is the producer-specific factor. Since the factor shares of each producer are constant, the log change in the producer-specific factor share is the same as the log change in the sales share of that producer. Therefore, we can replace $d\log p_i / d\log x$ in the covariance with$(\gamma - 1)\left(\partial \log \lambda_i / \partial \log x\right) d\log x$ (the other components of price changes are common to all producers and drop out of the covariance). Plugging this into the expression above and solving yields a closed-form expression for $\partial \lambda_i / \partial \log x$. This allows us to apply Proposition 5 to get the difference between welfare and real GDP up to a second order approximation:

$$EV^M - \Delta\mathrm{log}Y \approx \frac{1}{2} \frac{Cov_{\Omega^{(0)}}(\Delta \log x, \Delta \log A)}{1 + (\theta_0 - 1)(1-\gamma)}.$$

Note that the denominator disappears when we have constant-returns to scale ($\gamma = 1$) *or* the elasticity of substitution across goods is one ($\theta_0 = 1$). Outside of these cases, complementarities ($\theta_0 < 1$) amplify the impact of preference shocks under decreasing returns to scale ($\gamma < 1$). Intuitively, if preferences shift in favor of some good, the price of that good rises due to decreasing returns to scale. The fact that the price of the good increases raises the sales share of that good due to complementarities, which creates a feedback loop, raising prices of the good further, and causing additional substitution. In other words, in the decreasing returns to scale model with complementarities, sales shares respond more strongly to demand shocks. Given that sales shares respond more strongly to demand shocks, the necessary adjustment to correct real GDP is larger.

Example 4 (Macro vs. Micro Welfare Change)**.** Finally, we demonstrate the difference between macro and micro welfare changes using the previous example. The economy in the previous example has multiple factors and unstable preferences. Therefore, macro and micro notions of welfare are different since the PPF is no longer linear.

To illustrate this difference, suppose that only preference shocks are active (there are no supply shocks $\Delta \log A = 0$ and $\Delta \log L = 0$). By Corollary 2, real GDP changes are zero. Since the PPF is being held constant, macro-welfare changes are also zero. Micro-welfare changes, on the other hand, are not equal to zero. Specifically, by Proposition 2, micro welfare improves $EV^m > 0$ if preference shocks negatively covary with price changes. By equation (41), changes in prices are

$$d \log p_i = \sum_{f \in F} \Psi^F_{if} \frac{\partial \log \Lambda_f}{\partial \log x} d \log x.$$

Using the derivations above, for each $i \in N$, we obtain

$$d \log p_i = \frac{(1-\gamma)}{(1+(\theta_0 - 1)(1-\gamma))} \frac{1}{\lambda_i} Cov_{\Omega^{(0)}} \left(d \log x, I_{(i)} \right),$$

where $I_{(i)}$ is the ith column of the identity matrix. If there are decreasing returns, $\gamma < 1$, then a positive demand shock for i raises the price of i. The change in the price is amplified if goods are complements and mitigated if goods are substitutes. We can now apply Proposition 3 to obtain the change in micro welfare, up to a second order,

$$EV^m \approx -\frac{1}{2} \frac{(1-\gamma)}{(1+(\theta_0 - 1)(1-\gamma))} Var_{\Omega^{(0)}} (\Delta \log x) \neq 0 = EV^M.$$

With decreasing returns to scale ($\gamma < 1$), micro welfare decreases since the demand shock increases the prices of goods the consumer now values more. From a micro perspective, where the agent takes the budget sets as given, the agent is worse off.

On the other hand, when the economy has increasing returns to scale ($\gamma > 1$), micro welfare increases in response to demand shocks. Intuitively, in this case, increased demand for a good lowers the price of that good, which makes the consumer better off.[34] Of course, from a societal perspective, welfare has not changed, since the production possibility set of the economy has not changed.

5 Applications

In this section, we consider three applications of our results. The first application is to the problem of long-run growth and the difference between welfare-relevant and measured aggregate productivity growth in the presence of income effects and demand instability. The second application shows that correlated firm-level supply and demand shocks drive a wedge between measured real GDP and welfare even in the short-run. Our last application considers how demand instability can make measured real GDP an unreliable metric for changes in production, and we illustrate this point for the Covid-19 crisis.

5.1 Long-Run Growth and Structural Transformation

Baumol (1967) showed that, as economies grow, sectors with lower relative productivity growth rates expand (in terms of sales and value-added) relative to sectors with faster productivity growth. This means that over time, aggregate productivity growth is increasingly determined by those sectors whose productivity growth is slowest. This phenomenon is oftentimes called Baumol's cost disease.

The causes of changes in industry mix over time, called structural transformation, are the subject of a large literature. As discussed by Buera and Kaboski (2009) and Herrendorf et al. (2013), there are two primary reasons the literature offers for why economies undergo structural transformation. The first reason is complementarities, either in demand or in production, that mean that industries with faster productivity growth, whose relative prices fall over time, shrink as a share of overall expenditures. The second reason is non-homotheticities or instability in final demand, whereby marginal propensities to consume across different sectors change as the economy grows, resulting in more expenditures on sectors with lower productivity growth (i.e. services).

Either way, changes in sales shares will affect aggregate productivity growth. Following Nordhaus et al. (2008), aggregate productivity growth can be decomposed into two

[34] If the economy has increasing returns to scale, then the decentralized equilibrium is potentially inefficient. However, the propositions regarding micro welfare changes, which take changes in prices as given, do not require that the decentralized equilibrium be efficient.

terms:

$$\Delta \log TFP = \sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{i,t_0} \Delta \log A_{i,t} + \underbrace{\sum_{t=t_0}^{t_1} \sum_{i \in N} \left(\lambda_{i,t} - \lambda_{i,t_0} \right) \Delta \log A_{i,t}}_{\text{Baumol Adjustment}},$$

where $\lambda_{i,t}$ is the sales shares of industry i in period t and $\Delta \log A_{i,t}$ is the growth in gross-output productivity over that time period.[35] The first term captures changes in aggregate TFP if industry-structure had remained fixed, and the second term is the adjustment attributed to the fact that sales shares change over time. The second-term captures the importance of Baumol's cost disease.

Proposition 4 implies that, for the purposes of welfare, changes in sales shares due to income effects or demand instability must be treated differently to changes in sales shares due to complementarities. In particular, the welfare-relevant measure of the change in TFP is

$$\Delta \log TFP^{w} = \sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{i,t_0} \Delta \log A_{i,t} + \underbrace{\sum_{t=t_0}^{t_1} \sum_{i \in N} \left(\lambda_{i,t} - \lambda_{i,t_0} \right) \Delta \log A_{i,t}}_{\text{Baumol Adjustment}} + \underbrace{\sum_{t=t_0}^{t_1} \sum_{i \in N} \left(\lambda_{i,t}^{ev} - \lambda_{i,t} \right) \Delta \log A_{i,t}}_{\text{Welfare Adjustment}},$$

where λ^{ev} is the hypothetical sales-shares of each industry holding fixed final preferences and income-level — that is, sales shares after they have been purged from changes due to factors other than changes in relative prices.

Two polar extremes. Computing these terms requires an explicit structural model of the economy. However, there are two polar cases in which the TFP adjustment term can be calculated without specifying the detailed model. On the one hand, demand is stable and homothetic, and changes in sales shares are due only to relative price changes (complementarities). On the other hand, there are no complementarities (as in a Cobb-Douglas economy) and changes in sales shares are only due to income effects or demand instability. If structural transformation is driven by a combination of complementarities and non-homotheticities or demand instability, then the change in welfare TFP will be somewhere in between these two cases. The following corollary of Proposition 4 summarizes the change in welfare-TFP in these two polar cases.

[35] Technically, this is an approximation, since we define aggregate TFP in continuous time but the data is measured in discrete time (at annual frequency). However, this approximation error, resulting from the fact that the Riemann sum is not exactly equal to the integral is likely to be negligible in practice. At our level of disaggregation, long run TFP growth is very similar if we weight sectors using sales shares at time t or time t and $t+1$ averages.

Corollary 3. *If changes in sales shares are due only due only to complementarities, then*

$$\Delta \log TFP^{w} = \Delta \log TFP = \sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{i,t} \Delta \log A_{i,t}.$$

If changes in sales shares are due only to non-homotheticity or instability of demand, then

$$\Delta \log TFP^{w} = \Delta \log TFP + \sum_{t=t_0}^{t_1} \sum_{i \in N} \left(\lambda_{i,t_1} - \lambda_{i,t}\right) \Delta \log A_{i,t} = \sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{i,t_1} \Delta \log A_{i,t}.$$

In the first case, since preferences are homothetic and stable, welfare-TFP is equal to TFP in the data. In the second case, since there are no complementarities in production or demand, sales shares do not respond to productivity changes. In order to hold utility and preferences fixed at their final value, we must compute welfare-TFP using terminal sales shares.

To quantify Corollary 3, we use US-KLEMS data on sales shares and TFP growth for 61 private-sector industries. The results are shown in Figure 2. We calculate changes in industry-level gross-output TFP following the methodology of Jorgenson et al. (2005) and Carvalho and Gabaix (2013). The constant-initial-sales-share term grows by around 58 log points (or 78%), whereas the chain-linked change in aggregate TFP grew by around 47 log points (or 60%). Hence, Baumol's cost-disease caused aggregate TFP to fall by -10 log points, reducing aggregate productivity growth by around 23 percent (from 78% to 60%).

If we assume that structural transformation is due solely to non-homotheticities or demand instability, then by Corollary 3 the growth in welfare-relevant TFP from 1948-2014 has been 37 log points (or 46%) instead of the measured 47 log points (or 60%) — that is, to say, a 23 percent additional reduction in the growth rate.[36]

Intuitively, welfare-based productivity increases less than TFP because, relative to 1948, preferences in 2014 favor low productivity growth sectors such as services (due to either income effects or demand instability). This means that, at 1948 prices, households require less income growth to be indifferent between their budget constraint in 1948 and the one in 2014. This is because sectors with high income elasticities, like services, were cheaper compared to manufacturing in 1948 than in 2014. This intuition is

[36]The gap between constant-initial-sales shares TFP and (chained-linked) aggregate TFP growth and the gap between aggregate TFP and welfare-relevant TFP growth are driven by two forces. First, reallocation of sales towards sectors with lower relative productivity growth rates (the standard Baumol's cost-disease mechanism). Second, reallocation in sales towards sectors with lower intermediate input use (see equation 17). In our quantitative results, the second force accounts for roughly 18% of the gap between constant-initial-sales shares TFP and aggregate TFP growth, and 35% of the gap between aggregate TFP growth and welfare-relevant TFP growth.

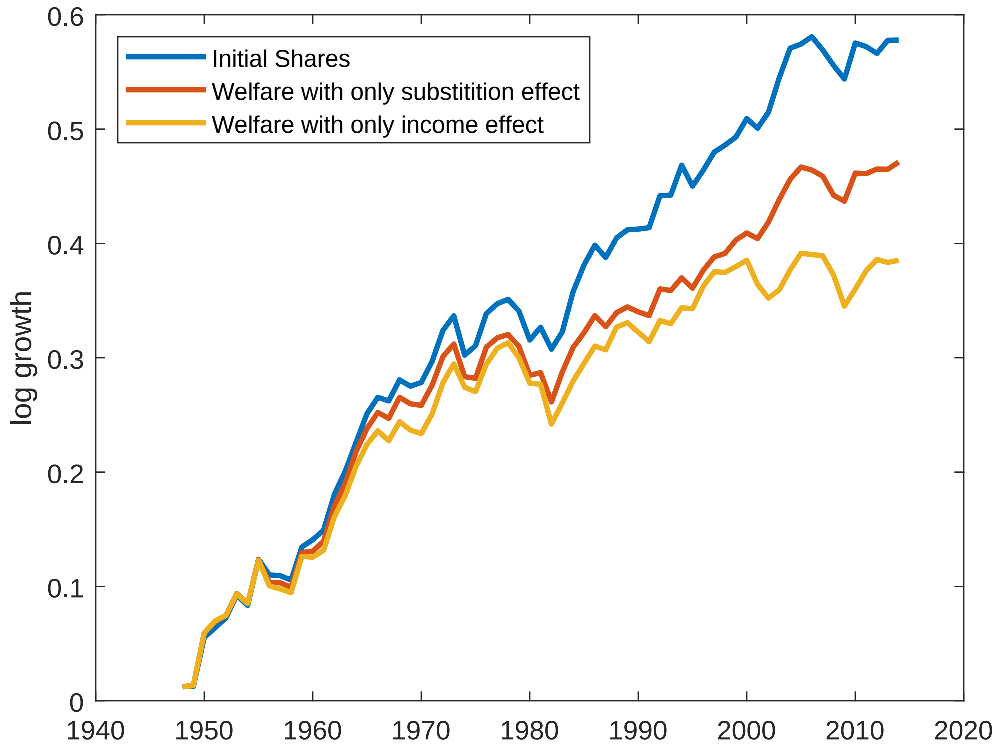

Figure 2: Growth in welfare-relevant TFP (in logs) from 1948 to 2014 using US-KLEMS. The blue line uses initial shares (in 1948) to calculate TFP changes. The red and yellow line measure the increase in welfare-relevant TFP in each year relative to 1948 under alternative assumptions about income and substitution elasticities. The red line assumes that sales shares change only due to substitution effects (welfare-relevant TFP is equal to measured chained-aggregate TFP). The yellow line assumes that sales shares change only due to income effects (or demand instability).

flipped for compensating variation. The reduction in income at 2014 prices necessary to make households indifferent between their budget constraint in 2014 and the one in 1948 is greater than TFP growth because, as households become poorer in 2014, they favor goods which are relatively cheap in 2014 such as manufacturing. Hence, we must reduce their income by more to make them indifferent between 2014 and 1948.[37]

To sum up, structural transformation caused by income effects or demand instability reduced welfare, E^M, by roughly twice as much as structural transformation caused by complementarities. To understand why the necessary adjustment is roughly twice as

[37]This means that if structural transformation is purely due to income effects, the welfare change based on CV is given by initial sale-share weighted productivity growth, $\sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{i,t_0} \Delta \log A_{i,t}$. In this case, the Baumol adjustment is not welfare-relevant for CV.

big, consider the second-order approximation in Proposition 4. Up to a second-order approximation, we can write

$$\Delta \log TFP^{\text{welfare}} \approx \Delta \log TFP + \frac{1}{2} \left[\sum_{i \in N} \frac{\partial \lambda_i}{\partial \log x} \Delta \log x + \frac{\partial \lambda_i}{\partial \log v} \Delta \log v \right] \Delta \log A_i, \tag{22}$$

where

$$\Delta \log TFP \approx \sum_{i \in N} \lambda_{i,t_0} \Delta \log A_i + \frac{1}{2} \sum_{i \in N} \Delta \lambda_i \Delta \log A_i.$$

If changes in sales shares are due entirely to demand-driven factors, then the term in square brackets in (22) is equal to $\sum_{i \in N} \Delta \lambda_i \Delta \log A_i$, so

$$\Delta \log TFP^{\text{welfare}} \approx \sum_{i \in N} \lambda_{i,t_0} \Delta \log A_i + \sum_{i \in N} \Delta \lambda_i \Delta \log A_i.$$

In other words, the adjustment to the initial sales shares must be roughly twice as large as the adjustment to the initial sales shares caused by complementarities. These second-order approximations are more accurate if changes in sales shares are well-approximated by linear time trends, and the surprising accuracy of the second-order approximation is a result of this fact.

The reasoning above implies that the welfare adjustment to bring measured TFP in line with welfare may be as high as 23 percent. However, it may also be as small as zero, if the data can be perfectly explained using only complementarities.

Quantitative illustration away from two polar extremes. In practice, both complementarities and non-homotheticities are likely to play an important role in explaining structural transformation. To dig deeper into the size of the welfare adjustment, we use a simplified version of the model introduced in Section 4 calibrated to the US economy, accounting for input-output linkages and complementarities, and use the model to quantify the size of the welfare-adjustment as a function of the elasticities of substitution.

Remarkably, Proposition 4 implies that to compute the welfare-relevant change in TFP, we must only supply the information necessary to compute λ^{ev}. That is, since we know sales shares in the terminal period 2014, we do not need to model the non-homotheticities or demand-shocks themselves, and the exercise requires no information on the functional form of non-homotheticities or the slope of Engel curves or magnitude of income elasticities *conditional* on knowing the elasticities of substitution.

We map the model to the data as follows. We assume that the constant-utility final demand aggregator has a nested-CES form. There is an elasticity θ_0 across the three groups of industries: primary, manufacturing, and service industries. The inner nest has elasticity of substitution θ_1 across industries within primary (2 industries), manufacturing (24 industries), and services (35 industries).[38] Production functions are also assumed to have nested-CES forms: there is an elasticity of substitution θ_2 between the bundle of intermediates and value-added, and an elasticity of substitution θ_3 across different types of intermediate inputs. For simplicity, we assume there is only one primary factor of production (a composite of capital and labor). We solve the non-linear model by repeated application of Proposition 7 in the fictional economy with stable and homothetic preferences.[39]

We calibrate the CES share parameters so that the model matches the 2014 input-output tables provided by the BEA. For different values of the elasticities of substitution $(\theta_0, \theta_1, \theta_2, \theta_3)$ we feed changes in industry-level TFP (going backwards, from 2014 to 1948) into the model and compute the resulting change in aggregate TFP. This number represents the welfare-relevant change in aggregate TFP. We report the results in Table 1.

Table 1: Percentage change in measured and welfare-relevant TFP in the US from 1948 to 2014.

$(\theta_0, \theta_1, \theta_2, \theta_3)$	(1,1,1,1)	(0.5,1,1,1)	(1,0.5,1,1)	(1,1,0.5,1)	(1,1,1,0.5)
Welfare TFP	46%	46%	54%	48%	55%
Measured TFP	60%	60%	60%	60%	60%

The first column in Table 1 shows the change in welfare-relevant TFP assuming that there are no complementarities (all production and consumption functions are Cobb-Douglas). In this case, all changes in sales shares in the data are driven by non-homotheticities or demand-instability, and hence welfare-relevant TFP has grown more slowly than measured TFP, exactly as discussed in the previous section. The other columns show how the results change given lower elasticities of substitution. As we increase the strength of complementarities, the implied non-homotheticities required to match changes in sales shares in the data are weaker. This in turn reduces the gap between measured and welfare-relevant productivity growth.

[38]In order to map this nested structure to our baseline model, good 0 is a composite of good 1-3, where good 1 is a composite of primary industries, good 2 is a composite of manufacturing industries, and good 3 is a composite of service industries. Goods 4-65 are the disaggregated industries. Finally, good 66 is the single factor of production.

[39]We calculate TFP shocks by industry in the data allowing for cross-industry variation in capital and labor shares, and then feed these shocks as a primitive into our one-factor model (alternatively, in order to allow for industry variation in capital intensity while keeping the one-factor simplicity, we must assume that labor and capital are perfect substitutes).

Table 1 also shows that not all elasticities of substitution are equally important. The results are much more sensitive to changes in the elasticity of substitution across more disaggregated categories, like materials, than aggregated categories, like agriculture, manufacturing, and services. In Appendix D, we show that this is an intuitive consequence of Proposition 6.

5.2 Aggregation Bias with Firm-Level Shocks

In the previous application, we considered a long-run industry-level application. Since industry-level sales shares are relatively stable over short-horizons, given industry-level data, biases in real GDP and consumption are likely to be modest at high frequency. However, this does not mean that these biases are necessarily absent from short-run data.

Whereas industry sales shares are stable at high frequency, firm or product-level sales shares are highly volatile even over the very short-run. If firms' or products' supply and demand shocks are correlated, then measured industry-level output is biased relative to what is relevant for welfare, even if all firms are infinitesimal within their industry. Furthermore, these biases in industry-level data are not diversified away as we aggregate, even if all industries are infinitesimal in the aggregate economy. If these firm-level supply and demand shocks are persistent, then this bias becomes even larger as we lengthen the horizon. Finally, if we start with industry-level (rather than firm-level) data, we are ruling out the existence of these biases by construction.

To make these points formally, we rely on a specification of our model with an explicit industrial structure.

Definition 6 (Industrial Structure)**.** An economy has an *industry structure* if the following conditions hold:

i. Each firm i belongs to one, and only one, industry I. Firms in the same industry share the same constant-returns-to-scale production function up to a firm-specific Hicks-neutral productivity shifter A_i.

ii. The representative household has homothetic preferences over industry-level goods, where the Ith industry-level consumption aggregator is

$$c_I = \left(\sum_{i \in I} \bar{b}_{iI} x_i c_i^{\frac{\zeta_I - 1}{\zeta_I}} \right)^{\frac{\zeta_I}{\zeta_I - 1}},$$

where c_i are consumption goods purchased by the household from firm i in industry I and x_i are firm-level demand shocks.

iii. Inputs purchased by any firm j from firms i in industry I are aggregated according to

$$m_{jI} = \left(\sum_{i \in I} \bar{s}_{iI} m_{ji}^{\frac{\sigma_I - 1}{\sigma_I}} \right)^{\frac{\sigma_I}{\sigma_I - 1}},$$

where m_{ji} are inputs purchased by firm j from firm i, and $\bar{s}_{iI}$ is a constant.

Input-output and production network models that are disciplined by industry-level data typically have an industry structure of the form defined above. For such economies, we can characterize the bias in real GDP relative to welfare easily.

Proposition 8 (Aggregation Bias)**.** *For models with an industry structure, in response to firm-level supply shocks $\Delta \log A$ and demand shocks $\Delta \log x$, we have*

$$\Delta \log EV^M \approx \Delta \log Y + \frac{1}{2} \sum_{i \in I} b_I Cov_{b_{(I)}}(\Delta \log x_{(i)}, \Delta \log A_{(i)}) + \Theta,$$

where b_I is industry I's share of final demand and $b_{(I)}$ is a vector whose ith element is b_i/b_I if i belongs to industry I and zero otherwise. The scalar Θ is the gap between real GDP and welfare in a version of the model with only industry-level shocks.[40]

In words, Proposition 8 implies that if firms' productivity and demand shocks are correlated with each other (but not necessarily across firms), then there is a gap between real GDP and welfare that does not appear in an industry-level specification of the model. Furthermore, this bias is, to a second-order, additive. That is, the overall bias is the sum of the industry-level bias (that we studied in the previous section) plus the additional bias driven by within-industry covariance of supply and demand shocks.

Identifying supply and demand shocks at the firm-level is notoriously difficult, and many papers impose that these shocks must be orthogonal in order to estimate them (or to estimate demand elasticities). However, a correlation between firm-level supply and demand shocks is a natural prediction in models with endogenous customer acquisition (Arkolakis, 2016; Foster et al., 2016; Fitzgerald et al.; 2016), as we discuss in Section 6.

[40]More specifically,

$$\Theta = \frac{1}{2} \sum_{i \in I} \sum_{j \in J} \left[\mathbb{E}_{b_I} \left(\Delta \log A_{iI} \right) \left(\frac{\partial b_I}{\partial \log x_j} \right) + \mathbb{E}_{s_I} \left(\Delta \log A_{iI} \right) \left(\frac{\partial \lambda_I^f}{\partial \log x_j} \right) \right] \Delta \log x_j,$$

where $\lambda_I^f = \lambda_I - b_I$ is the sales share of industry I excluding sales to the household. In words, Θ is proportional to the product of changes in industry-level sales shares caused by demand shocks times the industry-level (sales-weighted) productivity shocks.

A detailed quantitative investigation of this issue is beyond the scope of this paper, but Proposition 8 can be used for a back-of-the-envelope calculation.

As a numerical illustration, suppose that firm- or product-level demand shocks and TFP shocks at annual frequency have a standard deviation of 0.5 and 0.2 log points, and suppose that they have a correlation of 0.2. Proposition 8 implies that, up to a second-order approximation and at annual frequency, the gap between welfare-relevant industry output and measured industry output is around $\frac{1}{2} \times 0.2 \times 0.5 \times 0.2 = 1\%$.[41] If these industry-level covariances are the same across all industries, then the gap between real GDP and welfare induced by firm-level shocks is also 1% at annual frequency. This is a one-off steady-state adjustment that does not compound over time (every year, the change in welfare would be 1% higher than real GDP) if the shocks are independent over time.

To continue this thought experiment, note that if supply and demand shocks at the firm level are persistent, then the bias grows over time. For example, suppose that firm-level supply and demand shocks are both AR(1) processes with persistence parameters $\rho^A \in [0,1]$ and $\rho^x \in [0,1]$. Then the gap between welfare and real GDP in period T relative to some period 0 due to firm-level shocks is (suppressing industry sub-indices)

$$Cov_b(\Delta \log x_T, \Delta \log A_T) = \frac{1-(\rho_x \rho_A)^T}{1-(\rho_x \rho_A)} Cov_b(\Delta \log x_1, \Delta \log A_1).$$

Hence, as we expand the horizon, the gap between real GDP and welfare becomes larger. In the limit where supply and demand shocks are correlated random walks, the gap expands without bound as we lengthen the horizon. If the persistence parameters are both equal to 0.9, then the asymptotic gap between welfare and real GDP is around 5 times larger than the one-period gap.

5.3 Path Dependence of Real GDP with an Application to Covid-19

In our last application, we switch from focusing on welfare to considering the properties of real GDP in the presence of unstable or non-homothetic preferences. We apply our analysis to one specific and noteworthy instance of demand instability: the onset of the Covid-19 crisis. As shown by Corollary 1 changes in real consumption measure changes in welfare accurately if, and only if, preferences are stable and homothetic. The question that begs to be asked is: *does the change in real GDP or real consumption measure the change in*

[41] The gap would be around 2% based on statistics for TFPQ and demand shocks reported in Table 2 of Eslava and Haltiwanger (2021) using Colombian data. We obtain very similar numbers if we compute the gap exactly (numerically) or if we use the second-order approximation in Proposition 8

any stable function of allocations when preferences are unstable or non-homothetic? The gradient theorem for line integrals provides an answer. In our context, this theorem implies that if sales shares are not derivatives of some stable scalar-valued function, then the change in real GDP between two allocations can be any arbitrary number depending on the path taken (see Hulten, 1973).

This means that whenever preferences are non-homothetic or unstable, changes in real consumption (or real GDP) generically depend on the path and timing of shocks, unlike changes in welfare, which depend only on the initial and final allocation. To see this, suppose final demand is homothetic but unstable, and the economy faces supply shocks $\Delta \log A$ and demand shocks $\Delta \log x$. Consider two different paths for how the shocks unfold: under path SD the supply shocks arrive first and then the demand shocks, and under DS the demand shocks arrive first and then the supply shocks. A corollary of Proposition 5 is then the following.

Corollary 4 (Path Dependence and Timing of Shocks)**.** *If preferences are homothetic, then*

$$\Delta \log Y^{SD} = EV^{M}(A_{t_0}, L_{t_0}, A_{t_1}, L_{t_1}; x_{t_0}), \text{ and } \Delta \log Y^{DS} = EV^{M}(A_{t_0}, L_{t_0}, A_{t_1}, L_{t_1}; x_{t_1}),$$

where $\Delta \log Y^{SD}$ *and* $\Delta \log Y^{DS}$ *denote real GDP computed under SD and DS, respectively.*

That is, the change in real GDP under the first path measures something different than the change in real GDP under the second path. In general, the increase in real GDP is larger if prices fall (or productivity rises) later for goods that are subject to earlier positive demand shocks.[42] If we go from the initial allocation to the final one along the path SD and then come back along DS, real GDP will be different even though we are back at the exact same allocation. Hence, how fast the economy recovers from a shock like the Covid-19 crisis can be affected by irrelevant details like the order in which supply and demand conditions revert to normal.

Disconcerting as this path dependence problem of real GDP (and other Divisia indices) may be, it is tempting to dismiss it as a theoretical curiosity. However, far from being an exotic theoretical possibility, path-dependence is a well-known problem for practitioners when dealing with high-frequency data. Most manuals on index number construction, for example OECD et al. (2004), warn against using chain-weighted indices for high frequency data.

The term *"chain drift"* bias in national accounting describes a situation when a chained index measures an overall change between t_0 and t_1 even though all prices and quantities

[42]Another implication of path-dependence is that, in the presence of demand instability, the order of reforms that increase sectoral productivity matters for the magnitude of the overall increase in real GDP

in t_0 and t_1 are identical. This is a specific manifestation of path dependence and, by the gradient theorem for line integrals, it must be driven by either demand instability, income effects, or approximation errors due to discreteness. Chain drift bias can thus appear when movements in prices and quantities are oscillatory, where changes that take place over some periods are reversed in subsequent periods. This happens frequently at sub-annual frequencies, so the usual advice given to practitioners is to avoid these problems by not chaining their data too frequently. Intuitively, if we only chain indices at annual or lower frequencies, then it will very rarely be the case that prices and quantities will revert back to their previous values (due to long run growth). Hence, we will not observe chain drift.

However, chaining at low frequencies does not solve the path dependence problem, it simply avoids having to confront it by making it less likely that we ever see multiple paths that lead to the same allocations. But the path dependence problem does not become less severe at lower frequencies, it simply becomes less noticeable. The approach we advocate in this paper, which corrects the index for instability and non-homotheticity, has the virtue of restoring path independence.

The Covid-19 recession provides an illustrative case in point. During the first two quarters of 2020, employment and production rapidly contracted in most economies around the world. This contraction was unusual in that it fell very unevenly across different industries. The sales shares of some sectors, like air transportation or entertainment venues, collapsed while those of other sectors, like retail trade, rose. These changes are difficult to justify on the basis of observed price changes since relative price movements across sectors were relatively modest. Given the fact that producing and consuming some goods requires much more face-to-face contact than others, it is natural to suppose that non-price factors likely played an important role in the observed changes in sales shares.[43] That is, final demand across sectors as a function of measured prices was unstable during this period.

We now consider how sensitive the change in real GDP can be to the exact timing of the supply and demand shocks. To answer this question, we use a modified version of

[43] These non-price factors could be changes in quality, changes in preferences, or changes in government policies during the pandemic. For our purposes in this application (i.e. path dependence of real GDP as measured in the data), it does not matter what is the source of demand instability at measured prices. Of course, in order to assess the welfare costs of Covid-19, the source of demand instability is important. However, since consumers do not make choices about whether or not they live in a world with Covid-19, the welfare costs of Covid-19 are not revealed by their choices. For example, choices between going to restaurants or cooking at home reveal preferences between these two goods conditional on living in a world with a pandemic. What is not revealed by these choices are preferences over restaurants with and without pandemics, and cooking at home with and without pandemics. For this reason, we refrain from analyzing the welfare costs of Covid-19 in this application.

the quantitative model introduced above. Since we are interested in a short-run application, we assume that factor markets are segmented by industry, so that labor and capital in each industry is inelastically supplied. We then calibrate share parameters to match the input-output table in 2018, and consider a range of elasticities of substitution. Following Baqaee and Farhi (2020), we hit the economy with a vector of primitive supply and demand shocks.[44] The supply shocks are reductions in factor endowments, calibrated to match peak-to-trough reductions in hours worked by sector from January, 2020 to May, 2020. The primitive demand shocks are shifters in final demand, calibrated to match the observed peak-to-trough reductions in personal consumption expenditures by sector from January, 2020 to May, 2020 (given the supply shocks).

We consider three different calibrations informed by empirical estimates from Atalay (2017) and Boehm et al. (2015): high complementarities, medium complementarities, and no complementarities (Cobb-Douglas). The high complementarity scenario sets the elasticity of substitution across consumption goods (outer and inner nests) to be 0.7, the one across intermediates to be 0.01, across value-added and materials to be 0.3, and the one between labor and capital to be 0.2. The medium complementarities case sets the elasticity of substitution across consumption goods to be 0.95, the one across intermediates to be 0.01, across value-added and materials to be 0.5, and the one between labor and capital to be 0.5. The Cobb-Douglas calibration sets all elasticities of substitution equal to unity.

Table 2: The change in real GDP under alternative paths for the arrival of supply and demand shocks between February 2020 to May 2020. We assume that prices and quantities are measured in, and chaining is done, in continuous time. The initial and final value of the demand and supply shocks are the same across rows, but each row corresponds to a different arrival path of the shocks.

Elasticities	High compl.	Medium compl.	Cobb-Douglas
Supply then demand	-16.2%	-12.5%	-10.8%
Demand then supply	-10.1%	-9.4%	-9.0%
Simultaneous supply and demand	-12.1%	-10.6%	-9.8%

Table 3 shows that the reduction in real GDP is much larger when supply shocks arrive first before the demand shocks, even though, the difference in welfare and the difference in physical allocations in every scenario is the same. This is because changes in demand are positively correlated with changes in supply shocks. That is, the reduction in the out-

[44]Changes in labor by sector and expenditures, used to calibrate supply and demand shocks, are taken from Baqaee and Farhi (2020). For related analysis of Covid-19 induced supply shocks, see e.g. Bonadio et al. (2020) and Barrot et al. (2020). For related analysis of Covid-19 induced demand shocks, see Cakmakli et al. (2020).

put of a sector reduces GDP by less if the households have already reduced their spending on that sector's output. Hence, reductions in the supply of transportation services or food and beverage stores are less costly if demand for those sectors contracts first, compared to the case in which the fall in production is followed by a reduction in demand.

The wide gaps we document here suggest that if, during the recovery, the supply and demand shocks do not disappear in exactly the same way as they arrived, measured real GDP can be higher or lower than it was before the crisis, even if the economy returns exactly to its pre-Covid allocation. And this difference can be as much as 6% — equivalent to several years worth of economic growth. Hence, during episodes where final demand is unstable, real GDP becomes an increasingly unreliable guide for measuring output.

In Appendix E we calculate the change in micro welfare using initial and final preferences. We show that micro welfare falls by more at final preferences than at initial preferences. This is because the demand shocks associated with Covid-19 raised the prices of goods that consumers valued more. This finding is consistent with Cavallo (2020), who shows that goods whose expenditure share rose in 2020 had higher inflation, which implies that chained CPI inflation is higher than CPI inflation using initial shares. Our results imply that the chained index does not sufficiently account for changes in expenditure shares due to demand shocks. This means that true (welfare-based) inflation is even higher than what would be measured by the chained-CPI advocated in Cavallo (2020).[45]

6 Extensions

In this section, we briefly summarize how our results can be extended in different ways.

Extensive margin. If preference instability or non-homotheticity causes a consumer to begin purchasing a good in t_1 that she did not consume in t_0 (or to stop consuming a good that she was previously consuming), then our global and local formulas apply to that consumer without change.

To make this more explicit, consider a consumer whose preferences are represented by the utility function

$$\left(\int_0^{x^*} c(z)^{\frac{\sigma-1}{\sigma}} dz \right)^{\frac{\sigma}{\sigma-1}}, \tag{23}$$

where goods are indexed by $z \in [0,1]$ but the consumer only values goods $z \in [0, x^*]$. In this situation, x^* is a preference parameter, where goods $z \in (x^*, 1]$ are available at a

[45]This is all abstracting from the issue of disappearing goods. See Diewert and Fox (2020) for a discussion of this in the context of Covid-19.

finite price, but the consumer chooses not to consume them.

Consider how the welfare of the consumer changes accounting for the fact that x^* can change between t_0 and t_1. The following is an application of Lemma 1.

Proposition 9 (New Goods Due to Taste Shocks)**.** *Consider a household with preferences defined by* (23)*. Up to a second-order approximation,*

$$\Delta \log EV^m \approx \Delta \log Y + \frac{1}{2} b(x^*) \Delta x^* \left[\mathbb{E}_b \left[\Delta \log p \right] - \Delta \log p(x^*) \right].$$

In words, the gap between welfare and real GDP depends on product of sales shares at the cut-off $b(x^*)$, the change in the cut-off Δx^*, and the difference between inflation at the cut-off versus average inflation. If new goods are added $\Delta x^* > 0$, and the new goods experienced lower than average inflation, then welfare is higher than what is detected by real GDP. However, this adjustment is second-order (since it involves products of Δ), and to a first-order, real GDP is equal to the true change in the cost of living.

It is interesting to contrast Proposition 9 to the well-known new-goods adjustment due to Feenstra (1994), which, to a first-order approximation, is

$$\Delta \log EV^m = \Delta \log Y + \frac{1}{1-\sigma} \Delta \log \left[\int_{\mathcal{C}} b(z) dz \right], \tag{24}$$

where $\mathcal{C}$ is the set of continuing goods and the integral is their share in expenditures. The difference in these results is due to a difference in interpretation. Under the interpretation in Proposition 9, the change in the extensive margin is caused by a change in tastes — that is, the goods were previously available to the consumer in the initial period but the consumer chose not to consume them (or goods are available in the final period, but the consumer chose to stop consuming them). Therefore, when we calculate welfare changes, we simply need to adjust the price index so that it accounts for the price of goods that the consumer is choosing to consume in the final period. On the other hand, under (24), when we compute the change in welfare, we assume that the consumer is unable to consume the new goods in the past or can no longer consume the disappearing goods in the present. That is, under (24), when goods are not consumed they are valued by the consumer but the implicit price is infinity.

Therefore, if a good is available in t_0 and t_1, but the consumer does not consume the good in period t_0 and does consume the good in t_1 (due to, for example, advertising), an application of (24) is not innocuous. If the change in consumer behavior is due to a change in tastes, as opposed to a change in availability, then no adjustment is necessary to a first-order, and to a second-order, the relevant adjustment is the one in Proposition 9.

Endogenous separable arguments in the utility function (e.g. leisure or home-production). If there are goods in the utility function that are endogenously chosen but not measured, then an all-encompassing welfare measure must impute shadow prices for these goods (see Jones and Klenow, 2016). For example, suppose that leisure is the non-measured argument in the utility function. If these are separable from market goods, so that preferences over c are stable when the quantity of leisure changes, then our baseline results apply to the market-good component of welfare, even if leisure changes.

Endogenous non-separable arguments in the utility function (e.g. advertising). If the parameters of the utility function x are not separable from goods c, then our welfare questions ask how changes in constraints over c affect welfare holding fixed x. That is, we do not attempt to answer how a change in x itself affects welfare, which may or may not be a question that can be answered. A salient example of x can be advertising, which can change households ranking over different consumption bundles, and is obviously non-separable from market goods. In principle, advertising may have value to the consumer — that is, the consumer can have preferences over the amount of advertising they wish to be exposed to. If advertising is informative, the consumer's utility may be increasing in advertising, and equally plausibly, if advertising is manipulative, then utility may be decreasing in advertising. We do not attempt to answer the question of how much the household values advertising, instead, we hold fixed the amount of advertising[46] (or indeed the weather, chemicals in the brain, and whatever else that affects valuations over consumption bundles), and measure how changes in the availability of market goods affects welfare.

Interestingly, unlike random fluctuations in tastes, advertising is a purposeful economic activity, and therefore, models of advertising and consumer acquisition, for example Arkolakis (2016), explicitly predict that changes in tastes induced by advertising will be correlated with changes in physical productivity, whereby more productive firms will expend more resources on advertising. This positive correlation means that we should expect real GDP or real consumption measures to be systematically biased in situations where advertising plays a large role in consumption choices.

Beyond CES. Our results in Section 4 can be generalized beyond CES functional forms relatively easily. In Appendix F, we discuss how Proposition 7 must be adjusted to allow for non-CES production and utility functions.

[46] In this sense, our approach is related to Dixit and Norman (1978), who study the welfare implications of advertising at either pre- or post-advertising preferences. As argued by Fisher and McGowan (1979), this does not answer the question of what is the value of advertisement taking into account the change in tastes.

Heterogeneous agents. Our microeconomic welfare results can be applied in economies without representative agents (including open economies). Our macroeconomic welfare results relied on the assumption that there is a representative agent because, in this case, the competitive equilibrium is maximizing this agent's welfare.

In economies with heterogeneous agents and non-aggregable preferences, the social welfare function (SWF) is not unambiguous. In such economies, every competitive equilibrium maximizes a potentially different SWF (see Negishi, 1960). Whereas to a first-order, changes in real output will coincide with changes in the SWF that is being maximized in the initial competitive equilibrium, at higher orders this relation breaks down.

If one commits to a particular SWF, results like the ones we developed in Section 3 can be used to correct the gap between social welfare and real GDP. For example, consider the following utilitarian SWF which weights the welfare of each agent i by their initial share in expenditures,

$$W = \sum_{i \in I} \frac{\bar{p}_i \bar{Y}_i}{\sum_{j \in I} \bar{p}_j \bar{Y}_j} \left[\frac{EV_i^m}{E\bar{V}_i^m} . \right]$$

This SWF is maximized in the initial equilibrium. Since the parameters of this SWF change in response to shocks that redistribute income, it is as if the social planner has unstable preferences. This instability can create a gap between social welfare and GDP even if preferences of individual agents are stable and homothetic.

For this case, in Appendix G we show that under the assumptions in Section 3, to a second-order approximation,

$$\Delta \log W = \Delta \log Y + \frac{1}{2} Cov_{\chi} \left(-\mathbb{E}_b[\Delta \log p], \Delta \log I - \mathbb{E}_b[\Delta \log p] \right),$$

where χ is the initial distribution of expenditures by each agent, and $\mathbb{E}_b[\Delta \log p]$ is the vector of inflation rates and $\Delta \log I$ is the vector of nominal income changes for each household. The change in social welfare is greater than the change in real GDP if changes in real income negatively covary with inflation across households. There is no bias if preferences are aggregable, in which case $\mathbb{E}_b[\Delta \log p]$ is uniform across households, if real income growth is uniform or, more generally, if changes in household-level inflation rates are uncorrelated with changes in real income.

7 Conclusion

In this paper, we provide a characterization of the gap, due to substitution bias, between standard measures of consumption and welfare that appear when preferences are non-homothetic or unstable. We characterize this bias in both partial and general equilibrium, and show that it can be large over long horizons relevant for long-run growth as well as for short-horizons, where expenditure shares at the firm and product-level change rapidly, if demand-driven changes in expenditures covary with prices.

Although our motivation and applications have focused on shocks across time, our results can also be applied to compare welfare across locations in space. Variation in tastes is likely to be even more significant across space than across time (see, e.g. Deaton, 2003 and Argente et al., 2020). Applying our results in a spatial context is an interesting avenue for future work.

References

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

Acemoglu, D. and V. Guerrieri (2008). Capital deepening and nonbalanced economic growth. *Journal of political Economy 116*(3), 467–498.

Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2019). Missing growth from creative destruction. *American Economic Review 109*(8), 2795–2822.

Alder, S., T. Boppart, and A. Müller (2019). A theory of structural change that can fit the data.

Argente, D. O., C.-T. Hsieh, and M. Lee (2020, September). Measuring the cost of living in mexico and the us. Working Paper 27806, National Bureau of Economic Research.

Arkolakis, C. (2016). A unified theory of firm selection and growth. *The Quarterly Journal of Economics 131*(1), 89–155.

Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics (Forthcoming)*.

Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2020, March). Measuring Welfare and Inequality with Incomplete Price Information. NBER Working Papers 26890, National Bureau of Economic Research, Inc.

Auer, R., A. Burstein, S. Lein, and J. Vogel (2021). Unequal expenditure switching: Evidence from switzerland.

Balk, B. M. (1989). Changing consumer preferences and the cost-of-living index: theory and nonparametric expressions. *Journal of Economics 50*(2), 157–169.

Baqaee, D. and E. Farhi (2019a). Networks, barriers, and trade. Technical report, National Bureau of Economic Research.

Baqaee, D. and E. Farhi (2020). Supply and demand in disaggregated keynesian economies with an application to the covid-19 crisis. Technical report, National Bureau of Economic Research.

Baqaee, D. R. and E. Farhi (2019b). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics 135*(1), 105–163.

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

Barrot, J.-N., B. Grassi, and J. Sauvagnat (2020). Sectoral effects of social distancing. *Available at SSRN*.

Basu, S., L. Pascali, F. Schiantarelli, and L. Serven (2012). Productivity and the welfare of nations. Technical report, National Bureau of Economic Research.

Baumol, W. J. (1967). Macroeconomics of unbalanced growth: the anatomy of urban crisis. *The American economic review*, 415–426.

Boehm, C., A. Flaaen, and N. Pandalai-Nayar (2015). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake.

Bonadio, B., Z. Huo, A. A. Levchenko, and N. Pandalai-Nayar (2020). Global supply chains in the pandemic. Technical report.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica 82*(6), 2167–2196.

Buera, F. J. and J. P. Kaboski (2009, 04-05). Can Traditional Theories of Structural Change Fit The Data? *Journal of the European Economic Association 7*(2-3), 469–477.

Buera, F. J., J. P. Kaboski, and R. Rogerson (2015). Skill biased structural change. Technical report, National Bureau of Economic Research.

Cakmakli, C., S. Demiralp, S. Kalemli-Ozcan, S. Yesiltas, and M. A. Yildirim (2020, May). Covid-19 and emerging markets: An epidemiological multi-sector model for a small open economy with an application to turkey. Working Paper 27191, National Bureau of Economic Research.

Carvalho, V. and X. Gabaix (2013). The Great Diversification and its undoing. *The American Economic Review 103*(5), 1697–1727.

Carvalho, V. M. and A. Tahbaz-Salehi (2018). Production networks: A primer.

Cavallo, A. (2020). Inflation with covid consumption baskets. Technical report, National Bureau of Economic Research.

Caves, D. W., L. R. Christensen, and W. E. Diewert (1982). The economic theory of index numbers and the measurement of input, output, and productivity. *Econometrica: Journal of the Econometric Society*, 1393–1414.

Comin, D. A., D. Lashkari, and M. Mestieri (2015). Structural change with long-run income and price effects. Technical report, National Bureau of Economic Research.

Deaton, A. (2003). Adjusted indian poverty estimates for 1999-2000. *Economic and political Weekly*, 322–326.

Diewert, W. E. and K. J. Fox (2020). Measuring real consumption and cpi bias under lockdown conditions. Technical report, National Bureau of Economic Research.

Dixit, A. and V. Norman (1978). Advertising and welfare. *The Bell Journal of Economics*, 1–17.

Domar, E. D. (1961). On the measurement of technological change. *The Economic Journal 71*(284), 709–729.

Eslava, M. and J. C. Haltiwanger (2021). The life-cycle growth of plants: The role of productivity, demand and wedges. Technical report, National Bureau of Economic Research.

Fajgelbaum, P. D. and A. K. Khandelwal (2016). Measuring the unequal gains from trade. *The Quarterly Journal of Economics 131*(3), 1113–1180.

Fally, T. (2020). Integrability, generalized separability, and a new class of demand systems.

Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.

Feenstra, R. C., H. Ma, D. P. Rao, et al. (2009). Consistent comparisons of real incomes across time and space. *Macroeconomic Dynamics 13*(S2), 169–93.

Feenstra, R. C. and M. B. Reinsdorf (2007, October). Should Exact Index Numbers Have Standard Errors? Theory and Application to Asian Growth. In *Hard-to-Measure Goods and Services: Essays in Honor of Zvi*

Griliches, NBER Chapters, pp. 483–513. National Bureau of Economic Research, Inc.

Fisher, F. M. and J. J. McGowan (1979, Autumn). Advertising and Welfare: Comment. *Bell Journal of Economics 10*(2), 726–727.

Fisher, F. M. and K. Shell (1968). *Taste and quality change in the pure theory of the true cost of living index*. Chicago: Aldine Publishing.

Fitzgerald, D., S. Haller, and Y. Yedid-Levi (2016). How exporters grow. Technical report, National Bureau of Economic Research.

Foster, L., J. Haltiwanger, and C. Syverson (2016). The slow growth of new plants: learning about demand? *Economica 83*(329), 91–129.

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

Hamilton, B. W. (2001, June). Using engel's law to estimate cpi bias. *American Economic Review 91*(3), 619–630.

Hausman, J. (1981). Exact consumer's surplus and deadweight loss. *American Economic Review 71*(4), 662–76.

Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. In *The economics of new goods*, pp. 207–248. University of Chicago Press.

Herrendorf, B., R. Rogerson, and A. Valentinyi (2013). Two perspectives on preferences and structural transformation. *American Economic Review 103*(7), 2752–89.

Hicks, J. R. (1940). The valuation of the social income. *Economica 7*(26), 105–124.

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.

Jones, C. I. and P. J. Klenow (2016). Beyond gdp? welfare across countries and time. *American Economic Review 106*(9), 2426–57.

Jorgenson, D. W., M. S. Ho, K. J. Stiroh, et al. (2005). productivity, volume 3: information technology and the american growth resurgence. *MIT Press Books 3*.

Kongsamut, P., S. Rebelo, and D. Xie (2001). Beyond Balanced Growth. *Review of Economic Studies 68*(4), 869–882.

Martin, R. S. (2020). Revisiting taste change in cost-of-living measurement. Technical report, US Department of Labor, US Bureau of Labor Statistics, Office of Prices and Living Conditions.

Negishi, T. (1960). Welfare economics and existence of an equilibrium for a competitive economy. *Metroeconomica 12*(2-3), 92–97.

Ngai, L. R. and C. A. Pissarides (2007). Structural change in a multisector model of growth. *American economic review 97*(1), 429–443.

Nordhaus, W. D. et al. (2008). Baumol's diseases: A macroeconomic perspective. *The BE Journal of Macroeconomics 8*(1), 1–39.

OECD, T. W. Bank, T. U. N. E. C. for Europe, S. O. of the European Communities, and Luxembourg (2004). *Consumer price index manual: Theory and practice*. International Monetary Fund.

Redding, S. J. and D. E. Weinstein (2020). Measuring aggregate price indices with taste shocks: Theory and evidence for ces preferences. *The Quarterly Journal of Economics 135*(1), 503–560.

Samuelson, P. A. and S. Swamy (1974). Invariant economic index numbers and canonical duality: survey and synthesis. *The American Economic Review 64*(4), 566–593.

Syverson, C. (2017). Challenges to mismeasurement explanations for the us productivity slowdown. *Journal of Economic Perspectives 31*(2), 165–86.

Appendix A Proofs

Proof of Lemma 1. By definition,

$$\begin{aligned}
EV^m &= \log \frac{e\left(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1}\right)}{e\left(p_{t_0}, v(p_{t_0}, I_{t_0}; x_{t_1}); x_{t_1}\right)} \\
&= \log \frac{e\left(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1}\right)}{e\left(p_{t_0}, v(p_{t_0}, I_{t_0}; x_{t_1}); x_{t_1}\right)} \frac{e\left(p_{t_1}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1}\right)}{e\left(p_{t_1}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1}\right)} \\
&= \log \frac{e\left(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1}\right)}{I_{t_0}} \frac{I_{t_1}}{e\left(p_{t_1}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1}\right)}.
\end{aligned}$$

To finish, rewrite

$$\log \frac{e\left(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1}\right)}{e\left(p_{t_1}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1}\right)} = -\int_{t_0}^{t_1} \frac{\partial \log e(p, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})}{\partial \log p} \frac{d \log p}{dt} dt,$$

and use the Shephard's lemma to express the price elasticity of the expenditure function in terms of budget shares. If the path of prices between t_0 and t_1 is not differentiable, then construct a new a modified path of prices that is differentiable, and apply the integral to this modified path. Since the integral is path independent, it only depends on p_{t_0} and p_{t_1}. Therefore any path that connects p_{t_0} and p_{t_1} gives the same integral. □

Proof of Proposition 1. If the path of prices is continuously differentiable, we can combine Lemma 1 with the definition of real consumption. □

Proof of Corollary 1. From Proposition 1, given any change in prices, income, and preferences, $EV^m - \Delta \log Y = 0$ if, and only if, $b = b^{ev}$, which is the case if, and only if, budget shares do not depend on the level of utility and x – that is preferences are homothetic and stable. □

Proof of Lemma 2. Differentiate real consumption

$$\Delta \log Y = \int_{t_0}^{t_1} d \log I(t) - \sum_{i \in N} b_i(p(t), u(t); x(t)) \frac{d \log p_i}{dt} dt$$

twice with respect to t_1 and evaluate the derivative at $t_1 = t_0$. This yields the desired expression. □

Proof of Proposition 2. By Lemma 1:

$$EV = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} \frac{\partial \log e(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial \log p} \frac{d \log p}{d \log t} d \log t$$

Differentiate EV twice with respect to t_1 and evaluate the derivative at $t_1 = t_0$

$$\begin{aligned}
\frac{dEV}{dt_1} &= \frac{d \log I}{dt} - \sum_{i \in N} \frac{\partial \log e(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial \log p_i} \frac{d \log p_i}{d \log t} - \\
&\quad - \int_{t_0}^{t_1} \sum_{i \in N} d \log v \frac{\partial^2 \log e(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial \log u \partial \log p_i} d \log p_i \\
&\quad - \int_{t_0}^{t_1} \sum_{i \in N} d \log x \frac{\partial^2 \log e(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial \log x \partial \log p_i} d \log p_i \\
\frac{d^2 EV}{dt_1^2} &= - \sum_{i \in N} b_i \frac{d^2 \log p}{d \log t^2} - \sum_{i \in N} \sum_{j \in N} \frac{\partial^2 \log e(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial \log p_i \partial \log p_j} d \log p_i d \log p_j \\
&\quad - 2 \sum_{i \in N} d \log v \frac{\partial^2 \log e(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial \log p_i \partial \log u} d \log p_i - 2 \sum_{i \in N} d \log x' \frac{\partial \log e(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial \log p_i \partial \log x} d \log p_i \\
&= - \sum_{i \in N} \sum_{j \in N} \frac{\partial b_i}{\partial \log p_i} d \log p_i d \log p_j - \sum_{i \in N} b_i \frac{d^2 \log p}{d \log t^2} \\
&\quad - 2 \sum_{i \in N} d \log v \frac{\partial b_i}{\partial \log u} d \log p_i - 2 \sum_{i \in N} d \log x' \frac{\partial b_i}{\partial \log x} d \log p_i \\
&= - \sum_{i \in N} \left[\sum_{j \in N} \frac{\partial b_i}{\partial \log p_j} d \log p_j + d \log v \frac{\partial b_i}{\partial \log u} + d \log x' \frac{\partial b_i}{\partial \log x} \right] d \log p_i - \sum_{i \in N} b_i \frac{d^2 \log p}{d \log t^2} \\
&\quad - \sum_{i \in N} d \log v \frac{\partial b_i}{\partial \log u} d \log p_i - \sum_{i \in N} d \log x' \frac{\partial b_i}{\partial \log x} d \log p_i \\
&= - \sum_{i \in N} d b_i d \log p_i - \sum_{i \in N} b_i \frac{d^2 \log p}{d \log t^2} - \sum_{i \in N} d \log v \frac{\partial b_i}{\partial \log u} d \log p_i - \sum_{i \in N} d \log x' \frac{\partial b_i}{\partial \log x} d \log p_i
\end{aligned}$$

By Lemma 2, the first two terms are equal to the second-order expansion of $\Delta \log Y$, and the remaining terms are the bias. □

Proof of Proposition 3. By Lemma 2, we have

$$\Delta \log Y \approx \Delta \log I - \sum_i b_i \Delta \log p_i - \frac{1}{2} \sum_i \Delta b_i \Delta \log p_i.$$

Substitute (8) in place of Δb to get the desired expression. For the bias, note that Proposition 1 implies that

$$EV - \Delta \log Y \approx -\frac{1}{2}\sum_i \left[\Delta b_i - \sum_j \frac{\partial b_i^H}{\partial \log p_j}\Delta \log p_j\right]\Delta \log p_i$$

where b^H is the Hicksian budget share (holding fixed utility and demand shifters). Using (8) in place of Δb above and the fact that $\frac{\partial b_i^H}{\partial \log p_i} = (1-\theta_0)b_i(1-b_i)$ for $i = j$ and $\frac{\partial b_i^H}{\partial \log p_j} = \theta_0 b_i b_j$ for $i \neq j$, yields the following

$$\Delta \log EV - \Delta \log Y \approx -\frac{1}{2}\sum_{i\in N}\left[(\varepsilon_i - 1)b_i\left(d\log I - \sum_{j\in N} b_j \Delta \log p_j\right) + b_i \Delta \log x_i\right]\Delta \log p_i,$$

which can be rearranged to give the desired expression. □

Proof of Lemma 3. Setting nominal GDP to be the numeraire, we can write

$$\begin{aligned}
\Delta \log Y &= -\int_{t_0}^{t_1} b' d\log p \\
&= -\int_{t_0}^{t_1} b'\left[-\Psi d\log A - \Psi^F d\log L + \Psi^F d\log \Lambda\right] \\
&= \int_{t_0}^{t_1} b'\Psi d\log A - \int_{t_0}^{t_1} b'\Psi^F [d\log \Lambda - d\log L] \\
&= \int_{t_0}^{t_1} \lambda' d\log A + \int_{t_0}^{t_1} \Lambda' d\log L - \int_{t_0}^{t_1} \Lambda d\log \Lambda \\
&= \int_{t_0}^{t_1} \lambda' d\log A + \int_{t_0}^{t_1} \Lambda' d\log L
\end{aligned}$$

where the second line uses Proposition 7, and we use the fact that Using $\lambda' = b'\Psi$, $\Lambda' = b'\Psi^F$, and $b'\Psi^F d\log \Lambda = \Lambda' d\log \Lambda = 0$ because the factor shares always sum to one: $\sum_{f\in F} \Lambda_f = 1$. □

Proof of Proposition 4. Recall that the macro equivalent variation at final preferences is defined by $EV^M = \phi$, where

$$V\left(A_{t_0}, e^{\phi} L_{t_0}; x_{t_1}\right) = V\left(A_{t_1}, L_{t_1}; x_{t_1}\right)$$

Denote by $p\left(A, L, x\right)$ goods prices under technologies A, factor quantities L, and preferences x. Without loss of generality, we fix income at I. We have $p_{t_1} \equiv p\left(A_{t_1}, L_{t_1}, x_{t_1}\right)$ and

$$v_{t_1} \equiv v\left(p_{t_1}, I; x_{t_1}\right) = V\left(A_{t_1}, L_{t_1}; x_{t_1}\right).$$

Define a hypothetical economy with fictional households that have stable homothetic preferences defined by the expenditure function $e^{ev}(p,u) = e(p, v_{t_1}; x_{t_1})\frac{u}{v_{t_1}}$. Budget shares of this fictional consumer are $b_i^{ev}(p) \equiv \frac{\partial e^{ev}(p,u)}{\partial p_i} = \frac{\partial e\left(p, v_{t_1}; x_{t_1}\right)}{\partial p_i}$. Given any technology vector, in this hypothetical economy we denote the Leontief inverse matrix by Ψ^{ev} and sales shares by λ^{ev}. Given technologies A_t and factor quantities L_t, we denote prices in this hypothetical economy by p_t^{ev}. Changes in prices in this hypothetical economy satisfy

$$d\log p^{ev} = -\Psi^{ev} d\log A + \Psi^{evF} d\log \Lambda^{ev}, \tag{25}$$

where Ψ^{ev} is the fictitious Leontief inverse. Note that $p(A_{t_1}, L_{t_1}, x_{t_1}) = p^{ev}(A_{t_1}, L_{t_1})$ and $p\left(A_{t_0}, e^{\phi} L_{t_0}, x_{t_1}\right) = p^{ev}\left(A_{t_0}, e^{\phi} L_{t_0}\right)$, where we used the fact that $V\left(A_{t_0}, e^{\phi} L_{t_0}; x_{t_1}\right) = v_{t_1}$. We will use the property that, with constant returns to scale, homothetic preferences, and constant income I,

$$p^{ev}(A, aL) = \frac{1}{a} p^{ev}(A, L)$$

for every $a > 0$. Using the previous results,

$$\begin{aligned}
V\left(A_{t_0}, e^{\phi} L_{t_0}; x_{t_1}\right) &= v\left(p\left(A_{t_0}, e^{\phi} L_{t_0}, x_{t_1}\right), I; x_{t_1}\right) \\
&= v\left(p^{ev}\left(A_{t_0}, e^{\phi} L_{t_0}\right), I; x_{t_1}\right) \\
&= v\left(e^{-\phi} p^{ev}\left(A_{t_0}, L_{t_0}\right), I; x_{t_1}\right) \\
&= v\left(p^{ev}\left(A_{t_0}, L_{t_0}\right), e^{\phi} I; x_{t_1}\right),
\end{aligned}$$

where the last equality used the fact that the value function is homogeneous of degree 0 in prices and income. We thus have

$$v\left(p^{ev}\left(A_{t_0}, L_{t_0}\right), e^{\phi} I; x_{t_1}\right) = v\left(p^{ev}\left(A_{t_1}, L_{t_1}\right), I; x_{t_1}\right),$$

which can be re-expressed using the expenditure function as

$$EV^M = \log \frac{e\left(p^{ev}\left(A_{t_1}, L_{t_1}\right), v_{t_1}; x_{t_1}\right)}{e\left(p^{ev}\left(A_{t_0}, L_{t_0}\right), v_{t_1}; x_{t_1}\right)}.$$

As in the proof of Lemma 1, rewrite EV^M as

$$EV^M = -\int_{t_0}^{t_1} \sum_{i \in N+F} \frac{\partial \log e(p, v_{t_1})}{\partial \log p_i} d\log p_i^{ev} = -\int_{t_0}^{t_1} \sum_{i \in N+F} b_i^{ev} d\log p_i^{ev}.$$

Following the same steps as in the proof of Lemma 3 (for the hypothetical economy), we obtain

$$EV^M = \int_{t_0}^{t_1} \sum_{i \in N} \lambda_i^{ev} d\log A_i + \int_{t_0}^{t_1} \sum_{f \in F} \lambda_i^{ev} d\log L_f.$$

□

In general, macro and micro welfare changes are not the same when preferences are unstable and nonhomothetic. However, when the PPF is linear, the following proposition shows that they coincide.

Proposition 10 (Macro vs. Micro Welfare)**.** *Macro and micro welfare changes are equal ($EV^m = EV^M$) if preferences are stable and homothetic, or if factor income shares are constant (as in a one factor economy).*

Proof of Proposition 10. By the proof of Proposition 4, $EV^m = EV^M$ if and only if $p^{ev}(A_t, L_t) = p(A_t, L_t, x_t)$. This condition is immediate if preferences are homothetic and stable. Consider now the case in which preferences are non-homothetic and/or unstable but factor income shares, Λ, are constant. Then by Proposition 7, changes in prices in response to changes in A, L, and x are given by the following differential equation:

$$d\log p = -\Psi d\log A - \Psi^F d\log L.$$

Furthermore, note that changes in Ψ are determined by changes in Ω since $\Psi = (I - \Omega)^{-1}$. Since every $i \in N$ has constant returns to scale, changes in Ω_{ij} depend only on changes in relative prices for every $i \in N$. This means that changes in Ω only depend on changes in relative prices, therefore changes in Ψ depend only on changes in relative prices. Since x and utility v do not appear in any of these expressions, this means that prices and incomes $p(A, L, x)$ and $I(A, L, x)$, relative to the numeraire, do not depend on x and v. Thus, $p^{ev}(A_t, L_t) = p(A_t, L_t, x_t)$.

□

Proof of Lemma 4. Differentiate real GDP,

$$\Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} \lambda_i(A(t); x(t)) \frac{d\log A_i}{dt} dt,$$

twice with respect to t_1 and evaluate the derivative at $t_1 = t_0$. This yields the desired expression. □

Proof of Proposition 5. Following similar steps as in the proof of Proposition 3,

$$EV^M \approx \Delta \log Y + \frac{1}{2} \sum_{i \in N} \left[\Delta \lambda_i - \sum_{j \in N} \frac{\partial \lambda_i^{ev}}{\partial \log A_j} \Delta \log A_j \right] \Delta \log A_i.$$

The term in square brackets is the change in sales shares due to changes in utility and demand shifters. This expression can be written as

$$EV^M \approx \Delta \log Y + \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} + \Delta \log A' \frac{\partial \log v}{\partial \log A} \frac{\partial \lambda_i}{\partial \log v} \right] \Delta \log A_i. \tag{26}$$

□

Proof of Proposition 6. Normalize nominal GDP to one. Applying Proposition 7 to a one-factor model yields

$$d \log p = -\Psi d \log A,$$

so that relative prices do not respond to changes in demand or income.

To solve for $\Delta \log Y$, use Lemma 4 in combination with the expression for $d \log p$ and $d\lambda$ in Proposition 7 in the case of one factor. To solve for EV^M, by Proposition 10, $EV^M = EV^m$. Solve for $EV^m - \Delta \log Y$ by plugging the expression for $d \log p$ into Proposition 2 and noting that $b' = \Omega^{(0)}$. □

Proof of Proposition 7. We normalize nominal GDP to be the numeraire. Then Shephard's lemma implies that, for each $i \in N$

$$d \log p_i = -d \log A_i + \sum_j \Omega_{ij} d \log p_j.$$

Furthermore, for $i \in F$

$$d \log p_i = -d \log A_i + d \log \Lambda_i.$$

Combining these yields the desired expression for changes in prices

$$d \log p = -\Psi d \log A + \Psi^F d \log \Lambda.$$

To get changes in sales shares, note that

$$
\begin{aligned}
\lambda &= b'\Psi \\
d\lambda &= d(b'\Psi) \\
&= b'\Psi d\Omega\Psi + db'\Psi
\end{aligned}
$$

$$
\begin{aligned}
\Omega_{ij} d\log\Omega_{ij} &= (1-\theta_i)\Omega_{ij}(d\log p_j - \sum_k \Omega_{ik} d\log p_k) \\
d\Omega_{ij} &= (1-\theta_i) Cov_{\Omega^{(i)}}(d\log p, I_{(j)}) \\
\sum_j d\Omega_{ij}\Psi_{jk} &= (1-\theta_i) Cov_{\Omega^{(i)}}(d\log p, I_{(j)})\Psi_{jk} \\
&= (1-\theta_i)\sum_j Cov_{\Omega^{(i)}}(d\log p, \Psi_{jk} I_{(j)}) \\
[d\Omega\Psi]_{ik} &= (1-\theta_i) Cov_{\Omega^{(i)}}(d\log p, \Psi_{(k)})
\end{aligned}
$$

Meanwhile

$$
\begin{aligned}
d\log b_i &= (1-\theta_0)\left(d\log p_i - \sum_i b_i d\log p_i\right) + (\varepsilon_i - 1) d\log Y + d\log x_i \\
&= (1-\theta_0) Cov_{\Omega^{(0)}}\left(d\log p, I_{(i)}\right) + Cov_{\Omega^{(0)}}\left(\varepsilon, I_{(i)}\right) d\log Y + Cov_{\Omega^{(0)}}\left(d\log x, I_{(i)}\right) \\
\sum_i db_i \Psi_{ik} &= Cov_{\Omega^{(0)}}\left((1-\theta_0) d\log p + \varepsilon d\log Y + d\log x, \Psi_{(k)}\right)
\end{aligned}
$$

Hence,

$$
d\lambda' = \lambda' d\Omega\Psi + db'\Psi
$$

can be written as

$$
d\lambda_k = \sum_i \lambda_i (1-\theta_i) Cov_{\Omega^{(i)}}(d\log p, \Psi_{(k)}) + Cov_{\Omega^{(0)}}\left(\varepsilon, \Psi_{(k)}\right) d\log Y + Cov_{\Omega^{(0)}}\left(d\log x, \Psi_{(k)}\right).
$$

□

Proof of Proposition 8. Start by setting nominal GDP to be the numeraire. To model the industry-structure, for each industry I, add two new CES aggregators. One buys the good for the household and one buys the good for firms. Let the price of the household aggregator be given by p_I^c and the price of the non-household aggregator be p_I^f and let

p_I be the price of the original industry. Let firm $i's$ share of industry I from household expenditures be b_{iI}. Let the expenditure share of other firms on firm i be s_{iI}. We have

$$\sum_{i\in I} b_{iI} = 1$$
$$\sum_{i\in I} s_{iI} = 1.$$

Let λ_I^c and λ_I^f be sales of industry I to households and firms. Then we have

$$d\lambda_I = d\lambda_I^c + d\lambda_I^f.$$

The sales of an individual firm i in industry I is given by

$$\begin{aligned}
\lambda_i &= b_{iI}\lambda_I^c + s_{iI}\lambda_I^f \\
d\lambda_i &= db_{iI}\lambda_I^c + b_{iI}d\lambda_I^c + ds_{iI}\lambda_I^f + s_{iI}d\lambda_I^f \\
db_{iI} &= Cov_{b_I}(d\log x_{(iI)} + (1-\sigma_I)d\log A_{(iI)}, I_{(i)}) \\
ds_{iI} &= Cov_{b_I}((1-\sigma_I)d\log A_{(iI)}, I_{(i)}).
\end{aligned}$$

The gap between macro welfare and real GDP, $EV^M - \Delta\log Y$, is approximately given by

$$\frac{1}{2}d\log x\frac{\partial\lambda}{\partial\log x}d\log A = \frac{1}{2}\sum_{i\in N}\left[\sum_{j\in N} d\log x_j\frac{\partial\lambda_i}{\partial\log x_j}\right]d\log A_i,$$

where the sums can be re-written as

$$\begin{aligned}
\sum_{i\in N}\left[\sum_{j\in N} d\log x_j\frac{\partial\lambda_i}{\partial\log x_j}\right]d\log A_i = \sum_{i\in N}\Bigg[&d\log x\frac{\partial b_{iI}}{\partial\log x}\lambda_I^c d\log A_i + b_{iI}d\log x\frac{\partial\lambda_I^c}{\partial\log x}d\log A_i \\
&+d\log x\frac{\partial s_{iI}}{\partial\log x}\lambda_I^f d\log A_i + s_{iI}d\log x\frac{\partial\lambda_I^f}{\partial\log x}d\log A_i\Bigg].
\end{aligned}$$

The individual terms of this expression are:

$$\begin{aligned}
\sum_{i\in N}\left[d\log x\frac{\partial b_{iI}}{\partial\log x}\lambda_I^c d\log A_i\right] &= \sum_{i\in N} Cov_{b_I}(d\log x_{(iI)}, I_{(i)})\lambda_I^c d\log A_i \\
&= Cov_{b_I}(d\log x_{(iI)}, \sum_{i\in N} I_{(i)}d\log A_i)\lambda_I^c \\
&= Cov_{b_I}(d\log x_{(iI)}, d\log A_{iI})\lambda_I^c;
\end{aligned}$$

$$\sum_{i\in N}\left[b_{iI} d\log x \frac{\partial \lambda_I^c}{\partial \log x} d\log A_i\right] = \mathbb{E}_{b_I}\left(d\log A_{iI}\right) d\log x \frac{\partial \lambda_I^c}{\partial \log x};$$

$$\sum_{i\in N} d\log x \frac{\partial s_{iI}}{\partial \log x}\lambda_I^f d\log A_i = 0;$$

and

$$\sum_i s_{iI} d\log x \frac{\partial \lambda_I^f}{\partial \log x} d\log A_i = \mathbb{E}_{s_I}\left(d\log A_{iI}\right) d\log x \frac{\partial \lambda_I^f}{\partial \log x}.$$

Of the four terms, two depend on changes on industry-level sales shares (the sum of which is denoted by Θ in the proposition), one of them is zero, and the remaining one (the first term) is the within-industry covariance of supply and demand shocks in the proposition. □

Proof of Proposition 9. Consider a household with preferences given by

$$C = \left(\int_0^{x^*} c(x)^{\frac{\sigma-1}{\sigma}} dx\right)^{\frac{\sigma}{\sigma-1}}.$$

Note that budget shares are

$$\lambda^{ev}(x,t,t_1) = \frac{p(x,t)^{1-\sigma}}{\left(\int_0^{x^*(t_1)} p(x,t)^{1-\sigma} dx\right)}$$

$$\lambda(x,t) = \frac{p(x,t)^{1-\sigma}}{\left(\int_0^{x^*(t)} p(x)^{1-\sigma} dx\right)}$$

Hence

$$EV^m = \frac{I}{\left(\int_0^{x^*(t_1)} p(x)^{1-\sigma} dx\right)^{\frac{1}{1-\sigma}}}.$$

Next

$$\Delta \log EV^m = \Delta \log I - \int_{t_0}^{t_1}\int_0^{x^*(t_1)} \lambda^{ev}(x,t,t_1) \frac{d\log p(x,t)}{dt} dx dt.$$

Without loss of generality, let's normalize changes in nominal income to zero. Let $\partial_i \lambda^{ev}$ refer to the partial derivative of λ^{ev} with respect to its ith argument. Differentiating and evaluating at the initial point, we get

$$\begin{aligned}
\frac{d\log EV^m}{dt_1} &= -\int_{t_0}^{t_1}\int_0^{x^*(t_1)} \partial_3\lambda^{ev}(x,t,t_1)\frac{d\log p(x,t)}{dt}dxdt \\
&\quad -\int_{t_0}^{t_1}\lambda^{ev}(x^*(t_1),t,t_1)\frac{d\log p(x^*(t_1),t)}{dt}\frac{dx^*}{dt_1}dt - \int_0^{x^*(t_1)}\lambda^{ev}(x,t_1,t_1)\frac{d\log p(x,t_1)}{dt_1}dx \\
\frac{d^2\log EV^m}{dt_1^2} &= -\int_0^{x^*(t_1)}\partial_3\lambda^{ev}(x,t_1,t_1)\frac{d\log p(x,t_1)}{dt_1}dx - \lambda^{ev}(x^*(t_1),t_1,t_1)\frac{d\log p(x^*(t_1),t_1)}{dt}\frac{dx^*}{dt_1} \\
&\quad -\lambda^{ev}(x^*(t_1),t_1,t_1)\frac{dx^*}{dt_1}\frac{d\log p(x^*,t_1)}{dt_1} - \int_0^{x^*(t_1)}\frac{d\lambda^{ev}(x,t_1,t_1)}{dt_1}\frac{d\log p(x,t_1)}{dt_1}dx \\
&\quad -\int_0^{x^*(t_1)}\lambda^{ev}(x,t_1,t_1)\frac{d^2\log p(x,t_1)}{dt_1^2}dx
\end{aligned}$$

Evaluating at the initial point this simplifies to

$$\begin{aligned}
\frac{d\log EV^m}{dt_1} &= -\int_0^{x^*}\lambda(x)\frac{d\log p(x,t)}{dt}dx \\
\frac{d^2\log EV^m}{dt_1^2} &= -\int_0^{x^*(t_1)}\partial_3\lambda^{ev}(x,t_1,t_1)\frac{d\log p(x,t_1)}{dt_1}dx - \lambda^{ev}(x^*(t_1),t_1,t_1)\frac{d\log p(x^*(t_1),t_1)}{dt}\frac{dx^*}{dt_1} \\
&\quad -\lambda^{ev}(x^*(t_1),t_1,t_1)\frac{dx^*}{dt_1}\frac{d\log p(x^*,t_1)}{dt_1} - \int_0^{x^*(t_1)}\frac{d\lambda^{ev}(x,t_1,t_1)}{dt_1}\frac{d\log p(x,t_1)}{dt_1}dx \\
&\quad -\int_0^{x^*(t_1)}\lambda^{ev}(x,t_1,t_1)\frac{d^2\log p(x,t_1)}{dt_1^2}dx
\end{aligned}$$

We note that

$$\lambda^{ev}(x,t,t_1) = \frac{p(x,t)^{1-\sigma}}{\left(\int_0^{x^*(t_1)} p(x,t)^{1-\sigma}dx\right)}$$

$$\begin{aligned}
\frac{\partial\log\lambda^{ev}(x,t,t_1)}{\partial t} &= (1-\sigma)\left(\frac{d\log p(x,t)}{dt} - \int_0^{x^*(t_1)}\lambda(x,t)\frac{d\log p(x,t)}{dt}dx\right). \\
\partial_3\log\lambda^{ev}(x,t,t_1) = \frac{\partial\log\lambda^{ev}(x,t,t_1)}{\partial t_1} &= \left(-\lambda(x^*,t)\frac{dx^*}{dt_1}\right).
\end{aligned}$$

Meanwhile, real consumption changes are given by

$$\begin{aligned}
\log Y &= -\int_{t_0}^{t_1}\int_0^{x^*(t)} \lambda(x,t)\frac{d\log p}{dt}dxdt \\
\frac{d\log Y}{dt_1} &= -\int_0^{x^*(t_1)} \lambda(x,t_1)\frac{d\log p}{dt_1}dx \\
\frac{d^2\log Y}{dt_1^2} &= -\lambda(x^*(t_1),t_1)\frac{dx^*}{dt_1}\frac{d\log p}{dt_1} \\
&\quad -\int_0^{x^*(t_1)} \frac{d\lambda(x,t_1)}{dt_1}\frac{d\log p}{dt_1}dx - \int_0^{x^*(t_1)} \lambda(x,t_1)\frac{d^2\log p}{dt_1^2}dx
\end{aligned}$$

where

$$\frac{d\log\lambda(x,t_1)}{dt_1} = (1-\sigma)\left(\frac{d\log p(x,t)}{dt} - \int_0^{x^*(t_1)} \lambda(x,t)\frac{d\log p(x,t)}{dt}dx\right) - \lambda(x^*,t)\frac{dx^*}{dt}.$$

Hence

$$\begin{aligned}
\frac{d\log EV^m}{dt_1} &= \frac{d\log Y}{dt_1} \\
\frac{d^2\log EV^m}{dt_1^2} &= \frac{d^2\log Y}{dt_1^2} - \int_0^{x^*(t_1)} \partial_3\lambda^{ev}(x,t_1,t_1)\frac{d\log p(x,t_1)}{dt_1}dx - \lambda^{ev}(x^*(t_1),t_1,t_1)\frac{d\log p(x^*(t_1),t_1)}{dt}\frac{dx^*}{dt_1} \\
&= \frac{d^2\log Y}{dt_1^2} + \lambda(x^*)\frac{dx^*}{dt_1}\int_0^{x^*(t_1)} \lambda(x)\frac{d\log p(x,t_1)}{dt_1}dx - \lambda(x^*)\frac{d\log p(x^*(t_1),t_1)}{dt}\frac{dx^*}{dt_1} \\
&= \frac{d^2\log Y}{dt_1^2} + \lambda(x^*)\frac{dx^*}{dt_1}\left[\int_0^{x^*(t_1)} \lambda(x)\frac{d\log p(x,t_1)}{dt_1}dx - \frac{d\log p(x^*(t_1),t_1)}{dt}\right] \\
&= \frac{d^2\log Y}{dt_1^2} + \lambda(x^*)\frac{dx^*}{dt_1}\left[\mathbb{E}_\lambda\left[\frac{d\log p}{dt}\right] - \frac{d\log p(x^*)}{dt}\right].
\end{aligned}$$

□

Appendix B Extension to Other Welfare Measures

Our baseline measure of welfare changes is equivalent variation under final preferences. Alternatively, we could measure changes in welfare using compensating (instead of equivalent) variation, or by using initial (rather than final) preferences. We focus on equivalent variation with final preferences since it uses indifference curves in the final allocation to make welfare comparisons (that is, preferences "today" for growth-accounting purposes). In this appendix, we show that our methods generalize to the other welfare measures. If preferences are homothetic, then the expenditure function can be written as $e(p,u;x) = e(p;x)\,u$, so equivalent and compensating variation are equal. If preferences are stable, then the expenditure function can be written as $e(p,u;x) = e(p,u)$, so equivalent variation under initial and final preferences are equal (and the same is the case for compensating variation).

Recall that when preferences are homothetic, then the expenditure function can be written as $e(p,u;x) = e(p;x)\,u$. Hence, in this case, for any fixed x, compensating variation is equal to equivalent variation.

B.1 Micro welfare changes

We consider four alternative measures of micro welfare changes. For each measure, we present expression for global welfare changes and the approximate gap with real consumption.

The *compensating variation* with *initial preferences* is $CV^m(p_{t_0}, I_{t_0}, p_{t_1}, I_{t_1}; x_{t_0}) = \phi$, where ϕ solves

$$v(p_{t_1}, e^{-\phi} I_{t_1}; x_{t_0}) = v(p_{t_0}, I_{t_0}; x_{t_0}). \tag{27}$$

The analog to (7) in Lemma 1 is

$$CV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i^{cv} d \log p_i, \tag{28}$$

where $b_i^{cv}(p) \equiv b_i(p, v(p_{t_0}, I_{t_0}; x_{t_0}); x_{t_0})$.

Whereas EV^m weights price changes by hypothetical budget shares evaluated at current prices for fixed *final preferences and final utility*, CV^m uses budget shares evaluated at current prices for fixed *initial preferences and initial utility*. An alternative way of calculating CV^m is to reverse the flow of time (the final period corresponds to the initial period), calculate the baseline EV measure under this alternative timeline, and then set $CV^m = -EV^m$.

We now briefly describe how to calculate b^{cv} to apply (34). For ex-ante counterfactuals, where $b(t_0)$ is known, we can construct $b^{cv}(p)$ between t_0 and t_1 by iterating on (9) starting at t_0 and going forward to t_1. For ex-post counterfactuals, $b(t_0)$ can be obtained from past data, so we can construct $b^{cv}(p)$ by iterating on (9) starting at t_0 and going forward to t_1.

To a second-order approximation

$$\Delta \log CV^m \approx \Delta \log I - b' \Delta \log p - \frac{1}{2} \sum_{i \in N} \left[\Delta \log p' \frac{\partial b_i}{\partial \log p} \right] \Delta \log p \tag{29}$$

$$\approx \Delta \log Y + \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial b_i}{\partial \log x} + \Delta \log v \frac{\partial b_i}{\partial \log u} \right] \Delta \log p. \tag{30}$$

Recall that changes in budget shares due to non-price factors are multiplied by $1/2$ in real consumption. However, they are multiplied by 0 in CV^m, since CV^m is based on budget shares at initial preferences and initial utility.

Combining (10) and (29), we see that up to a second order approximation,

$$0.5 \left(EV^m + CV^m \right) \approx \Delta \log Y.$$

That is, locally (but not globally) changes in real consumption equal a simple average of equivalent variation under final preferences and compensating variation under initial preferences.

Alternatively, we can measure the change in welfare using the *micro equivalent variation* with *initial preferences*, $EV^m(p_{t_0}, I_{t_0}, p_{t_1}, I_{t_1}; x_{t_0}) = \phi$ where ϕ solves

$$v(p_{t_1}, I_{t_1}; x_{t_0}) = v(p_{t_0}, e^{\phi} I_{t_0}; x_{t_0}). \tag{31}$$

Globally, changes in welfare are

$$EV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i^{ev} d \log p_i, \tag{32}$$

where $b_i^{ev}(p) \equiv b_i(p, v(p_{t_1}, I_{t_1}; x_{t_0}); x_{t_0})$. The gap between changes in welfare and real consumption is, up to a first order approximation,

$$\Delta \log EV^m - \Delta \log Y \approx \frac{1}{2} \sum_{i \in N} \left[-\Delta \log x' \frac{\partial b_i}{\partial \log x} + \Delta \log v \frac{\partial b_i}{\partial \log u} \right] \Delta \log p.$$

Finally, the change in welfare measured using the *micro compensating variation* with *final preferences* is $CV^m(p_{t_0}, I_{t_0}, p_{t_1}, I_{t_1}; x_{t_1}) = \phi$ where ϕ solves

$$v(p_{t_1}, e^{-\phi} I_{t_1}; x_{t_1}) = v(p_{t_0}, I_{t_0}; x_{t_1}). \tag{33}$$

Globally, changes in welfare are given by

$$CV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i^{cv} d \log p_i, \tag{34}$$

where $b_i^{cv}(p) \equiv b_i(p, v(p_{t_0}, I_{t_0}; x_{t_0}); x_{t_1})$. The gap between changes in welfare and real consumption is, up to a first order approximation,

$$\Delta \log CV^m - \Delta \log Y \approx \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial b_i}{\partial \log x} - \Delta \log v \frac{\partial b_i}{\partial \log u} \right] \Delta \log p.$$

Note for EV with initial preferences or CV with final preferences, we must be able to separate demand instability from income effects. For this reason, to compute welfare changes, the elasticities of substitution are not sufficient — we must also know income elasticities or the demand shocks.

Finally, we note that real consumption can be interpreted as representing an alternative measure of welfare, defined as the sum of instantaneous welfare changes using current preferences at each point in time. In particular, real consumption can be written as

$$\Delta \log Y = \int_{t_0}^{t_1} \frac{\partial EV^m(p(t), I(t), p(t), I(t); x(t))}{\partial p_{t_1}} dp(t) + \frac{\partial EV^m(p(t), I(t), p(t), I(t); x(t))}{\partial I_{t_1}} dI(t),$$

where, at every $t \in [t_0, t_1]$, the integrand is the instantaneous welfare change in response to changes in prices and income, dp and dI, measured using equivalent variation with preferences $x(t)$. In contrast to our welfare measures, this measure does not represent the welfare change over a single preference ordering, and is path dependent (it does not only depend only on initial and final income and prices).

B.2 Macro welfare changes

For each alternative micro welfare measure there is a corresponding macro welfare measure. For example, the *macro compensating variation* with *initial preferences* is $CV^M(A_{t_0}, L_{t_0}, A_{t_1}, L_{t_1}; x_{t_0}) = \phi$, where ϕ solves

$$V(A_{t_0}, L_{t_0}; x_{t_0}) = V(A_{t_1}, e^{-\phi} L_{t_1}; x_{t_0}).$$

In words, CV^M is the proportional change in final factor endowments necessary to make a planner with preferences $\succeq_{x_{t_0}}$ indifferent between the initial PPF (A_{t_0}, L_{t_0}) and PPF defined by $(A_{t_1}, e^{-\phi} L_{t_1})$.

Equation (14) in Proposition 4 applies using $\lambda^{cv}(A)$, the sales shares in a fictional economy with the PPF A, L but where consumers have stable homothetic preferences represented by the expenditure function $e^{cv}(p,u) = e(p, v_{t_0}, x_{t_0})\frac{u}{v_{t_0}}$ where $v_{t_0} = v(p_{t_0}, I_{t_0}; x_{t_0})$. Growth accounting for welfare is based on hypothetical sales shares evaluated at current technology but for fixed initial preferences and initial utility. The only information on preferences we need to know is elasticities of substitution at the final allocation. As discussed above, CV^M is equal to $-EV^M$ if we reverse the flow of time.

The gap between changes in welfare and real GDP is, to a second-order approximation (the analog of equation 16 in Proposition 5) is

$$CV^M \approx \Delta \log Y - \frac{1}{2}\sum_{i\in N}\left[\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} + \Delta \log A' \frac{\partial \log v}{\partial \log A}\frac{\partial \lambda_i}{\partial \log v}\right]\Delta \log A_i. \tag{35}$$

We can also define *macro equivalent variation* with *initial preferences*, $EV^M(A_{t_0}, L_{t_0}, A_{t_1}, L_{t_1}; x_{t_0}) = \phi$, where ϕ solves

$$V(A_{t_1}, L_{t_1}; x_{t_0}) = V(A_{t_0}, e^{\phi}L_{t_0}; x_{t_0}).$$

Growth accounting for welfare is based on hypothetical sales shares evaluated at current technology for fixed initial preferences and final utility. In contrast to our previous measures, in order to implement this measure we must know initial demand shifters or income effects. The gap between changes in welfare and real GDP is

$$EV^M \approx \Delta \log Y + \frac{1}{2}\sum_{i\in N}\left[-\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} + \Delta \log A' \frac{\partial \log v}{\partial \log A}\frac{\partial \lambda_i}{\partial \log v}\right]\Delta \log A_i. \tag{36}$$

Finally, define *macro compensating variation* with *final preferences*, $CV^M(A_{t_0}, L_{t_0}, A_{t_1}, L_{t_1}; x_{t_1}) = \phi$, where ϕ solves

$$V(A_{t_0}, L_{t_0}; x_{t_1}) = V(A_{t_1}, e^{-\phi}L_{t_1}; x_{t_1}).$$

Growth accounting for welfare is based on hypothetical sales shares evaluated at current technology for fixed final preferences and initial utility, which requires information on demand shifters or income effects. The gap between changes in welfare and real GDP is

$$CV^M \approx \Delta \log Y + \frac{1}{2}\sum_{i\in N}\left[\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} - \Delta \log A' \frac{\partial \log v}{\partial \log A}\frac{\partial \lambda_i}{\partial \log v}\right]\Delta \log A_i. \tag{37}$$

Appendix C Non-homothetic CES preferences

This appendix provides a derivation of the log-linearized expression (8). Changes in Marshallian budget share are given by

$$\begin{aligned} d\log b_i^M &= d\log p_i - d\log I + \sum_j \varepsilon_{ij}^M d\log p_j + \varepsilon_i^w d\log I + d\log x_i, \\ &= d\log p_i - d\log I + \sum_j \left(\varepsilon_{ij}^H - \varepsilon_i^w b_j\right) d\log p_j + \varepsilon_i^w d\log I + d\log x_i, \end{aligned}$$

where ε^H and ε^M are the Hicksian and Marshallian price elasticities, $\varepsilon^w$are the income elasticities, and $d\log x_i$ is a residual that captures changes in shares not attributed to changes in prices or income. The third line is an application of Slutsky's equation. When preferences are non-homothetic CES, then the Hicksian demand curve can be written as

$$c_i = \gamma_i \left(\frac{p_i}{\sum_j p_j c_j}\right)^{-\theta_0} u^{\xi_i},$$

where γ_i and ξ_i are some parameters. The Hicksian price elasticity for $j \neq i$ is

$$\frac{\partial \log c_i}{\partial \log p_j} = \varepsilon_{ij}^H = \theta_0 \frac{p_j c_j}{I} = \theta_0 b_j.$$

Using this fact and the identity $\varepsilon_{ii}^H = -\sum_{j\neq i} \varepsilon_{ij}^H$, we can rewrite changes in budget shares as

$$\begin{aligned} d\log b_i^M &= \sum_j \left(\varepsilon_{ij}^H - \varepsilon_i^w b_j\right) d\log p_j + d\log p_i + \left(\varepsilon_i^w - 1\right) d\log I + d\log x \\ &= \left(1 - \sum_{j\neq i} \varepsilon_{ij}^H\right) d\log \frac{p_i}{I} + \sum_{j\neq i} \varepsilon_{ij}^H d\log \frac{p_j}{I} + \varepsilon_i^w \left[d\log I - \sum_j b_j d\log p_j\right] + d\log x_i \\ &= \left(1 - \sum_{j\neq i} \theta_0 b_j\right) d\log \frac{p_i}{I} + \sum_{j\neq i} \theta_0 b_j d\log \frac{p_j}{I} + \varepsilon_i^w \left[d\log I - \sum_j b_j d\log p_j\right] + d\log x_i \\ &= \left(1 - \theta_0(1 - b_i)\right) d\log \frac{p_i}{I} + \sum_{j\neq i} \theta_0 b_j d\log \frac{p_j}{I} + \varepsilon_i^w \left[d\log I - \sum_j b_j d\log p_j\right] + d\log x_i \\ &= (1-\theta_0) \left[d\log p_i - \sum_j b_j d\log p_j\right] + \left(\varepsilon_i^w - 1\right) \left[d\log I - \sum_j b_j d\log p_j\right] + d\log x_i. \end{aligned}$$

Appendix D The Role of Elasticities of Substitution for Changes in Welfare-Based Productivity

In this appendix, we show why the results in Table 1 are differentially sensitive to changes in different elasticities of substitution. Combine Propositions 6 and 10 to obtain the following second-order approximation:

$$\Delta \log TFP^{\text{welfare}} \approx \sum_i \lambda_i \Delta \log A_i + \frac{1}{2} \sum_{j \in \{0\}+\mathcal{N}} (\theta_j - 1) \lambda_j Var_{\Omega^{(j)}} \left(\sum_{k \in \mathcal{N}} \Psi_{(k)} \Delta \log A_i \right). \quad (38)$$

The second term is half the sum of changes in Domar weights due to substitution effects (i.e. changes in welfare-relevant sales shares) times the change in productivities. Note that changes in these welfare-relevant sales shares are linear in the microeconomic elasticities of substitution. The importance of some elasticity θ depends on

$$\sum_j \lambda_j Var_{\Omega^{(j)}} \left(\sum_{k \in \mathcal{N}} \Psi_{(k)} \Delta \log A_i \right),$$

where the index j sums over all CES nests whose elasticity of substitution is equal to θ (i.e. all j such that $\theta_j = \theta$). Therefore, elasticities of substitution are relatively more potent if: (1) they control substitution over many nests with high sales shares, or (2) if the nests corresponding to those elasticities are heterogeneously exposed to the productivity shocks.

We compute the coefficients in (38) for our model's various elasticities using the IO table at the end of the sample. The coefficient on $(\theta_0 - 1)$, the elasticity of substitution between agriculture, manufacturing, and services in consumption is only 0.01. This explains why the results in Table 1 are not very sensitive to this elasticity. On the other hand, the coefficient on $(\theta_1 - 1)$, the elasticity across disaggregated consumption goods, is much higher at 0.21. The coefficient on $(\theta_2 - 1)$, the elasticity between materials and value-added bundles is 0.07. Finally, the coefficient on $(\theta_3 - 1)$, the elasticity between disaggregated categories of materials is 0.25. This underscores the fact that elasticities of substitution are more important if they control substitution in CES nests which are very heterogeneously exposed to productivity shocks — that is, nests that have more disaggregated inputs.

According to equation (38), setting $\theta_1 = \theta_2 = \theta_3 = 1$ (which is similar to abstracting from heterogeneity within the three broader sectors and heterogeneity within intermediate inputs), then θ_0 is the only parameter that can generate complementarities in the model. This may help understand why more aggregated models of structural transformation (e.g. Buera et al., 2015 and Alder et al., 2019) require low values of θ_0 to account for the extent of sectoral reallocation in the data.

Appendix E Micro and Macro Welfare in the Covid-19 Application

Table 3 displays welfare changes between January 2020 and May 2020 in the calibrated model of section 5.3. We report separately micro and macro welfare based on initial and final preferences. Recall that micro and macro welfare are not equal in this economy because the PPF is nonlinear (because there are multiple factors). For comparison, we also report the change in real consumption assuming supply and demand shocks arrive simultaneously (as in the last row of Table 3).

The numbers for macro welfare coincide with the changes in real GDP reported in Table 3 under different assumptions on the timing of the demand and supply shocks. Recall that, according to Corollary 4, welfare at initial preferences is equal to real GDP when supply shocks arrive first, and welfare at final preferences is equal to real GDP when demand shocks arrive first and then the supply shocks. Since supply and demand shocks are positively correlated, the decline in welfare is larger under initial preferences than under final preferences.

On the other hand Table 3 shows that the drop in micro welfare is larger under final preferences than under initial preferences. This is because, as shown in our analytic example 4, demand shocks reduce welfare in the presence of decreasing returns to scale (since demand shocks increase the price of goods that consumers value more over time).

Focusing on final preferences, which are more relevant, we see that chained real consumption under-measures welfare losses for the microeconomic change in welfare (comparing initial and final budget sets) and it over-measures welfare losses for the macroeconomic change in welfare (comparing initial and final PPFs). This example also illustrates that micro and macro welfare answer different questions, and the answers to these questions can be quantitatively very different.

Table 3: The change in micro and macro welfare with initial and final preferences given the supply and demand shocks between February 2020 to May 2020. Chained real consumption is computed assuming supply and demand shocks arrive simultaneously.

Elasticities	High compl.	Medium compl.	Cobb-Douglas
Micro initial preferences	-11.7%	-9.1%	-8.7%
Micro final preferences	-13.2%	-12.3%	-10.9%
Macro initial preferences	-16.2%	-12.5%	-10.8%
Macro final preferences	-10.1%	-9.4%	-9.0%
Chained real consumption	-12.1%	-10.6%	-9.8%

Appendix F Non-CES Functional Forms

In this appendix, we generalize Proposition 7 beyond CES functional forms. To do this, for each producer k with cost function $\mathbf{C}_k$, we define the Allen-Uzawa elasticity of substitution between inputs x and y as

$$\theta_k(x,y) = \frac{\mathbf{C}_k d^2\mathbf{C}_k/(dp_x dp_y)}{(d\mathbf{C}_k/dp_x)(d\mathbf{C}_k/dp_y)} = \frac{\epsilon_k(x,y)}{\Omega_{ky}},$$

where $\epsilon_k(x,y)$ is the elasticity of the demand by producer k for input x with respect to the price p_y of input y, and Ω_{ky} is the expenditure share in cost of input y. For the household $k=0$, we use the household's expenditure function in place of the cost function (where the Allen-Uzawa elasticities are disciplined by Hicksian cross-price elasticities and expenditure shares).

Following Baqaee and Farhi (2019c), define the *input-output substitution operator* for producer k as

$$\Phi_k(\Psi_{(i)}, \Psi_{(j)}) = -\sum_{1\leq x,y\leq N+1+F} \Omega_{kx}[\delta_{xy} + \Omega_{ky}(\theta_k(x,y)-1)]\Psi_{xi}\Psi_{yj}, \tag{39}$$

$$= \frac{1}{2}E_{\Omega^{(k)}}\left((\theta_k(x,y)-1)(\Psi_i(x)-\Psi_i(y))(\Psi_j(x)-\Psi_j(y))\right), \tag{40}$$

where δ_{xy} is the Kronecker delta, $\Psi_i(x) = \Psi_{xi}$ and $\Psi_j(x) = \Psi_{xj}$, and the expectation on the second line is over x and y. The second line can be obtained from the first using the symmetry of Allen-Uzawa elasticities of substitution and the homogeneity identity.

Then, Proposition 7 generalizes as follows:

Proposition 11. *Consider some perturbation in final demand $d\log x$ and technology $d\log A$.*

Then changes in prices of goods and factors are

$$d \log p_i = -\sum_{j \in N} \Psi_{ij} d \log A_j + \sum_{f \in F} \Psi^F_{if} d \log \lambda_f. \tag{41}$$

Changes in sales shares for goods and factors are

$$\begin{aligned} \lambda_i d \log \lambda_i = & \sum_{j \in \{0\}+N} \lambda_j \Phi_j \left(-d \log p, \Psi_{(i)}\right) \\ & + Cov_{\Omega^{(0)}} \left(d \log x, \Psi_{(i)}\right) + Cov_{\Omega^{(0)}}(\varepsilon, \Psi_{(i)}) \left(\sum_{k \in N} \lambda_k d \log A_k\right). \end{aligned} \tag{42}$$

Since Φ_j shares many of the same properties as a covariance (it is bilinear and symmetric in its arguments, and is equal to zero whenever one of the arguments is a constant), the intuition for Proposition 11 is very similar to that of Proposition 7. Computing the equilibrium response in Proposition 11 requires solving a linear system exactly as in Proposition 7.

Appendix G Heterogeneous Agents

Consider an economy with a set of agents indexed by h, where each agent has stable and homothetic preferences. This means that real consumption for agent h, $\Delta \log Y_h$, is equal to EV_h. The social welfare function is given by

$$W = \sum_h \frac{\bar{p}_h \bar{Y}_h}{\sum_j \bar{p}_j \bar{Y}_j} \left[\frac{Y_h}{\bar{Y}_h}\right],$$

where bars denote values at the initial equilibrium. We show that to a second order approximation

$$\Delta \log W \approx \Delta \log Y - \frac{1}{2} Cov_\chi(\mathbb{E}_{b_h}[\Delta \log p], \Delta \log Y_h),$$

where χ is the vector of initial expenditure shares for each agent as a share of aggregate spending, and the covariance is applied across individuals.

To see this,

$$\begin{aligned}
dW &= \sum_h \frac{\bar{p}_h Y_h}{\sum_j \bar{p}_j \bar{Y}_j} d\log Y_h, \\
d\log W &= \frac{1}{W} \sum_h \frac{\bar{p}_h Y_h}{\sum_j \bar{p}_j \bar{Y}_j} d\log Y_h, \\
d^2 \log W &= \frac{1}{W} \sum_h \frac{\bar{p}_h Y_h}{\sum_j \bar{p}_j \bar{Y}_j} d^2 \log Y_h + \frac{1}{W} \sum_h \frac{\bar{p}_h Y_h}{\sum_j \bar{p}_j \bar{Y}_j} \left(d\log Y_h\right)^2 - d\log W^2, \\
&= \sum_h \chi_h d^2 \log Y_h + \sum_h \chi_h \left(d\log Y_h\right)^2 - d\log Y^2,
\end{aligned}$$

where the last line uses the fact that $d\log W = d\log Y$ at the initial point.

Next consider the change in real GDP:

$$\log Y = \int_{t_0}^{t_1} \sum_i \frac{p_i(t) q_i(t)}{\sum_j p_j(t) q_j(t)} d\log q_i(t),$$

and

$$p_i q_i = \sum_h \chi_h b_{hi} GDP,$$

where b_{hi} is the budget share of agent h on good i and GDP is nominal GDP. We have

$$\begin{aligned}
d\log q_i &= \sum_h \frac{q_{hi}}{q_i} d\log q_{hi} \\
&= \sum_h \frac{\chi_h b_{hi}}{\sum_g \chi_g b_{gi}} d\log q_{hi} \\
\log Y &= \int^{t_1} \sum_h \chi_h \sum_i b_{hi} d\log q_{hi} \\
&= \int^{t_1} \sum_h \chi_h d\log Y_h. \\
d\log Y &= \chi \cdot d\log Y_{(h)}.
\end{aligned}$$

Therefore, to a second order

$$\begin{aligned}
d^2 \log Y &= d\chi \cdot d\log Y_{(h)} + \chi \cdot d^2 \log Y_{(h)} \\
&= d\chi \cdot d\log Y_{(h)} + d^2 \log W - \sum_h \chi_h \left(d\log Y_h\right)^2 + d\log Y^2.
\end{aligned}$$

Using the fact that

$$d\chi_h = \chi_h d\log\chi_h = \chi_h \left[\mathbb{E}_{b_h}[d\log p] + d\log Y_h - \sum_j \chi_j \left[\mathbb{E}_{b_j}[d\log p] + d\log Y_j\right]\right],$$

we have

$$\begin{aligned}
d^2\log Y &= \sum_h \chi_h \left[\mathbb{E}_{b_h}[d\log p] + d\log Y_h - \sum_j \chi_j \left[\mathbb{E}_{b_j}[d\log p] + d\log Y_j\right]\right] d\log Y_h \\
&\quad + d^2\log W - \sum_h \chi_h \left(d\log Y_h\right)^2 + d\log Y^2 \\
&= \sum_h \chi_h \mathbb{E}_{b_h}[d\log p] d\log Y_h + \sum_h \chi_h d\log Y_h d\log Y_h - \sum_h \chi_h d\log Y_h \sum_j \chi_j \left[\mathbb{E}_{b_j}[d\log p] + d\log Y_j\right] \\
&\quad + d^2\log W - \sum_h \chi_h \left(d\log Y_h\right)^2 + d\log Y^2 \\
&= Cov_\chi(\mathbb{E}_{b_h}[d\log p], d\log Y_{(h)}) + d^2\log W.
\end{aligned}$$

This implies that

$$d^2\log W = d^2\log Y - Cov_\chi(\mathbb{E}_{b_h}[d\log p], d\log Y_{(h)}).$$

To a second order approximation changes in the social welfare function are given by

$$d\log W + \frac{1}{2}d^2\log W = d\log Y + \frac{1}{2}d^2\log Y - \frac{1}{2}Cov_\chi(\mathbb{E}_{b_h}[d\log p], d\log Y_{(h)}).$$