

NBER WORKING PAPER SERIES

WELFARE AND OUTPUT WITH INCOME EFFECTS AND TASTE SHOCKS

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 August 2021

We thank Conor Foley and Sihwan Yang for outstanding research assistance. We thank Andy Atkeson, Natalie Bau, Javier Cravino, Pierre Sarte, David Weinstein, 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 conclusions and analysis are our own, calculated in part on data from Nielsen Consumer LLC and provided through the NielsenIQ Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. NielsenIQ is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by David Baqaee and Ariel Burstein. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

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

ABSTRACT

We characterize how welfare responds to changes in budget sets and technologies when preferences are non-homothetic or subject to shocks, in both partial and general equilibrium. We generalize Hulten's theorem, the basis for constructing aggregate quantity indices, to this context. We show that calculating the response of welfare to a shock only requires knowledge of expenditure shares and elasticities of substitution and (given these elasticities) not of income elasticities and taste shocks. We also characterize the gap between welfare and chain-weighted indices. We apply our results to long- and short-run phenomena. In the long-run, we show that if structural transformation is caused by income effects or changes in tastes, rather than substitution effects, then Baumol's cost disease is twice as important for welfare. In the short-run, we show that standard deflators understate welfare-relevant inflation because product-level demand shocks are positively correlated with price changes. Finally, using the Covid-19 recession we illustrate the differences between partial and general equilibrium notions of welfare, and show that real consumption and real GDP are unreliable metrics for measuring welfare or production.

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

1 Introduction

In this paper, we study how a change in the economic environment affects welfare. For example, how does an individual's welfare change when her budget constraint changes, or how does national welfare change when technologies change? Under some strong assumptions, a chain-weighted index of real consumption, as measured by statistical agencies, answers both of these questions.¹ Two of these assumptions are stability and homotheticity of preferences. Both assumptions are highly convenient, but highly counterfactual. Homotheticity requires that the income elasticity of demand equal one for every good. Stability requires that consumers only change spending between goods in response to changes in incomes and relative prices.²

In this paper, we relax both assumptions and characterize changes in welfare, changes in chained-weighted consumption, and the gap between the two in terms of measurable sufficient statistics. Our baseline welfare measure is the equivalent variation at fixed final preferences, which answers the question: "*holding fixed preferences, how much must consumers' initial endowment change to make them indifferent between their choice sets at t_0 and t_1 ?*" where t can refer to, for example, time or space.

We first study this problem in partial equilibrium, where choice sets are defined in terms of budget sets (prices are exogenous) and the endowment is income. Here, our welfare measure answers a microeconomic question, comparing two budget sets for an infinitesimal agent who does not alter market-level prices through her choices. We then study this problem in general equilibrium, where choice sets are defined in terms of production possibility frontiers (technologies are exogenous) and the endowment is a bundle of factor inputs.³ In this case, our welfare measure answers a macroeconomic question comparing two technologies for a collection of agents whose collective decisions alter market-level

¹To aggregate data on prices and quantities over multiple goods, chain-weighted indices use good-specific weights that are updated every period. As compared to fixed-weight indices, chain-weighted indices account for substitution by consumers. The continuous time analog to a chain-weighted index is called a *Divisia* index. Chained-weighted indices are used to calculate most measures of real economic activity and price deflators, ranging from aggregates like output (real GDP), total factor productivity, private consumption and investment, to less aggregated objects like industry-level measures of production and inflation. The fact that under suitable assumptions these indices approximately measure changes in welfare and production justifies their recommended use in the United Nations' System of National Accounts (see e.g. Chapters 15 and 17 in IMF, 2004).

²As 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. aging, illness, advertising, and 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 chain-weighted measures of real consumption. Our results are relevant when either of these forces is active.

³For the macro problem, we consider neoclassical economies with representative (or homogeneous) agents. We show how to generalize all of our results to economies with heterogeneous agents in a companion paper, Baqaee and Burstein (2021).

prices. When preferences are homothetic and stable, the change in macro and microeconomic welfare are both equal to changes in chain-weighted real consumption. However, when preferences are non-homothetic or unstable, the change in macro and micro welfare caused by the same primitive shock are not equal to each other and neither of them is measured by chain-weighted real consumption. Intuitively, macro and micro welfare differ in general because 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.

We provide exact and approximate characterizations of the change in micro and macro welfare, and generalize chain-weighted (or more precisely Divisia) indices to measure both. In particular, we extend Hulten (1978) to environments with non-homotheticities and taste shocks. In contrast to a chain-weighted consumption index, which weighs changes in prices or technologies using observed expenditure shares, welfare-relevant indices weigh changes in prices or technologies using counterfactual expenditure shares calculated under the final indifference curve. We show that for this reason, compared to welfare, chain-weighted consumption *undercounts* expenditure-switching due to income effects or taste shocks (but not substitution effects).

To understand why, consider the following example. Over the post-war period, spending on healthcare grew relative to manufacturing. Suppose this was caused by consumers getting older and richer, and the fact that older and richer consumers spend more on healthcare. If this were the case, then a chain-weighted consumption index would not correctly account for expenditure-switching by consumers. Intuitively, when we compare the past to the present, we must use demand curves that are relevant for the older and richer consumers of today, and not demand curves that were relevant in the past. Whereas a chained deflator weighs changes in prices that happened during the 1950's using demand from the 1950's, a welfare-relevant index uses demand from today to weigh changes in prices throughout the sample. We show that a chain-weighted consumption index differs from a welfare-relevant index if income- or taste-driven expenditure-switching is correlated with changes in relative prices.

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 does not require direct knowledge of the taste shocks or income elasticities. Instead, what we must know are expenditure shares and elasticities of substitution at the final allocation. For micro welfare, these are the household's expenditure shares and elasticities of substitution. For macro welfare, these are the input-output table and elasticities of substitution in both production and consumption.

Our results can be used both for ex-post accounting and ex-ante counterfactuals.

For very simple economies with one factor, constant returns to scale, and no intermediates, the difference between welfare and chain-weighted consumption is approximately half the covariance of supply and demand shocks. We generalize this formula to more complex economies and 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 chain-weighted consumption. The discrepancies between welfare and chain-weighted consumption that we emphasize do not get “aggregated” away. In fact, the more we disaggregate, the more important these discrepancies are likely to become. 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, 2019b).

We illustrate the relevance of our results for understanding long-run and short-run phenomena with three applications.

- i. **Long-run application:** Since Baumol (1967), an enduring stylized fact is 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 1947 to 2014, aggregate TFP in the US grew by 60%. If the US economy had kept its original 1947 industrial structure, then aggregate TFP would have grown by 78% instead. We show that if this transformation is caused solely by income effects and demand instability, then welfare-relevant TFP grew by only 47% instead of 60%. We also find a similar pattern in consumption data.
- ii. **Short-run application:** While our first application focuses on long-run patterns, our second application shows that gaps between real consumption and welfare are likely present at high frequencies too. Whereas industry-level sales shares are relatively stable over short-horizons, firm or product-level sales shares are not. In 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. These biases do not disappear as we aggregate up to the level of real GDP even if firms and industries are infinitesimal. When we work with industry-level (rather than disaggregated firm- or product-level) data, we rule out the existence of these biases by assumption. We quantify these biases at the industry level using product-level non-durable consumer goods data between 2004 and 2019. We find that standard price indices, like the Sato-Vartia index

and chained-weighted price index, understate the welfare-relevant inflation rate by around 0.5% between 2018 and 2019, and this bias grows to 4% over the whole sample.

- iii. **Business-cycle application:** Our final application draws on the Covid-19 recession to illustrate the difference between macroeconomic and microeconomic notions of welfare. During this recession, household expenditures switched to favor certain sectors at the same time that those sectors experienced higher inflation. We show that this implies that microeconomic welfare, taking changes in prices as given, fell by more than macroeconomic welfare, taking into account the fact that changes in prices are themselves caused by demand shocks. Furthermore, real consumption failed to measure either object. This is because in episodes where spending patterns are partly driven by taste shocks, as in the Covid-19 recession, changes in real consumption generically depend on irrelevant details like the order in which supply and demand shocks hit the economy. In these circumstances, the change in real consumption between two time periods is not a function of only prices and quantities in those two periods. Real consumption can be different between the initial and final periods even if initial and final prices and quantities are the same. The same logic applies to real GDP, which means that real GDP or TFP are unreliable metrics for measuring changes in productive capacity.

Of course, there are other reasons, besides instability and non-homotheticity of preferences, why chained indices fail to accurately measure welfare. Many of these well-known reasons can be thought of as being due to missing prices and quantities. For example, it is well-known that real consumption 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); real consumption does not properly account for changes in the quality of goods (see Syverson, 2017); and, real consumption fails to properly account for changes in non-market components of welfare, like changes in the user-cost of durable consumption or leisure and mortality (see Jones and Klenow, 2016). 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. The biases caused by non-homotheticities and taste shocks are different in the sense that they are not caused by mismeasurement of market prices. For this reason, we abstract from these important mismeasurement 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.

Relatedly, taste shocks 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, even holding prices and incomes fixed. However, while they have similar implications for changes in observed 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*. In the case of taste shocks (or non-homotheticities), the gap between welfare and real consumption is caused by a difference between measured and welfare-relevant *expenditure shares*.

Other related literature. Measuring changes in welfare using a money metric when there are income effects is standard in microeconomic theory (see, e.g. chapter 7 in Deaton and Muellbauer, 1980).⁴ We characterize the gap between this notion of welfare and real consumption with non-homotheticities and taste shocks. We then extend this to a general equilibrium context, and develop welfare-relevant growth accounting. A standard assumption in growth accounting is the existence of a stable and homothetic final aggregator. We generalize welfare-relevant growth-accounting to environments where preferences are neither homothetic nor stable, and provide exact and approximate characterizations of how welfare responds to shocks in general equilibrium (extending Domar 1961, and Hulten 1978). This is an issue of central importance in the literature on disaggregated and production network models (see, for example, Carvalho and Tahbaz-Salehi, 2018 and the references therein).⁵

A recent and related paper is Redding and Weinstein (2020), who show that variations in sales are difficult to explain via shifts in supply curves alone, and shifts in demand curves (i.e. taste shocks) are an important source of variation in the data. Their approach to evaluate welfare changes in the presence of taste shocks contrasts with ours because, 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 de-

⁴An alternative to the money-metric approach, which we do not pursue, is axiomatic index theory, which postulates axioms that an ideal index should satisfy, and then finds functional forms that satisfy those axioms (see Chapter 16 of IMF, 2004).

⁵The biases we identify, and the failure of Hulten's theorem, are not caused by inefficiencies (e.g. markups, wedges, taxes). Baqaee and Farhi (2019a) 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.

termines 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 that 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 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. Moreover, as mentioned above, in practical terms the adjustments we derive require the use of counterfactual expenditure shares and not counterfactual *taste-adjusted* prices. We compare the two approaches in more detail in Appendix E.⁶

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 shifts in demand curves caused by income effects— 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), or taste shocks— households spend more of their income on some goods as they become older (Cravino et al., 2019). Our results suggest that structural transformation driven by relative price changes has different welfare implications than structural transformation driven by non-homotheticity or taste shocks.⁷

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. In Section 3, we set up the macroeconomic general

⁶Given 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 Balk (1989) who discusses various ways one can define changes in the cost of living, 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, Tornqvist-based indices correspond to a geometric average of welfare changes under initial and final preferences. We characterize welfare (in partial and general equilibrium) at either initial or final preferences and using either EV or CV. In contrast to Tornqvist and Sato-Vartia, studied by Caves et al. (1982) and Feenstra and Reinsdorf (2007), Divisia indices cannot generically be interpreted as corresponding to any mixture of well-defined preferences. This is because, as we discuss in Section 5, Divisia (or chained) indices are path-dependent, so they can violate basic properties like assigning a higher value to a strictly larger budget set.

⁷For 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).

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. This section helps build intuition for Section 3, where we model the equilibrium determination of prices.

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 $c \in \mathbb{R}^N$, where N is the number 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 $u(c; x)$. Since the consumer makes no choices over x , preferences over x , if they exist, are not revealed by choices. Hence, whereas $u(c; x) > u(c'; x)$ implies that \succeq_x prefers c to c' , a comparison of $u(c; x')$ and $u(c; x)$ is not necessarily meaningful and may not encode any information.⁸

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

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

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

⁸In 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 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). The vector p includes all relevant prices, and if \succeq_x is intertemporal, then p includes the path of current and future prices.⁹

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, but it could equally refer to space. 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

$$v(p_{t_1}, I_{t_1}; x_{t_1}) = v(p_{t_0}, e^\phi I_{t_0}; x_{t_1}). \quad (1)$$

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 (p_{t_1}, I_{t_1}) . The new budget set is preferred to the initial one, if and only if, EV^m is positive. We focus on $\succeq_{x_{t_1}}$ since, for intertemporal comparisons (as opposed to interregional comparisons), today's preferences are more relevant than the those of the past. Hence, if t_1 represents the present and t_0 represents the past, then $\succeq_{x_{t_1}}$ is more relevant than $\succeq_{x_{t_0}}$. See Remark 2 and Appendix C where we show that our methods readily extend to $\succeq_{x_{t_0}}$ and compensating (as opposed to equivalent) variation. Finally, the superscript m in EV^m represents the fact that this is the *micro* equivalent variation, since we take prices as given.

⁹We explicitly discuss how to apply our results in dynamic economies in Section 4.3.

Discussion of our welfare criterion. Following Fisher and Shell (1968), the welfare criterion in Definition 3 measures the change in welfare by presenting the consumer with a hypothetical choice holding fixed their preferences. 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 watching cartoons as a child to the amount of utils that individual derives from drinking coffee as an adult. 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.” However, if we fix the consumer’s age x , we can meaningfully compare the consumer’s choices about budget sets they faced at different points in their life or 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. Quality adjustments are applicable to situations where the consumer can conceivably make choices between the good at differing 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. Our approach of holding x constant allows us to study welfare in situations where, either for practical or philosophical reasons, it is not possible to model preferences over x itself.

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 .

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

$$\Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} b_{it} \frac{d \log c_{it}}{dt} dt = \int_{t_0}^{t_1} \sum_{i \in N} b_i d \log c_i, \quad (2)$$

where $b_{it} \equiv p_{it}c_{it}/I_t$ is the budget share of good i given prices, income, and preferences at time t .¹⁰

The last equation on the right-hand side suppresses dependence on t in the integral. We sometimes use this convention to simplify notation. Equation (2) is called a *Divisia* quantity

¹⁰For any variable z , we denote by dz its change over infinitesimal time intervals, so that $\Delta z = \int_{t_0}^{t_1} dz$.

index. In practice, since perfect data is not available in continuous time, statistical agencies approximate this integral via a (Riemann) sum using chained indices. We abstract from the imperfections of these discrete time approximations in this paper.¹¹ Moreover, we assume that the data on prices and quantities is *perfect* — completely accurate, comprehensive, and adjusted for any necessary quality changes. This is because the important and well-studied 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.

Define the expenditure function for any value of x by

$$e(p, u; x) = \min_c \left\{ \sum_{i \in N} p_i c_i : u(c; x) = u \right\}.$$

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

$$b_i(p, u; x) \equiv \frac{p_i c_i(p, u; x)}{e(p, u; x)} = \frac{\partial \log e(p, u; x)}{\partial \log p_i}, \quad (3)$$

where the second equality is Shephard's lemma. Using the budget constraint, real consumption in (2) can be expressed in terms of changes in nominal income deflated by price changes:

$$\Delta \log Y = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i(p_t, u_t, x_t) \frac{d \log p_{it}}{dt} dt. \quad (4)$$

In 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 preferences 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.

¹¹In discrete time, one can approximate this Riemann integral in different ways. For example, we can use left-Riemann sums (Chained Laspeyres), right-Riemann sums (Chained Paasche), or mid-point Riemann sums (Chained Tornqvist or Fisher). In continuous time, all of these procedures are equivalent and yield the same answer.

Global results. We start by expressing changes in welfare in terms of changes in prices and expenditure shares.

Lemma 1 (Micro Welfare). *For any smooth path of prices, income, and tastes that unfold as a function of time t , micro welfare changes are given by*

$$EV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i(p_t, u_{t_1}, x_{t_1}) \frac{d \log p_{it}}{dt} dt, \quad (5)$$

where $b_i(p_t, u_{t_1}, x_{t_1}) = b_i(p_t, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})$ denotes budget shares at prices p_t , but fixing final preferences x_{t_1} and final utility $u_{t_1} = v(p_{t_1}, I_{t_1}; x_{t_1})$.

Comparing (5) and (4) clarifies the differences between welfare and real consumption. Real consumption weighs changes in prices at time t by observed budget shares at time t . In contrast, welfare weighs changes in prices at time t by some hypothetical budget shares (at fixed utility and tastes). Intuitively, EV^m depends only on the terminal value of x_{t_1} because this is the only value of x that appears in the definition of EV^m . Moreover, EV^m depends on budget shares evaluated at final utility, u_{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 . For instance, 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 .¹²

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 (5).¹³

¹²When there are no taste shocks, real consumption, defined by (2), is a multi-good version of the change in consumer surplus, which is the area under the Marshallian demand curve. Similarly, by equation (5), welfare is the area under a Hicksian demand curve. Hence, in a partial equilibrium context with stable preferences, the gap between real consumption and welfare is also the gap between consumer surplus and welfare, studied by Hausman (1981) and McKenzie and Pearce (1982) amongst others. This equivalence does not hold when preferences are unstable (since Marshallian consumer surplus is not the same as chained real consumption) or in general equilibrium (since micro and macro welfare are not the same, as we discuss in Section 4).

¹³By definition, EV^m only depends on initial and final prices and income, given t_1 preferences. By the gradient theorem for line integrals, the integral in (5) 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 path of prices over time, which as described in the text is assumed to be available in continuous time.

To make the notation more compact, denote the *welfare-relevant* budget shares in Lemma 1 by

$$b_i^{ev}(p_t) \equiv b_i(p_t, u_{t_1}, x_{t_1}).$$

We can reinterpret these hypothetical budget shares $b^{ev}(p)$ as corresponding to those of a fictional consumer with homothetic and stable preferences with expenditure function $e^{ev}(p, u) = e(p, u_{t_1}; x_{t_1}) \frac{u}{u_{t_1}}$, where $u_{t_1} = v(p_{t_1}, I_{t_1}; x_{t_1})$. 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 this fictional consumer has homothetic and stable preferences, so 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 the terminal elasticities of substitution, as discussed in the following remark.

Remark 1 (Non-homothetic CES preferences). To illustrate how Lemma 1 can be used, consider a non-homothetic CES example as in Comin et al. (2015) or Fally (2020). For this demand system, the following equation pins down changes in budget shares at time t :

$$d \log b_{it} = \underbrace{[1 - \theta_0] [d \log p_{it} - \mathbb{E}_{b_t}(d \log p_t)]}_{\text{substitution effects}} + \underbrace{[\varepsilon_{it} - 1] [d \log I_t - \mathbb{E}_{b_t}(d \log p_t)]}_{\text{income effects}} + \underbrace{d \log x_{it}}_{\text{taste shock}}, \quad (6)$$

where $\mathbb{E}_b(\cdot)$ is a weighted average using budget shares in place of “probability” weights. The elasticity θ_0 is the (constant utility) elasticity of substitution across goods and ε_{it} is the income elasticity of good i . The term $d \log x_{it}$ is a demand shifter (i.e. a taste shock), a residual that captures changes in expenditure shares not attributable to changes in income or prices. Note that when ε_{it} is equal to 1 for every i and t , final demand is homothetic, and when x_{it} is constant for all i and t , final demand is stable.¹⁴

If we know b_{t_1} , we can construct the welfare-relevant budget shares $b^{ev}(p)$ between t_0 and t_1 by iterating on

$$d \log b_{it}^{ev} = [1 - \theta_0] [d \log p_{it} - \mathbb{E}_{b_t^{ev}}(d \log p_t)], \quad (7)$$

starting at t_1 with initial value $b_{t_1}^{ev} = b_{t_1}$ and going back to t_0 . These are changes in budget shares which are only due to substitution effects, and hence omit the last two terms in equation (6). Given the path of b_t^{ev} , we can then apply Lemma 1. For non-homothetic CES,

¹⁴Since b_{it} are expenditure shares that always add up to one, it must be that $\mathbb{E}_{b_t}(d \log x_t) = 0$ and $\mathbb{E}_{b_t}(\varepsilon_t) = 1$. See Appendix D for a derivation and mapping between ε and $d \log x$ and primitive preference parameters.

the integral in Lemma 1 has a closed form solution¹⁵

$$EV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i^{ev}(p_t) \frac{d \log p_{it}}{dt} dt = \Delta \log I + \log \left(\sum_i b_{it_1} \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}. \quad (8)$$

This shows that the income elasticities and taste shocks are not directly required.^{16,17}

If b_{t_1} is not known, we first have to predict b_{t_1} by iterating on equation (6) from t_0 to t_1 to obtain b_{t_1} . This first step requires full knowledge of demand shocks and income elasticities over time (see Appendix D for more details). Once in possession of b_{t_1} , apply (8) to get the change in welfare.

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 initial budget shares and elasticities of substitution, whereas equivalent variation requires knowledge of final budget shares and elasticities of substitution. This means that EV^m is more convenient for ex-post comparisons and CV^m (at initial preferences) is more convenient for ex-ante comparisons or counterfactuals. This is because in these cases, we use “today’s” budget shares and elasticities of substitution to undertake the welfare comparisons (without needing knowledge of taste shocks or income elasticities). See Appendix C for more details.¹⁸

¹⁵In Appendix D.2, we show that when preferences are non-homothetic CES, changes in the utility index are not the same as changes in equivalent (or compensating) variation. Hence, the non-homothetic CES utility index is not a money-metric for welfare.

¹⁶In practice, estimating the elasticity of substitution θ_0 may require knowing the income elasticities (via Slutsky’s equation). 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 price elasticities and apply Lemma 1 to measure the heterogeneous welfare effects of changes in foreign prices in the presence of demand non-homotheticities.

¹⁷The result that only terminal elasticities of substitution are necessary to calculate EV^m is true for arbitrary non-CES functional forms, but since the intuition for the more general case is similar to the CES case, we leave the more general non-parametric results in Appendix I. We use non-homothetic CES in our worked-out examples since it provides a clean separation between substitution elasticities (necessary for computing welfare) and other parameters of the utility function, in contrast to other commonly used non-homothetic demand systems such as PIGL and AIDS. Furthermore, substitution elasticities are symmetric and constant for CES, which also helps keep the examples intuitive.

¹⁸In Appendix C 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.

We now contrast changes in real consumption and welfare.

Proposition 1 (Consumption vs. Welfare). *Given a smooth path of prices, income, and tastes that unfold as a function of time t , the difference between welfare changes and real consumption is*

$$EV^m - \Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} (b_{it} - b_{it}^{ev}) \frac{d \log p_{it}}{dt} dt = (t_1 - t_0) \mathbb{E}_t \text{Cov}(b_t - b_t^{ev}, d \log p_t),$$

where the covariance is calculated across goods at a point in time, and the average 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, $b_{it}^{ev} = b_{it}$ for every path of shocks and every t .

To gain more intuition for the gap between welfare and real consumption, we use a second-order approximation.

Local results. Consider local approximations of the objects of interest as the time period goes to zero, $t_1 - t_0 = \Delta t \rightarrow 0$.¹⁹ 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.

Differentiating (4) twice, and evaluating at t_0 , implies that the change in real consumption around t_0 is approximately

$$\Delta \log Y \approx \underbrace{\Delta \log I - \mathbb{E}_b(\Delta \log p)}_{\text{first-order}} - \underbrace{\frac{1}{2} \text{Cov}_b(\Delta \log b, \Delta \log p)}_{\text{second-order}}, \quad (9)$$

where $\mathbb{E}_b(\cdot)$ and $\text{Cov}_b(\cdot)$ are evaluated using budget shares at t_0 as probability weights. The first-order term is just the change in nominal income deflated by average prices (where the average uses budget shares at the point of linearization). The second-order terms depend on how expenditures change in response to the shock, and these changes in expenditures can be driven by either substitution effects, income effects, or taste shocks.

¹⁹For our local approximations, we assume that the exogenous parameters (prices, income, and taste shifters) are smooth functions of t and that the expenditure function is a smooth function of x .

To make the relationship between real consumption and welfare more concrete, we use the non-homothetic CES aggregator introduced in Remark 1.²⁰

Proposition 2 (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*

$$\begin{aligned} \Delta \log Y \approx & \Delta \log I - \mathbb{E}_b(\Delta \log p) - \underbrace{\frac{1}{2}(1 - \theta_0) \text{Var}_b(\Delta \log p)}_{\substack{\text{expenditure-switching} \\ \text{due to substitution effect}}} \\ & - \underbrace{\frac{1}{2} \text{Cov}_b(\Delta \log x, \Delta \log p)}_{\substack{\text{expenditure-switching} \\ \text{due to taste shock}}} - \underbrace{\frac{1}{2} [\Delta \log I - \mathbb{E}_b(\Delta \log p)] \text{Cov}_b(\varepsilon, \Delta \log p)}_{\substack{\text{expenditure-switching} \\ \text{due to income effect}}}, \end{aligned} \quad (10)$$

and the change in welfare is

$$\begin{aligned} EV^m \approx & \Delta \log I - \mathbb{E}_b(\Delta \log p) - \frac{1}{2}(1 - \theta_0) \text{Var}_b(\Delta \log p) \\ & - \text{Cov}_b(\Delta \log x, \Delta \log p) - [\Delta \log I - \mathbb{E}_b(\Delta \log p)] \text{Cov}_b(\varepsilon, \Delta \log p), \end{aligned} \quad (11)$$

where $\mathbb{E}_b(\cdot)$, $\text{Var}_b(\cdot)$, and $\text{Cov}_b(\cdot)$ are evaluated using budget shares at t_0 as probability weights.

We begin by considering the change in real consumption in (10), which rewrites the nonlinear terms in (9) in terms of primitives. 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 welfare is convex in prices and variance in price changes boosts welfare by raising the expenditure share of goods that become relatively cheap. The second line of (10) captures the effect of income effects and demand shocks. If the composition of demand shifts in favor of goods that happen to become relatively cheap, either due to income effects $\text{Cov}_b(\varepsilon, \Delta \log p) (\Delta \log I - \mathbb{E}_b(\Delta \log p)) < 0$ or demand shocks $\text{Cov}_b(\Delta \log x, \Delta \log p) < 0$, then real consumption increases.

Now consider changes in welfare in (11). The first-order terms are identical to real consumption, but discrepancies are present at the second order. In particular, welfare places a larger weight on changes in expenditure shares that occurred due to income effects and taste shocks. Whereas changes in real consumption only take into consideration changes in expenditures as the changes unfold over time, changes in welfare account for changes in expenditure shares due to income effects and taste shocks from the start. Therefore,

²⁰For a more elaborate discussion of Proposition 2 without imposing non-homothetic CES, see Proposition 10 in Appendix A. The intuition remains similar.

changes in budget shares due to income and tastes are multiplied by 1/2 in real consumption, but they are multiplied by 1 in welfare. This implies that, for example, the increase in welfare from a price reduction in a good i with increasing demand (due to an increase in x_i or a relatively high ε_i) is not fully reflected in real consumption, implying $EV^m > \Delta \log Y$. If preferences are stable and homothetic, then welfare changes coincide with changes in real consumption. Furthermore, even if preferences are unstable or non-homothetic, real consumption strays from welfare only when price changes covary with non-price changes in demand.²¹

In Appendix E we extend Proposition 2 to incorporate unobserved changes in quality. We show that the biases caused by non-homotheticities and taste shocks are very different to the ones caused by quality changes. We also explicitly compare the biases that we study to the ones discussed in Redding and Weinstein (2020).

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 are assumed to be exogenous. 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).

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.

²¹In Remark 1 we pointed out that, starting at b_{t_1} , computing welfare does not require knowledge of income elasticities or taste shocks if we know the elasticities of substitution. However, the approximation in (11) depends on income elasticities and taste shocks. The reason is because this approximation is around initial budget shares b_{t_0} . If we start with budget shares at t_1 , we get

$$EV^m \approx \Delta \log I - \mathbb{E}_b(\Delta \log p) - \frac{1}{2}(1 - \theta_0)Var_b(\Delta \log p),$$

where $\mathbb{E}_b(\cdot)$ and $Var_b(\cdot)$ are evaluated using budget shares at t_1 . Hence, starting at the terminal budget shares, EV^m only depends on substitution effects as in Remark 1. Both expressions are valid second-order approximations and in either case, real consumption undercounts expenditure-switching caused by income effects or taste shocks.

3.1 Definition of Welfare and Real GDP

Consider a perfectly competitive neoclassical closed economy with a representative agent.²² 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 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 .

Let A be the $N \times 1$ vector of technology shifters and L be the $F \times 1$ vector of primary (exogenously given) 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 . Given our assumption that production functions have constant returns to scale, the PPF is linear if there is only one factor of production.

For each A , L , and x , we denote equilibrium prices and aggregate income by $p(A, L, x)$ 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 solution to the following planning problem:

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

This is the maximum amount of utility the economy can deliver given technologies (A, L) and preferences \succeq_x . 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 only 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.

Consider shifts in the PPF as technologies and factor endowments change from $(A_{t_0},$

²²For brevity, we use the representative (or homogeneous) agent assumption in this paper since social welfare is unambiguous in this case. A companion paper, Baqaee and Burstein (2021), shows how to generalize the results in this paper to models with heterogeneous agents. See Section 6 for more details.

L_{t_0}) to (A_{t_1}, L_{t_1}) , along with changes in preferences from x_{t_0} to x_{t_1} . We generalize our microeconomic 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, EV^M is the proportional change in initial factor endowments required to make the 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, defined by (A_{t_1}, L_{t_1}) . 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.²³ Moreover, as we show below, measuring welfare changes in terms of factor endowments results in a Hulten-type expression for EV^M . As we discuss below, EV^M and EV^m always coincide if the PPF is linear or if preferences are homothetic and stable.

Difference between macroeconomic and microeconomic welfare. Macro EV^M and micro EV^m welfare changes are different because they answer different questions. Consider a situation where preferences change between t_0 and t_1 but technologies and factor endowments do not. Since the PPF is unchanged, the change in macro welfare is zero by construction. However, if the PPF is nonlinear, the relative price of goods does change between t_0 and t_1 : prices rise for those goods that become more desirable. In this case, microeconomic welfare, at final preferences, falls even though the PPF is unchanged. Hence, micro welfare EV^m is a poor guide for measuring technological change for a society. For more details, see Example 4 below.

Now consider an example where both preferences and technologies change. More concretely, suppose we are interested in measuring technological progress in an economy with growth and aging. Households are richer and older in t_1 than in t_0 , so they prefer to spend more of their income on healthcare. Microeconomic welfare is measured by the endowment a single consumer, living in t_1 , would have to be given at t_0 to make her willing to go back to the economy with t_0 prices. However, since in t_0 households were poor and young, healthcare services are relatively cheap in t_0 . This makes t_0 prices seem very attractive to the consumer in t_1 . But this is not because technologies in t_0 are any better. If the older and

²³When preferences are stable and homothetic, EV^M is the same as consumption equivalents, but we do not define welfare changes in terms of consumption equivalents because when preferences are non-homothetic or unstable, households' desired consumption bundle is not stable.

wealthy consumers were transported to the t_0 economy, the fact that they demand more healthcare would raise healthcare prices in general equilibrium, 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 technology is deceptive, since the initial budget set reflects both the technologies and demand in t_0 . 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 endowments 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 base prices.

When relative prices do not respond to consumers' choices (i.e. the PPF is linear, which is the case in our model if there is a single factor of production), then for a given primitive shock, macro and micro welfare are always the same. Alternatively, if preferences are homothetic and stable, then macro and micro welfare are the same (regardless of the shape of the PPF). Proposition 11 in Appendix B formalizes this result. For a quantitative illustration of the difference between micro and macro welfare see the Covid-19 case study in Section 5.

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.

As in Section 2, to study this problem we index the path of technologies, factor endowments, and preferences by time t . The definition of $\Delta \log Y$ is the same as before: $\Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} b_{it} d \log c_{it}$. In the general equilibrium model and (its applications), we refer to $\Delta \log Y$ as real GDP.

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 $\mathbf{1}$ is an indicator function, and 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, the Domar weight of factors sum to one: $\sum_{i \in F} \lambda_i = 1$.

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.

Proposition 3 (Real GDP). *Given a path of technologies, factor quantities, and tastes that unfold as a function of time t , the change in real GDP is*

$$\Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} \lambda_i(A_t, L_t, x_t) \frac{d \log A_{it}}{dt} dt + \int_{t_0}^{t_1} \sum_{i \in F} \lambda_i(A_t, L_t, x_t) \frac{d \log L_{it}}{dt} dt. \quad (12)$$

In (12), 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 endowments. Lemma 3 shows that changes in real GDP are equal to sales-weighted changes in technology and factor inputs, where sales shares at t are pinned down by the PPF (A_t, L_t) and x_t . 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 $e^{ev}(p, u) = e(p, u_{t_1}, x_{t_1}) \frac{u}{u_{t_1}}$ where $u_{t_1} = v(p_{t_1}, I_{t_1}; x_{t_1})$, similar to Section 2. We call λ^{ev} the *welfare-relevant* sales share.

Proposition 4 (Macro Welfare). *For any smooth path of technologies, factor quantities, and tastes that unfold as a function of time t , changes in macro welfare are*

$$EV^M = \int_{t_0}^{t_1} \sum_{i \in N} \lambda_i^{ev}(A_t, L_t) \frac{d \log A_{it}}{dt} dt + \int_{t_0}^{t_1} \sum_{i \in F} \lambda_i^{ev}(A_t, L_t) \frac{d \log L_{it}}{dt} dt. \quad (13)$$

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 (12), which uses sales shares evaluated at current technology and current preferences. As with real GDP, the first N summands of (13) 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.²⁴

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).

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 1 (Demand Shocks Only). *In response to changes in preferences, x , that keep the PPF, A and L , unchanged between t_0 and 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 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 $\Delta \log Y$ and macroeconomic welfare changes EV^M are both zero so long as we stay on the surface of a given PPF, Proposition 4 implies that the two are not equal when the PPF shifts. This is because real GDP is based on a path of sales shares λ that treats income, substitution, and taste shocks symmetrically, whereas welfare is based on a path of sales shares λ^{ev} that accounts for income and taste shocks from the start, but only takes substitution into account along the path. Therefore, if productivity rises for goods for which there are negative taste or income effects, 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.

²⁴Following 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.

Local Results. We characterize, up to a second order approximation (as $t_1 - t_0 = \Delta t \rightarrow 0$), 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 $\Delta \log L = 0$.²⁵ For the following proposition, we make explicit that equilibrium sales shares depend on the consumer's expenditure function $\lambda(A, x) = \lambda(A, e(p(A, x), V(A, x); x))$, since the expenditure function determines final demand.

Proposition 5 (Approximate Macro Welfare and Real GDP). *Up to a second order approximation, the change in real GDP is*

$$\Delta \log Y \approx \lambda' \Delta \log A + \underbrace{\frac{1}{2} \sum_{i \in N} \left[\Delta \log A' \frac{\partial \lambda_i}{\partial \log A} + \Delta \log x' \frac{\partial \lambda_i}{\partial \log x} \right]}_{\text{total expenditure-switching, } \Delta \lambda' \Delta \log A} \Delta \log A_i, \quad (14)$$

and the change in welfare is

$$EV^M \approx \Delta \log Y + \underbrace{\frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} + \Delta \log A' \frac{\partial \log V}{\partial \log A} \frac{\partial \lambda_i}{\partial \log V} \right]}_{\text{expenditure-switching due to taste shocks and income effects}} \Delta \log A_i, \quad (15)$$

where $\frac{\partial \lambda_i}{\partial \log A}$ and $\frac{\partial \lambda_i}{\partial \log x}$ denote the total derivative of sales shares with respect to productivities and taste shocks, and $\frac{\partial \lambda_i}{\partial \log V}$ denotes the partial derivative of sales shares with respect to utility, and all derivatives are evaluated at t_0 .

We discuss (14) and (15) in turn. The first term in (14) is the Hulten-Domar formula. The remaining terms capture nonlinearities due to changes in sales shares (since these are second-order, they are multiplied by 1/2). Intuitively, if sales shares decrease for those goods with higher productivity growth, then real GDP growth slows down. This type of effect, known as Baumol's cost disease, is an important driver of the slow-down in aggregate productivity growth.

Equation (15) shows that the gap between macro welfare and real GDP is similar to that for our micro results (the signs are flipped because a positive productivity shock reduces prices). 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,

²⁵Shocks 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.

welfare measures treat changes in shares due to technology shocks differently than changes in shares due to demand shocks or non-homotheticities. That is, real GDP “undercorrects” for changes in shares caused by non-homotheticities or changes in preferences. These terms are multiplied by 1/2 in real GDP, but they are multiplied by 1 in welfare.

Proposition 5 is different to its micro counterpart (Proposition 2) in some important ways. For micro welfare, Proposition 2 shows that if all price changes are the same, there can be no gap between micro welfare EV^m and real consumption. The general equilibrium counterpart of this statement is not true. That is, 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 $\Delta \log A_i = \Delta \log A > 0$ 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

$$EV^M - \Delta \log Y \approx \frac{1}{2} \left[\Delta \log x' \frac{\partial \lambda^{sum}}{\partial \log x} + \Delta \log A' \frac{\partial \log V}{\partial \log A} \frac{\partial \lambda^{sum}}{\partial \log V} \right] \Delta \log A, \quad (16)$$

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 $EV^M < \Delta \log Y$. 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 Section 3 are reduced-form in the sense that they take changes in observed and welfare-relevant sales shares as given. In this section, we solve for changes in these endogenous objects in terms of observable sufficient statistics. For clarity, we restrict attention to nested-CES economies. The general case is in Appendix I, and the intuition is very similar. After providing a characterization to solve for changes in prices and shares in general equilibrium, we work out some analytical examples to provide more intuition. We also discuss how our results can be applied in dynamics economies.

Nested-CES economies. Household preferences are represented by a non-homothetic CES aggregator, which imply that budget shares vary according to (6). 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

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

where ω_{ij} and ω_{if} are constant parameters. 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, 2019b).

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 $i \notin N$).²⁶ 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

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

where Id is the identity matrix. The Leontief inverse matrix $\Psi \geq Id$ records the *direct and indirect* exposures through the supply chains in the production network. We partition Ψ in the following way:

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

²⁶We 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 $i \notin N$.

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 General characterization for nested-CES economies

According to Lemma 3 and Proposition 5, changes in real GDP and welfare can be computed by weighing technology shocks by observed and welfare-relevant sales shares (λ and λ^{ev} , respectively) and cumulating the results. The following proposition pins down λ and λ^{ev} along the transition path. This proposition can then be used in combination with Lemma 3 and Proposition 5 to calculate (globally) changes in real GDP and welfare.

To simplify notation, we again assume away shocks to factor endowments and consider some path of taste shocks x and technology shocks A between t_0 and t_1 .

Proposition 6 (Characterization for nested-CES economies). *At any point in time t , changes in observed prices and the Leontief inverse are pinned down by the following equations:*

$$\begin{aligned}
 d \log p_{it} &= - \underbrace{\sum_j \Psi_{ijt} d \log A_{jt}}_{\text{upstream TFP changes}} + \underbrace{\sum_{f \in F} \Psi_{ift} d \log \lambda_{ft}}_{\text{upstream factor price changes}} \quad , \\
 d \Psi_{ilt} &= \underbrace{\sum_j \Psi_{ijt} (\theta_j - 1) \text{Cov}_{\Omega_{(j,:),t}} \left(-d \log p, \Psi_{(:,l),t} \right)}_{\text{substitution effect}} + \underbrace{\mathbf{1}_{\{i=0\}} \text{Cov}_{\Omega_{(0,:),t}} \left(d \log x_t + \varepsilon_t d \log Y_t, \Psi_{(:,l),t} \right)}_{\text{taste shocks and income effect}} \quad , \\
 & \hspace{15em} (17)
 \end{aligned}$$

where changes in observed sales shares are given by $d\lambda_{it} = d\Psi_{0it}$ and changes in real GDP are given by $d \log Y_t = \sum_i \lambda_{it} d \log A_{it}$.

On the other hand, changes in welfare-relevant variables are pinned down by the following system of differential equations

$$d \log p_{it}^{ev} = - \sum_j \Psi_{ijt}^{ev} d \log A_{jt} + \sum_{f \in F} \Psi_{ift}^{ev} d \log \lambda_{ft}^{ev}$$

$$d\Psi_{ilt}^{ev} = \sum_j \Psi_{ijt}^{ev} (\theta_j - 1) \text{Cov}_{\Omega_{(j,:),t}^{ev}} \left(-d \log p_t^{ev}, \Psi_{(:,l),t}^{ev} \right), \quad (18)$$

where changes in welfare-relevant sales shares are given by $d\lambda_{it}^{ev} = d\Psi_{0it}^{ev}$ and changes in welfare are given by (13).

For all of these expressions, the summations are evaluated over all goods and factors, so that i and $j \in \{0\} + N + F$, $\text{Cov}_{\Omega_{(j,:),t}(\cdot)}$ is the covariance using the j th row of Ω at time t as the probability weights, and $\Psi_{(:,i),t}$ is the i th column of the Leontief inverse at time t .

These differential equations can be solved by repeated iteration.²⁷ Once in possession of these paths, the change in real GDP and welfare are straightforward to calculate by cumulating the λ and λ^{ev} -weighted sum of technology shocks. For this iterative procedure, the boundary condition of the differential equations are that prices satisfy $p_{t_1} = p_{t_1}^{ev} = 1$ and the Leontief inverse matches $\Psi_{t_1}^{ev} = \Psi_{t_1}$.

For ex-post welfare questions, where the Leontief inverse Ψ is observed at t_1 , we can calculate Ψ^{ev} between t_0 and t_1 by starting (18) at t_1 and going backwards to t_0 . This process does not require knowledge of either the income elasticities ε nor the taste shocks $\Delta \log x$ since they do not appear in either the equation for $d \log p^{ev}$ nor the equation for $d\Psi^{ev}$.

Each term in the differential equations in Proposition 6 has a clear interpretation. We start by discussing the equation determining prices $d \log p$. This equation captures the fact that the price of each good $d \log p_i$ is determined by its (direct and indirect) exposure to the price of inputs j and factors f (captured by Ψ_{ijt} and Ψ_{ift} at time t). The equation for $d \log \Psi_{ilt}$, in turn, shows that changes in the Leontief inverse are determined by substitutions by j , if j is an intermediary between i and l , as well as income and substitution effects if i is the household ($i = 0$). Finally, the welfare-relevant versions of these equations, $d \log p^{ev}$ and $d\Psi^{ev}$ are identical except that they do not account for expenditure-switching due to income effects or taste shocks.²⁸

Remark 3 (Micro Welfare). Proposition 6 can also be used to compute changes in microeconomic welfare EV^m . To do this, we compute the actual path of prices $d \log p$ using Proposition 6 and then plug these price changes into (8). Unlike macroeconomic welfare EV^M , calculating $d \log p$ generically requires knowledge of both income elasticities and taste shocks.

Remark 4 (Compensating Variation). Proposition 6 can also be used to compute changes in compensating variation at initial preferences (instead of equivalent variation at final preferences). To do this, we would still solve the differential equations for $d \log p^{ev}$ and $d\Psi^{ev}$,

²⁷When evaluating (17) between t_0 and t_1 , we must take into account that income elasticities ε change with budget shares, as described in Appendix D.1.

²⁸Proposition 6 generalizes Baqaee and Farhi (2019b) to economies with income effects and taste shocks.

however, we would use different boundary conditions. The boundary conditions for compensating variation at initial preferences would match the data at t_0 instead of t_1 , setting $p_{t_0}^{ev} = \mathbf{1}$ and $\Psi_{t_0}^{ev} = \Psi_{t_0}$. We would then solve the differential equations forward from t_0 to t_1 and use the resulting welfare-relevant shares to weight the technology shocks. This procedure effectively makes use of the fact that compensating variation at initial preferences going from t_0 to t_1 is equal to equivalent variation at final preferences if we go from t_1 to t_0 . As with EV^M , conditional on the boundary conditions, we do not need to know the income elasticities or the taste shocks.

To build more intuition, we first focus on economies with only a single factor of production. In this case, the differential equations for $d \log p$ and $d \log p^{ev}$ are decoupled from the ones for $d\Psi$ and $d\Psi^{ev}$. This follows from the fact that the economy's single primary factor must have a sales share of unity. In other words, the following set of equations always hold when the economy has a single factor: $\lambda_f = \lambda_f^{ev} = \Psi_{0f} = \Psi_{0f}^{ev} = 1$ for $f \in F$. This allows for a simple closed-form characterization of both welfare and real GDP up to a second-order approximation.

Proposition 7 (Approximate Macro Welfare vs GDP: Single Factor). *Consider some perturbation in technology, $\Delta \log A$, and tastes, $\Delta \log x$. When the economy has one factor of production, the change in real GDP is*

$$\begin{aligned} \Delta \log Y \approx & \sum_{i \in N} \lambda_i \Delta \log A_i + \frac{1}{2} \sum_j \lambda_j (\theta_j - 1) \text{Var}_{\Omega_{(j,:)}} \left(\sum_{i \in N} \Psi_{(:,i)} \Delta \log A_i \right) \\ & + \frac{1}{2} \text{Cov}_{\Omega_{(0,:)}} \left(\Delta \log x + \left(\sum_{i \in N} \lambda_i \Delta \log A_i \right) \varepsilon, \sum_{i \in N} \Psi_{(:,i)} \Delta \log A_i \right), \end{aligned} \quad (19)$$

and the difference between welfare and GDP is

$$EV^M - \Delta \log Y \approx \frac{1}{2} \text{Cov}_{\Omega_{(0,:)}} \left(\Delta \log x + \left(\sum_{i \in N} \lambda_i \Delta \log A_i \right) \varepsilon, \sum_{i \in N} \Psi_{(:,i)} \Delta \log A_i \right), \quad (20)$$

where λ , Ω , Ψ , and ε are evaluated at t_0 .

Proposition 7 is a general equilibrium counterpart to Proposition 2. We discuss (19) and (20), starting with (19). The first term in Equation (19) 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 are emphasized by Baqaee and Farhi, 2019b). The terms on the second line of (19), which are

the ones we focus on in this paper, capture changes in sales shares due to changes in tastes or non-homotheticities.

Equation (20) shows that while real GDP correctly accounts for substitution due to supply shocks, in order to measure welfare, it needs to be corrected for expenditure-switching due to changes in final demand caused by taste shocks or income effects. Whereas in partial equilibrium, the gap between welfare and real GDP is proportional to the covariance of supply and demand shocks (see Proposition 2), equation (20) shows that in general equilibrium, the relevant statistic is the covariance of demand shocks with a network-adjusted notion of supply shocks, and not supply shocks per-se. Furthermore, Proposition 7 shows that the elasticities of substitution are irrelevant for the gap between welfare and real GDP in one-factor models. This is because, in response to demand-driven forces, relative prices do not change as the equilibrium moves along a linear PPF. 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.

4.2 Analytical Examples

We now work through some simple examples to illustrate the forces that drive a gap between λ and λ^{ev} and, by extension, real GDP and welfare.

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 $\lambda_i = b_i = \Omega_{0i}$. Therefore, Proposition 7 simplifies to

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

Welfare changes are greater than the change in real GDP if productivity and demand shocks (i.e. shifts in demand curves) are positively correlated. This could happen either because preferences exogenously change to favor high productivity growth goods, $Cov_b(\Delta \log x, \Delta \log A) > 0$, or income effects favor high productivity growth goods, $Cov_b(\varepsilon, \Delta \log A) \Delta \log Y > 0$. 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 $\Psi \geq Id$ and $\lambda_i \geq b_i$; (2) the production network “mixes” the shocks, and this may reduce the correlation of sup-

ply 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 7. 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.

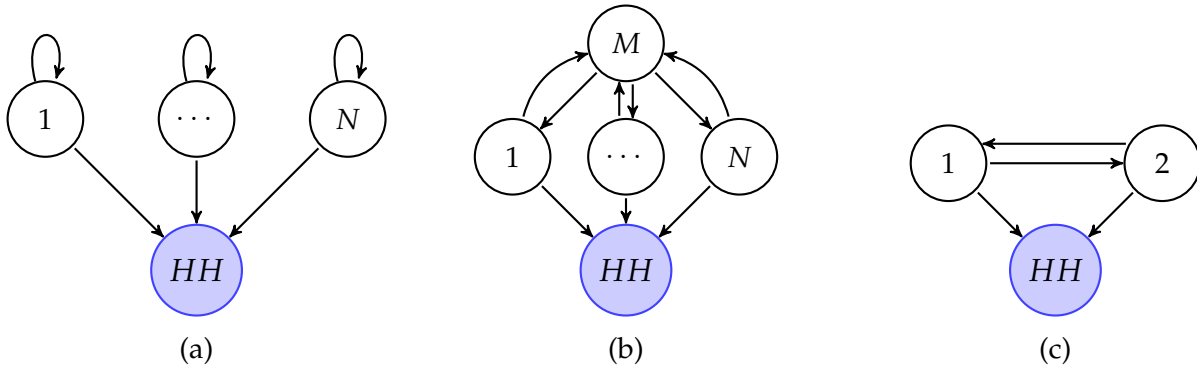


Figure 1: Three different kinds of round-about economy. The arrows represent the flow of goods. The only factor is labor which is not depicted in the diagram.

For Panel 1a, we get

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

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 sales shares is now larger due to the roundabout nature of production.²⁹

²⁹As discussed after Equation (16), 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.

On the other hand, for Panel 1b, we get³⁰

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

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.³¹

Finally, consider Panel 1c. For clarity, we 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

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

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 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.

To recap, in the first, second, and third example the gap between welfare and real consumption increases, is independent of, and decreases in the intermediate input share. Hence, the effect of input-output networks on the adjustment are potent but theoretically ambiguous.

We now work through some simple examples with multiple factors of production to illustrate how nonlinear PPFs affect the previous results.

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

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

³⁰For 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 $i \in N$.

³¹As indicated in Footnote 29, 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.

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 propositions 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 6 to compute changes in sales shares due to demand shocks and then plug this into Proposition 5 to get the difference between welfare and real GDP up to a second order approximation:

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

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 1, 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. Using Proposition 6, changes in prices are

$$d \log p_i = \frac{(1 - \gamma)}{(1 + (\theta_0 - 1)(1 - \gamma))} \frac{1}{\lambda_i} Cov_b \left(d \log x, Id_{(:,i)} \right),$$

where $Id_{(:,i)}$ is the i th 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 2 to obtain micro welfare, up to a second order,

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

With decreasing returns to scale ($\gamma < 1$), micro welfare is negative 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.³² Of course, from a societal perspective, welfare has not changed, since the production possibility set of the economy has not changed.

4.3 Dynamic Economies

As mentioned earlier, at an abstract level, all of our results can be applied to dynamic economies by using the Arrow-Debreu formalism. In particular, we can index goods by period of time and state of nature and apply our results to these economies (see e.g. Basu et al., 2012). In a dynamic economy the utility function is intertemporal and capital accumulation must be treated as an intertemporal intermediate good, as advocated by Barro (2021). Proposition 4 implies that, in such a model, macro welfare can be computed using the final (intertemporal) indifference curve of the representative agent.

In this subsection, we show how Proposition 4 can be used to make steady-state to steady-state welfare comparisons in models with unstable and non-homothetic preferences without requiring us to solve a dynamic model. For simplicity, we restrict our discussion to non-homothetic CES preferences, though this logic can be extended.

Consider a perfectly competitive dynamic economy indexed by the initial period t with a representative agent whose intertemporal preferences are given by

$$U_t = \sum_{s=t}^{\infty} \beta^{s-t} u(C_s), \quad \sum_i \omega_{i0} x_{it} \left(\frac{C_{is}}{C_s^{\theta_i}} \right)^{\frac{\theta_0-1}{\theta_0}} = 1,$$

³²If 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.

where C_s is a non-homothetic (and potentially unstable) CES aggregator. The economy has the same set of goods every period, and every good i in period s is produced according to constant returns production technology

$$y_{is} = A_{is} G_i \left(\{m_{ijs}\}_{j \in N}, H(l_{is}, k_{is}) \right),$$

where A_{is} is a productivity shifter, l_{is} are capital and labor inputs, and H is constant returns to scale.

Labor L_s in each period is inelastically supplied, and capital is accumulated according to a capital accumulation technology

$$K_{s+1} = (1 - \delta) (K_s + X_s),$$

where X_s is aggregate investment. Investment goods are produced according to a constant returns technology

$$X_s = A_{Is} X \left(\{m_{Ijs}\}_{j \in N}, H(l_{Is}, k_{Is}) \right).$$

The intertemporal PPF of economy t is defined by an initial capital stock inherited from the past, a path of future labor endowments, and a path of vectors of productivities: $(K_t, \{L_s\}_{s=t}^{\infty}, \{A_s\}_{s=t}^{\infty})$. This economy has infinitely many factors: the initial capital stock and the path of labor endowments $(K_t, \{L_s\}_{s=t}^{\infty})$. The welfare change between t_0 and t_1 is the proportional change in factor endowments of the t_0 economy required to make the household indifferent between that and the t_1 economy. We say that economy t is in *steady-state* if the vector of productivities A_s , labor endowments L_s , per-period utility $u(C_s)$, and capital stocks K_s are constant over time.

The following proposition shows that computing the welfare change between t_0 and t_1 is straightforward if the economy is in steady-state in both t_0 and t_1 .

Proposition 8 (Dynamic Welfare Change). *Consider two dynamic economies, denoted t_0 and t_1 , that are in steady-state. The change in macro welfare is given by*

$$EV^M = \log \left(\frac{\sum_i p_{it_1} c_{it_1}}{\sum_i p_{it_0} c_{it_0}} \right) + \log \left(\sum_i b_{it_1} \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}.$$

In words, the change in nominal consumption deflated by the ideal price index associated with the t_1 indifference curve is equal to the macroeconomic change in welfare.

5 Applications

In this section we consider three applications of our results. The first application is about long-run growth, quantifying the difference between welfare-relevant and measured aggregate productivity growth as well as the difference between welfare and measured real consumption, in the presence of income effects and demand instability. The second application shows that correlated product-level supply and demand shocks within industries drive a wedge between measured real consumption and welfare even in the short-run. Our final application is a business cycle event study, where we use the Covid-19 recession to demonstrate the difference between macroeconomic and microeconomic welfare and how demand instability can make measured real GDP an unreliable metric for changes in production.

5.1 Long-Run Growth and Structural Transformation

As economies grow, sectors with low productivity growth tend to expand compared 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.

Following Nordhaus et al. (2008), aggregate productivity growth between t_0 and t_1 can be decomposed into two terms:

$$\Delta \log TFP = \sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{it_0} \Delta \log A_{it} + \underbrace{\sum_{t=t_0}^{t_1} \sum_{i \in N} (\lambda_{it} - \lambda_{it_0}) \Delta \log A_{it}}_{\text{Baumol Adjustment}}$$

where λ_{it} is the sales shares of industry i in period t and $\Delta \log A_{it}$ is the growth in gross-output productivity over each time period.³³ 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.³⁴

³³Technically, 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.

³⁴For this exercise, we abstract from investment decisions and apply our formulas statically. This means that we assume a reduced-form representation whereby preference relations are defined over all final goods in a given period (including government spending, net exports, and investment) and calculate welfare changes between two time periods taking preferences, technologies, and factor quantities as given. When

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 substitution effects. In particular, the welfare-relevant measure of the change in TFP is

$$\Delta \log TFP^w = \sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{it_0} \Delta \log A_{it} + \underbrace{\sum_{t=t_0}^{t_1} \sum_{i \in N} (\lambda_{it} - \lambda_{it_0}) \Delta \log A_{it}}_{\text{Baumol Adjustment}} + \underbrace{\sum_{t=t_0}^{t_1} \sum_{i \in N} (\lambda_{it}^{ev} - \lambda_{it}) \Delta \log A_{it}}_{\text{Welfare Adjustment}}$$

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 the welfare adjustment term to obtain $\Delta \log TFP^w$ requires an explicit structural model of the economy. However, there are two polar cases in which $\Delta \log TFP^w$ can be calculated without specifying the detailed model. The first extreme is when demand is stable and homothetic, and changes in sales shares are due only to relative price changes (substitution effects). The second extreme is when there are no substitution effects (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 substitution effects and non-homotheticities or demand instability, then the change in welfare TFP will be somewhere in between these two cases, as discussed in Appendix F. The following corollary of Proposition 4 summarizes the change in welfare-TFP in these two polar cases.

Corollary 2. *If changes in sales shares are due only due only to substitution effects, then*

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

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

$$\Delta \log TFP^w = \Delta \log TFP + \sum_{t=t_0}^{t_1} \sum_{i \in N} (\lambda_{it_1} - \lambda_{it}) \Delta \log A_{it} = \sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{it_1} \Delta \log A_{it}.$$

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 substitution effects in production or calculating welfare using consumption data below, we apply the steady-state results of a dynamic economy implied by Proposition 8.

demand, sales shares do not respond to productivity changes. By focusing on the change in welfare-relevant aggregate productivity, we do not need to take a stand on the number of primary factors in the economy.

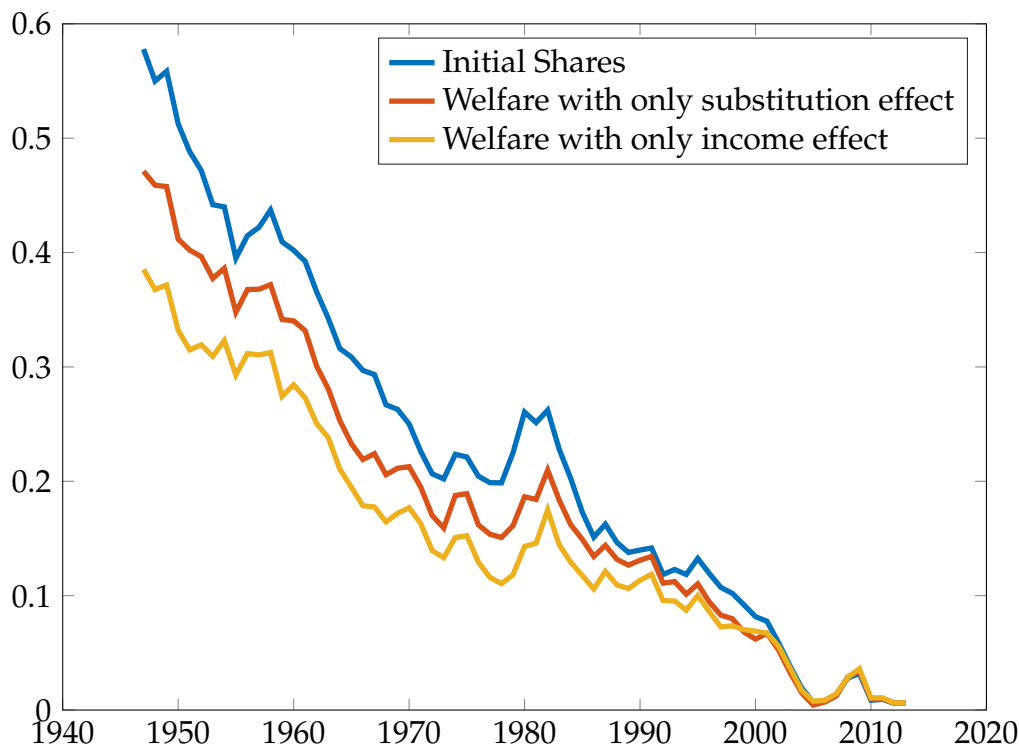


Figure 2: Growth in welfare-relevant TFP (in logs) from 1947 to 2014 using US-KLEMS. The blue line uses initial shares (in each year t between 1947 and 2014) to calculate TFP changes. The red and yellow line measure the increase in welfare-relevant TFP between t and 2014 under alternative assumptions 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).

To quantify Corollary 2, we use US-KLEMS data on sales shares and TFP growth for 61 private-sector industries between 1947 and 2014. We calculate changes in industry-level gross-output TFP following the methodology of Jorgenson et al. (2005) and Carvalho and Gabaix (2013).³⁵

Figure 2 plots EV^M comparing 2014 to previous years under alternative assumptions about substitution and income elasticities. For comparisons that are relatively close to 2014, the change in welfare is not very sensitive to our assumptions about elasticities. This is because at high frequency, the shocks are small and the sales shares are reasonably stable.

³⁵For each industry, the change in TFP is itself a chain-weighted index calculated as output growth minus share-weighted input growth. Inputs are industry-level measures of materials, labor, and capital services.

However, the assumptions about substitution and income elasticities do start to play a role as we roll the comparison back farther in time. Comparing 1947 to 2014, 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, and reduced aggregate productivity growth by around 23 percent (from 78% to 60%).

If we assume that structural transformation is due solely to income effects and taste shocks, then by Corollary 2 the growth in welfare-relevant TFP from 1947-2014 was 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 1947, preferences in 2014 favor low productivity growth sectors such as services (due to either income effects or demand instability). This means that, at 1947 prices, households require less income growth to be indifferent between their budget constraint in 1947 and the one in 2014. This is because sectors with high income elasticities or that consumers prefer in 2014, like services, were cheaper compared to manufacturing in 1947 than in 2014.³⁷

To sum up, structural transformation caused by income effects or demand instability reduced welfare, EV^M , by roughly twice as much as structural transformation caused by substitution effects. The necessary adjustment for taste and income effects is roughly twice as big as those for substitution effects because, in the second-order approximation in Proposition 5, the former are multiplied by 1 whereas the latter are multiplied by 1/2. We elaborate on this point in Appendix F and also provide some quantitative illustrations away from the two polar extremes we discussed above. In this appendix, we compute welfare changes for different values of elasticities of substitution in consumption and production using Proposition 6. Recall that this does not require taking a stance on the income elasticities. We show that welfare-relevant TFP is closer to measured TFP if the elasticity of

³⁶The 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 16). 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.

³⁷This intuition is flipped for compensating variation. As households become poorer in 2014 to be made equally well-off as under their budget constraint in 1947, they favor goods which are relatively cheap in 2014 such as manufacturing, so their income must be reduced by more. More generally, if structural transformation is purely due to income effects or preference instability, then welfare-based productivity growth using CV at initial preferences is given by initial sale-share weighted productivity growth, $\sum_{t=t_0}^{t_1} \sum_{i \in N} \lambda_{it_0} \Delta \log A_{it}$ (which corresponds to the Initial Shares line in Figure 2), so for this exercise, the Baumol adjustment is not welfare-relevant.

substitution across disaggregated industries (in consumption or production) is lower than one.

The Baumol effect in real consumption. So far in this application we have examined aggregate productivity. We now show that similar conclusions apply if we measure welfare changes using data on consumer prices and budget shares across goods in the US between 1947 and 2019.

Specifically, we measure the change in microeconomic welfare using Lemma 1 under alternative assumptions about income and substitution elasticities. We apply this formula statically and calculate welfare changes between two time periods taking as given changes in prices and nominal expenditures. Under the assumptions of Proposition 8, these static numbers also represent the change in macroeconomic welfare in a dynamic model between two steady-states.

If changes in budget shares are driven by substitution effects only, then welfare changes are equal to growth in real consumption per capita. If changes in budget shares are driven by income effects and demand instability only, then welfare changes between any year and 2019 are given by changes in nominal expenditures deflated by a price index using 2019 budget shares.

Figure 3 shows that for comparisons that are close to 2019, the change in welfare is not very sensitive to the assumptions on demand instability and income effects versus substitution effects because, at high frequency, the shocks are small and the sales shares at our level of aggregation (66 goods and services) are stable. On the other hand, for longer time periods, welfare growth is smaller if changes in budget shares took place due to income effects (or demand instability) rather than substitution effects. That is, comparing 1947 and 2019, the change in welfare per capita was 145 log points if preferences are homothetic and stable, but it was only 126 log points if changes in budget shares were entirely due to demand shocks and income effects. As before, structural transformation in consumption caused by demand shocks and income effects is roughly twice as important for welfare as structural transformation caused by substitution effects.

5.2 Aggregation Bias with Firm-Level Taste 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, the gaps between welfare and real GDP or consumption due to demand instability or income effects are usually modest at high frequency. However, this does not mean that these biases are necessarily absent from short-run data.

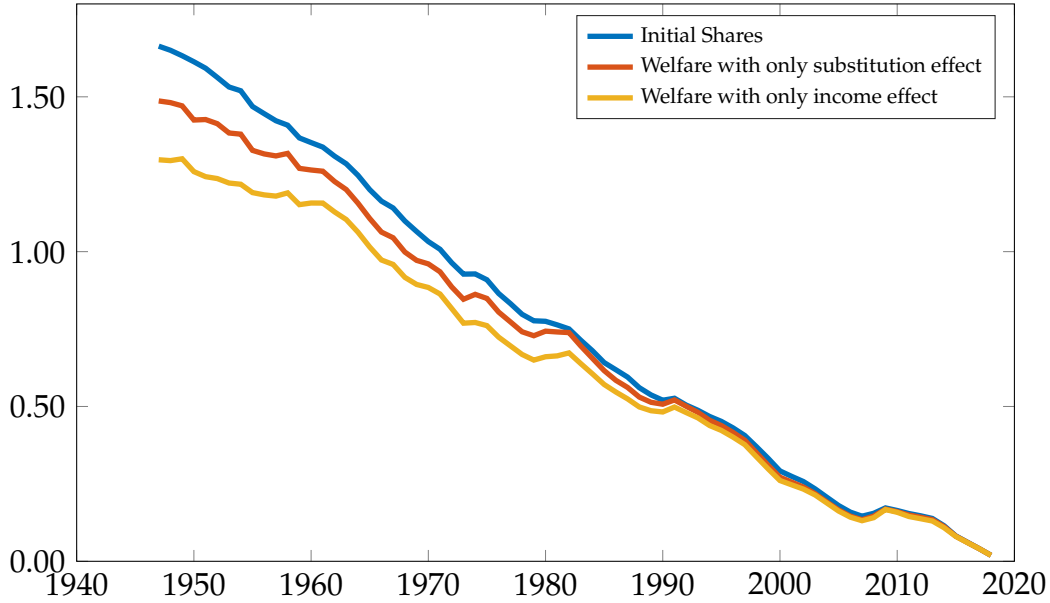


Figure 3: Change in welfare per capita from 1947 to 2019 using Personal Consumption Expenditure (PCE) prices and expenditures for 66 goods and services from the BEA. The blue line uses initial shares (in each year t between 1947 and 2019) to calculate the deflator. The red and yellow line measure the increase in welfare between t and 2019 under alternative assumptions about income and substitution elasticities. The red line assumes that budget shares change only due to substitution effects (welfare is equal to measured chained-real consumption). The yellow line assumes that budget shares change only due to income effects (or demand instability).

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. In Appendix G, we formally show that the biases in industry-level data are not diversified away as we aggregate, even if all firms are infinitesimal in their industry and all industries are infinitesimal in the aggregate economy. Furthermore, we show that the within-industry biases are, to a second-order, linearly separable from the across-industry biases. That is, the overall bias is the sum of the cross-industry bias (that we studied in the previous section) plus additional biases driven by within-industry covariance of supply and demand shocks. If supply and demand shocks at the firm level are persistent, then the covariance of supply and demand shocks, and hence the bias, is larger over longer horizons.

We provide an empirical illustration of the magnitude of the biases caused by taste shocks in product-level data using the Nielsen Consumer Panel database. In the body of the paper, we only briefly describe the dataset, and refer readers to Appendix H for more details. The Nielsen Consumer Panel tracks the purchasing behavior of about 40,000

to 60,000 panelists every year from 2004 to 2019 as they shop in a wide variety of non-durable consumer goods (food, non-food groceries, general merchandise, etc.). A product in the data is defined by its unique Universal Product Code (UPC), and each product is assigned to a module. Our balanced sample covers roughly 820 modules. Panelists in the sample are assigned weights, allowing purchases by the panel to be projected to a nationally representative sample.

We model national demand for UPCs in a given module using a homothetic but unstable CES functional form. We set $t_1 = 2019$ and then for each $t_0 < 2019$, we calculate a welfare-relevant deflator by module for continuing goods using preferences in 2019 following equation (8). The price of each UPC in each year is calculated as the ratio of national expenditures on that UPC over units sold over the whole year. For each t_0 , we include only UPCs purchased in each quarter of each year between t_0 and t_1 (abstracting from product entry and exit). For the same set of UPCs, we also compute the change in inflation as measured by a chained Tornqvist index (a discrete time approximation to the Divisia index) as well as the commonly used Sato-Vartia (SV) index. We then combine these module-level inflation rates into a single number by weighing each module according to its share in overall expenditures in the year 2019, which corresponds to assuming demand across modules is Cobb-Douglas.

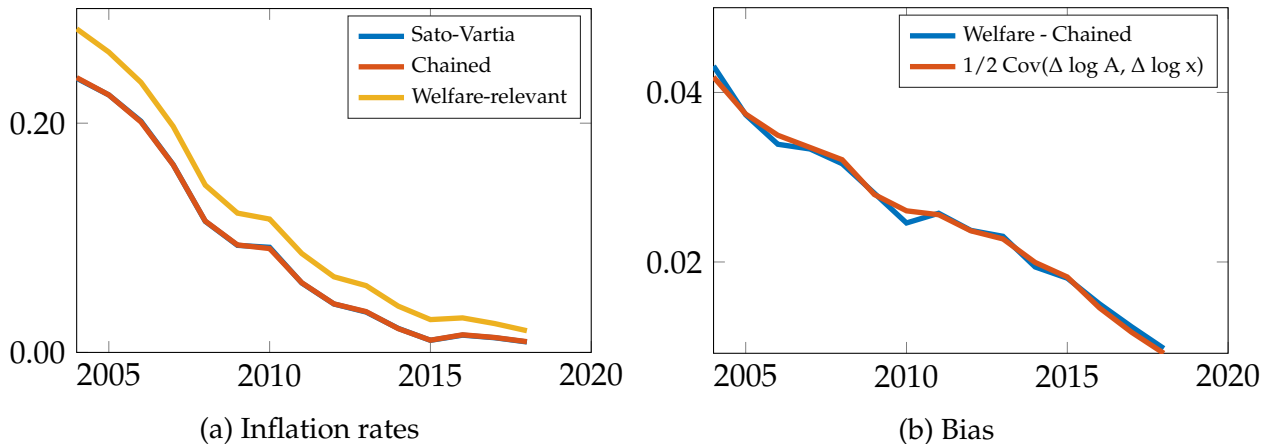


Figure 4: Welfare-relevant, chain-weighted, and Sato-Vartia inflation rate for continuing products. The chain-weighted and Sato-Vartia index are almost the same. The welfare-relevant rate is computed assuming that the elasticity of substitution across UPCs in the same module is 4.5.

Figure 4 displays the results assuming that the elasticity of substitution across UPCs in the same module is 4.5.³⁸ Panel 4a shows the welfare-relevant, chained, and SV inflation

³⁸Estimating the elasticity of substitution is beyond the scope of this paper, therefore, for our empirical illustration we draw on estimates from the literature. An elasticity of 4.5 is at the lower range of estimates

rates for each t_0 . The chained and SV index move closely with each other, but the welfare-relevant inflation rate is higher. Starting in 2018, welfare-relevant inflation is around 0.5% higher than the chained and SV indices. Intuitively, this is because changes in prices and changes in demand residuals are positively correlated, and hence, following the logic of Proposition 2, the chained index understates inflation. The gap increases as we go back further in time and over the whole sample, the gap widens to 4.3%.

Panel 4b plots the difference between the welfare-relevant and chained indices and the expected bias term from Proposition 7 (or Proposition 13 in Appendix G which explicitly models the industry-firm structure). To calculate UPC-level productivity shocks $\Delta \log A$ we use the change in each UPC's log price relative to the average log price change for that module. We construct Δx as the difference between the observed change in expenditure shares and the change in expenditure share implied by a CES aggregator with elasticity $\theta_0 = 4.5$.

Panel 4b shows that the second-order approximation performs well, as the gap between the welfare-relevant and chained index is approximately equal to half the covariance between supply and demand shifters. Furthermore, the gap widens as we extend the time horizon because the covariance of supply and demand shifters increases. This is natural if supply and demand shocks are persistent. Hence, these disaggregated product-level biases, which are assumed away if one begins with more aggregated data, are non-negligible.

We report robustness with respect to the elasticity of substitution parameter in Appendix H. In this appendix, we show that the size of the bias gets smaller as we get closer to Cobb-Douglas. This is because in the data changes in prices and changes in expenditure shares are approximately uncorrelated. When demand is Cobb-Douglas, changes in expenditure shares are driven only by taste shocks, and so taste shocks are roughly uncorrelated with price changes. Hence, following the logic of Proposition 2, the bias is smaller in the Cobb-Douglas case.

5.3 Case Study: the Covid-19 Recession

Our final application examines how real GDP, microeconomic welfare, and macroeconomic welfare were affected during the Covid-19 recession. The Covid-19 recession is an interesting case study since sectoral expenditure shares changed substantially during this time, these changes were not explainable via changes in observed prices alone, and the movements in demand curves were correlated with movements in supply curves. These are

reported by Redding and Weinstein (2020). In Appendix H, we report results for higher and lower elasticities. We find that the size of the bias is increasing in the elasticity of substitution. In this sense, the results in Figure 4 are relatively conservative.

exactly the conditions under which micro welfare, macro welfare, and real GDP can diverge from each other.

Cavallo (2020) argues that, during this episode, the fact that price indices were not being chained at high enough frequency led to “biases” in official measures of inflation. However, since final demand was unstable during this period, chaining is not theoretically justified. As we have argued, chaining is only theoretically valid if expenditure-switching is caused by substitution effects, and not if expenditure-switching is caused by shocks to demand. Furthermore, if changes in prices are themselves caused by changes in demand (due to decreasing returns to scale), then microeconomic welfare and macroeconomic welfare changes are different.

In this section, we do not attempt to measure the welfare costs of Covid-19 itself. This is because households do not make choices over whether or not they live in a world with Covid-19. Therefore, their preferences about Covid-19 itself are not revealed by their choices. Instead, we ask a more modest question: how does the household value changes in prices (micro welfare) and changes in production (macro welfare), holding *fixed* the presence of Covid-19.

To study this episode, we use a modified version of the quantitative model introduced in Section 4. 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 calibrate share parameters to match the 71 industry US input-output table in 2018 (we exclude government sectors) from the BEA, and consider a range of elasticities of substitution. Following Baqaee and Farhi (2020), we model the Covid-19 recession as a combination of negative sectoral employment shocks and sectoral taste shifters. We hit the economy with a vector of primitive supply and demand shocks. The reductions in sectoral employment are calibrated to match peak-to-trough reductions in hours worked by sector from January, 2020 to May, 2020. The primitive demand shifters are calibrated to match the observed peak-to-trough reductions in personal consumption expenditures by sector from January, 2020 to May, 2020 (conditional on the supply shocks and the elasticities of substitution).³⁹

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 to be 0.7, the one across intermediates to

³⁹Changes in labor by sector and personal consumption 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).

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 1 displays welfare changes between January 2020 and May 2020 in the calibrated model. We report separately micro and macro welfare based on pre-Covid (initial, Q1-2018) and post-Covid (final, Q2-2020) preferences. Recall that micro and macro welfare are not equal in this economy because the PPF is nonlinear. For comparison, we also report the change in real consumption assuming supply and demand shocks arrive simultaneously over time.

Table 1: The change in micro and macro welfare with pre-Covid and post-Covid 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 pre-Covid preferences	-11.7%	-9.1%	-8.7%
Micro post-Covid preferences	-13.2%	-12.3%	-10.9%
Macro pre-Covid preferences	-16.2%	-12.5%	-10.8%
Macro post-Covid preferences	-10.1%	-9.4%	-9.0%
Chained real consumption	-12.1%	-10.6%	-9.8%

Table 1 shows that the drop in micro welfare is larger under post-Covid preferences than under pre-Covid preferences. This is because, as shown in Example 4, demand shocks reduce micro welfare in the presence of decreasing returns to scale. Intuitively, if there are decreasing returns to scale, then demand shocks increase the price of goods that consumers value more and this causes micro welfare to drop (since whatever households value becomes more expensive relative to the past).

This pattern is exactly reversed for macro welfare. Macro welfare is higher at post-Covid preferences than at pre-Covid preferences. This is because the negative supply shocks were biggest in those sectors where demand also fell more drastically (e.g. transportation and energy). Hence, the reduction in welfare is smaller with post-Covid preferences because those goods that the economy is less capable of producing are less desirable. This illustrates that micro and macro welfare answer different questions, and the answers to these questions can be quantitatively very different. Furthermore, comparing

columns of Table 1 shows that the magnitude of these differences depend on the details of the production structure like the extent of complementarities in production. As we raise the elasticities of substitution in production closer to unity (Cobb-Douglas), the differences between macro and micro notions become less dramatic. This is because the PPF becomes less curved.

In Table 1, we also compute real consumption assuming supply and demand shocks arrive simultaneously and linearly over time. Interestingly, chained real consumption in Table 1 does not exactly measure any of the different welfare notions. This is because supply and demand shocks are not orthogonal along the path. In fact, if we change the order or path of supply and demand shocks, real consumption changes value (even though the initial and final allocation are not changing). For example, if the supply shocks arrive before the demand shocks, then real consumption equals macro welfare changes at pre-Covid preferences. On the other hand, if demand shocks arrive before the supply shocks, then real consumption equals macro welfare changes at post-Covid preferences.⁴⁰

Hence, if the supply and demand shocks do not disappear in exactly the same way as they arrived, measured real consumption (or GDP) after the recovery can be higher or lower than it was before the crisis, even if the economy returns exactly to its pre-Covid allocation. If in the downturn, demand shocks arrive before supply shocks (so real consumption falls by roughly 10% in the high complementarities case, according to Table 1) and, in the recovery, demand shocks disappear before the supply shocks (so real consumption rises by roughly 16%), then real consumption is as much as 6% higher when comparing pre-shock real consumption to post-recovery real consumption. This is despite the fact that every price and quantity is the same when comparing the pre-shock allocation to the post-recovery allocation. Hence, during episodes where final demand is unstable, chained real GDP and consumption are unreliable guides for measuring output or welfare, even if we chain in continuous time.⁴¹

This example highlights that when comparing chain-weighted measures in the data against chain-weighted measures constructed using model-generated data, the specific path of shocks between two periods must be taken into account and matters if the model features non-homotheticities and/or taste shifters.

⁴⁰These two observations follow from the fact that demand shocks alone have no impact on real GDP (see Corollary 1) and, conditional on fixed and homothetic preferences, real GDP equals macro welfare.

⁴¹This is related to a problem known as “*chain drift*” bias in national accounting. Chain drift occurs when a chained index registers an overall change between t_0 and t_1 even though all prices and quantities in t_0 and t_1 are identical. This is a specific manifestation of path dependence of chained indices (see Hulten, 1973) 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. Welfare changes do not exhibit chain drift since, by definition, they depend only on t_0 and t_1 variables.

6 Extensions

In this section, we briefly summarize how our theoretical 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

$$u(c; x^*) = \left(\int_0^{x^*} c(z)^{\frac{\sigma-1}{\sigma}} dz \right)^{\frac{\sigma}{\sigma-1}}, \quad (21)$$

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

Consider how welfare of the consumer (taking prices as given) 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 and Taste Shocks). *Consider a household with preferences defined by (21). Up to a second-order approximation,*

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

In words, the gap between welfare and real consumption 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 consumption. However, this adjustment is second-order (since it involves products of Δ), and to a first-order, real consumption 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

$$\Delta \log EV^m \approx \Delta \log Y + \frac{1}{1-\sigma} \Delta \log \left[\int_{\mathcal{C}} b(z) dz \right], \quad (22)$$

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 (22), 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 (22), 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 (22) 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. By fixing x when measuring welfare, we do not attempt to answer the question of how much the household values advertising, instead, we hold fixed the amount of advertising (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.⁴²

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 as measures of welfare 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 I, we discuss how Proposition 6 must be adjusted to allow for non-CES production and utility functions.

Heterogeneous agents. Our microeconomic welfare results can be applied to individual households in economies without representative agents, but for brevity, we have abstracted from preference heterogeneity when defining and characterizing macroeconomic welfare. When households have non-aggregable preferences, societal welfare is ambiguous to define. In Baqaee and Burstein (2021), we generalize the results in this paper to environments where there is no representative agent, using the Kaldor-Hicks compensation principle to measure societal welfare. Specifically, to measure the change in welfare from between t_0 and t_1 , we ask: “*what is the minimum amount endowments in t_0 must change so that it is possible to make every consumer indifferent between t_0 and t_1 ?*” We show that all the results in this paper generalize to economies without representative agents if we define social welfare in this way, and this definition collapses to the definition we use in this paper when there is a representative agent.

7 Conclusion

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

⁴²In 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.

Although our motivation and applications have focused on shocks across time, our results can also be applied to compare welfare across locations in space (see e.g. Deaton, 2003, and Argente et al., 2020). Variation in tastes and income effects are likely to be even more significant across space than across time. The micro welfare change between two countries is the income consumers of country A must be given in country B to make them indifferent between staying in country A and moving to country B . This can be calculated using equation (5), where the welfare-relevant budget shares (as a function of prices) use country A 's demand.

The distinction between macroeconomic and microeconomic welfare is also important in a spatial context. Comparing budget constraints in one location to another may be misleading as a way to compare the technologies of two economies. This is because, even if PPFs in both locations are exactly the same, the relative price of goods households value more in one location will be lower in the other location. 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).
- 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.
- Baqae, D. and A. Burstein (2021). Aggregate welfare and output with heterogeneous agents.
- Baqae, 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.
- Baqae, D. R. and E. Farhi (2019a). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics* 135(1), 105–163.
- Baqae, D. R. and E. Farhi (2019b). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. *Econometrica* 87(4), 1155–1203.
- Barro, R. J. (2021, 01). Double Counting of Investment. *The Economic Journal* 131(638), 2333–2356.

- 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.
- Cravino, J., A. A. Levchenko, and M. Rojas (2019, September). Population Aging and Structural Transformation. NBER Working Papers 26327, National Bureau of Economic Research, Inc.
- Deaton, A. (2003). Adjusted indian poverty estimates for 1999-2000. *Economic and political Weekly*, 322–326.
- Deaton, A. and J. Muellbauer (1980). *Economics and consumer behavior*. Cambridge university press.
- 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.
- 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 Eco-*

- nomics* 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.
- Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica* 79(3), 733–772.
- 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.
- IMF (2004). *Consumer Price Index Manual: Theory and Practice*. USA: International Monetary Fund.
- 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.
- McKenzie, G. W. and I. F. Pearce (1982). Welfare measurement: a synthesis. *The American Economic Review* 72(4), 669–682.
- 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.
- 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.

Online Appendix

A	Non-Parametric Comparison of EV^m and Real Consumption	A1
B	Proofs	A2
C	Extension to Other Welfare Measures	A13
D	Non-homothetic CES preferences	A16
	D.1 Derivation of Marshallian budget shares	A16
	D.2 Comparison of welfare and changes in utility index	A18
E	Quality changes	A21
F	Additional details on the Baumol application	A23
	F.1 Intuition for size of welfare-adjustment	A23
	F.2 Welfare-TFP outside of polar extremes	A24
G	Within-Industry Supply and Demand Shocks	A26
H	Additional Details on Nielsen Application	A30
I	Non-CES Functional Forms	A32

Appendix A Non-Parametric Comparison of EV^m and Real Consumption

In this appendix we provide and discuss an alternative version of Proposition 2 based on Hicksian demands without imposing non-homothetic CES preferences. The proof is in Appendix B.

Proposition 10 (Approximate Micro using Hicksian Demand). *To a second-order approximation, the change in real consumption is⁴³*

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

and the change in welfare is

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

where the derivatives are evaluated at t_0 .

Comparing the expression for real consumption $\Delta \log Y$ and welfare EV^m shows that to a first order, they are the same. Discrepancies between the two arise starting at the second-order and involve how expenditure-switching is treated. Real consumption accounts for changes in budget shares in the same way regardless of their cause. The first term in the square brackets reflects changes in budget shares due to changes in relative prices (substitution effects) and the next two terms correspond to changes in budget share due to non-price factors (taste shocks and income effects).

The second line shows that welfare treats changes in budget shares due to substitution effects differently to changes in budget shares due to taste shocks or income effects. To understand the 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 taste shocks as the shock unfolds over time, changes in welfare 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

⁴³The terms $\Delta \log x$ and $\Delta \log v$ 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, $\Delta \log p \approx (\partial \log p / \partial t) \Delta t + 1/2 (\partial^2 \log p / \partial t^2) \Delta t^2$.

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 $EV^m < \Delta \log Y$.

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 ($EV^m < \Delta \log Y$).

Appendix B Proofs

Proof of Lemma 1. By definition,

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

To finish, rewrite

$$\log \frac{e(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})}{e(p_{t_1}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})} = - \int_{t_0}^{t_1} \sum_{i \in N} \frac{\partial \log e(p_t, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})}{\partial \log p_{it}} \frac{d \log p_{it}}{dt} dt,$$

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 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 smooth path that connects p_{t_0} and p_{t_1} gives the same integral. \square

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

Proof of Proposition 10. For real consumption, differentiate real consumption

$$\Delta \log Y = \int_{t_0}^{t_1} \left(\frac{d \log I_t}{dt} dt - \sum_{i \in N} b_i(p_t, u_t; x_t) \frac{d \log p_{it}}{dt} \right) dt$$

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

For EV^m , using Lemma 1, we can write

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

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

$$\begin{aligned} \frac{dEV^m}{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_{it}}{dt} \\ &\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_{it} \\ &\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_{it} \\ \frac{d^2 EV^m}{dt_1^2} &= \frac{d^2 \log I}{dt^2} - \sum_{i \in N} b_i \frac{d^2 \log p_{it}}{dt^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_{it} d \log p_{jt} \\ &\quad - 2 \sum_{i \in N} d \log v_t \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_{it} - 2 \sum_{i \in N} d \log x'_t \frac{\partial \log e(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial \log p_{it} \partial \log x} d \log p_{it} \\ &= \frac{d^2 \log I}{dt^2} - \sum_{i \in N} \sum_{j \in N} \frac{\partial b_i}{\partial \log p_j} d \log p_{it} d \log p_{jt} \\ &\quad - \sum_{i \in N} b_i \frac{d^2 \log p_{it}}{dt^2} - 2 \sum_{i \in N} d \log v \frac{\partial b_i}{\partial \log u} d \log p_{it} - 2 \sum_{i \in N} d \log x'_t \frac{\partial b_i}{\partial \log x} d \log p_{it} \\ &= \frac{d^2 \log I}{dt^2} - \sum_{i \in N} \left[\sum_{j \in N} \frac{\partial b_i}{\partial \log p_j} d \log p_{jt} + d \log v_t \frac{\partial b_i}{\partial \log u} + d \log x'_t \frac{\partial b_i}{\partial \log x} \right] d \log p_{it} \\ &\quad - \sum_{i \in N} b_i \frac{d^2 \log p_{it}}{dt^2} - \sum_{i \in N} d \log v_t \frac{\partial b_i}{\partial \log u} d \log p_{it} - \sum_{i \in N} d \log x'_t \frac{\partial b_i}{\partial \log x} d \log p_{it} \\ &= \frac{d^2 \log I}{dt^2} - \sum_{i \in N} db_{it} d \log p_{it} - \sum_{i \in N} b_i \frac{d^2 \log p_{it}}{dt^2} \\ &\quad - \sum_{i \in N} d \log v \frac{\partial b_i}{\partial \log u} d \log p_{it} - \sum_{i \in N} d \log x'_t \frac{\partial b_i}{\partial \log x} d \log p_{it} \end{aligned}$$

The first three terms are equal to the second-order expansion of $\Delta \log Y$, and the remaining terms are the bias. \square

Proof of Proposition 2. By Proposition 10, 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 (6) in place of Δb to get the desired expression. For the gap between real consumption and EV^m , note that Proposition 1 implies that

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

where b^H is the Hicksian budget share (holding fixed utility and demand shifters). Using (6) 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

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

which can be rearranged to give the desired expression. \square

Proof of Lemma 3. For this proof, we use notation introduced in Section 4. Setting nominal GDP to be the numeraire, we can write (simplifying notation by suppressing dependence on t in the integral)

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

where the second line uses Proposition 6, and we use the fact that $\lambda' = b' \Psi$, $\Lambda' = b' \Psi^F$, and $b' \Psi^F d \log \Lambda = \Lambda' d \log \Lambda = 0$ because the factor shares always sum to one: $\sum_{f \in F} \Lambda_f = 1$. \square

Proof of Proposition 4. Recall that the macro equivalent variation at final preferences is de-

fined by $EV^M = \phi$, where

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

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

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

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 $b_i^{ev}(p) \equiv \frac{\partial \log e^{ev}(p, u)}{\partial \log p_i} = \frac{\partial \log e(p, v_{t_1}; x_{t_1})}{\partial \log p_i}$. 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

$$d \log p^{ev} = -\Psi^{ev} d \log A + \Psi^{evF} d \log \Lambda^{ev}, \quad (23)$$

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 $p(A_{t_0}, e^\phi L_{t_0}, x_{t_1}) = p^{ev}(A_{t_0}, e^\phi L_{t_0})$, where we used the fact that $V(A_{t_0}, e^\phi L_{t_0}; x_{t_1}) = v_{t_1}$. 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,

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

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

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

which can be re-expressed using the expenditure function as

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

This observation is a key step in the proof. Macro welfare changes can be re-expressed as micro welfare changes given changes in equilibrium prices in a fictional economy with preferences represented by $e^{ev}(p, u)$. As in the proof of Lemma 1, rewrite EV^M as (suppressing dependence on t in the integral)

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

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

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

□

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 11 (Macro vs. Micro Welfare). *Consider changes in technologies A , factor quantities L , and tastes x . 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 11. By the proof of Proposition 4, $EV^m = EV^M$ if and only if $p^{ev}(A_t, L_t) = p(A_t, L_t, x_t)$ for all 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 6, 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 Proposition 5. Differentiate real GDP (abstracting from changes in factor endowments),

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

twice with respect to t_1 and evaluate the derivative at $t_1 = t_0$. This yields the desired expression. Following similar steps as in the proof of Proposition 2,

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

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

$$EV^M \approx \Delta \log Y + \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} + \Delta \log A' \frac{\partial \log V}{\partial \log A} \frac{\partial \lambda_i}{\partial \log V} \right] \Delta \log A_i. \quad (24)$$

□

Proof of Proposition 6. We normalize nominal GDP to be the numeraire. Then Shephard's lemma implies that, for each $i \in N$

$$d \log p_i = -d \log A_i + \sum_{j \in N+F} \Omega_{ij} d \log p_j,$$

where $d \log p_j$ is the change in the price of $j \in N + F$. 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

$$\begin{aligned}
\lambda &= b'\Psi \\
d\lambda &= d(b'\Psi) \\
&= b'\Psi d\Omega\Psi + db'\Psi \\
\Omega_{ij}d\log\Omega_{ij} &= (1-\theta_i)\Omega_{ij}(d\log p_j - \sum_k \Omega_{ik}d\log p_k) \\
d\Omega_{ij} &= (1-\theta_i)\text{Cov}_{\Omega_{(i,:)}}(d\log p, Id_{(:,j)}) \\
\sum_j d\Omega_{ij}\Psi_{jk} &= (1-\theta_i)\text{Cov}_{\Omega_{(i,:)}}(d\log p, Id_{(:,j)})\Psi_{jk} \\
&= (1-\theta_i)\sum_j \text{Cov}_{\Omega_{(i,:)}}(d\log p, \Psi_{jk}Id_{(:,j)}) \\
[d\Omega\Psi]_{ik} &= (1-\theta_i)\text{Cov}_{\Omega_{(i,:)}}(d\log p, \Psi_{(:,k)})
\end{aligned}$$

Meanwhile

$$\begin{aligned}
d\log b_i &= (1-\theta_0)\left(d\log p_i - \sum_i b_i d\log p_i\right) + (\varepsilon_i - 1)d\log Y + d\log x_i \\
&= (1-\theta_0)\text{Cov}_{\Omega_{(0,:)}}(d\log p, Id_{(:,i)}) + \text{Cov}_{\Omega_{(0,:)}}(\varepsilon, Id_{(:,i)})d\log Y + \text{Cov}_{\Omega_{(0,:)}}(d\log x, Id_{(:,i)}) \\
\sum_i db_i\Psi_{ik} &= \text{Cov}_{\Omega_{(0,:)}}\left((1-\theta_0)d\log p + \varepsilon d\log Y + d\log x, \Psi_{(:,k)}\right)
\end{aligned}$$

Hence, $d\lambda' = \lambda'd\Omega\Psi + db'\Psi$ can be written as

$$d\lambda_k = \sum_i \lambda_i(1-\theta_i)\text{Cov}_{\Omega_{(i,:)}}(d\log p, \Psi_{(:,k)}) + \text{Cov}_{\Omega_{(0,:)}}(\varepsilon, \Psi_{(:,k)})d\log Y + \text{Cov}_{\Omega_{(0,:)}}(d\log x, \Psi_{(:,k)}).$$

□

Proof of Proposition 7. Normalize nominal GDP to one. Applying Proposition 6 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 Proposition 5 in combination with the expression for $d\log p$ and $d\lambda$ in Proposition 6 in the case of one factor. To solve for EV^M , by Proposition 11, $EV^M = EV^m$. Solve for $EV^m - \Delta\log Y$ by plugging the expression for $d\log p$ into Proposition 10 and noting that $b' = \Omega_{(0,:)}$. □

Proof of Proposition 8. Consider intertemporal preferences

$$V(\mathbf{A}, \mathbf{L}, K_0) = \sum_{s=t}^{\infty} \beta^{s-t} u(C_s).$$

Comparing economies t and t' , macro EV solves the following equation:

$$V(\mathbf{A}, \phi\mathbf{L}, \phi K_0) = \sum_{s=t}^{\infty} \beta^{s-t} u(C_s(\mathbf{A}, \phi\mathbf{L}, \phi K_0)) = \sum_{s=t'}^{\infty} \beta^{s-t'} u(C_s(\mathbf{A}', \mathbf{L}', K'_0)) = V(\mathbf{A}', \mathbf{L}', K'_0).$$

Since the economy t' is in steady-state, we are looking for

$$\sum_{s=t}^{\infty} \beta^{s-t} u(C_s(\mathbf{A}, \phi\mathbf{L}, \phi K_0)) = \frac{1}{1-\beta} u(C(\mathbf{A}', \mathbf{L}', K'_0)).$$

Furthermore, since $(\mathbf{A}, \phi\mathbf{L}, \phi K_0)$ is also a steady-state (by Lemma 2 below), we are searching for

$$u(C(\mathbf{A}, \phi\mathbf{L}, \phi K_0)) = u(C(\mathbf{A}', \mathbf{L}', K'_0))$$

or

$$C(\mathbf{A}, \phi\mathbf{L}, \phi K_0) = C(\mathbf{A}', \mathbf{L}', K'_0).$$

Let $v(p, I)$ be the static indirect utility function. Then we know that we are searching for

$$v(p(\mathbf{A}, \phi\mathbf{L}, \phi K_0), m) = v(p(\mathbf{A}, \mathbf{L}, K_0), \phi m) = v(p(\mathbf{A}', \mathbf{L}', K'_0), m'),$$

where the first equality uses the fact within period relative goods prices do not depend on within period preferences (since the static PPF is linear). Hence,

$$\begin{aligned} \phi &= \frac{e(p(\mathbf{A}, \mathbf{L}, K_0), v_{t_1})}{e(p(\mathbf{A}, \mathbf{L}, K_0), v_{t_0})} = \frac{e(p(\mathbf{A}, \mathbf{L}, K_0), v_{t_1})}{e(p(\mathbf{A}, \mathbf{L}, K_0), v_{t_0})} \frac{e(p(\mathbf{A}', \mathbf{L}', K'_0), v_{t_1})}{e(p(\mathbf{A}', \mathbf{L}', K'_0), v_{t_1})} \\ &= \frac{e(p(\mathbf{A}', \mathbf{L}', K'_0), v_{t_1})}{e(p(\mathbf{A}, \mathbf{L}, K_0), v_{t_0})} \frac{e(p(\mathbf{A}, \mathbf{L}, K_0), v_{t_1})}{e(p(\mathbf{A}', \mathbf{L}', K'_0), v_{t_1})} \\ &= \exp EV^m. \end{aligned}$$

Hence, we can use micro EV^m to calculate the change in macro welfare. □

Lemma 2. *The steady-state choice of capital (and investment) is the same for any homothetic and stable within-period preferences.*

Proof. Suppose intertemporal welfare is given by

$$U_t = \sum_{s=t}^{\infty} \beta^{s-t} u(C_s),$$

where C_s is some homothetic aggregator of within-period consumption goods. Since all goods are produced with constant-returns to scale and every good uses the same homothetic bundle of capital and labor, we can write the consumption aggregator as depending on

$$C_s = G(L_{cs}, K_{cs})$$

for some function constant-returns-to-scale function G . Similarly, investment goods are created according to some constant returns to scale function

$$X_s = X(L_{Xs}, K_{Xs}),$$

and the capital accumulation equation is

$$K_{s+1} = (1 - \delta)(K_s + X_s).$$

The Lagrangean is

$$\begin{aligned} \mathcal{L} = \sum_{s=t}^{\infty} \beta^{s-t} [& u(C_s) + \mu_s(G(L_{cs}, K_{cs}) - C_s) + \kappa_s(K_{s+1} - (1 - \delta)(K_s + X(L_{Xs}, K_{Xs}))) \\ & + \rho_s(L_s - L_{cs} - L_{Xs}) + \psi_t(K_s - K_{cs} - K_{Xs})] \end{aligned}$$

The first order conditions are

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial C_s} : u'(C_s) &= \mu_s \\ \frac{\partial \mathcal{L}}{\partial K_{s+1}} : \kappa_s - \beta \kappa_{s+1}(1 - \delta) + \beta \psi_{s+1} &= 0 \\ \frac{\partial \mathcal{L}}{\partial K_{Xs}} : -\kappa_s(1 - \delta) \frac{\partial X_s}{\partial K_{Xs}} &= \psi_s = \mu_s \frac{\partial G}{\partial K_{cs}} \\ \frac{\partial \mathcal{L}}{\partial K_{cs}} : \mu_s \frac{\partial G}{\partial K_{cs}} &= \psi_s \\ \frac{\partial \mathcal{L}}{\partial L_{cs}} : \mu_s \frac{\partial G}{\partial L_{cs}} &= \rho_s \\ \frac{\partial \mathcal{L}}{\partial L_{Xs}} : -\kappa_s(1 - \delta) \frac{\partial X}{\partial L_{Xs}} &= \rho_s. \end{aligned}$$

Hence

$$-\kappa_s(1 - \delta) = \mu_s \frac{\partial G / \partial K_{cs}}{\partial X_s / \partial K_{Xs}}$$

$$\begin{aligned} \kappa_s &= \beta \kappa_{s+1}(1 - \delta) - \beta \psi_{s+1} \\ u'(C_s) &= \beta(1 - \delta)u'(C_{s+1}) \frac{\partial G / \partial K_{cs+1}}{\partial G / \partial K_{cs}} \partial X_s / \partial K_{Xs} \left[(\partial X_s / \partial K_{Xs+1})^{-1} + 1 \right]. \end{aligned}$$

In steady state we have

$$1 = \beta(1 - \delta) [1 + \partial X_s / \partial K_{Xs}].$$

Hence, the capital stock and investment in steady-state are pinned down by the following 5 equations in 5 unknowns (K_C, K_X, K, L_C, L_I):

$$\begin{aligned} 1 &= \beta(1 - \delta) [1 + \partial X / \partial K_X], \\ \frac{K_C}{L_C} &= \frac{K_X}{L_X}, \\ K &= K_C + K_X, \\ L &= L_C + L_X, \\ \delta K &= (1 - \delta)X(L_X, K_X). \end{aligned}$$

Since G does not appear in any of these equations, the steady-state investment and capital stock do not depend on the shape of the within-period utility function G . \square

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

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

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 i th argument. Differentiating and evaluating at the initial point, we get

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

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

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

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 C 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, it uses 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 for any x 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).

Micro welfare changes We consider four alternative measures of micro welfare changes.

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

$$v(p_{t_1}, e^{-\phi} I_{t_1}; x_{t_0}) = v(p_{t_0}, I_{t_0}; x_{t_0}). \quad (25)$$

The analog to (5) in Lemma 1 is (suppressing dependence on t in the integral)

$$CV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_{it}^{cv} d \log p_{it}, \quad (26)$$

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

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 (26). 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 (7) (with b^{cv} in place of b^{ev}) 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 (7) starting at t_0 and going forward to t_1 . For non-homothetic CES,

$$CV^m = \Delta \log I - \log \left(\sum_i b_{it_0} \left(\frac{p_{it_1}}{p_{it_0}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}. \quad (27)$$

To a second-order approximation around t_0 (as in Proposition 10)

$$\begin{aligned} \Delta \log CV^m &\approx \Delta \log I - b' \Delta \log p - \frac{1}{2} \sum_{i \in N} \left[\Delta \log p' \frac{\partial b_i}{\partial \log p} \right] \Delta \log p \\ &\approx \Delta \log Y + \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial b_i}{\partial \log x} + \Delta \log v \frac{\partial b_i}{\partial \log u} \right] \Delta \log p. \end{aligned} \quad (28)$$

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 Proposition 10 and (28), we see that up to a second order approximation,

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

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

$$v(p_{t_1}, I_{t_1}; x_{t_0}) = v(p_{t_0}, e^\phi I_{t_0}; x_{t_0}). \quad (29)$$

Globally, changes in welfare are given by (5) where $b_{it}^{ev} \equiv b_i(p_t, v(p_{t_1}, I_{t_1}; x_{t_0}); x_{t_0})$. 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

$$v(p_{t_1}, e^{-\phi} I_{t_1}; x_{t_1}) = v(p_{t_0}, I_{t_0}; x_{t_1}). \quad (30)$$

Globally, changes in welfare are given by (26) where $b_{it}^{cv} \equiv b_i(p_t, v(p_{t_0}, I_{t_0}; x_{t_0}); x_{t_1})$. Note for EV with initial preferences or CV with final preferences, we must be able to separate taste shocks from income effects.

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 the 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 (13) in Proposition 4 applies using $\lambda^{cv}(A, L)$, 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 pref-

erence parameters 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 that in Proposition 5) is

$$CV^M \approx \Delta \log Y - \frac{1}{2} \sum_{i \in N} \left[\Delta \log x' \frac{\partial \lambda_i}{\partial \log x} + \Delta \log A' \frac{\partial \log V}{\partial \log A} \frac{\partial \lambda_i}{\partial \log V} \right] \Delta \log A_i. \quad (31)$$

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. 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. In contrast to our first two measures, in order to implement these last two measure we must know initial demand shifters or income effects.

Appendix D Non-homothetic CES preferences

In this appendix, we derive (6). We also compare EV^m with the utility index (under a popular cardinalization) in non-homothetic CES preferences and show that changes in the utility index are not equal to changes in equivalent or compensating variation.

D.1 Derivation of Marshallian budget shares

This appendix provides a derivation of the log-linearized expression (6). Changes in Marshallian budget share (before imposing functional forms) are given by

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

where ε^H and ε^M are the Hicksian and Marshallian price elasticities, ε^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 second equality is an application of Slutsky's equation. When preferences are non-homothetic CES, then the Hicksian demand curve can be written as

$$c_i = \omega_i \left(\frac{p_i}{\sum_j p_j c_j} \right)^{-\theta_0} u^{\zeta_i}, \quad (32)$$

where ω_i and ζ_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

$$\begin{aligned} d \log b_i &= \sum_j \left(\varepsilon_{ij}^H - \varepsilon_i^w b_j \right) d \log p_j + d \log p_i + (\varepsilon_i^w - 1) d \log I + d \log x \\ &= \left(1 - \sum_{j \neq i} \varepsilon_{ij}^H \right) d \log \frac{p_i}{I} + \sum_{j \neq i} \varepsilon_{ij}^H d \log \frac{p_j}{I} + \varepsilon_i^w \left[d \log I - \sum_j b_j d \log p_j \right] + d \log x_i \\ &= \left(1 - \sum_{j \neq i} \theta_0 b_j \right) d \log \frac{p_i}{I} + \sum_{j \neq i} \theta_0 b_j d \log \frac{p_j}{I} + \varepsilon_i^w \left[d \log I - \sum_j b_j d \log p_j \right] + d \log x_i \\ &= (1 - \theta_0(1 - b_i)) d \log \frac{p_i}{I} + \sum_{j \neq i} \theta_0 b_j d \log \frac{p_j}{I} + \varepsilon_i^w \left[d \log I - \sum_j b_j d \log p_j \right] + d \log x_i \\ &= (1 - \theta_0) \left[d \log p_i - \sum_j b_j d \log p_j \right] + (\varepsilon_i^w - 1) \left[d \log I - \sum_j b_j d \log p_j \right] + d \log x_i. \end{aligned} \quad (33)$$

In the body of the paper we use ε_i in place of ε_i^w . This completes the derivation of (6) where $d \log x$ could be any perturbation to budget shares not explained by income and substitution effects.

For ex-ante counterfactuals, we can also use (33) as a differential equation to solve for budget shares in the future. To do this, we must put some structure on $d \log x$. For example, if $d \log x$ is being driven by changes in taste parameters $\omega(t)$, and θ_0 and ζ_i are constant, then we can show that

$$d \log x_{it} = \left[d \log \omega_{it} - \frac{\zeta_i}{\sum_i b_{it} \zeta_i} \sum_i b_{it} d \log \omega_{it} \right]. \quad (34)$$

Next, we use the identities

$$\frac{\partial \log e}{\partial \log u} = \frac{\sum_i b_i \tilde{\zeta}_i}{1 - \theta_0}$$

and

$$\frac{\partial \log b_i}{\partial \log u} = \tilde{\zeta}_i - (1 - \theta_0) \frac{\partial \log e}{\partial \log u}$$

to express income elasticities as

$$\varepsilon_{it}^w - 1 = (1 - \theta_0) \left(\frac{\tilde{\zeta}_i}{\mathbb{E}_{b_t}[\tilde{\zeta}]} - 1 \right). \quad (35)$$

We can thus rewrite (6) and (33) as

$$d \log b_{it} = (1 - \theta_0) [d \log p_{it} - \mathbb{E}_{b_t}[d \log p_t]] + (1 - \theta_0) \left(\frac{\tilde{\zeta}_i}{\mathbb{E}_{b_t}[\tilde{\zeta}]} - 1 \right) [d \log I_t - \mathbb{E}_{b_t}[d \log p_t]] + d \log x_{it},$$

where $d \log x$ is given by (34). This is a differential equation that pins down budget shares b as a function of prices, incomes and primitives ω . Hence, it can be solved numerically for any path of prices, incomes, and ω to arrive at b .

D.2 Comparison of welfare and changes in utility index

In this appendix, we discuss the difference between changes in welfare as measured by the equivalent variation and changes in the utility index in non-homothetic CES preferences (under the most popular cardinalization). We show that there is no normalization of the parameters such that the equivalent variation is equal to changes in the utility index unless preferences are homothetic. The same result applies for compensating variation.

The expenditure function corresponding to non-homothetic CES (with Hicksian demand given by 32) can be expressed as

$$e(p, u) = \left(\sum_i \omega_i p_i^{1-\theta_0} u^{\tilde{\zeta}_i} \right)^{\frac{1}{1-\theta_0}}$$

where for brevity we have assumed away taste shocks. The micro equivalent variation is given by

$$EV^m = \log \frac{e(p_0, v(p_1, I_1))}{e(p_0, v(p_0, I_0))},$$

where $v(p, I)$ is the indirect utility function, initial prices and income are p_0 and I_0 , and final prices and income are p_1 and I_1 .

Consider now the change in the utility index u in the expenditure function introduced

above, which is given by

$$U \equiv \log \frac{v(p_1, I_1)}{v(p_0, I_0)}.$$

As this definition makes clear, EV and U are not generically the same. In particular, whereas EV can be defined in terms of a hypothetical choice and is independent of the utility function chosen to represent preferences (how much income would the household need to be given to make them indifferent), U will depend on the cardinal properties of the utility function.

If preferences are homothetic ($\xi_{ii} = \bar{\xi}$ for all i), then $e(p, u) = \left(\sum_i \omega_i p_i^{1-\theta_0}\right)^{\frac{1}{1-\theta_0}} u^{\frac{\bar{\xi}}{1-\theta_0}}$ and we can write

$$EV^m = \frac{\bar{\xi}}{1-\theta_0} U.$$

So, when preferences are homothetic, in order for $EV^m = U$ we must cardinalize utility in such a way that the expenditure function is homogeneous of degree 1 in u ($d \log e / d \log u = 1$). In other words, although there are infinitely many utility functions that represent these preferences, when preferences are homothetic, there is one representation where $EV^m = U$.

We now consider the non-homothetic case, and we characterize the difference between EV^m and U to a first and second order. We write these results in terms of primitive shocks (that is, changes in income and prices) rather than in terms of changes in endogenous objects like budget shares.

Define the budget share of good i by

$$b_i \equiv \frac{p_i c_i}{I} = \frac{\partial \log e}{\partial \log p_i}.$$

Furthermore, let ε_i^w be the income elasticity of demand for good i .

Using Proposition 2, we have that to a first-order EV^m is

$$dEV^m = d \log e - b d \log p = d \log Y,$$

where $d \log Y$ is the first-order change in real consumption as measured by Tornqvist. Hence, to a first order, Tornqvist and EV are the same. The second-order change in EV^m is, by Proposition 2, equal to

$$\begin{aligned} d^2 EV^m &= d^2 \log e - d b d \log p - (d \log e - b d \log p) \text{Cov}_b(\varepsilon^w, d \log p) \\ &= d^2 \log Y - (d \log e - b d \log p) \text{Cov}_b(\varepsilon^w, d \log p), \end{aligned}$$

where ε^w is the vector of income elasticities and $d^2 \log Y$ is the change in real consumption

as measured by a Tornqvist or Divisia index (to a second-order, they are equivalent). On the other hand, the first and second-order changes in the utility index are given by (derivations are available upon request)

$$dU = \frac{1 - \theta_0}{\sum_i b_i \bar{\zeta}} (d \log e - bd \log p),$$

$$d^2U = \frac{1 - \theta_0}{\sum_i b_i \bar{\zeta}_i} \left[d^2 \log e - dbd \log p - (d \log e - bd \log p) \sum_i b_i (\varepsilon_i^w - 1) d \log p_i \right. \\ \left. - \frac{1}{1 - \theta_0} \sum_i b_i \varepsilon_i^w ((\varepsilon_i^w - 1)) (d \log e - bd \log p)^2 \right]$$

The derivatives EV^m and U are in general different. Whereas EV^m is a function of observables, U depends on a normalization since $\sum_i b_i \bar{\zeta}_i$ is not pinned down by observables (scaling all $\bar{\zeta}$ proportionally does not change any observable, but it does change how U responds).

Now we compare the first and second-order derivatives in turn. The first order difference is

$$dU - dEV^m = \left(\frac{1 - \theta_0}{\sum_i b_i \bar{\zeta}_i} - 1 \right) (d \log e - bd \log p).$$

If we impose a normalization on utility parameters such that, in the initial point,

$$\frac{1 - \theta_0}{\sum_i b_i \bar{\zeta}_i} = 1,$$

we have that $dU = dEV^m = d \log Y$. This normalization is effectively ensuring that $\partial \log e / \partial \log u = 1$.

Now let's consider the second-order difference and let's impose the same normalization

$$d^2U - d^2EV^m = -\frac{1}{1 - \theta_0} \sum_i b_i \varepsilon_i^w (\varepsilon_i^w - 1) (d \log e - bd \log p)^2 \\ - (d \log e - bd \log p) \left[\sum_i b_i (\varepsilon_i^w - 1) \sum_i b_i d \log p_i \right] \\ = -\frac{1}{1 - \theta_0} \sum_i b_i \varepsilon_i^w (\varepsilon_i^w - 1) (d \log e - bd \log p)^2$$

$$= -\frac{1}{1-\theta_0} \text{Var}_b(\varepsilon_i^w) (d \log e - b d \log p)^2 \neq 0,$$

where we used $\sum_i b_i \varepsilon_i^w = 1$. Hence, unless preferences are homothetic (in which case $\varepsilon_i^w = 1$ for every i), the change in U and EV^m are not the same even under the normalization. This is not to mention that globally, we cannot ensure that the normalization

$$\frac{1-\theta_0}{\sum_i b_i \zeta_i} = 1$$

always holds. This means that the gap between EV^m and U , which exists at the initial equilibrium, only gets more severe if, once we commit to a specific normalization of utility, $\frac{1-\theta_0}{\sum_i b_i \zeta_i}$ starts to change from 1.

Recall from Appendix C that changes in real consumption are equal to an average of equivalent and compensating variation, up to a second order approximation. Since changes in the utility index are not equal to a Tornqvist real consumption index, it follows that the utility index is not equal to an average of EV and CV .

Appendix E Quality changes

In this appendix, we discuss how our results can be extended to environments with quality changes that are not reflected in market prices. In this appendix, we also contrast the bias we identify with the “taste shock bias” discussed by Redding and Weinstein (2020).

If quality changes are equivalent to changes in prices, then we can write the quality-adjusted price as $p = \tilde{p}/q$, where \tilde{p} is the market price and q is the quality adjustment. In other words, changes in quality-adjusted prices are given by $\Delta \log p = \Delta \log \tilde{p} - \Delta \log q$. Substituting this into our various propositions allows us to isolate the way quality changes affect our results. For example, Proposition 2 becomes the following. For brevity, we assume homothetic preferences below.

Proposition 12 (Approximate Micro with Quality Change). *Consider some perturbation in demand $\Delta \log x$, market prices $\Delta \log \tilde{p}$, quality $\Delta \log q$, and income $\Delta \log I$. Then, to a second-order approximation, the change in real consumption is*

$$\begin{aligned} \Delta \log Y \approx & \Delta \log I - \mathbb{E}_b(\Delta \log \tilde{p}) - \frac{1}{2}(1-\theta_0) \text{Var}_b(d \log \tilde{p}) \\ & + \frac{1}{2}(1-\theta_0) \text{Cov}_b(d \log q, d \log \tilde{p}) - \frac{1}{2} \text{Cov}_b(d \log x, d \log \tilde{p}), \end{aligned}$$

and the change in welfare is

$$EV^m \approx \Delta \log I - \mathbb{E}_b(\Delta \log \tilde{p} - \Delta \log q) - \frac{1}{2}(1 - \theta_0)Var_b(\Delta \log \tilde{p}) \\ - \frac{1}{2}(1 - \theta_0)Var_b(\Delta \log q) + (1 - \theta_0)Cov_b(\Delta \log \tilde{p}, \Delta \log q) - Cov_b(\Delta \log x, \Delta \log p),$$

where $\mathbb{E}_b(\cdot)$, $Var_b(\cdot)$, and $Cov_b(\cdot)$ are evaluated using budget shares at t_0 as probability weights.

Hence, by subtracting these two expressions, we can derive the gap between real consumption and welfare up to a second order approximation as

$$EV^m - \Delta \log Y \approx \underbrace{\mathbb{E}_b(\Delta \log q)}_{\text{average quality}} + \frac{1}{2} \underbrace{(\theta_0 - 1)Var_b(\Delta \log q)}_{\text{dispersion in quality}} + \frac{1}{2} \underbrace{(1 - \theta_0)Cov_b(\Delta \log \tilde{p}, \Delta \log q)}_{\text{covariance of price and quality}} \\ - \frac{1}{2} \underbrace{Cov_b(\Delta \log x, \Delta \log \tilde{p})}_{\text{covariance of taste and price}} + \underbrace{Cov_b(\Delta \log x, \Delta \log q)}_{\text{covariance of taste and quality}}. \quad (36)$$

The first term on the right-hand side captures how the average increase in quality raises welfare relative to real consumption. The second term captures the fact that dispersion in quality raises welfare if the elasticity of substitution is greater than one (since the consumer substitutes towards goods with relatively higher quality, but quality is not captured by market prices in real consumption). The third term is an interaction (cross-partial) effect that raises welfare if market prices fall for goods whose quality rose, as long as the elasticity of substitution is greater than one. The fourth term is the bias we have been emphasizing in the paper so far. The final term is the interaction between quality and taste changes — welfare is higher, at final preferences, if tastes increase for goods whose quality also increased.

In our analysis, we assume that prices have already been adjusted for quality, so the only non-zero term is the fourth one. Welfare is higher than real consumption if the covariance between taste shocks and prices is negative. This is independent of the value of the elasticity of substitution. In other words, in the body of the paper, we assume that $\Delta \log q = 0$, which means that (36) simplifies to

$$EV^m - \Delta \log Y \approx -\frac{1}{2}Cov_b(\Delta \log x, \Delta \log \tilde{p}). \quad (37)$$

Comparison to Redding and Weinstein (2020). We can use (36) to contrast our results with those of Redding and Weinstein (2020).⁴⁴ The “taste shifters” in that paper are equivalent to quality shocks $\Delta \log q \neq 0$, whereas for them $\Delta \log x = 0$. Hence, in their case, (36)

⁴⁴We abstract from product creation and destruction, which Redding and Weinstein (2020) also consider.

simplifies

$$EV^m - \Delta \log Y \approx \mathbb{E}_b(\Delta \log q) + \frac{1}{2}(\theta_0 - 1)Var_b(\Delta \log q) - \frac{1}{2}(\theta_0 - 1)Cov_b(\Delta \log \tilde{p}, \Delta \log q). \quad (38)$$

Comparing (37) to (38) elucidates the differences. First, the average level of $\Delta \log q$ affects welfare but the average level of $\Delta \log x$ does not. Redding and Weinstein (2020) neutralize this effect by assuming it is zero. Second, for shocks to $\Delta \log q$, even when they are mean zero, dispersion in q can raise or lower welfare depending on the elasticity of substitution. Hence, shocks to q on their own can change welfare, even if choice sets have not changed, and the sign of this effect depends on the elasticity of substitution. This is in contrast to shocks to x which do not change welfare on their own. Third, in both (37) and (38), the covariance of taste shifters and market prices matters, however, in (38) the sign of the covariance depends on whether the elasticity of substitution is greater than or less than one, whereas in (37), the sign is always the same.

Appendix F Additional details on the Baumol application

In this appendix, we provide some intuition for why, from a welfare perspective, structural transformation caused by income effects or taste shocks is roughly twice as important as that caused by substitution effects. We also use a structural nested-CES model to explore the change in welfare-relevant TFP outside of the two polar extremes in Section 5.

F.1 Intuition for size of welfare-adjustment

According to our results in Section 5, structural transformation caused by income effects or demand instability reduced welfare by roughly twice as much as structural transformation caused by substitution effects. To understand why the necessary adjustment is roughly twice as big, consider the second-order approximation in Proposition 5:

$$\Delta \log TFP^{\text{welfare}} \approx \Delta \log TFP + \frac{1}{2} \sum_{i \in N} \left[\frac{\partial \lambda_i}{\partial \log x} \Delta \log x + \frac{\partial \lambda_i}{\partial \log V} \Delta \log V \right] \Delta \log A_i, \quad (39)$$

where

$$\Delta \log TFP \approx \sum_{i \in N} \lambda_{it_0} \Delta \log A_i + \frac{1}{2} \sum_{i \in N} \Delta \lambda_i \Delta \log A_i.$$

If changes in sales shares are due entirely to demand-driven factors, then the term in square brackets in (39) is equal to $\Delta\lambda_i$, so

$$\Delta \log TFP^{\text{welfare}} \approx \sum_{i \in N} \lambda_{it_0} \Delta \log A_i + \sum_{i \in N} \Delta \lambda_i \Delta \log A_i.$$

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 substitution effects.⁴⁵

F.2 Welfare-TFP outside of polar extremes

In practice, both substitution effects and non-homotheticities are likely to play an important role in explaining structural transformation. To dig deeper into the size of the welfare adjustment outside our two polar cases, 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. We calculate TFP by industry in the data allowing for cross-industry variation in capital and labor shares. For simplicity, we feed these TFP shocks as primitive shocks into a one-factor model.

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).⁴⁶ 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 inter-

⁴⁵These 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.

⁴⁶In 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.

mediate 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 6 in the fictional economy with stable and homothetic preferences.

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 1947) 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 2.

Table 2: Percentage change in measured and welfare-relevant TFP in the US from 1947 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 2 shows the change in welfare-relevant TFP assuming that there are no substitution effects (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 (so that substitution effects are active), 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.

Table 2 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.

To see why the results in Table 2 are differentially sensitive to changes in different elasticities of substitution, we use Proposition 7 to obtain the following second-order approximation:

$$\Delta \log TFP^{\text{welfare}} \approx \sum_i \lambda_i \Delta \log A_i + \frac{1}{2} \sum_{j \in \{0\} + \mathcal{N}} (\theta_j - 1) \lambda_j \text{Var}_{\Omega_{(j,:)}} \left(\sum_{k \in \mathcal{N}} \Psi_{(:,k)} \Delta \log A_k \right), \quad (40)$$

where λ, Ω , and Ψ are evaluated at t_1 . 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

$$\sum_j \lambda_j \text{Var}_{\Omega(j,:)} \left(\sum_{k \in \mathcal{N}} \Psi_{(:,k)} \Delta \log A_k \right),$$

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 λ_j , or (2) if the nests corresponding to those elasticities are heterogeneously exposed to the productivity shocks.

We compute the coefficients in (40) 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 2 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 (40), 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 substitution effects 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 G Within-Industry Supply and Demand Shocks

In this appendix, we introduce a specification of our model with an explicit firm-industry structure. We show that within-industry supply and demand shocks can also drive a wedge between welfare and real GDP, and we show that this gap is linearly separable (to a second-order) from across-industry biases. For simplicity, we abstract from non-homotheticities.

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 I th 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, the following proposition characterizes the bias in real GDP relative to welfare.

Proposition 13 (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*

$$\Delta \log EV^M \approx \Delta \log Y + \frac{1}{2} \sum_I b_I \text{Cov}_{b_{(I)}}(\Delta \log x, \Delta \log A) + \Theta,$$

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

In words, Proposition 13 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. Note that if supply and demand shocks at the firm level are correlated and persistent, then the bias grows over time, as in our product-level data discussed below.

Proof of Proposition 13. 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 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

$$\begin{aligned}\sum_{i \in I} b_{iI} &= 1 \\ \sum_{i \in I} s_{iI} &= 1.\end{aligned}$$

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

$$\lambda_i = b_{iI}\lambda_I^c + s_{iI}\lambda_I^f, \quad (41)$$

$$d\lambda_i = db_{iI}\lambda_I^c + b_{iI}d\lambda_I^c + ds_{iI}\lambda_I^f + s_{iI}d\lambda_I^f, \quad (42)$$

$$db_{iI} = Cov_{b_I}(d \log x + (1 - \zeta_I)d \log A, Id_{(:,i)}),$$

$$ds_{iI} = Cov_{s_I}((1 - \sigma_I)d \log A, Id_{(:,i)}),$$

where $Id_{(:,i)}$ is a vector of all zeros except for its i th element which is equal to one, b_I is a vector of market shares in final sales of industry I , and s_I is a vector of market shares in non-final sales of industry I .

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

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

Using (42), the sums can be re-written as

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

where now the subscript I indicates the industry that the firm i belongs to.

The individual terms can be written out as

$$\sum_{i \in N} \left[d \log x \frac{\partial b_{iI}}{\partial \log x} \lambda_I^c d \log A_i \right] = \sum_{i \in N} \text{Cov}_{b_I}(d \log x, Id_{(:,i)}) \lambda_I^c d \log A_i \\ = \text{Cov}_{b_I}(d \log x, \sum_{i \in N} Id_{(:,i)} d \log A_i) \lambda_I^c \\ = \text{Cov}_{b_I}(d \log x, d \log A) \lambda_I^c;$$

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

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

and

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

Of the four terms, two depend on changes on industry-level sales shares, one of them is zero, and the remaining one (the first term) is the within-industry covariance of supply and demand shocks that is highlighted in the statement of the proposition. Hence, the remaining terms in the statement of the proposition are

$$\Theta = \sum_I \left[\mathbb{E}_{s_I}(d \log A) d \log x \frac{\partial \lambda_I^f}{\partial \log x} + \mathbb{E}_{b_I}(d \log A) d \log x \frac{\partial \lambda_I^c}{\partial \log x} \right].$$

□

Appendix H Additional Details on Nielsen Application

Details on the construction of Figure 4 The Nielsen Consumer Panel data are provided under subscription through the Kilts Center for Marketing at the University of Chicago. A first file provides quantity and expenditures net of discount by UPC (universal product code) for each shopping trip recorded by roughly 60,000 households in the panel.⁴⁷ Additional files record the date of each shopping trip and describe household characteristics, including the Nielsen-defined market in which each household resides. Nielsen provides a set of weights so that each household in the panel can be understood to represent a certain number of households in their market for a given panel year. Nielsen also provides a file with descriptions of each product, including a set of Nielsen-defined product categories. The lowest level of product categorization in this scheme is known as a module. The Kilts Center tracks UPCs over time, assigning UPC version numbers that record if characteristics associated with a given barcode change over time. Thus, a UPC-version has a fixed set of product characteristics over time, which makes it less likely that its quality changes over time. As pointed out by Redding and Weinstein (2020), this is because firms prefer to use different barcodes for products with different observable characteristics for inventory and stock control purposes. We use this stable-characteristic notion of UPCs.

We construct our sample as follows. After dropping trips with non-positive quantity or non-positive expenditure net of discounts, we collapse household-trip-UPC observations by summing to household-quarter-UPC observations. For each household-quarter-UPC, we calculate the average unit value (expenditures/quantity) and drop observations that are more than three times or less than one third the median unit value for observations in the same market-quarter-UPC, as well as those for which the quantity purchased is more than 24 times the median within the same market-quarter-UPC.

In turn, we collapse the household-quarter-UPC data to a year-UPC panel by summing (scaled by the Nielsen household projection factor) quantities and expenditures by UPC and by year. Annual price is defined as the ratio of annual expenditures and annual quantity.

We calculate the growth rate of each good's price and expenditure between adjacent years (e.g. 2013 price / 2012 price), and identify observations with "extreme growth rates" as instances where the price and/or expenditure growth rate are outside the 1st and 99th percentiles among all year-to-year price and expenditure growth rates for goods with non-zero expenditures in all 8 quarters in adjacent years.

We set $t_1 = 2019$, and $t_0 = 2004, \dots, 2018$. For each t_0 we construct a balanced sample

⁴⁷40,000 households in panel years 2004 to 2006.

of UPCs with non-extreme growth rates and non-zero expenditures in every quarter between t_0 and 2019. In addition, we impose a balanced panel of modules that have at least two unique UPCs available in every quarter from 2004 to 2019. As noted in the main text, this panel of modules also excludes so-called magnet series and "unclassified" module categories. For $t_0 = 2018$, the balanced sample includes 822 modules and 247,611 products (average of 301 products per module, median of 137 products per module). For $t_0 = 2004$, the balanced sample includes the same 822 modules and 32,030 products (average of 39 products per module, median of 17 products per module).

For each t_0 (x-axis in the figure) we construct chained-Tornqvist, long-difference Sato-Vartia, and "welfare-relevant" (equivalent variation at $t_1 = 2019$ preferences) price indices for each module including only those goods in the corresponding t_0 balanced sample. These module price indices are combined into a single aggregate index by weighting each module's price index by its share of expenditure in 2019 among modules in the balanced panel (i.e. every set of t_0 module price indices is aggregated using the same weights, and weights sum to unity for the 822 modules in the sample). For the chained-Tornqvist, for each module we construct year-by-year Tornqvist price indices and cumulate them between t_0 and t_1 . For the long-difference Sato-Vartia, we apply the standard formula to a single change in prices and expenditures between t_0 and t_1 (in contrast to the chained-Tornqvist, which uses prices and expenditures in every year between t_0 and t_1). For welfare, we assume for each module a homothetic-CES aggregator with elasticity of substitution $\theta_0 = 4.5$ (we report robustness to lower and higher values of θ_0). The welfare-relevant price index based on $t_1 = 2019$ preferences, given price changes between t_0 and t_1 , is

$$\log \left(\sum_i b_{it_1} \left(\frac{p_{it_1}}{p_{it_0}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}$$

where b_{it_1} denotes the t_1 budget share of good i within its module among goods in the t_0 -continuing goods sample.

The first panel of Figure 4 reports all three price indices for $t_0 = 2004, \dots, 2018$. Note that, for each t_0 , all three price indices are based on the same sample of products but the sample varies with t_0 due to product entry and exit.⁴⁸

⁴⁸We consider an alternative chained-price index which constructs year-by-year Tornqvist price indices using products with non-extreme growth rates and nonzero expenditures in all 8 quarters in these two adjacent years (but not balanced in the overall period between t_0 and t_1) which we then cumulate between t_0 and t_1 . The resulting inflation is lower than using our balanced Tornqvist index, implying an even larger gap between chained and welfare-relevant inflation. We choose the balanced Tornqvist index as a baseline so that all three indices are based on the same set of observations for any given t_0 , and because in our welfare measure we abstract from entry and exit.

The second panel of Figure 4 compares the gap between the “welfare-relevant” series and the “chained” series to the gap implied by the approximation formula in Proposition 13. Specifically, we report the term $\frac{1}{2} \sum_I b_I Cov_{b(I)} (\Delta \log x, \Delta \log A)$, again weighting modules by their spending share in 2019. We construct $\Delta \log A$ as the difference between each good’s own log-price change and the module’s average log-price change (for goods in the corresponding balanced sample). We construct Δx as the difference between the observed change in expenditure shares and the change in expenditure share implied by a CES aggregator with elasticity $\theta_0 = 4.5$ given observed price changes.

Robustness Figures 5 and 6 replicate Figure 4 using lower and higher values for the elasticity of substitution. The size of the bias gets smaller as we get closer to Cobb-Douglas. This is because in the data changes in prices and changes in expenditure shares are approximately uncorrelated. When demand is Cobb-Douglas, changes in expenditure shares are taste shocks, and since taste shocks are uncorrelated with price changes, following the logic of Proposition 2, the bias is small.

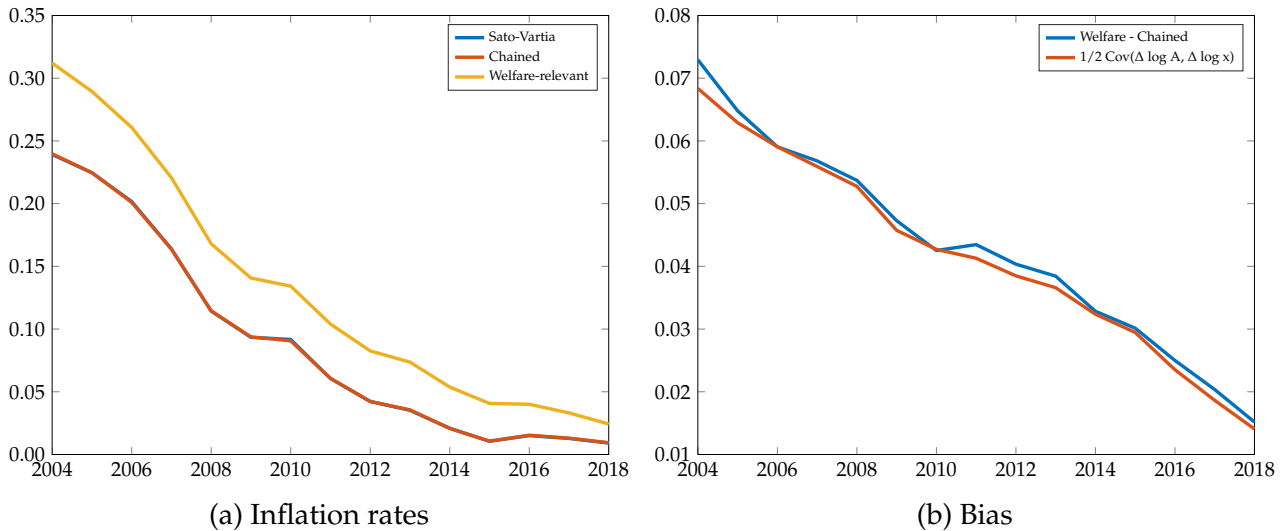


Figure 5: Welfare-relevant, chain-weighted, and Sato-Vartia inflation rate for continuing products. The welfare-relevant rate is computed assuming that the elasticity of substitution across UPCs in the same module is 6.5.

Appendix I Non-CES Functional Forms

In this appendix, we generalize Proposition 6 beyond CES functional forms. To do this, for each producer k with cost function C_k , we define the Allen-Uzawa elasticity of substitution

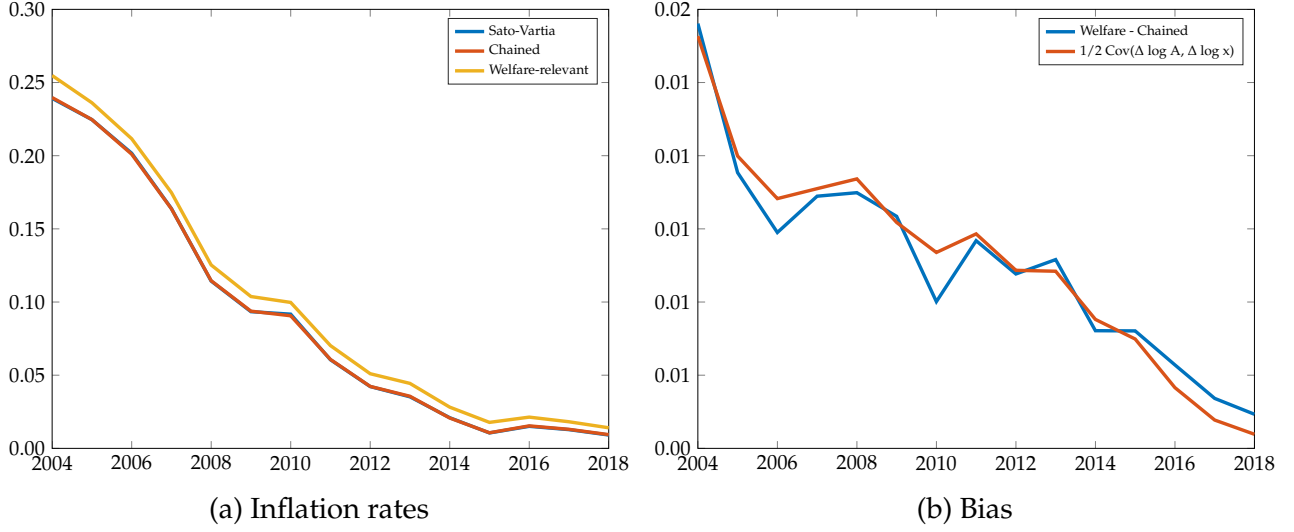


Figure 6: Welfare-relevant, chain-weighted, and Sato-Vartia inflation rate for continuing products. The welfare-relevant rate is computed assuming that the elasticity of substitution across UPCs in the same module is 2.5.

between inputs x and y as

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

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 (2019b), define the *input-output substitution operator* for producer k as

$$\Phi_k(\Psi_{(i)}, \Psi_{(j)}) = - \sum_{1 \leq x, y \leq N+1+F} \Omega_{kx} [\delta_{xy} + \Omega_{ky} (\theta_k(x, y) - 1)] \Psi_{xi} \Psi_{yj}, \quad (43)$$

$$(44)$$

where δ_{xy} is the Kronecker delta. Then, Proposition 6 generalizes as follows:

Proposition 14. *At any point in time t , changes in the relevant variables are pinned down by the following system of equations*

$$d \log p_{it} = - \sum_{j \in N} \Psi_{ijt} d \log A_{jt} + \sum_{f \in F} \Psi_{ift}^F d \log \lambda_{ft}. \quad (45)$$

Changes in sales shares for goods and factors are

$$\begin{aligned} \lambda_{it} d \log \lambda_{it} = & \sum_{j \in \{0\} + N} \lambda_{jt} \Phi_j \left(-d \log p_t, \Psi_{(:,t),t} \right) \\ & + \text{Cov}_{\Omega_{(0,:),t}} \left(d \log x_t, \Psi_{(:,i),t} \right) + \text{Cov}_{\Omega_{(0,:),t}} \left(\varepsilon_t, \Psi_{(:,i),t} \right) \left(\sum_{k \in N} \lambda_{kt} d \log A_{kt} \right). \end{aligned} \quad (46)$$

Changes in welfare-relevant variables are pinned down by the same set of differential equations above where the second line of (46) is set to zero and the boundary conditions are that $\Omega = \Omega_{t_1}$ and $\Psi = \Psi_{t_1}$.

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 14 is very similar to that of Proposition 6. Computing the equilibrium response in Proposition 14 requires solving a linear system exactly as in Proposition 6.