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 2022

First version: April 2021. We thank Conor Foley, Yasutaka Koike-Mori, and Sihwan Yang for outstanding research assistance. We thank Fernando Alvarez, Andy Atkeson, David Atkin, Natalie Bau, Paco Buera, Javier Cravino, Joe Kaboski, Greg Kaplan, Stephen Redding, Esteban Rossi-Hansberg, 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.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 2022
JEL No. E0,E01,E1,E21,E23,E3

ABSTRACT

We present a unified treatment of how welfare responds to changes in budget sets or technologies with taste shocks and non-homothetic preferences. We propose a welfare metric that ranks production possibility frontiers that differs from one that ranks budget sets, and characterize it using a general equilibrium generalization of Hicksian demand. This extends Hulten's theorem, the basis for constructing aggregate quantity indices, to environments with non-homothetic and unstable preferences. We illustrate our results using both long- and short-run applications. 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 our preferred measure of welfare. In the short run, we show that standard chain-weighted deflators understate welfare-relevant inflation for current tastes. Finally, using the Covid-19 recession we illustrate that chain-weighted real consumption and real GDP are unreliable metrics for measuring welfare or production when there are taste shocks.

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? If preferences are homothetic and there are no taste shocks, there is a well-known formula that answers these questions: the chain-weighted (Divisia) index of real consumption.¹ Assuming that there are no income effects or changes in tastes is highly convenient, but also highly inconsistent with reality. In this paper, we relax both of these assumptions and characterize changes in welfare, changes in chained-weighted consumption, and the gap between the two in terms of measurable sufficient statistics.

To compare two different choice sets, indexed by t_0 and t_1 , for a consumer with preferences $\succeq$, we ask: "how much must the t_0 endowment change to make $\succeq$ indifferent between the two choice sets?" If the preference relation at t_0 and t_1 are different, due to changes in tastes, then to answer this question one must take a stance on which preference relation is used for the comparison because the answer depends on this choice.² In this paper, our baseline measure of welfare is the equivalent variation using t_1 preferences, though we also characterize and report results for other welfare questions (i.e. compensating variation and t_0 preferences) and explain how they differ.

We begin by studying this problem in partial equilibrium, taking prices as given, where our welfare metric compares and ranks different budget constraints. This welfare metric is appropriate for an infinitesimal agent whose choices do not alter prices. We call this a *micro* welfare metric. We then propose a generalization of money-metric measures that ranks production possibility frontiers (PPFs) and takes into account that prices are endogenous to collective choices. We call this a *macro* welfare metric. We show that comparing two PPFs using their corresponding equilibrium budget constraints, which is what standard micro money-metrics do, can be misleading if preferences are non-homothetic or unstable.

Calculating micro welfare requires integrating Hicksian (or compensated) demand curves with respect to prices. We prove that there are general equilibrium counterparts to Hicksian demand curves, which we call Hicksian *sales*, and show that integrating them with respect to technologies yields macro welfare. This general equilibrium integral also gener-

---

¹Chain-weighted indices weigh changes in prices or quantities by good-specific weights that are updated every period. The continuous time analog is called a *Divisia* index. Chained-weighted indices are the standard way to measure real economic activity (real GDP, productivity, real consumption, etc.), and their use is justified under homotheticity and stability of preferences (see e.g. Chapters 15 in the UN's System of National Accounts Manual, United Nations 2009, and Chapters 15 and 17 in the CPI Manual, IMF 2004).

²For background on how to account for taste shocks in welfare measures, see Fisher and Shell (1968) and Samuelson and Swamy (1974).

alizes Hulten’s theorem to measure welfare in environments featuring taste shocks and non-homotheticities. In this paper, we focus on neoclassical economies with homogeneous agents, but a companion paper (Baqae and Burstein, 2021) generalizes our results to economies with heterogeneous agents and distortions.

We provide exact and approximate characterizations of the change in micro and macro welfare. In contrast to the standard chain-weighted consumption index formula, which weighs changes in prices or technologies using observed shares, welfare-relevant indices weigh changes in prices or technologies using Hicksian shares. We show that this implies that, compared to our baseline welfare measure, chain-weighted consumption *undercounts* expenditure switching due to income effects or taste shocks (but not substitution effects).

To understand why chain-weighted indices undercount expenditure-switching due to income effects and taste shocks, 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, because older and richer consumers spend more on healthcare. In this case, a chain-weighted consumption index does 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 the chain-weighted consumption index is higher than the welfare-relevant index if income- or taste-driven expenditure-switching is positively correlated with changes in 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 in consumption. 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 and can 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 may become.

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

- i. **Long-run application:** Baumol (1967) argues that aggregate productivity growth slows down if industries with relatively low productivity growth become larger as a share of the economy over time. 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, rather than substitution effects, then our baseline measure of welfare-relevant TFP grew by only 47% instead of 60%.
- ii. **Short-run application:** Our second application looks at shorter horizons but uses more disaggregated (and volatile) data. We compute changes in the welfare-relevant price index and compare these to the chained index using product-level non-durable consumer goods data between 2004 and 2019. We find that the chained-weighted price index understates inflation rates if we use 2019 preferences. This is because goods that became more popular over time, due to taste shocks, have higher inflation rates. At annual frequency, the gap is around 1 percentage point, and it grows to around 4 percentage points over the whole sample. This also means that for past preferences, rather than current preferences, the welfare-relevant inflation rate is lower than the chained index because goods that became less popular experienced lower inflation.
- iii. **Business-cycle application:** Our final application draws on the Covid-19 recession to illustrate the difference between macro and micro 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 micro welfare using mid-2020 preferences, taking changes in prices as given, fell by more than macro welfare, taking into account the fact that changes in prices are themselves caused by demand shocks. Furthermore, unlike macro welfare, real consumption, real GDP, and aggregate TFP are unreliable metrics for measuring changes in productivity capacity, because they depend on irrelevant details like the order in which

supply and demand shocks hit the economy.

In addition to preference stability and homotheticity, a chained index accurately measures welfare only if prices and quantities change continuously and are measured correctly. Many of the well-known reasons why chained indices fail to measure welfare are due to violations of these measurement assumptions. For example, it is well-known that real consumption fails to account for product creation and destruction if we do not 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. Non-homotheticities and taste shocks are different from mismeasured prices because they violate the maintained assumptions about preferences, not prices, and correcting the index requires the use of welfare-relevant, rather than observed, expenditure shares. For this reason, we abstract from these 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.

Other related literature. Measuring changes in welfare using money-metrics is standard in microeconomic theory (see, e.g. chapter 7 in Deaton and Muellbauer, 1980). We characterize the gap between this notion of welfare and real consumption with non-homotheticities and taste shocks. Our general equilibrium results relax the standard assumption in growth accounting that there exists a stable and homothetic final good aggregator (extending Domar 1961, and Hulten 1978). This is also an important maintained assumption in the literature on disaggregated and production network models of the economy (see, for example, Carvalho and Tahbaz-Salehi, 2018 and the references therein).

Our approach focusing on the money-metric at final preferences can be contrasted with common practice in the literature on index numbers, which focuses on Konüs price indices for intermediate levels of utility or tastes between t_0 and t_1 (see, for example, Diewert 1976, Caves et al. 1982, and Feenstra and Reinsdorf 2007). These papers show that, under some assumptions (i.e. translog or CES), commonly used indices like Tornqvist and Sato-Vartia do answer a meaningful question. The advantage of this approach is that constructing these indices requires far less information; the disadvantage is that, unlike our preferred

welfare measure, these indices are not money-metrics that can be used for policy or counterfactual analysis, and they do not provide specific information about the reference indifference curve being used (what budget level, price vector, and preferences it corresponds to). Furthermore, in practice, most index numbers are constructed by chaining, and the aforementioned results do not apply to chained indices. An additional contribution of our paper is to characterize how welfare measures differ from chained (Divisia) indices.³ Finally, relative to this literature, we also provide a unified analysis of non-homotheticity and taste shocks, and we define and characterize a welfare measure for comparing PPFs rather than budget sets (taking into account that prices are endogenous to choices).

A recent and related paper is Redding and Weinstein (2020), who also study welfare changes with both taste shocks and non-homotheticities. Their approach depends on untestable assumptions about cardinal properties of utility functions. On the other hand, we use a money-metric approach that relies only on ordinal preference relations and not on the way utility is cardinalized. We compare our approach to that of Redding and Weinstein (2020) in Section 2.4.

Our approach to calculate ex-post welfare changes requires knowing price changes. When information on prices is incomplete, if preferences are non-homothetic an alternative approach is to infer ex-post changes in welfare by relying on changes in prices, expenditures, price elasticities, and Engel curve slopes for only a subset of goods, given assumptions on separability and stability in preferences (see e.g. Hamilton, 2001 and, more recently, Atkin et al., 2020). In a different vein, Jaravel and Lashkari (2022) provide a procedure to measure microeconomic welfare changes without direct knowledge of elasticities of substitution in the absence of taste shocks and under some additional assumptions. In a separate note (Baqae et al., 2022) inspired by their paper, we show how our results can also be used to accomplish this task under the same assumptions. Unlike the aforementioned papers, in this paper we are not only interested in ex-post welfare measurement with non-homotheticities but also consider taste shocks, counterfactuals, and general equilibrium.

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

---

³Under the assumption that the path of prices is linear, Feenstra and Reinsdorf (2000) shows the equivalence between Divisia and a Konüs price index for an intermediate utility level under AIDS preferences. In practice, price paths tend to be nonlinear (for evidence using scanner-level data, see Ivancic et al., 2011). Therefore, in contrast to Tornqvist and Sato-Vartia, chained indices cannot generically be interpreted as welfare measures corresponding to any 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 choice set. Oulton (2008) discusses how Konüs price indices resolve the path-dependency problem of Divisia indices.

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 or taste shocks— households spend more of their income on some goods as they become richer (e.g. Kongsamut et al., 2001; Boppart, 2014; Comin et al., 2021; Alder et al., 2019) or 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 changes.

The structure of the paper is as follows. In Section 2, we set up the microeconomic problem and provide exact and approximate characterizations of welfare and chain-weighted real consumption. In Section 3, we set up the macroeconomic general equilibrium model and provide exact and approximate characterizations of welfare and chain-weighted 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 how to use our results in dynamic settings and how to treat new goods in Section 6, and conclude in Section 7.

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 this with chain-weighted real consumption. This section builds intuition for Section 3, where we model the equilibrium determination of prices.

2.1 Environment and Definitions

Consider a set of preference relations, $\{\succeq_x: x \in X\}$, over bundles of goods $c \in \mathbb{R}^N$ where N is the number of goods. The vector c includes all relevant goods, and if $\succeq_x$ is intertemporal, then c is a path of current and future consumption bundles.

These preferences are indexed by some vector of parameters, denoted by x , that the consumer does not make choices about but that can affect preference rankings over bundles of goods. For example, x could include calendar time, age, exposure to fads, or state of nature. For every $x \in X$, 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 $U(c; x) > U(c'; x)$ if, and only if, $c \succeq_x c'$, a comparison of $U(c; x')$ and $U(c; x)$ is not meaningful and does not encode any information,

because it is affected by how $U(\cdot; x)$ and $U(\cdot; x')$ are cardinalized.

Note that x should not be interpreted as quality because consumers have preferences and choices over quality but not over parameters of their own utility function (tastes). Hence, the welfare implications of quality changes are different to those of taste changes. Moreover, quality has interpretable cardinal units (e.g. GHz for the CPU of a computer) whereas tastes do not. Quality characteristics must be included as a part of the description of the consumption bundle c , not in x . See Appendix D for a more detailed discussion.

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

Definition 1. Preferences are *homothetic* if whenever $c \sim_x c'$ then $\alpha c \sim_x \alpha c'$ for every $\alpha > 0$.

When $\succeq_x$ is homothetic, we can write $U(c; x)$ so that for every $\alpha > 0$, $U(\alpha c; x) = \alpha U(c; x)$.

Definition 2. Preferences are *stable* if $\succeq_x$ is the same as $\succeq_{x'}$ for every x and x' in X .

If preferences are stable, then the utility function $U(c; x)$ is separable in c and x .

The indirect utility function, for any value of x , is

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

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

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

Since changes in utility do not have meaningful units, we use a money-metric to measure how the consumer values different budget sets. 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 (p_{t_1}, I_{t_1}) is preferred to the initial one (p_{t_0}, I_{t_0}) if, and only if, EV^m is positive.

As Fisher and Shell (1968) point out, this is different to the following question that one may wish to answer: “how much better off is the consumer in t_1 compared to the consumer in t_0 ?” This question is unanswerable using only ordinal choice information. As discussed above, a comparison of two different utility functions, $U(\cdot; x_{t_0})$ and $U(\cdot; x_{t_1})$, depends on how each utility function cardinalizes the underlying preference relation, which is arbitrary because utility functions are defined only up to monotone transformations. Thus, a welfare measure based only on the ordinal preference relation must hold preference parameters x constant in the comparison. For a discussion of papers that adopt a cardinal approach, see Section 2.4.

Other than keeping x constant, we make two other choices in Definition 3. First, we focus on *final* rather than *initial* preferences, and second, we use *equivalent* rather than *compensating* variation. In principle, one could study initial preferences and compensating variation instead. In general, since these alternative measures ask different questions they give different answers, unless preferences happen to be both stable and homothetic. Our methods can be also used to understand these alternative welfare measures, but to streamline the exposition we leave some of the additional details for these alternative measures in Appendix B.

In our baseline, we focus on final preferences $\succeq_{x_{t_1}}$, as opposed to initial preferences $\succeq_{x_{t_0}}$, because for temporal comparisons, the asymmetry of time makes current preferences more relevant than preferences in the past. As Fisher and Shell (1968, page 5) write, “...every practical question which one wants the cost of living index to answer is answered with reference to current, not base-year tastes.” Of course, in some contexts one may be interested in using past preferences to value choice sets, and for counterfactuals one may want to use current rather than future preferences.⁴

We also focus on equivalent variation as our benchmark, rather than compensating variation. We focus on equivalent variation because, unlike the compensating variation, the equivalent variation is a money-metric (see McKenzie and Pearce, 1982). Specifically, the equivalent variation is itself an index of utility which transforms the utility value of different outcomes into dollar values under a common price system (prices in t_0). This is not true for compensating variation, as discussed in Section 2.4.

We now define real consumption based on a chain-weighted index. This is the stan-

---

⁴For optimal policy questions in societies with heterogeneous and/or changing preferences, one must specify a social welfare function which makes the necessary interpersonal utility comparisons. However, our paper is about assigning value to choices given some preferences, and not about how different preferences should be weighed against each other.

standard procedure used by national income accountants and statistical agencies to construct aggregate quantities 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 .⁵

Equation (2) is called a *Divisia* quantity index. In practice, since perfect data is not available in continuous time, statistical agencies approximate this integral via a (Riemann) sum, which they call a *chained* index. We abstract from the imperfections of these discrete time approximations in this paper.⁶ 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 between measured and welfare-relevant objects associated with imperfections in the data, like the lack of quality adjustment, missing prices, or infrequent measurement, are different to the biases we study.

2.2 Characterization of Welfare

To characterize changes in welfare, define the expenditure function for any value of x to be

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

The compensated or Hicksian 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 $c_i(p, u; x)$ is Hicksian demand for good i and preferences x , and the second equality is Shephard's lemma.

---

⁵For any variable z , we use the shorthand dz to denote its change over infinitesimal time intervals, $dz \equiv \frac{dz}{dt} dt$, so that $\Delta z = \int_{t_0}^{t_1} dz$. The last term on the right-hand side of (2) suppresses dependence on t in the integral. We sometimes use this convention to simplify notation.

⁶In 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.

The following lemma expresses changes in welfare given changes in prices and income. It is a key result that we use throughout the paper.

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*

$$EV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i(p_t, u_{t_1}, x_{t_1}) \frac{d \log p_{it}}{dt} dt, \quad (4)$$

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

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

$$EV^m = \log \frac{e(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})}{e(p_{t_0}, v(p_{t_0}, I_{t_0}; x_{t_1}); x_{t_1})} = \Delta \log I - \log \frac{e(p_{t_1}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})}{e(p_{t_0}, v(p_{t_1}, I_{t_1}; x_{t_1}); x_{t_1})}, \quad (5)$$

and recognizing that the second term can be written as the integral in (4).^{7,8}

The budget shares $b_i(p_t, u_{t_1}, x_{t_1})$ that are used to calculate welfare in equation (4) can be reinterpreted as those under prices p_t of a fictional consumer with homothetic and stable preferences with expenditure function $\tilde{e}(p, u) = e(p, u_{t_1}; x_{t_1}) u$. This implies that to compute $b_i(p_t, u_{t_1}, x_{t_1})$ and EV^m , we need to know the budget shares and elasticities of substitution at t_1 but not the income elasticities or the demand shocks. This holds generally, but we illustrate it using the convenient non-homothetic CES functional form below.

Example 1 (Non-homothetic CES.). To illustrate how Lemma 1 can be used, consider a non-homothetic CES example as in Hanoch (1975), Comin et al. (2021), Matsuyama (2019), and Fally (2020). The expenditure function can be written as

$$e(p, u; x) = \left(\sum_{i \in N} x_i p_i^{1-\theta_0} u^{\xi_i} \right)^{\frac{1}{1-\theta_0}}, \quad (6)$$

---

⁷The ratio of expenditure functions, holding fixed utility, is called a Konüs (1939) price index. Equation (5) shows that EV^m requires deflating nominal income by the t_1 utility and preferences Konüs price index. This is in contrast to common practice in the index number theory literature, say Diewert (1976), that uses an intermediate level of utility or preferences between u_{t_0} and u_{t_1} . We discuss how our results relate to this alternative approach in Appendix C.

⁸By definition, EV^m depends only on initial and final prices and income, given t_1 preferences. Therefore, the integral in (4) must be path-independent and yields the same answer under every continuously differentiable path of prices that go from p_{t_0} to p_{t_1} .

with budget shares

$$b_i(p, u, x) = \frac{x_i u^{\xi_i} p_i^{1-\theta_0}}{\sum_{j \in N} x_j u^{\xi_j} p_j^{1-\theta_0}}. \quad (7)$$

The parameter θ_0 is the elasticity of substitution across goods and ξ_i governs the utility elasticity of good i . When the utility elasticities ξ_i are equal for every i , preferences are homothetic. When the set of preference parameters $\{x_i\}$ only scale proportionately, preferences are stable. Note that the underlying preference relation only pins down relative ξ_i 's and relative x_i 's, but not their overall levels. That is, scaling all x_i by the same positive number and adding the same constant to every ξ_i has no effect on choices, but alters the utility associated with each choice.

Welfare-relevant budget shares $b(p, u_{t_1}, x_{t_1})$ as functions of prices p can be calculated as

$$b_i(p, u_{t_1}; x_{t_1}) = \frac{b_{it_1} \times (p_i / p_{it_1})^{1-\theta_0}}{\sum_{j \in N} b_{jt_1} \times (p_j / p_{jt_1})^{1-\theta_0}}. \quad (8)$$

Using (8) in the integral in Lemma 1 yields

$$EV^m = \Delta \log I + \log \left(\sum_{i \in N} b_{it_1} \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}. \quad (9)$$

Hence, to compute $b(p, u_{t_1}, x_{t_1})$ and EV^m , we need to know the final budget shares b_{t_1} and the elasticity of substitution θ_0 , but not the utility elasticities ξ or taste shocks x .⁹ Equation (9) shows when CES preferences are non-homothetic or unstable, the price deflator for EV^m is the exact hat-algebra CES deflator (sometimes referred to as the Lloyd-Moulton index) from t_1 to t_0 rather than the Sato-Vartia index that must be used.

If b_{t_0} is known but b_{t_1} is not, we first have to calculate b_{t_1} before applying (9). Appendix E shows how this can be done given knowledge of taste shocks and income elasticities of demand.

---

⁹In practice, estimating the elasticity of substitution θ_0 must take into account the possibility of demand shocks and income effects. For example, Auer et al. (2021) estimate compensated price elasticities and apply Lemma 1 to measure the heterogeneous welfare effects of changes in foreign prices in the presence of demand non-homotheticities.

2.3 Comparing Welfare and Real Consumption

Using the budget constraint, real consumption in (2) can be expressed in terms of changes in nominal income deflated by price changes:

$$\Delta \log Y = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i(p_t, u_t, x_t) \frac{d \log p_{it}}{dt} dt, \quad (10)$$

where u_t is $v(p_t, I_t; x_t)$ and $b(p_t, u_t, x_t)$ are observed budget shares, b_{it} , at t . In words, changes in real consumption are equal to changes in income minus changes in the consumption price deflator. Changes in real consumption (and 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.

Combining (4) and (10) implies the following proposition.

Proposition 1 (Micro Welfare and Real Consumption). *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*

$$EV^m - \Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} (b_i(p_t, u_t, x_t) - b_i(p_t, u_{t_1}, x_{t_1})) \frac{d \log p_{it}}{dt} dt. \quad (11)$$

This proposition clarifies the differences between welfare and real consumption. Real consumption weighs changes in prices at time t by observed budget shares at time t , taking into account expenditure-switching as it happens. In contrast, welfare takes into account expenditure-switching due to income effects and taste shocks from the beginning, weighing changes in prices at time t by $b_i(p_t, u_{t_1}, x_{t_1})$. Intuitively, EV^m depends on budget shares evaluated at final utility (u_{t_1}) and tastes (x_{t_1}), since EV^m adjusts the level of income in t_0 to make consumers with t_1 preferences as well off as they are in t_1 . For instance, if the consumer prefers the t_1 to the t_0 budget set, then she must be given more income in t_0 to make her indifferent between t_0 and t_1 . As we give her more income in t_0 , the shape of her 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 .¹⁰

---

¹⁰When 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 (4), 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 in this case Marshallian consumer surplus is not the same as chained real consumption.

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 on utility u over time. That is, when preferences are homothetic and stable, $b_{it} = b_i(p_t, u_t, x_t) = b_i(p_t, u_{t_1}, x_{t_1})$ for every path of shocks and every t . Real consumption and welfare are also the same if the difference between the actual and Hicksian budget shares in equation (11) are orthogonal to price changes.

To gain more intuition, we provide a second-order approximation as the time period goes to zero, $t_1 - t_0 = \Delta t \rightarrow 0$. Formally, for any variable z , we write $\Delta \log z \approx (\partial \log z / \partial t) \Delta t + 1/2(\partial^2 \log z / \partial t^2)(\Delta t)^2$, where the remainder in the approximation is of order $(\Delta t)^3$.¹¹ Furthermore, throughout we use the shorthand $\mathbb{E}_z(y)$ for $\sum_i z_i y_i$ for any pair of vectors z and y where $z_i \geq 0$ and $\sum_i z_i = 1$. Similarly, $\text{Cov}_z(w, y)$ is shorthand for $\mathbb{E}_z(wy) - \mathbb{E}_z(w)\mathbb{E}_z(y)$.

Proposition 2 (Approximate Real Consumption and Micro Welfare). *To a second-order approximation, the change in real consumption is*

$$\Delta \log Y \approx \Delta \log I - \mathbb{E}_{b_{t_0}}(\Delta \log p) - \underbrace{\frac{1}{2} \text{Cov}_{b_{t_0}}(\Delta \log b, \Delta \log p)}_{\text{expenditure-switching}}, \quad (12)$$

and the change in welfare is

$$EV^m \approx \Delta \log I - \mathbb{E}_{b_{t_0}}(\Delta \log p) - \text{Cov}_{b_{t_0}}(\Delta \log b, \Delta \log p) + \underbrace{\frac{1}{2} \sum_j \Delta \log p_j \frac{\partial b_i}{\partial \log p_j} \Delta \log p_i}_{\substack{\text{expenditure-switching} \\ \text{due to substitution effect}}}, \quad (13)$$

where $\partial b / \partial \log p$ is the partial derivative of the Hicksian budget share with respect to prices p evaluated at t_0 .

We discuss real consumption and welfare in turn. The first two terms in (12) are just the change in nominal income deflated by average price changes. The last term (multiplied by 1/2) captures how expenditure-switching affects real consumption. If expenditures rise for goods whose prices are rising, then this lowers real consumption. In principle, the expenditure shares can change for three different reasons: substitution effects, income effects, and taste shocks.

Equation (13) shows that the first-order terms in welfare are identical to real consumption, but discrepancies are present at the second order. Specifically, welfare treats changes

---

¹¹For local approximations throughout the paper, we assume that the exogenous parameters are smooth functions of t and that the expenditure function is a smooth function of preference parameters x .

in budget shares due to substitution effects differently to changes in budget shares due to taste shocks or income effects. The adjustment term in (13) implies that welfare puts more weight on expenditure-switching if it is due to income effects or taste shocks than if it is due to substitution effects. To better understand the expenditure-switching terms, consider the non-homothetic CES demand system. To apply Proposition 2, we need to approximate changes in expenditure shares only to the first-order. This is because $\Delta \log b$ is always multiplied by $\Delta \log p$, so any nonlinear terms in $\Delta \log b$ will have order higher than two once they are multiplied by $\Delta \log p$. Log-linearizing the budget shares (7) with respect to time yields¹²

$$d \log b_i = \underbrace{[1 - \theta_0] [d \log p_i - \mathbb{E}_b(d \log p)]}_{\text{substitution effects}} + \underbrace{[\varepsilon_i - 1] [d \log I - \mathbb{E}_b(d \log p)]}_{\text{income effects}} + \underbrace{d \log x_i}_{\text{taste shifter}}, \quad (14)$$

where ε is the vector of income elasticities of demand for each good, and $d \log x$ is the vector of demand shifters due to taste shocks.¹³ Substituting (14) into Proposition 2 results in the following proposition.

Proposition 3 (Approximation with Non-Homothetic CES). *For a consumer with non-homothetic preferences, defined by (6), to a second-order approximation, the change in real consumption is*

$$\begin{aligned} \Delta \log Y \approx & \Delta \log I - \mathbb{E}_{b_{t_0}}(\Delta \log p) - \underbrace{\frac{1}{2}(1 - \theta_0) \text{Var}_{b_{t_0}}(\Delta \log p)}_{\substack{\text{expenditure-switching} \\ \text{due to substitution effect}}} \\ & - \frac{1}{2} \underbrace{\text{Cov}_{b_{t_0}}(\Delta \log x, \Delta \log p)}_{\substack{\text{expenditure-switching} \\ \text{due to taste shock}}} - \frac{1}{2} \underbrace{\left[\Delta \log I - \mathbb{E}_{b_{t_0}}(\Delta \log p) \right] \text{Cov}_{b_{t_0}}(\varepsilon, \Delta \log p)}_{\substack{\text{expenditure-switching} \\ \text{due to income effect}}}, \end{aligned} \quad (15)$$

and the change in welfare is

$$\begin{aligned} EV^m \approx & \Delta \log I - \mathbb{E}_{b_{t_0}}(\Delta \log p) - \frac{1}{2}(1 - \theta_0) \text{Var}_{b_{t_0}}(\Delta \log p) \\ & - \text{Cov}_{b_{t_0}}(\Delta \log x, \Delta \log p) - \left[\Delta \log I - \mathbb{E}_{b_{t_0}}(\Delta \log p) \right] \text{Cov}_{b_{t_0}}(\varepsilon, \Delta \log p). \end{aligned} \quad (16)$$

Consider the expenditure-switching terms in (15) one-by-one. If goods are substitutes, $\theta_0 > 1$, then welfare is convex in prices and variance in price changes boosts welfare by

¹²See Appendix E for a more detailed derivation. We assume here that all goods have positive expenditure shares. We discuss new goods in Section 6.

¹³In terms of primitives of the utility function, $\varepsilon_i = 1 + (1 - \theta_0) (\xi_i / \mathbb{E}_b(\xi) - 1)$ and $d \log x_i = d \log x_i - (\xi_i / \mathbb{E}_b(\xi)) \mathbb{E}_b(d \log x)$. As expected, only relative ξ 's and relative x 's are identified but not the overall levels $\mathbb{E}_b(\xi)$ and $\mathbb{E}_b(d \log x)$, since these are not pinned down by the underlying preference relation.

allow substitution towards goods that become relatively cheap. The second line of (15) captures the effect of taste shocks and income effects. If the composition of demand shifts in favor of goods that become relatively cheap, either due to taste shocks $Cov_b(\Delta \log x, \Delta \log p) < 0$ or income effects $Cov_b(\varepsilon, \Delta \log p)(\Delta \log I - \mathbb{E}_b(\Delta \log p)) < 0$, then real consumption increases.

Having understood (15), consider now expenditure-switching terms in (16). Welfare places the same weight on substitution effects as real consumption does. However, it places a larger weight on expenditure-switching due to income effects and taste shocks. Whereas $\Delta \log Y$ only takes into account expenditure-switching as it occurs over time, EV^m accounts for expenditure-switching due to income effects and taste shocks from the start. Therefore, expenditure-switching due to income and tastes are multiplied by $1/2$ for $\Delta \log Y$ and 1 for EV^m . 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 $EV^m > \Delta \log Y$. Even if preferences are unstable or non-homothetic, real consumption strays from welfare only if income effects or taste shocks covary with price changes.¹⁴

2.4 Alternative Welfare Measures

In this section, we discuss two alternative welfare measures that are widely used in the literature: changes in utility indices and compensating variation.

Utility Index. A common approach in the literature is to use u_t as a measure of welfare and define the associated "ideal" price index $P_t \equiv e(p_t, u_t; x_t)/u_t$ by analogy to the homothetic and stable case. This is particularly common with CES preferences with either taste shocks or non-homotheticities (e.g. Comin et al., 2021; Redding and Weinstein, 2020; Ehrlich et al., 2019; Argente et al., 2021). If the utility function is not changing, then u_t obviously ranks choices according to the underlying preference relation, but the magnitude of the changes in u_t and the associated price index P_t are not interpretable. This is because

---

¹⁴In Section 2.2, 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 (16) 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 and take an approximation as t_0 moves back in time, 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 Example 1. Both expressions are valid second-order approximations and in either case, real consumption undercounts expenditure-switching caused by income effects or taste shocks.

they depend on the overall level of utility elasticities ξ , which are arbitrary. Furthermore, when preferences are non-homothetic, there is no normalization of the ξ parameters such that changes in u_t are equal to EV^m (see Appendix E.2).

This issue is even more severe when there are taste shocks. In this case, changes in u_t and P_t depend on the overall change in taste shifters x_t . Once again, the overall level of x_t in (6) is not pinned down by choice behavior. Therefore, to compute u_t one needs to assume a value for the overall level of x_t 's period-by-period. This means u_t no longer corresponds to the ranking of choices according to any preference relation, but is rather determined by the untestable assumption that determines the overall level of x_t 's in each period. For example, Redding and Weinstein (2020) set the average of x_t 's to be constant over time. However, the same data is also consistent with the average of x_t 's rising or falling in arbitrary ways through time. Each assumption results in a different answer. As Muellbauer (1975) forcefully argues: "It is not possible to set up a market experiment which would unambiguously indicate whether the consumer's efficiency as a utility machine has changed." As he points out, "no such problems attend an ordinal cost-of-living index" like the ones we characterize.¹⁵

Compensating variation. Our baseline measure of welfare changes is equivalent variation. An alternative is the compensating variation, which is the reduction in t_1 income that makes the consumer indifferent to their t_0 budget set. As mentioned earlier, EV^m is a money-metric that ranks all choices, whereas CV^m is not. To see this, consider a household in t_0 choosing between (p_{t_1}, I_{t_1}) and (p'_{t_1}, I'_{t_1}) . It follows that $u'_{t_1} = v(p'_{t_1}, I'_{t_1}; x) > v(p_{t_1}, I_{t_1}; x) = u_{t_1}$ if, and only if, $e(p_{t_0}, u'_{t_1}; x)/I_{t_0} > e(p_{t_0}, u_{t_1}; x)/I_{t_0}$. That is, the household prefers whichever choice has the highest EV^m . Hence, EV^m is itself an index of utility. On the other hand, CV^m cannot be used to compare t_1 and t'_1 because CV^m for t_1 and t'_1 are expressed in terms of income adjustments in t_1 and t'_1 prices respectively. Compensating variation can be used to compare t_0 to t_1 and t_0 to t'_1 but it cannot be used to compare t_1 to t'_1 .

Nevertheless, CV^m can be characterized using similar methods, as shown in Appendix B. For example, every result in the paper can be translated into compensating variation under initial preferences simply by reversing the flow of time. That is,

$$CV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i(p_t, u_{t_0}, x_{t_0}) \frac{d \log p_{it}}{dt} dt. \quad (17)$$

In particular, EV^m preserves the shape of the indifference curve at t_1 , whereas CV^m pre-

---

¹⁵See also Balk (1989) and Martin (2020) for related discussion of the cardinal and ordinal approaches.

serves the shape of the indifference curve at t_0 .¹⁶ To a second order starting at t_0 , for a consumer with non-homothetic CES preferences,

$$CV^m \approx \Delta \log I - \mathbb{E}_{b_{t_0}} (\Delta \log p) - \frac{1}{2}(1 - \theta_0) \text{Var}_{b_{t_0}} (\Delta \log p).$$

Comparing this with (16) reveals that EV^m accounts for expenditure-switching due to income effects whereas CV^m does not. Furthermore, up to a second-order approximation, $CV^m - \Delta \log Y = \Delta \log Y - EV^m$. Hence, the average of EV at final preferences and CV at initial preferences is equal to real consumption $\Delta \log Y$, up to a second order. See Appendix B for a generalization beyond non-homothetic CES.

3 Macroeconomic Changes in Welfare and Consumption

In the previous section we showed how to value changes in budget sets for a consumer who takes prices as given. For a whole society however, prices are endogenous to collective choices. In this section, we extend our analysis to show how to assign value to different production possibility frontiers (PPFs) rather than budget sets. The *macro*, or society-level, notion of welfare we define is not the same as micro welfare evaluated in general equilibrium.

We first introduce a closed, representative-agent, neoclassical economy with non-homothetic and unstable preferences. We generalize our definitions of welfare, now at the macroeconomic level, and present exact and approximate expressions for real GDP and welfare in terms of endogenous sufficient statistics. In the next section, Section 4, we solve for these endogenous sufficient statistics in terms of observable primitives. Without non-homotheticities and taste shocks, our economy is the same as the one studied by Baqaee and Farhi (2019). We extend their approximation formulas to capture real GDP and welfare when there are income effects and taste shocks.

---

¹⁶Hence, calculating CV^m requires knowledge of initial budget shares and elasticities of substitution, whereas equivalent variation EV^m requires knowledge of final budget shares and elasticities of substitution. This means that EV^m at final preferences is more convenient for ex-post comparisons and CV^m at initial preferences is more convenient for counterfactuals.

3.1 Environment and Definitions

Consider a perfectly competitive neoclassical closed economy with a representative agent.¹⁷ 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.¹⁸ 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.¹⁹ 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 (for both goods and factors) by $p(A, L, x)$ and aggregate income, or nominal GDP, by $I(A, L, x)$. 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

¹⁷In Baqaee and Burstein (2021) we show that our results can be generalized to economies with heterogeneous agents and distortions. To measure the change in welfare from t_0 to 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 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. We capture distortions by including wedges as part of the primitives of the economy.

¹⁸Note that unlike utility, physical output of each individual good has interpretable units and can be measured directly, without the need to use a money-metric.

¹⁹Allowing for endogenous labor-leisure choice requires including the time endowment in L and leisure in the consumption bundle c .

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}, 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

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

The superscript M in EV^M represents the fact that this is the *macro* equivalent variation, in contrast to EV^m for the *micro* welfare measure. In words, EV^M is the proportional change in initial factor endowments required to make the consumer 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. That is, EV^M is itself an index of utility and $(A, L) \succeq_x (A', L')$ if, and only if, $EV^M(A_{t_0}, L_{t_0}, A, L; x) \geq EV^M(A_{t_0}, L_{t_0}, A', L', x)$. The macro equivalent variation, EV^M , is a useful metric because it ranks PPFs without reference to endogenous prices.

Definition 5 shows that EV^M is a generalization of EV^m in the sense that the two coincide when the PPF is the same as the budget constraint. However, EV^M is also a generalization of *consumption-equivalents* commonly used to measure welfare in macroeconomics (e.g. Lucas, 1987). To see this, note that EV^M can also be defined in the following way.

Definition 6 (Alternative Definition of Macro Welfare). Let $\mathcal{C}_t$ denote the production possibility set associated with (A_t, L_t) . Then $EV^M = \phi$, where ϕ solves

$$V(\mathcal{C}_{t_1}; x_{t_1}) = V(e^\phi \mathcal{C}_{t_0}; x_{t_1}),$$

and $e^\phi \mathcal{C}_{t_0} \equiv \{e^\phi c : c \in \mathcal{C}_{t_0}\}$.

In words, EV^M also measures the proportional shift in $\mathcal{C}_{t_0}$ necessary to make $\succeq_{x_{t_1}}$ indifferent between $e^\phi \mathcal{C}_{t_0}$ and $\mathcal{C}_{t_1}$. This definition is isomorphic to Definition 5. This is because scaling all factor endowments by a constant shifts out the PPF by the same constant under our assumptions.

For a graphical representation of EV^m and EV^M , see Figure 1. Macro and micro welfare answer different questions. To see the difference, consider a situation where households age between t_0 and t_1 but technologies and factor endowments stay the same. Since the

Figure 1: The left panel shows the change in t_0 income that makes the household with final preferences indifferent between the budget constraint in t_0 and t_1 . The right panel shows the change in t_0 endowments that makes the household with final preferences indifferent between the PPF in t_0 and t_1 . EV^m and EV^M are equal if the PPF is linear, or if preferences are homothetic and stable.

PPF is unchanged, the change in macro welfare is zero by construction. However, if the PPF is nonlinear, the relative price of goods changes between t_0 and t_1 : prices rise for goods that become more desirable. In this case, EV^m necessarily falls even though the PPF is unchanged. Hence, EV^m does not rank PPFs for a society.

The following examples further illustrate why micro welfare can be a misleading measure of the welfare impact of technological change when preferences are non-homothetic or unstable.

Example 2 (Micro versus macro welfare and growth). Consider economies that are endowed with cows, and suppose each cow produces one unit of steak and one unit of offal. At t_0 the economy is endowed with one cow. We consider what happens if we double the number of cows at t_1 . For any stable and homothetic preferences, macro and micro welfare doubles ($EV^M = EV^m = \log 2$).

Panel (a) in Figure 2 considers a case with non-homothetic but stable preferences. For this example, suppose that the consumer at t_0 considers steak and offal to be perfect substitutes because the consumer is poor and only cares about total calories. Hence, the relative price of offal and steak is one and the consumption bundle is one unit of steak and one unit offal in t_0 . At t_1 , the consumer consumes two units of steak and two units of offal. If steak is an extreme luxury good, then at t_1 , the relative price of steak is much higher than offal. Macro welfare measures the distance between the original PPF and the new PPF, which is twice as big since the economy in t_1 produces double the number of both steaks and offal. However, in this example, micro welfare only rises by an arbitrarily small number ε . That is, even though the economy is twice as good, micro welfare has increased by almost zero. Intuitively, in the initial equilibrium, steak was relatively cheap (because consumers were poor), so initial income in t_0 only needs to be raised by ε to allow the consumer to reach the

t_1 indifference curve.

In other words, micro welfare measures the fact that an infinitesimal consumer with an endowment of $1 + \varepsilon$ cows could purchase a total of $2 + \varepsilon$ steaks. However, if all consumers' endowments were raised by ε , and all consumers tried to sell their offal to buy steak, then the relative price of steak would quickly rise because the aggregate supply of steaks has only risen by ε . This means that, in practice, consumers with $1 + \varepsilon$ cows are not as well off as they thought they would be. Hence, to make them indifferent, we would have to give them more endowment — doing this, would change relative prices again. Macro welfare is the fixed point to this problem, where the increase in the endowment is exactly enough to make consumers indifferent taking into account the fact that relative prices are endogenous.²⁰

Figure 2: The left panel shows an example with non-homothetic preferences in which the PPF doubles but micro welfare increases by an arbitrarily small ε . The right panel shows an example in which the PPF expands by $2 - \varepsilon$ but micro welfare falls. Indifference curves are blue, PPFs are green.

Panel (b) in Figure 2 considers a case with homothetic but unstable preferences instead. Starting with the same PPF as before in t_0 , suppose that in t_1 the number of cows increases by a factor $2 - \varepsilon$. Due to a concurrent advertising campaign, consumer preferences shift

²⁰Surprisingly, even though macro welfare is the solution to this seemingly complicated fixed point problem, our analytical results will show that it is in fact simpler to characterize than micro welfare is in general equilibrium.

so that in t_1 steak becomes more popular than offal (the indifference curves in Figure 2 (b) become flatter). In this example, since the number of cows has increased by a factor $2 - \varepsilon$, the macro welfare change is $\log(2 - \varepsilon)$. However, micro welfare declines because the consumer with t_1 preferences strictly prefers the t_0 budget set. Intuitively, a consumer with t_1 preferences could sell the offal in t_0 and purchase a total of 2 steaks, which is preferred to $2 - \varepsilon$ steaks she can get in t_1 . That is, even though the economy has almost doubled in terms of its productive capabilities, micro welfare in general equilibrium has fallen. Hence, micro welfare in general equilibrium is not a good measure of how changes in choice sets affect societal welfare.²¹

The previous examples illustrate that using the initial budget set to represent the initial PPF 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.

When relative prices do not respond to consumers' choices (i.e. the PPF is linear), 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. The following proposition formalizes this:

Proposition 4 (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 the PPF is linear.*

We provide a quantitative illustration of the difference between micro and macro welfare using the Covid-19 crisis in Section 5.

3.2 Characterization of Welfare and Real GDP

To characterize macro welfare, we define a Hicksian (or compensated) economy indexed by (A, L, u, x) , where (A, L) determines the PPF and (u, x) specifies an indifference curve corresponding to $U(c; x) = u$.

Definition 7 (Hicksian economy). Consider a fictitious consumer whose preferences are represented by the expenditure function $\tilde{e}(p, \tilde{u}) \equiv e(p, u; x)\tilde{u}$, for fixed values of u and x . In words, $\tilde{e}(p, \tilde{u})$ maps prices p and utility values $\tilde{u}$ into required expenditures of the fictitious consumer. The *Hicksian* economy, indexed by (A, L, u, x) , is an economy with PPF (A, L) and expenditure function $\tilde{e}(p, \tilde{u})$.

---

²¹ Although EV^m and EV^M can have different signs when preferences are unstable, they have the same sign when preferences are stable. This is because u_t is equal in the micro and macro problem (by the first welfare theorem), and $EV^m > 0$ and $EV^M > 0$ if, and only if, $u_{t_1} > u_{t_0}$.

In the definition above, u and x are parameters that define the fictional consumer's preferences. When the true consumer's preferences, represented by the expenditure function $e(\cdot)$, are stable and homothetic, then they coincide with the fictional consumer's preferences, represented by $\tilde{e}(\cdot)$. In this case, the Hicksian economy is the same as the uncompensated original economy. We refer to variables, like prices and sales, in the Hicksian economy as Hicksian variables and index them by (A, L, u, x) . We refer to variables in the uncompensated original economy as *observed* variables and index them by t .

Denote the sales shares relative to GDP of each good or factor i by

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

where $\mathbf{1}$ is an indicator function, and N and F are the set of goods and factors. The sales share λ_i is often referred to as the *Domar* weight of i . Note that referring to λ_i as a “share” is an abuse of language since $\sum_{i \in N} \lambda_i > 1$ whenever there are intermediate inputs.

To characterize macro welfare and compare it to real GDP, we use *Hicksian* sales shares, $\lambda(A, L, u, x)$, which are the sales shares in a Hicksian economy. The Hicksian sale shares are the general equilibrium counterpart to partial equilibrium Hicksian budget shares. The following proposition shows that changes in macro welfare is the integral of Hicksian sales shares with respect to changes in technologies and factor quantities.

Proposition 5 (Macro Welfare). *For any smooth path of technologies and factor quantities that unfold as a function of time t , changes in macro welfare are*

$$EV^M = \underbrace{\int_{t_0}^{t_1} \sum_{i \in N} \lambda_i(A_t, L_t, u_{t_1}, x_{t_1}) \frac{d \log A_{it}}{dt} dt}_{\text{welfare-relevant TFP growth}} + \underbrace{\int_{t_0}^{t_1} \sum_{i \in F} \lambda_i(A_t, L_t, u_{t_1}, x_{t_1}) \frac{d \log L_{it}}{dt} dt}_{\text{welfare-relevant factor growth}}. \quad (18)$$

In words, growth accounting for welfare should be based on compensated or Hicksian sales shares evaluated at current technology but final preferences and utility. We define the first N summands of (18) to be changes in *welfare-relevant TFP* and the last F summands are changes in welfare due to changes in factor inputs. As we show below, this is analogous to how measured TFP is defined using real GDP.

To apply Proposition 5, the only thing we need to know about preferences is how $e(p, u_{t_1}; x_{t_1})$ varies as a function of prices for given u_{t_1} and x_{t_1} . This uniquely pins down the fictional consumers' preferences that generate the Hicksian shares $\lambda(A_t, L_t, u_{t_1}, x_{t_1})$. In other words, we need to know the budget shares and elasticities of substitution at the final allocation but not income elasticities or taste shocks.²²

²²Following the observation made in Section 2.4, for compensating variation at initial preferences, we need

We now consider changes in real GDP, defined using the Divisia index for final goods as $\Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} b_{it} d \log c_{it}$. The following result shows that changes in real GDP can be calculated by integrating observed sales shares with respect to technology changes.

Proposition 6 (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*

$$\Delta \log Y = \underbrace{\int_{t_0}^{t_1} \sum_{i \in N} \lambda_i(A_t, L_t, u_t, x_t) \frac{d \log A_{it}}{dt} dt}_{TFP \text{ growth}} + \underbrace{\int_{t_0}^{t_1} \sum_{i \in F} \lambda_i(A_t, L_t, u_t, x_t) \frac{d \log L_{it}}{dt} dt}_{\text{factor growth}}, \quad (19)$$

where $\lambda(A_t, L_t, u_t, x_t)$ are observed sales shares λ_{it} at t .

In (19), 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. Proposition 6 is a slight generalization of Hulten (1978) to environments with unstable and non-homothetic final demand. In general equilibrium, the sales shares play the role that budget shares played in partial equilibrium. Whereas in partial equilibrium, integrating budget shares with respect to prices yielded real consumption, in general equilibrium integrating sales shares with respect to technologies and factors yields real GDP. If the PPF is not moving, that is, technology A and endowments L are constant, then pure preference shocks have no effect on real GDP even though they can change allocations.

Propositions 5 and 6 show that the difference between macro welfare and real GDP can be understood in terms of the distinction between Hicksian and observed sales shares. If preferences are homothetic and stable, then $\lambda(A, L, u_{t_1}, x_{t_1}) = \lambda_{it}$ and the change in welfare is equal to real GDP.

To get more intuition for the difference between welfare and real GDP, we use a second-order approximation of the response of real GDP and welfare to technology and preference shocks as $t_1 - t_0 = \Delta t \rightarrow 0$. To make the formulas more compact and without loss of generality, when we write local approximations we abstract from shocks to factor endowments ($\Delta \log L = 0$).

Proposition 7 (Approximate Real GDP and Macro Welfare). *Up to a second order approximation, the change in real GDP is*

$$\Delta \log Y \approx \sum_{i \in N} \lambda_i \Delta \log A_i + \frac{1}{2} \sum_{i \in N} \Delta \lambda_i \Delta \log A_i, \quad (20)$$

---

to know elasticities of substitution at the initial allocation instead of the final one.

and the change in welfare is

$$EV^M \approx \sum_{i \in N} \lambda_i \Delta \log A_i + \sum_{i \in N} \Delta \lambda_i \Delta \log A_i - \underbrace{\frac{1}{2} \sum_{j \in N} \Delta \log A_j \frac{\partial \lambda_i}{\partial \log A_j} \Delta \log A_i}_{\text{expenditure-switching due to substitution effects}} , \quad (21)$$

where $\partial \lambda / \partial \log A$ is the partial derivative of the Hicksian sales share with respect to technology, and all terms are evaluated at t_0 .

We discuss (20) and (21) in turn. The first term in (20) is the Hulten-Domar formula. The second term in (20), emphasized by Baqaee and Farhi (2019), captures nonlinearities due to changes in sales shares. Intuitively, if sales shares decrease for those goods with higher productivity growth, then real GDP growth slows down.

Equation (21) 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. Equation (21) shows that welfare puts more weight on changes in sales shares that are due to demand shocks and non-homotheticities than those due to substitution effects. That is, real GDP “undercorrects” for changes in shares caused by non-homotheticities or changes in preferences, similar to the partial equilibrium counterparts in Proposition 2 and 3.

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 compensated sales shares as given. In this section, we solve for changes in these endogenous objects in terms of microeconomic sufficient statistics. For clarity, we restrict attention to nested-CES economies. The general case is in Appendix J, and the intuition is very similar. This section also contains some worked-out analytical examples.

Nested-CES economies. Household preferences are non-homothetic CES, with expenditure function (6), where once again θ_0 is the elasticity of substitution across consumption goods. 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, 2019).

Input-output matrix. We stack the expenditure shares of the 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 household, we index it 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 $i \notin N$). 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:

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

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' . The next N rows and columns correspond to goods, and the last F rows and columns correspond to the factors. The right F columns

of Ψ represent the network-adjusted factor intensities of each good. The sum of network-adjusted factor intensities for every 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 identity that $\lambda' = b'\Psi$.

4.1 General characterization for nested-CES economies

According to Propositions 5 and 6, changes in real GDP and welfare can be computed by integrating observed and Hicksian sales shares with respect to technology shocks. The following proposition pins down sales shares as a function of primitives. This proposition can then be used in combination with Propositions 5 and 6 to calculate exact changes in real GDP and welfare. For readability, we again assume away shocks to factor endowments.

Proposition 8 (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 differential equations:*

$$d \log p_{it} = - \underbrace{\sum_j \Psi_{ijt} d \log A_{jt}}_{\text{upstream TFP changes}} + \underbrace{\sum_{f \in F} \Psi_{ift} d \log \lambda_{ft}}_{\text{upstream factor price changes}}, \quad (22)$$

$$d\Psi_{ilt} = \underbrace{\sum_j \Psi_{ijt} (\theta_j - 1) \text{Cov}_{\Omega_{(j,:),t}} \left(-d \log p_t, \Psi_{(:,l),t} \right)}_{\text{substitution effect}} + \underbrace{\mathbf{1}_{\{i=0\}} \text{Cov}_{\Omega_{(0,:),t}} \left(d \log \mathbf{x}_t + \varepsilon_t d \log Y_t, \Psi_{(:,l),t} \right)}_{\text{taste shocks and income effect}}, \quad (23)$$

where changes in sales shares are $d\lambda_{it} = d\Psi_{0it}$ and changes in real GDP are $d \log Y_t = \sum_i \lambda_{it} d \log A_{it}$. As before, the demand shifter for good i is $d \log \mathbf{x}_{it} = d \log x_{it} - (\xi_i / \mathbb{E}_{b_t}[\xi]) \mathbb{E}_{b_t}[d \log x]$ and the income elasticity of demand for good i is $\varepsilon_{it} = 1 + (1 - \theta_0) (\xi_i / \mathbb{E}_{b_t}[\xi] - 1)$.²³

To compute changes in real GDP using Proposition 6, we need to know $\lambda_{it} = \lambda(A_t, L_t, u_t, x_t)$ for $t \in [t_0, t_1]$. These are solutions to the differential equations above given some boundary condition that pins down Ψ at some point in time $t \in [t_0, t_1]$.

On the other hand, to compute macro welfare, we need Hicksian sales shares $\lambda(A_t, L_t, u_{t_1}, x_{t_1})$ as a function of t . These are also solutions to the differential equations above, except that the terms involving taste shocks and income effects in (23) are set to zero. In this case, the boundary condition is that the Leontief inverse at t_1 is equal to the observed Leontief

²³For 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 .

inverse Ψ_{t_1} at t_1 . Therefore, if Ψ_{t_1} is observed, we can calculate Hicksian sales shares between t_0 and t_1 by starting (23) 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 $\Delta \log x$.

Each term in the differential equations in Proposition 8 has a clear interpretation. We start by discussing the equation determining prices (22). 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).

On the other hand, the equation for $d \log \Psi_{ilt}$ 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 Hicksian version of this equation is identical except that it does not adjust expenditures due to income effects and taste shocks along the transition.

Remark 1 (Micro Welfare). Proposition 8 can also be used to compute changes in microeconomic welfare EV^m in a general equilibrium model. To do this, we compute $p(A_{t_0}, L_{t_0}, u_{t_0}, x_{t_0})$ and $p(A_{t_1}, L_{t_1}, u_{t_1}, x_{t_1})$ using Proposition 8 and then plug these price changes into (9). Unlike macroeconomic welfare, microeconomic welfare changes generically require knowledge of both income elasticities and taste shocks since they are needed to calculate the initial and or final prices.

To build more intuition about the difference between real GDP and macro welfare, consider economies with only a single factor of production. In this case the system is simplified since we know that the economy's single primary factor always has a sales share equal to one. This allows for a simple characterization of both welfare and real GDP up to a second-order approximation.

Proposition 9 (Approximate Macro Welfare vs GDP: Single Factor). *When the economy has one factor of production, up to a second order approximation, the change in real GDP is*

$$\begin{aligned} \Delta \log Y \approx & \sum_{i \in N} \lambda_i \Delta \log A_i + \frac{1}{2} \sum_j \lambda_j (\theta_j - 1) \text{Var}_{\Omega_{(j,:)}} \left(\sum_{i \in N} \Psi_{(:,i)} \Delta \log A_i \right) \\ & + \frac{1}{2} \text{Cov}_{\Omega_{(0,:)}} \left(\Delta \log x + \left(\sum_{i \in N} \lambda_i \Delta \log A_i \right) \varepsilon, \sum_{i \in N} \Psi_{(:,i)} \Delta \log A_i \right), \end{aligned} \quad (24)$$

and welfare is

$$EV^M \approx \Delta \log Y + \frac{1}{2} \text{Cov}_{\Omega_{(0,:)}} \left(\Delta \log x + \left(\sum_{i \in N} \lambda_i \Delta \log A_i \right) \varepsilon, \sum_{i \in N} \Psi_{(:,i)} \Delta \log A_i \right), \quad (25)$$

where $\lambda, \Omega, \Psi, \varepsilon$, and $\Delta \log x$ are evaluated at t_0 .

Proposition 9 is a general equilibrium counterpart to Proposition 3. We first discuss real GDP. The first term in Equation (24) 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 respond to changes in relative prices caused by technology shocks (these effects are emphasized by Baqaee and Farhi, 2019). The terms on the second line of (24), which are the ones we focus on in this paper, capture changes in sales shares due to changes in tastes or non-homotheticities.

Equation (25) 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 price and demand shocks (see Proposition 3), equation (25) shows that in general equilibrium, the relevant statistic is the covariance of demand shocks with a network-adjusted notion of supply shocks.

4.2 Analytical Examples

We now work through some simple examples to illustrate the forces that drive a gap between real GDP and welfare.

Example 3 (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 9 simplifies to

$$EV^M - \Delta \log Y \approx \frac{1}{2} (Cov_b(\Delta \log x, \Delta \log A) + Cov_b(\varepsilon, \Delta \log A) \mathbb{E}_b[\Delta \log A]). \quad (26)$$

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. By Proposition 4, in this example, macro welfare EV^M is the same as micro welfare EV^m .

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

Example 4 (Decreasing Returns to Scale). Consider the one-sector model without intermediate inputs in Example 3 but now suppose that production functions are non-constant-returns-to-scale. Specifically, the production for good i is

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

where L_i is labor and γ need not equal 1. 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.

For simplicity, suppose that preferences are homothetic ($\varepsilon_i = 1$ for every i), but potentially unstable ($\Delta \log x \neq 0$). We apply Proposition 7 to compute the difference between welfare and real GDP. To do this, we first use Proposition 8 to compute changes in observed and Hicksian sales shares due to demand shocks and then plug this into Proposition 7 to get the difference between welfare and real GDP up to a second order approximation:

$$EV^M - \Delta \log Y \approx \frac{1}{2} \frac{\text{Cov}_b(\Delta \log x, \Delta \log A)}{1 + (\theta_0 - 1)(1 - \gamma)}. \quad (27)$$

This expression simplifies when we have constant-returns to scale ($\gamma = 1$) or when the elasticity of substitution is unity ($\theta_0 = 1$). 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 5 (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 $\Delta \log L = 0$). 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 3, micro welfare improves $EV^m > 0$ if preference shocks negatively covary with price changes. Using Proposition 8, changes in prices are

$$d \log p_i = \frac{(1 - \gamma)}{(1 + (\theta_0 - 1)(1 - \gamma))} (d \log x_i - \mathbb{E}_b[d \log \mathbf{x}]).$$

If there are decreasing returns, $\gamma < 1$, then a positive demand shock for i raises the price of i . This initial change in the price is amplified by general equilibrium forces if goods are complements and mitigated if goods are substitutes (this is the multiplier in the denominator). We can now apply Proposition 3 to obtain micro welfare, up to a second order,

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

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. Of course, from a societal perspective, welfare has not changed, since the production possibility set of the economy has not changed.

Appendix F contains additional examples showing how input-output connections can amplify or mitigate the gap between macro welfare and real GDP.

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 in the presence of income effects and demand instability. The second application is about short-run fluctuations, showing 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 using a quantitative model. We also use this model to show that demand instability can make real GDP an unreliable metric for changes in production.

5.1 Application I: 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.²⁴ The first term captures changes in aggregate TFP if industrial structure had remained fixed, and the second term is the adjustment attributed to the fact that sales shares change over time. The second-term captures the importance of Baumol's cost disease.

Proposition 5 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, for EV^M , the relevant measure of the change in TFP is

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

where $\lambda^{ev} \equiv \lambda(A_t, L_t, u_{t_1}, x_{t_1})$ is the Hicksian sales shares of each industry.

Two polar extremes. Computing the Hicksian sales shares to obtain $\Delta \log TFP^{ev}$ requires an explicit structural model of the economy. However, there are two polar cases in which $\Delta \log TFP^{ev}$ can be calculated without specifying the detailed model. The first extreme is when demand is stable and homothetic, so that changes in sales shares are due only to relative price changes (substitution effects). The second extreme is when there are no substitution effects in sales shares (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

²⁴Technically, 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.

discussed in Appendix G. The following corollary of Proposition 5 summarizes the change in welfare-TFP in these two polar cases.

Corollary 1. *If changes in sales shares are due only due only to substitution effects, then*

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

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

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

In the first case, since preferences are homothetic and stable, $\lambda(A_t, L_t, u_{t_1}, x_{t_1}) = \lambda_{it}$, so welfare-TFP is equal to TFP in the data. In the second case, since there are no substitution effects in production or demand, compensated sales shares do not respond to productivity changes, so $\lambda(A_t, L_t, u_{t_1}, x_{t_1}) = \lambda_{it_1}$.

Figure 3: Growth in welfare-relevant TFP (in logs) from 1947 to 2014 using US-KLEMS. The blue line uses initial 1947 shares to calculate TFP changes. The red and yellow line measure the increase in welfare-relevant TFP between 1947 and each year t between 1948 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 1, 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).²⁵ Figure 3 plots EV^M comparing 1947 to subsequent years under alternative assumptions about substitution and income elasticities. For comparisons that are close to 1947, the change in welfare is not very sensitive to our assumptions about elasticities. This is because at higher frequencies, the shocks are small and industry sales shares are reasonably stable. However, the assumptions about substitution and income elasticities play a larger role over longer horizons. 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 1 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.

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). These sectors were cheaper compared to manufacturing in 1947 than in 2014. This means that at 1947 prices, when services were relatively cheap, households require less income growth to be indifferent between their budget constraint in 1947 and the one in 2014.²⁶

In Appendix G, we provide some quantitative illustrations away from the two polar extremes discussed above. In this appendix, we compute welfare changes for different values of elasticities of substitution in consumption and production using Proposition 8. We show that welfare-relevant TFP is closer to measured TFP if the elasticity of substitution across industries (in consumption or production) is significantly lower than one.

²⁵For 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.

²⁶This intuition is flipped for compensating variation at 1947 preferences. Preferences in 1947 favor manufacturing over services. Therefore, at 2014 prices, when manufacturing is relatively cheap, households require a larger reduction in 2014 income to make them indifferent to 1947. 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 3), so for this exercise, the Baumol adjustment is not welfare-relevant. The fact that EV at final preferences and CV at initial preferences are different stems from the fact that they answer different questions, so EV uses demand from 2014 whereas CV uses demand from 1947.

5.2 Application II: Inflation with Product-Level Taste Shocks

In the previous application, we considered a long-run industry-level application. With industry-level data, the gaps between welfare and chain-weighted indices are usually modest over the short-run because industry-level sales shares are relatively stable at high frequency. However, this does not mean that these issues are absent in short-run data.

Whereas industry sales shares are stable at high frequency, firm or product-level sales shares are highly volatile even over the very short-run. If firms' or products' supply and demand shocks are correlated, then measured industry-level output is biased relative to what is relevant for welfare.²⁷ 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 I 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 for 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 year $t_0 < 2019$, we calculate welfare-relevant deflators for t_0 and t_1 preferences module-by-module for continuing products using (9). 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 . In other words, we abstract from product entry and exit by focusing on the continuing-goods price index (see Section 6 for how to deal with product entry-exit when preferences are unstable). To combine module-level inflation rates into a single number, we assume preferences in both t_0 and t_1 across modules are Cobb-Douglas. 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).

Figure 4 displays the welfare-relevant and chained inflation rates for each t_0 assuming

---

²⁷In Appendix H, we formally show that the biases in industry-level data are not diversified away as we aggregate, even if all products are infinitesimal in their industry and all industries are infinitesimal in the aggregate economy. Furthermore, we provide conditions under which 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.

Figure 4: Welfare-relevant and chain-weighted inflation rates for continuing products. The welfare-relevant rates are computed assuming that the elasticity of substitution across UPCs in the same module is 4.5 and the elasticity of substitution across modules is one.

that the elasticity of substitution across UPCs in the same module is 4.5.²⁸ Starting in 2018, inflation for 2019 preferences is around 1 percentage point higher than the chained index. Intuitively, this is because changes in prices and changes in demand shocks between t_0 and t_1 are positively correlated. That is, goods that are more popular in 2019 had relatively higher inflation rates. Following the logic of Proposition 3, this means that the chained index understates inflation for final preferences. The gap increases as we go farther back in time because the shocks are persistent and cumulate. Over the whole sample, the gap widens to 4.3 percentage points.

On the other hand, for initial preferences, the chained index overstates the inflation rate. This is because goods that are relatively more popular in t_0 compared to t_1 (i.e. goods that became less popular over time) experience lower inflation rates. Since goods that become more popular are also the goods that have higher inflation rates, it is natural that inflation with final tastes is higher than inflation with initial tastes. The two series are different

²⁸Estimating 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 reported by Redding and Weinstein (2020) and Jaravel (2019). In Appendix I, 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.

because they answer different questions.^{29,30}

We report robustness with respect to the elasticity of substitution parameter in Appendix I. In this appendix, we show that the size of the bias gets smaller as θ gets closer to one. 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 3, the bias is smaller in the Cobb-Douglas case and larger if the calibrated elasticity is greater than 4.5.

5.3 Application III: 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 exactly the conditions under which micro welfare, macro welfare, and real GDP can diverge from each other.³¹

In this application, 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.

---

²⁹If we were to average initial and final tastes for each t_0 , then we would get a series that is fairly close to the chained index in this example. Since the chained index is, to a second-order, equal to the average of the series with initial and final tastes, this suggests that the second-order approximation is relatively accurate for this example. In Appendix I, we also report results using monthly, rather than annual, data. In that case, the chained measure is not as close to the average of initial and final tastes because the second-order approximation is less accurate.

³⁰The “taste-adjusted inflation rate” in Redding and Weinstein (2020) is lower than the chained-weighted inflation. This may appear related to our finding that inflation for initial tastes is lower than chained inflation. However, for initial tastes, our welfare-relevant inflation is simply exact hat-algebra with initial expenditure shares, discarding any information on subsequent taste shocks. This is different to Redding and Weinstein (2020) whose measure does account for changing tastes, but models an increase in tastes as equivalent to a reduction in price and pins down the overall level of taste shifters by assuming that their average stays constant through time. For a better understanding of why our approach is fundamentally different to that of Redding and Weinstein (2020), see Appendix D.

³¹Cavallo (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, 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.

To study this episode, we use a modified version of the quantitative model introduced in Section 4. Since this is 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 in household demand are calibrated to match the observed peak-to-trough change in personal consumption expenditures by sector from January, 2020 to May, 2020 (conditional on the supply shocks and the elasticities of substitution).³²

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 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 ($t_0 = \text{Q1-2018}$) and Covid ($t_1 = \text{Q2-2020}$) preferences. Recall that micro and macro welfare are not equal in this economy because the PPF is nonlinear.

Table 1 shows that micro welfare falls by more than macro welfare under Covid preferences. This is because, as shown in Example 5, when demand rises for goods in the presence of decreasing returns to scale, micro welfare falls more than macro welfare. Intuitively, if household demand for, say, toilet paper rises, then this raises the price of toilet paper, and reduces micro welfare compared to macro welfare. That is, a single consumer with Covid preferences, who values toilet paper and takes prices as given, is made worse off by the fact that demand for toilet paper rose. However, society as a whole does not take

---

³²Changes 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).

Table 1: The change in micro and macro welfare with pre-Covid and Covid preferences given the supply and demand shocks between February 2020 to May 2020. Chained real GDP 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 Covid preferences	-13.2%	-12.3%	-10.9%
Macro pre-Covid preferences	-16.2%	-12.5%	-10.8%
Macro Covid preferences	-10.1%	-9.4%	-9.0%
Chained real GDP	-12.1%	-10.6%	-9.8%

prices as given. Therefore, when comparing the t_0 PPF to the t_1 PPF, the fact that toilet paper was cheaper in t_0 due to lower demand is irrelevant.

This pattern is reversed for pre-Covid preferences. Macro welfare at pre-Covid preferences fell by a lot since supply contracted in sectors that were popular pre-Covid (e.g. airline travel). Micro welfare losses at pre-Covid preferences are smaller since demand also fell in those sectors (e.g. airline travel) and this reduction in demand reduced their equilibrium prices. Hence, the consumer with pre-Covid preferences, who still likes airline travel, is not as worse off as the society which still likes airline travel.

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 GDP assuming supply and demand shocks arrive simultaneously and linearly over time. Under this assumption, real GDP is somewhere in between the welfare measures at t_0 and t_1 preferences for both macro and micro. However, this is only because of our assumption that the supply and demand shocks arrive simultaneously and at the same speed. If we change the path of supply and demand shocks, real GDP 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 GDP equals macro welfare changes at t_0 preferences. On the other hand, if demand shocks arrive before the supply shocks, then real GDP equals macro welfare changes at t_1 preferences.

Hence, if the supply and demand shocks do not disappear in exactly the same way

as they arrived, measured real 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. That is, if in the downturn, demand shocks arrive before supply shocks (so real GDP 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 GDP rises by roughly 16%), then real GDP is as much as 6% higher when comparing pre-shock real GDP to post-recovery real GDP. 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.³³

6 Extensions

In this section, we briefly summarize how our theoretical results can be used in dynamics settings and extended to account for new goods.

Dynamic Economies

At an abstract level, our results can be applied to dynamic economies by using the Arrow-Debreu formalism of indexing goods by period of time and state of nature (e.g. as in 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). In Appendix K, we specialize our results further to show how Proposition 5 can be used to make welfare comparisons in a dynamic multi-sector model with production of consumption goods and investment goods, with unstable and non-homothetic preferences. We show that steady-state to steady-state macro welfare is equal to the change in nominal consumption expenditures deflated by the exact-algebra CES price index associated with the t_1 indifference curve, exactly as for the partial equilibrium microeconomic welfare in expression (9), despite the fact that the dynamic general equilibrium economy has infinitely many factors.

---

³³This 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 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.

Extensive margin.

In our discussion, we do not explicitly deal with the new goods problem (i.e. goods discontinuously appearing or disappearing). The classic treatment assumes that consumers start and stop consuming goods because the price falls from or goes to infinity. Using stable CES preferences, Feenstra (1994) provides a simple-to-implement formula to adjust price indices for entry and exit of goods. However, when preferences are unstable, it is possible that consumers start or stop consuming a good not because of a change in its availability but because of a change in their tastes. In this section, using the commonly used CES specification, we show that the welfare implications of this can be profound.

Consider a consumer whose preferences are homothetic CES with taste shifters x and elasticity of substitution $\theta_0 > 1$. Good i is *unavailable* if its price is infinite, $p_i = \infty$, and *available* otherwise. Demand for good i may be zero either because the good is unavailable ($p_i = \infty$) or because the consumer does not value the good ($x_i = 0$).

Split the set of goods that consumers value at t_1 into three sets: (1) $\mathcal{C}$: continuing goods consumed in both periods, $b_{it_0} > 0$ and $b_{it_1} > 0$; (2) $\mathcal{N}$: newly consumed goods that either were unavailable at t_1 ($p_{it_0} = \infty$ and $p_{it_1} < \infty$) or that were available ($p_{it_0} < \infty$) but are valued at t_1 and not at t_0 ($x_{it_0} = 0$ and $x_{it_1} > 0$); (3) $\mathcal{X}$: exiting goods that become unavailable at t_1 ($p_{it_0} < \infty$ and $p_{it_1} = \infty$).

The following proposition derives the change in welfare.

Proposition 10 (Welfare with Product Entry/Exit). *The change in equivalent variation at t_1 tastes is given by*

$$EV^m = \log \frac{I_{t_1}}{I_{t_0}} - \log \frac{P_{t_1}}{P_{t_0}}, \quad (28)$$

where

$$\frac{P_{t_1}}{P_{t_0}} = \left(\frac{b_{t_1}^c \left(\frac{P_{t_1}^c}{P_{t_0}^c} \right)^{\theta_0-1} + (1 - b_{t_1}^c) \left(\frac{P_{t_1}^n}{P_{t_0}^n} \right)^{\theta_0-1}}{1 - b_{t_0}^{x_{t_1}}} \right)^{\frac{1}{\theta_0-1}}. \quad (29)$$

In this expression $b_{t_1}^c$ is the expenditure share on continuing goods in t_1 , and $P_{t_1}^c / P_{t_0}^c$ and $P_{t_1}^n / P_{t_0}^n$ is the change in a CES price index for continuing, $\mathcal{C}$, and new, $\mathcal{N}$, goods under t_1 tastes. Finally, $b_{t_0}^{x_{t_1}}$ is the expenditure share on exiting goods under t_1 tastes and p_{t_0} prices.

Applying (29) requires three pieces of information:

- i. The t_1 share of continuing goods, $b_{t_1}^c$, and changes in the t_1 price index for continuing goods $P_{t_1}^c / P_{t_0}^c$.

- ii. The price index for newly consumed goods $P_{t_1}^n / P_{t_0}^n$, which combines newly available goods, for which $p_{i_{t_1}} / p_{i_{t_0}} = 0$, and goods that were available in both periods but consumers did not have tastes for at t_0 , for which $p_{i_{t_1}} / p_{i_{t_0}} \neq 0$.
- iii. The counterfactual share of exiting goods at t_0 prices but t_1 preferences $b_{t_0}^{x_{t_1}}$. It is reasonable to think that $b_{t_0}^{x_{t_1}} \in [0, b_{t_0}^{x_{t_0}}]$. The lower bound takes the position that exiting goods are no longer valuable to the consumer with t_1 tastes, and the upper bound takes the position that exiting are not relatively more valuable to the consumer with t_1 tastes than the consumer with t_0 tastes (i.e. demand curves for those goods that exited did not shift out).

By making different assumptions about the bounds, we can derive three noteworthy special cases of expression (29):

No extensive margin. Assume that all goods are continuing ($b_{t_1}^c = 1$). Then the price deflator is the same as in expression (9)

$$\frac{P_{t_1}}{P_{t_0}} = \frac{P_{t_1}^c}{P_{t_0}^c} = \left(\sum_{i \in \mathcal{C}} b_{it_1} \left(\frac{p_{it_1}}{p_{it_0}} \right)^{\theta_0 - 1} \right)^{\frac{1}{\theta_0 - 1}}. \quad (30)$$

This is the assumption we make in Section 5.2, since we only compute the price index for continuing goods.

Feenstra (1994) with taste shocks. Suppose that newly consumed goods at t_1 were unavailable at t_0 ($P_{t_1}^n / P_{t_0}^n = 0$). Furthermore, suppose that, under t_1 tastes but t_0 prices, the share of expenditures on exiting goods equals the observed share of exiting goods at t_0 : ($b_{t_0}^{x_{t_1}} = b_{t_0}^{x_{t_0}} = 1 - b_{t_0}^c$). In this case, (29) reduces to

$$\frac{P_{t_1}}{P_{t_0}} = \left(\frac{b_{t_1}^c}{b_{t_0}^c} \right)^{\frac{1}{\theta_0 - 1}} \frac{P_{t_1}^c}{P_{t_0}^c}.$$

The term in front of the continuing price index coincides with the well-known new-goods adjustment with CES preferences due to Feenstra (1994). Note that P_{t_1} / P_{t_0} is still not the same as Feenstra (1994) because $P_{t_1}^c / P_{t_0}^c$ is calculated for fixed x_{t_1} tastes and hence is not the Sato-Vartia price index. In other words, $P_{t_1}^c / P_{t_0}^c$ is the price index in (30).

Entry/exit only due to taste shocks. Suppose that all goods were available in t_0 and t_1 , but some goods are consumed in t_1 and not t_0 because of changes in tastes. In this case,

$b_{t_0}^{x_{t_1}} = 0$, so

$$\frac{P_{t_1}}{P_{t_0}} = \left(b_{t_1}^c \left(\frac{P_{t_1}^c}{P_{t_0}^c} \right)^{\theta_0-1} + (1 - b_{t_1}^c) \left(\frac{P_{t_1}^n}{P_{t_0}^n} \right)^{\theta_0-1} \right)^{\frac{1}{\theta_0-1}}. \quad (31)$$

In this case, the price index is a weighted CES aggregator of the price index for continuing goods and the price index for newly consumed goods. The price index increases less than the price index for continuing goods if inflation for newly consumed goods is less than inflation for continuing goods.³⁴

7 Conclusion

In this paper, we provide a toolkit for studying how welfare changes in response to changes in budget sets or production possibility sets allowing for preference instability and non-homotheticity. In contrast to the case of stable and homothetic preferences, there is no model-free statistic like chain-weighted real consumption that simultaneously provides answers to different welfare questions (i.e. equivalent or compensating variation, initial or final preferences, partial or general equilibrium). When preferences are non-homothetic and unstable, there are different welfare questions one may ask, and they have different answers. In this paper, we characterize these measures and show the difference between them and chain-weighted real consumption.

Although our motivation and applications have focused on shocks across time, an interesting avenue to explore is welfare comparisons across locations (see e.g. Deaton, 2003, and Argente et al., 2021). Whereas in a temporal context, the preferences of today are typically more relevant than the preferences of yesterday, in a spatial context, both locations' preferences are equally interesting. 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.³⁵

³⁴For example, the model in Arkolakis (2016) predicts 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. In this case price of newly consumed goods increase on average less than continuing goods.

³⁵This implies that consumers from one location prefer the budget set in the other location, which resembles the Gerschenkron (1951) effect that the relative GDP of a country is higher when evaluated at another country's prices (the grass is greener in the other side). However, while the Gerschenkron effect is a spatial version of the discrepancy between Laspeyres and Paasche indices, our result is driven by the endogeneity of relative prices to demand forces and not by substitution bias.

References

- 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 (2021). Measuring the cost of living in mexico and the us. *American Economic Journal: Macroeconomics*. Forthcoming.
- 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* 9(4), 254–80.
- Atkin, D., B. Faber, T. Fally, and M. Gonzalez-Navarro (2020, March). Measuring Welfare and Inequality with Incomplete Price Information. NBER Working Papers 26890, National Bureau of Economic Research, Inc.
- Auer, R., A. Burstein, S. Lein, and J. Vogel (2021). Unequal expenditure switching: Evidence from switzerland.
- Balk, B. M. (1989). Changing consumer preferences and the cost-of-living index: theory and nonparametric expressions. *Journal of Economics* 50(2), 157–169.
- Baqae, D. and A. Burstein (2021). Aggregate welfare and output with heterogeneous agents.
- Baqae, D., A. Burstein, and Y. Koike-Mori (2022). Measuring welfare with income effects using cross-sectional data.
- 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 (2019). 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 Re-*

- view* 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., D. Lashkari, and M. Mestieri (2021). Structural change with long-run income and price effects. *Econometrica* 89(1), 311–374.
- 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.
- Diewert, W. E. (1976). Exact and superlative index numbers. *Journal of econometrics* 4(2), 115–145.
- Domar, E. D. (1961). On the measurement of technological change. *The Economic Journal* 71(284), 709–729.
- Duernecker, G., B. Herrendorf, and A. Valentinyi (2021). The productivity growth slowdown and kaldor growth facts. *Journal of Economic Dynamics and Control* 130, 104200.
- Durán, J. and O. Licandro (2018). Is the output growth rate in nipa a welfare measure?
- Ehrlich, G., J. Haltiwanger, R. Jarmin, D. Johnson, and M. D. Shapiro (2019). Minding your ps and qs: Going from micro to macro in measuring prices and quantities. In *AEA Papers and Proceedings*, Volume 109, pp. 438–43.
- 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. and M. B. Reinsdorf (2000). An exact price index for the almost ideal demand system. *Economics Letters* 66(2), 159–162.
- 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 K. Shell (1968). *Taste and quality change in the pure theory of the true cost of living index*. Value, Capital and Growth. Edinburgh University Press.
- Gerschenkron, A. (1951). *A dollar index of Soviet machinery output, 1927–28 to 1937*, Volume 197. Rand Corporation.
- Hamilton, B. W. (2001, June). Using engel’s law to estimate cpi bias. *American Economic Review* 91(3), 619–630.
- Hanoch, G. (1975). Production and demand models with direct or indirect implicit additivity. *Econometrica: Journal of the Econometric Society*, 395–419.
- 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.

- Ivancic, L., W. E. Diewert, and K. J. Fox (2011). Scanner data, time aggregation and the construction of price indexes. *Journal of Econometrics* 161(1), 24–35.
- Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the us retail sector. *The Quarterly Journal of Economics* 134(2), 715–783.
- Jaravel, X. and D. Lashkari (2022). Nonparametric measurement of long-run growth in consumer welfare.
- 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.
- Konüs, A. A. (1939). The problem of the true index of the cost of living. *Econometrica: Journal of the Econometric Society*, 10–29.
- Licandro, O., J. Ruiz-Castillo, and J. Durán (2002). The measurement of growth under embodied technical change. *Recherches Économiques de Louvain/Louvain Economic Review* 68(1-2), 7–19.
- Lucas, R. E. (1987). *Models of business cycles*, Volume 26. Basil Blackwell Oxford.
- 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.
- Matsuyama, K. (2019). Engel’s law in the global economy: Demand-induced patterns of structural change, innovation, and trade. *Econometrica* 87(2), 497–528.
- McKenzie, G. W. and I. F. Pearce (1982). Welfare measurement: a synthesis. *The American Economic Review* 72(4), 669–682.
- Muellbauer, J. (1975). The cost of living and taste and quality change. *Journal of Economic Theory* 10(3), 269–283.
- 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.
- Oulton, N. (2008). Chain indices of the cost-of-living and the path-dependence problem: An empirical solution. *Journal of Econometrics* 144(1), 306–324.
- 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.
- United Nations (2009). *System of National Accounts 2008*. USA: United Nations.

Online Appendix

A	Proofs	A1
B	Extension to Other Welfare Measures	A9
C	Relation to Konüs Price Indices	A12
D	Comparison of Quality and Taste Changes	A14
E	Non-homothetic CES preferences	A16
	E.1 Derivation of Marshallian budget shares	A17
	E.2 Comparison of welfare and changes in utility index	A18
F	Analytical Examples with Input-Output Connections	A21
G	Additional details on Application I	A23
H	Within-Industry Supply and Demand Shocks	A26
I	Additional Details on Application II	A29
J	Non-CES Functional Forms	A32
K	Dynamic Economies	A34

Appendix A Proofs

We do not present the proofs in the same order as the propositions in the paper because the proofs of earlier propositions rely on notation and logic from later propositions.

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 2. Differentiate real consumption

$$\Delta \log Y = \int_{t_0}^{t_1} \left(\frac{d \log I_s}{ds} - \sum_{i \in N} b_{is} \frac{d \log p_{is}}{ds} \right) ds$$

twice with respect to t_1

$$\begin{aligned} \frac{d \log Y}{dt_1} &= \frac{d \log I}{dt} - \sum_{i \in N} b_i \frac{d \log p_{it}}{dt} \\ \frac{d^2 \log Y}{dt_1^2} &= \frac{d^2 \log I}{dt^2} - \sum_{i \in N} b_i \frac{d^2 \log p_i}{dt^2} - \sum_{i \in N} \frac{db_i}{dt} \frac{d \log p_i}{dt}, \end{aligned}$$

where all right-hand-side terms are evaluated at t_1 . Therefore,

$$\begin{aligned}\Delta \log Y &\approx \frac{d \log Y}{dt_1} \Delta t + \frac{1}{2} \frac{d^2 \log Y}{dt_1^2} (\Delta t)^2 \\ &= \left[\frac{d \log I}{dt} - \sum_{i \in N} b_i \frac{d \log p_{it}}{dt} \right] \Delta t \\ &\quad + \frac{1}{2} \left[\frac{d^2 \log I}{dt^2} - \sum_{i \in N} b_i \frac{d^2 \log p_{it}}{dt^2} - \sum_{i \in N} \frac{db_{it}}{dt} \frac{d \log p_{it}}{dt} \right] (\Delta t)^2.\end{aligned}$$

Thus,

$$\Delta \log Y \approx \Delta \log I - \sum_{i \in N} b_i \Delta \log p_i - \frac{1}{2} \sum_{i \in N} \Delta b_i \Delta \log p_i, \quad (32)$$

where recall that for any variable z , $\Delta \log z \approx (\partial \log z / \partial t) \Delta t + 1/2 (\partial^2 \log z / \partial t^2) (\Delta t)^2$ and the remainder terms are of order $(\Delta t)^3$. Evaluating (32) at $t_1 = t_0$ yields (12).

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

$$EV^m = \Delta \log I - \int_{t_0}^{t_1} \sum_{i \in N} b_i(p, v(p_{t_1}, x_{t_1}), x_{t_1}) \frac{d \log p_{is}}{ds} ds$$

Differentiate EV^m with respect to t_1

$$\begin{aligned}\frac{dEV^m}{dt_1} &= \frac{d \log I}{dt} - \sum_{i \in N} b_i(p, v(p_{t_1}, x_{t_1}), x_{t_1}) \frac{d \log p_{it}}{dt} \\ &\quad - \int_{t_0}^{t_1} \sum_{i \in N} \frac{du_t}{dt} \frac{\partial b_i(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial u} \frac{d \log p_{is}}{ds} ds - \int_{t_0}^{t_1} \sum_{i \in N} \frac{dx_t}{dt} \frac{\partial b_i(p, v(p_{t_1}, x_{t_1}), x_{t_1})}{\partial x} \frac{d \log p_{is}}{ds} ds,\end{aligned}$$

where the shares and derivatives are evaluated at t_1 , and then differentiate again and eval-

uate at $t_1 = t_0$ to get

$$\begin{aligned}
\frac{d^2 EV^m}{dt_1^2} &= \frac{d^2 \log I}{dt^2} - \sum_{i \in N} \sum_{j \in N} \frac{\partial b_i}{\partial \log p_j} \frac{d \log p_{it}}{dt} \frac{d \log p_{jt}}{dt} \\
&\quad - \sum_{i \in N} b_i \frac{d^2 \log p_{it}}{dt^2} - 2 \sum_{i \in N} \frac{du_t}{dt} \frac{\partial b_i}{\partial u} \frac{d \log p_{it}}{dt} - 2 \sum_{i \in N} \frac{dx_t}{dt} \frac{\partial b_i}{\partial x} \frac{d \log p_{it}}{dt} \\
&= \frac{d^2 \log I}{dt^2} - \sum_{i \in N} \left[\sum_{j \in N} \frac{\partial b_i}{\partial \log p_j} \frac{d \log p_{jt}}{dt} + \frac{du_t}{dt} \frac{\partial b_i}{\partial u} + \frac{d \log x_t}{dt} \frac{\partial b_i}{\partial x} \right] \frac{d \log p_{it}}{dt} \\
&\quad - \sum_{i \in N} b_i \frac{d^2 \log p_{it}}{dt^2} - \sum_{i \in N} \frac{du_t}{dt} \frac{\partial b_i}{\partial u} \frac{d \log p_{it}}{dt} - \sum_{i \in N} \frac{dx_t}{dt} \frac{\partial b_i}{\partial x} \frac{d \log p_{it}}{dt} \\
&= \frac{d^2 \log I}{dt^2} - \sum_{i \in N} \frac{db_{it}}{dt} \frac{d \log p_{it}}{dt} - \sum_{i \in N} b_i \frac{d^2 \log p_{it}}{dt^2} \\
&\quad - \sum_{i \in N} \frac{du_t}{dt} \frac{\partial b_i}{\partial u} \frac{d \log p_{it}}{dt} - \sum_{i \in N} \frac{dx_t}{dt} \frac{\partial b_i}{\partial x} \frac{d \log p_{it}}{dt},
\end{aligned}$$

where $\partial b_i / \partial \log p_j$ holds fixed utility values u and taste shifters x (without loss of generality, represented by a scalar). The first three terms are equal to the second-order expansion of $\Delta \log Y$, and the remaining terms are the gap between $\Delta \log Y$ and EV^m . Finally, we use $db_i = \sum_j d \log p_j \frac{\partial b_i}{\partial \log p_j} + du \frac{\partial b_i}{\partial u} + dx \frac{\partial b_i}{\partial x}$ to obtain (13). $\square$

Proof of Proposition 3. Substitute (14) in place of Δb in (32) to get the expression for $\Delta \log Y$. For the gap between real consumption and EV^m , note that Proposition 2 implies that

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

Using (14) in place of Δb above and the fact that $\partial b_i / \partial \log p_i = (1 - \theta_0) b_i (1 - b_i)$ for $i = j$ and $\partial b_i / \partial \log p_j = \theta_0 b_i b_j$ 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 Proposition 6. For this proof, we use notation introduced in Section 4. Let nominal

GDP be the numeraire. Start by noting that from Shephard's lemma

$$d \log p_i = -d \log A_i - d \log L_i \mathbf{1}_{\{i \in F\}} + \sum_{j \in N+F} \Omega_{ij} d \log p_j,$$

where $d \log A_i$ is a factor-quality shock if i is a factor. Inverting this linear system yields

$$d \log p_i = -\Psi d \log A + \Psi^F (d \log \Lambda - d \log L),$$

where Ψ^F is the $(1 + N + F) \times F$ matrix consisting of the right F columns of Ψ and Λ is the vector of factor income shares.

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 substitutes in our solution for $d \log p$, and the following lines 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 5. Recall that the macro equivalent variation at final preferences is defined 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

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

Define a hypothetical Hicksian economy with fictional households that have stable homothetic preferences defined by the expenditure function $e^h(p, u) = e(p, u_{t_1}; x_{t_1}) u$ where e is the expenditure function of the consumer in the primitive economy. Budget shares of this

fictional consumer are $b_i^h(p) \equiv \frac{\partial \log e^h(p, u)}{\partial \log p_i} = \frac{\partial \log e(p, u_{t_1}; x_{t_1})}{\partial \log p_i}$. Given any technology vector, in this hypothetical Hicksian economy we denote the Leontief inverse matrix by Ψ^h and sales shares by λ^h . Given technologies A_t and factor quantities L_t , we denote prices in this hypothetical economy by p_t^h . Changes in prices in this hypothetical economy satisfy

$$d \log p^h = -\Psi^h d \log A + \Psi^{hF} d \log \Lambda^h, \quad (33)$$

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

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

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

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

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

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

which can be re-expressed using the expenditure function as

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

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 Hicksian economy with preferences represented by $e^h(p, u)$. As in the proof of Lemma 1, rewrite EV^M as (suppressing dependence on t in the integral)

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

Following the same steps as in the proof of Proposition 6 (for the hypothetical economy), we obtain

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

□

Proof of Proposition 4. By the proof of Proposition 5, $EV^m = EV^M$ if, and only if, $p^h(A_t, L_t) = p(A_t, L_t, x_t)$ 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 (using notation introduced in Section 4 and in the proof of Proposition 6). Then by Proposition 8, 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 u 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 u . Thus, $p^h(A_t, L_t) = p(A_t, L_t, x_t)$.

□

Proof of Proposition 7. The proof follows similar steps as in the proof of Proposition 3. Write changes in real GDP (abstracting from changes in factor endowments) as,

$$\Delta \log Y = \int_{t_0}^{t_1} \sum_{i \in N} \lambda_{is} \frac{d \log A_i}{ds} ds,$$

differentiate twice with respect to t_1 and evaluate the derivative at $t_1 = t_0$. Write EV^M as

$$EV^M = \int_{t_0}^{t_1} \sum_{i \in N} \lambda_i(A_s, L_s, u_{t_1}, x_{t_1}) \frac{d \log A_i}{ds} ds,$$

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

Proof of Proposition 8. 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 $i \in F$, and recalling that we assume that factor endowments are unchanged, we have

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

where $d \log A_i$ when i is a factor represents a factor-quality shock. Combining these yields the desired expression for changes in prices

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

To get changes in sales shares, note that

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

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

Meanwhile

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

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 9. Normalize nominal GDP to one. Applying Proposition 8 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 7 in combination with the expression for $d \log p$ and $d\lambda$ in Proposition 8 in the case of one factor. To solve for EV^M , by Proposition 4, $EV^M = EV^m$. Solve for $EV^m - \Delta \log Y$ by plugging the expression for $d \log p$ into Proposition 2 and noting that $b' = \Omega_{(0,:)}$. □

Proof of Proposition 10. By Lemma 1, welfare changes measured as the equivalent variation at t_1 preferences (which is equal to compensating variation at t_1 tastes since preferences are homothetic) is given by

$$EV^m = \log \frac{I_{t_1}}{I_{t_0}} - \log \frac{P_{t_1}}{P_{t_0}}, \quad (34)$$

where

$$\frac{P_{t_1}}{P_{t_0}} = \left(\frac{\sum_i x_{it_1} p_{it_0}^{1-\theta_0}}{\sum_j x_{jt_1} p_{jt_1}^{1-\theta_0}} \right)^{\frac{1}{\theta_0-1}} = \left(\frac{\sum_i x_{it_1} p_{it_1}^{1-\theta_0} \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0}}{\sum_j x_{jt_1} p_{jt_1}^{1-\theta_0}} \right)^{\frac{1}{\theta_0-1}}. \quad (35)$$

We re-express (35) as

$$\begin{aligned} \left(\frac{P_{t_1}}{P_{t_0}} \right)^{\theta_0-1} &= \frac{\sum_{i \in \mathcal{C}} x_{it_1} p_{it_1}^{1-\theta_0} \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} + \sum_{i \in \mathcal{N}} x_{it_1} p_{it_1}^{1-\theta_0} \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} + \sum_{i \in \mathcal{X}} x_{it_1} p_{it_0}^{1-\theta_0}}{\sum_j x_{jt_1} p_{jt_1}^{1-\theta_0}} \\ &= b_{t_1}^c \sum_{i \in \mathcal{C}} b_{it_1}^c \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} + (1 - b_{t_1}^c) \sum_{i \in \mathcal{N}} b_{it_1}^n \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} + b_{t_0}^{x_{t_1}} \left(\frac{P_{t_1}}{P_{t_0}} \right)^{\theta_0-1} \end{aligned} \quad (36)$$

where $b_{t_1}^c \equiv \sum_{i \in \mathcal{C}} b_{it_1}$ denotes the expenditure share on continuing goods at t_1 ,

$$b_{it_1}^c \equiv \frac{x_{it_1} p_{it_1}^{1-\theta_0}}{\sum_{i \in \mathcal{C}} x_{it_1} p_{it_1}^{1-\theta_0}}, \quad b_{it_1}^n \equiv \frac{x_{it_1} p_{it_1}^{1-\theta_0}}{\sum_{i \in \mathcal{N}} x_{it_1} p_{it_1}^{1-\theta_0}},$$

and

$$b_{t_0}^{x_{t_1}} \equiv \frac{\sum_{i \in \mathcal{X}} x_{it_1} p_{it_0}^{1-\theta_0}}{\sum_j x_{jt_1} p_{jt_0}^{1-\theta_0}}$$

is the (unobserved) share of exiting goods at t_0 prices under t_1 preferences. Defining a price index for continuing goods under t_1 preferences,

$$\frac{P_{t_1}^c}{P_{t_0}^c} = \left(\sum_{i \in \mathcal{C}} b_{it_1}^c \left(\frac{p_{it_1}}{p_{it_0}} \right)^{\theta_0-1} \right)^{\frac{1}{\theta_0-1}}$$

and a price index for new goods under t_1 preferences,

$$\frac{P_{t_1}^n}{P_{t_0}^n} = \left(\sum_{i \in \mathcal{N}} b_{it_1}^n \left(\frac{p_{it_1}}{p_{it_0}} \right)^{\theta_0-1} \right)^{\frac{1}{\theta_0-1}},$$

we can express the change in the welfare-relevant price index as

$$\frac{P_{t_1}}{P_{t_0}} = \left(\frac{b_{t_1}^c \left(\frac{P_{t_1}^c}{P_{t_0}^c} \right)^{\theta_0-1} + (1 - b_{t_1}^c) \left(\frac{P_{t_1}^n}{P_{t_0}^n} \right)^{\theta_0-1}}{1 - b_{t_0}^{x_{t_1}}} \right)^{\frac{1}{\theta_0-1}} \quad (37)$$

□

Appendix B Extension to Other Welfare Measures

Our baseline measure of welfare changes is equivalent variation under final preferences. Alternatively, we could measure changes in welfare using compensating (instead of equivalent) variation, or by using initial (rather than final) preferences. In this appendix, we show how to calculate the alternative welfare measures. Note that 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*, which we discussed in Section 2.4, is

$CV^m(p_{t_0}, I_{t_0}, p_{t_1}, I_{t_1}; x_{t_0}) = \phi$, where ϕ solves

$$v(p_{t_1}, e^{-\phi} I_{t_1}; x_{t_0}) = v(p_{t_0}, I_{t_0}; x_{t_0}). \quad (38)$$

The analog to (4) in Lemma 1 is given in equation (17). 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 $CV^m = -EV^m$.

For non-homothetic CES, CV^m is equal to the exact hat-algebra price index with initial shares b_{t_0} :

$$CV^m = \Delta \log I - \log \left(\sum_i b_{it_0} \left(\frac{p_{it_1}}{p_{it_0}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}. \quad (39)$$

To a second-order approximation around t_0 (without imposing non-homothetic CES, as in Proposition 2)

$$\Delta \log CV^m = \Delta \log I - \mathbb{E}_{b_{t_0}}(\Delta \log p) - \frac{1}{2} \sum_{i \in N} \sum_j \Delta \log p_j \frac{\partial b_i}{\partial \log p_j} \Delta \log p_i \quad (40)$$

Recall that changes in budget shares due to non-price factors are multiplied by 1/2 in real consumption and by 1 in EV^m . 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 2 and equation (40), we see that up to a second order approximation (without imposing any specific form of preferences),

$$0.5 \times (EV^m + CV^m) \approx \Delta \log Y. \quad (41)$$

That is, locally 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

$$v(p_{t_1}, I_{t_1}; x_{t_0}) = v(p_{t_0}, e^{\phi} I_{t_0}; x_{t_0}). \quad (42)$$

Globally, changes in welfare are given by (4) using Hicksian budget shares $b_i(p_t, v(p_{t_1}, I_{t_1}; x_{t_0}); x_{t_0})$.

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

$$v(p_{t_1}, e^{-\phi} I_{t_1}; x_{t_1}) = v(p_{t_0}, I_{t_0}; x_{t_1}). \quad (43)$$

Globally, changes in welfare are given by (17) using Hicksian budget shares $b_i(p_t, v(p_{t_0}, I_{t_0}; x_{t_0}); x_{t_1})$. Note that to calculate 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 (18) in Proposition 5 applies using Hicksian budget shares $\lambda(A, L, u_{t_0}, x_{t_0})$, the sales shares in a fictional economy with the PPF A, L but where consumers have stable homothetic preferences represented by the expenditure function $e(p, u) = e(p, u_{t_0}, x_{t_0})u$ where $u_{t_0} = v(p_{t_0}, I_{t_0}; x_{t_0})$. Growth accounting for welfare is based on Hicksian sales shares evaluated at current technology but for fixed initial preferences and initial utility. The only information on preference parameters we need to know is elasticities of substitution at the initial allocation.

Changes in welfare are, to a second-order approximation (the analogue of that in Proposition 7)

$$CV^M = \sum_{i \in N} \lambda_i \Delta \log A_i + \frac{1}{2} \sum_{i \in N} \sum_{j \in N} \Delta \log A_j \frac{\partial \lambda_i}{\partial \log A_j} \Delta \log A_i, \quad (44)$$

where $\partial \lambda / \partial \log A$ is the partial derivative of the Hicksian sales share with respect to technology, and all terms are evaluated at t_0 .

Proposition 8 can be used to compute CV^M (instead of EV^M). To do this, we need Hicksian sales shares $\lambda(A_t, L_t, u_{t_0}, x_{t_0})$ as a function of t . These are solutions to the differential equations in Proposition 8 with the terms involving taste shocks and income effects in (23) set to zero. In this case, the boundary condition is that the Leontief inverse at t_0 is equal

to the observed Leontief inverse Ψ_{t_0} at t_0 . Therefore, if Ψ_{t_0} is observed, we can calculate Hicksian sales shares between t_0 and t_1 by starting (23) at t_0 and going forward to t_0 . This process does not require knowledge of either the income elasticities ε nor the taste shocks $\Delta \log x$.

Appendix C Relation to Konüs Price Indices

A Konüs price index is defined as the ratio of the expenditure function at two different price systems holding fixed utility and preferences:

$$\frac{P_{t_1}(u, x)}{P_{t_0}(u, x)} = \frac{e(p_{t_1}, u; x)}{e(p_{t_0}, u; x)}.$$

Lemma 1 shows that EV^m can be calculated by deflating nominal income changes by the Konüs price index corresponding to final preferences and final utility (i.e. the final indifference curve).³⁶

In the index number theory literature, it is common to work with Konüs price indices for some intermediate preferences or utility levels. For example, Diewert (1976), Caves et al. (1982), and Feenstra and Reinsdorf (2007). The advantage of this approach is that it requires far less information. For example, Diewert (1976) shows that a Tornqvist index of t_0 and t_1 measures the Konüs price index for a consumer with stable translog preferences with utility level $(u_{t_0} u_{t_1})^{\frac{1}{2}}$; Caves et al. (1982) and Feenstra and Reinsdorf (2007) prove a similar result for homothetic but unstable CES or translog preferences. In contrast to EV^m , these indices can all be computed *without* knowledge of any elasticities. In particular, these papers show that, under some assumptions (translog or CES), commonly used indices like Tornqvist and Sato-Vartia do answer an economically meaningful question.

However, whilst these papers provide an interpretation for these commonly used indices, these indices do not measure EV or CV , which are of interest per se in applied micro and macro welfare analysis. Furthermore, these indices are not money metrics, as we show below. Our contribution, relative to common practice in the index number theory literature, is to instead characterize and analyze EV and CV , both for ex-post accounting and ex-ante counterfactuals. Furthermore, we provide a unified analysis of non-homotheticity and taste shocks, whereas the literature has tended to focus on one at a time under parametric assumptions or second-order approximations. We also show how to also develop a general equilibrium measure of welfare.

---

³⁶ CV^m can be calculated by deflating nominal income changes by the Konüs price index corresponding to the *initial* indifference curve.

To relate the aforementioned results to ours, consider the economic question that changes in nominal income between t_1 and t_0 deflated by a Konüs price index evaluated at some intermediate level of utility answers. For any base period t_b (which does not need to lie between t_0 and t_1) we can write

$$\log \frac{I_{t_1}}{I_{t_0}} - \log \frac{P_{t_1}(u_b, x_b)}{P_{t_0}(u_b, x_b)} = \left(\log \frac{I_{t_b}}{I_{t_0}} - \log \frac{P_{t_b}(u_b, x_b)}{P_{t_0}(u_b, x_b)} \right) + \left(\log \frac{I_{t_1}}{I_{t_b}} - \log \frac{P_{t_1}(u_b, x_b)}{P_{t_b}(u_b, x_b)} \right),$$

or

$$\log \frac{I_{t_1}}{I_{t_0}} - \log \frac{P_{t_1}(u_b, x_b)}{P_{t_0}(u_b, x_b)} = \log \frac{e(p_{t_0}, u_{t_b}; x_{t_b})}{e(p_{t_0}, u_{t_0}; x_{t_b})} - \log \frac{e(p_{t_1}, u_{t_b}; x_{t_b})}{e(p_{t_1}, u_{t_1}; x_{t_b})}.$$

The first summand on the right-hand side is $EV^m(p_{t_0}, I_{t_0}, p_{t_b}, I_{t_b}; x_b)$ and the second summand is $-EV^m(p_{t_1}, I_{t_1}, p_{t_b}, I_{t_b}; x_b)$. In words, $\log \frac{I_{t_1}}{I_{t_0}} - \log \frac{P_{t_1}(u_b, x_b)}{P_{t_0}(u_b, x_b)}$ answers the question “For a consumer with preferences $\succeq_{x_{t_b}}$, what is the change in the t_0 endowment that makes her indifferent between her choice set at t_0 and t_b minus the change in the t_1 endowment that makes her indifferent between her choice set at t_1 and t_b ?” In particular, note that the first term is in units of t_0 prices whereas the second one is in units of t_1 prices. Therefore, this is not a money metric that can be used to compare all choices. In sum, although our approach has stronger information requirements, it characterizes a widely-studied and fundamentally different object (i.e. a money metric) than what has commonly been studied in the index number theory literature.

Deaton and Muellbauer (1980) write in reference to the Konüs at intermediate utility result:

If we were willing to accept the reference indifference curve labelled by u^* (note: the geometric average of a u_{t_0} and u_{t_1}) as the relevant one, this property of the Tornqvist index is attractive since the quadratic specification can provide a second-order approximation to any arbitrary cost function. Without knowing the parameters of the cost function, we lack more specific information about the reference indifferent curve (such as what budget level and price vector correspond to it), and the result is of no help in constructing a constant utility cost-of-living index series with more than two elements [elements refer to time periods]. A chained series of pairwise Tornqvist indices can always be constructed, but this has a different reference indifference curve for every link in the chain. (Deaton and Muellbauer, 1980, page 174)

That is, in practice most index numbers are constructed by chaining, but the intermediate utility result does not apply to chained indices unless the path of prices is linear (Feenstra and Reinsdorf, 2000). In our paper we characterize how equivalent variation at final preferences or compensating variation at initial preferences differ from chained (Divisia) indices

under arbitrary price and income paths.

Appendix D Comparison of Quality and Taste Changes

In this appendix, we discuss how our welfare results can be extended to environments with unobserved quality changes. We also contrast the bias we identify with the “taste shock bias” discussed by Redding and Weinstein (2020).

The standard approach to modeling quality is hedonics, where goods are bundles of characteristics and consumers have preferences over characteristics. For example, for computers, CPU speed is a characteristic that consumers value. If a computer increases its CPU speed, the consumer can consume more of this characteristic. Choices made by consumers over computers with different CPU speeds reveal how consumers value this characteristic. Note that there is no reason to normalize the level of quality across goods because the units of characteristics are observable (e.g. GHz). However, even after all the quality-adjustments have been done, demand curves can still shift. We model such residual shifts in demand curves as changes in tastes x and hold x constant in the comparison because consumer preferences over x , if they exist, are by definition unobservable. Of course, if consumers have some preferences over x and we can measure x , then x must be included as part of the description of the commodity space rather than treated as a taste shifter.

To make this more concrete, suppose that consumers have CES preferences (indexed by tastes x) over $q_i c_i$ where i indexes a variety, and c_i and q_i are the quantity and quality of each i . For example, each i is a different variety of chocolate, c_i is the number of boxes of chocolate, and q_i is the weight of each box of chocolate i . So the characteristic that consumers have preferences over is the total weight of chocolate they purchase of each type, and consumers do not care about how many boxes their chocolate came in.

Under these assumptions, quality changes are equivalent to changes in prices, so we can write the quality-adjusted price of good i as $p_i = \tilde{p}_i / q_i$, where $\tilde{p}_i$ is the observed market price of good i . In our example, $\tilde{p}_i$ is the observed price per box and $\tilde{p}_i / q_i$ is the price per ounce. Changes in quality-adjusted prices are given by $\Delta \log p_i = \Delta \log \tilde{p}_i - \Delta \log q_i$. On the other hand, changes in x indicate changes in preferences for the different varieties of chocolate. Unlike changes in q_i , which are measured in ounces per box (or any other observable cardinal units such as grams per box), changes in x do not have interpretable units and the effects on the utility index depend on the choice of cardinalization.

Substituting this into our various propositions allows us to isolate the way quality changes affect our results and how they compare with changes in tastes. For example, Proposition 3 becomes the following (for brevity, we assume homothetic CES preferences):

Proposition 11 (Approximate Micro with Quality Change). *Consider some perturbation in demand $\Delta \log \mathbf{x}$, 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*

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

and the change in welfare is

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

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.

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

$$\begin{aligned} EV^m - \Delta \log Y \approx & \underbrace{\mathbb{E}_b(\Delta \log q)}_{\text{average quality}} + \underbrace{\frac{1}{2}(\theta_0 - 1) \text{Var}_b(\Delta \log q)}_{\text{dispersion in quality}} + \underbrace{\frac{1}{2}(1 - \theta_0) \text{Cov}_b(\Delta \log \tilde{p}, \Delta \log q)}_{\text{covariance of price and quality}} \\ & - \underbrace{\frac{1}{2} \text{Cov}_b(\Delta \log \mathbf{x}, \Delta \log \tilde{p})}_{\text{covariance of taste and price}} + \underbrace{\text{Cov}_b(\Delta \log \mathbf{x}, \Delta \log q)}_{\text{covariance of taste and quality}}. \end{aligned} \quad (45)$$

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 increase.

In our analysis, we assume that prices have already been adjusted for quality so the only non-zero term is the fourth one. In other words, in the body of the paper, we assume

that $\Delta \log q = 0$, which means that (45) simplifies to

$$EV^m - \Delta \log Y \approx -\frac{1}{2} \text{Cov}_b (\Delta \log x, \Delta \log \tilde{p}). \quad (46)$$

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.

Comparison to Redding and Weinstein (2020). We can use (45) to contrast our approach to that of Redding and Weinstein (2020). The “taste shifters” in that paper are mathematically equivalent to quality shocks ($\Delta \log q \neq 0$), and preferences are stable over “taste-adjusted consumption” ($\Delta \log x = 0$). Equation (45) simplifies to

$$EV^m - \Delta \log Y \approx \mathbb{E}_b (\Delta \log q) + \frac{1}{2}(\theta_0 - 1) \text{Var}_b (\Delta \log q) - \frac{1}{2}(\theta_0 - 1) \text{Cov}_b (\Delta \log \tilde{p}, \Delta \log q). \quad (47)$$

Comparing (46) to (47) 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) assume that unweighted average of $\Delta \log q$ is zero.³⁷ 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, holding prices and income constant, and the sign of this effect depends on the elasticity of substitution. This is in contrast to shocks to x which cannot change money-metric welfare on their own if prices and income are held constant. Third, in both (46) and (47), the covariance of taste shifters and market prices matters, however, in (47) the sign of the covariance depends on whether the elasticity of substitution is greater than or less than one, whereas in (46), the sign is always the same.

Appendix E Non-homothetic CES preferences

In this appendix, we derive (14). 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.

---

³⁷As discussed earlier, if $\Delta \log q$ is interpreted as a taste shock rather than a quality shock, then there is nothing in the data that pins down the average level of $\Delta \log q$ since it is not a primitive of the ordinal preference relation.

E.1 Derivation of Marshallian budget shares

This appendix provides a derivation of the log-linearized expression (14). When preferences are non-homothetic CES, the expenditure function can be written as

$$e(p, u; x) = \left(\sum_{i \in N} x_i p_i^{1-\theta_0} u^{\xi_i} \right)^{\frac{1}{1-\theta_0}}, \quad (48)$$

with Hicksian demand

$$c_i = x_i \left(\frac{p_i}{\sum_j p_j c_j} \right)^{-\theta_0} u^{\xi_i}, \quad (49)$$

and budget shares

$$b_i(p, x, u) \equiv \frac{p_i c_i}{\sum_j p_j c_j} = x_i u^{\xi_i} \left(\frac{p_i}{\sum_j p_j c_j} \right)^{1-\theta_0} = \frac{x_i u^{\xi_i} p_i^{1-\theta_0}}{\sum_{j \in N} x_j u^{\xi_j} p_j^{1-\theta_0}} \quad (50)$$

Differentiating (48) and (50) at any point t ,

$$d \log b_{it} = d \log x_{it} + (1 - \theta_0) (d \log p_{it} - d \log I_t) + \xi_i d \log u_t, \quad (51)$$

and

$$d \log u_t = \frac{1 - \theta}{\sum_j b_{jt} \xi_j} \left[d \log I_t - \sum_j b_{jt} d \log p_{jt} \right] - \frac{1}{\sum_j b_{jt} \xi_j} \sum_j b_{jt} d \log x_{jt}. \quad (52)$$

Substituting (52) into (51),

$$d \log b_{it} = (1 - \theta_0) \left[d \log p_{it} - \sum_j b_{jt} d \log p_{jt} \right] + (\varepsilon_{it} - 1) \left[d \log I_t - \sum_j b_{jt} d \log p_{jt} \right] + d \log x_{it},$$

with demand shifters

$$d \log x_{it} = d \log x_{it} - \frac{\xi_i}{\sum_j b_{jt} \xi_j} \sum_j b_{jt} d \log x_{jt}, \quad (53)$$

and income elasticities

$$\varepsilon_{it} - 1 = (1 - \theta_0) \left(\frac{\xi_i}{\sum_j b_{jt} \xi_j} - 1 \right). \quad (54)$$

This is a differential equation that pins down budget shares b as a function of prices, incomes, and primitives x , given budget shares and income elasticities at some point in time.

E.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. This utility index is used in Section IIIA of Redding and Weinstein (2020) as a welfare measure. We show that there is no normalization of the parameters such that the equivalent variation (or the compensating variation) is equal to changes in the utility index unless preferences are homothetic and stable.

In this section, for brevity we assume away taste shocks (for taste shocks, see Appendix D). The micro equivalent variation is given by

$$EV^m = \log \frac{e(p_{t_0}, v(p_{t_1}, I_{t_1}))}{e(p_{t_0}, v(p_{t_0}, I_{t_0}))},$$

where $v(p, I)$ is the indirect utility function, initial prices and income are p_{t_0} and I_{t_0} , and final prices and income are p_{t_1} and I_{t_1} .

The utility index u at t is equal to $v(p_t, I_t)$, and can be calculated by solving for u in $I_t = e(p_t, u)$. Equivalently, one can calculate changes in u_t using the price index $P_t \equiv e(p_t, u_t)/u_t$. The change in the utility index between t_0 and t_1 is given by

$$U \equiv \log \frac{v(p_{t_1}, I_{t_1})}{v(p_{t_0}, I_{t_0})}.$$

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.

Consider the expenditure function in equation (48). If preferences are homothetic ($\zeta_i = \bar{\zeta}$ for all i), then $e(p, u) = \left(\sum_i \omega_i p_i^{1-\theta_0}\right)^{\frac{1}{1-\theta_0}} u^{\frac{\bar{\zeta}}{1-\theta_0}}$ and we can write

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

So, when preferences are homothetic, in order for $EV^m = U$ we must cardinalize utility by setting $\bar{\zeta} = 1 - \theta_0$ so 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.

Using Proposition 3, we have that to a first-order EV^m is

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

where $d \log Y$ is the first-order change in real consumption as measured by Tornqvist or Divisa (to a first-order, they are equivalent). Hence, to a first order, Tornqvist and EV are the same. The second-order change in EV^m is, by Proposition 3, equal to

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

where ε is the vector of income elasticities and $d^2 \log Y$ is the change in real consumption as measured by a Tornqvist or Divisa 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)

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

$$\begin{aligned} d^2 U &= \frac{1 - \theta_0}{\sum_i b_i \bar{\xi}_i} \left[d^2 \log e - dbd \log p - (d \log e - bd \log p) \sum_i b_i (\varepsilon_i - 1) d \log p_i \right. \\ &\quad \left. - \frac{1}{1 - \theta_0} \sum_i b_i \varepsilon_i ((\varepsilon_i - 1)) (d \log e - bd \log p)^2 \right] \end{aligned}$$

The derivatives EV^m and U are in general different. Whereas EV^m is a function of observables, U depends on the cardinalization of the utility function. In particular $\sum_i b_i \bar{\xi}_i$ affects the response of U but is not a primitive parameter of the ordinal preference relation, and hence is not pinned down by observables, as discussed in Section E.1. A standard approach in the literature to pin down $\sum_i b_i \bar{\xi}_i$ is to set one of the $\bar{\xi}$ to 1.

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

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

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

$$\frac{1 - \theta_0}{\sum_i b_i \xi_i} = 1,$$

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

$$\begin{aligned} d^2U - d^2EV^m &= -\frac{1}{1 - \theta_0} \sum_i b_i \varepsilon_i (\varepsilon_i - 1) (d \log e - bd \log p)^2 \\ &\quad - (d \log e - bd \log p) \left[\sum_i b_i (\varepsilon_i - 1) \sum_i b_i d \log p_i \right] \\ &= -\frac{1}{1 - \theta_0} \sum_i b_i \varepsilon_i (\varepsilon_i - 1) (d \log e - bd \log p)^2 \\ &= -\frac{1}{1 - \theta_0} \text{Var}_b(\varepsilon_i) (d \log e - bd \log p)^2 \neq 0, \end{aligned}$$

where we used $\sum_i b_i \varepsilon_i = 1$. Hence, unless preferences are homothetic (in which case $\varepsilon_i = 1$ 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

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

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, $\frac{1 - \theta_0}{\sum_i b_i \xi_i}$ starts to change from 1.

Recall from Appendix B 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 F Analytical Examples with Input-Output Connections

We first discuss some differences between Proposition 3 and Proposition 7 in the presence of intermediate inputs. Proposition 3 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 7 implies that the gap between real GDP and welfare is

$$EV^M - \Delta \log Y \approx \frac{1}{2} \left[\Delta \log x' \frac{\partial \lambda^{sum}}{\partial \log x} + \Delta \log V \frac{\partial \lambda^{sum}}{\partial \log u} \right] \Delta \log A, \quad (55)$$

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.

In our quantitative results in Application 1 (section 5.1), the reallocation in sales towards sectors with lower intermediate input use 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.

We now extend the analytic examples in Section 4.2 to show how input-output connections can amplify or mitigate the gap between macro welfare EV^M and real GDP $\Delta \log Y$. 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 supply and demand shocks by making the technology shocks more uniform. However, since it is the covariance (not the correlation) of the shocks that matters, this means the effects are, at least theoretically, ambiguous.

To see these two forces, consider the three economies depicted in Figure 5. Each of these economies has a roundabout structure. Panel 5a depicts a situation where each producer uses only its own output as an input, Panel 5b a situation where all producers use the same basket of goods (denoted by M) as an intermediate input, and Panel 5c 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 9. 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 5: 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 5a, 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.³⁸

On the other hand, for Panel 5b, we get³⁹

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

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.⁴⁰

³⁸Even 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.

³⁹For 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 $i \in N$.

⁴⁰As indicated in Footnote 38, if the labor share is heterogeneous across producers, there is an additional

Finally, consider Panel 5c. 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

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

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.

Appendix G Additional details on Application I

In this appendix, we use a structural nested-CES model to explore the change in welfare-relevant TFP outside of the two polar extremes in Section 5.

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.

Proposition 5 implies that to compute the welfare-relevant change in TFP, we must only supply the information necessary to compute Hicksian sales shares at the terminal indifference curve. 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

---

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.

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).⁴¹ Production functions are also assumed to have nested-CES forms: there is an elasticity of substitution θ_2 between the bundle of intermediates and value-added, and an elasticity of substitution θ_3 across different types of intermediate inputs. For simplicity, we assume there is only one primary factor of production (a composite of capital and labor). We solve the non-linear model by repeated application of Proposition 8 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

⁴¹In 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.

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 9 to obtain the following second-order approximation:

$$\Delta \log TFP^{\text{ev}} \approx \sum_i \lambda_i \Delta \log A_i + \frac{1}{2} \sum_{j \in \{0\} + \mathcal{N}} (\theta_j - 1) \lambda_j \text{Var}_{\Omega(j,:)} \left(\sum_{k \in \mathcal{N}} \Psi_{(:,k)} \Delta \log A_k \right), \quad (56)$$

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 (56) 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 (56), 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 H 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 8 (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 12 (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 12 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 12. 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

$$\lambda_i = b_{iI}\lambda_I^c + s_{iI}\lambda_I^f, \quad (57)$$

$$d\lambda_i = db_{iI}\lambda_I^c + b_{iI}d\lambda_I^c + ds_{iI}\lambda_I^f + s_{iI}d\lambda_I^f, \quad (58)$$

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

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

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 (58), the sums can be re-written as

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

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

The individual terms can be written out as

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

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

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

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 I Additional Details on Application II

In this appendix, we provide additional details on how we treat the Nielsen data when constructing Figure 4, and we perform some robustness exercises with respect to the elasticity of substitution.

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.⁴² 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, and we use this stable-characteristic notion of UPCs.

This makes it unlikely that the good undergoes quality changes over time. First, 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. Second, even if a product keeps the same barcode but undergoes a change in one of the observable characteristics tracked by the Kilts Center, then it is not treated as the same product in our sample.

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 quan-

---

⁴²40,000 households in panel years 2004 to 2006.

tity.

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 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. 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 and "welfare-relevant" (equivalent variation at $t_1 = 2019$ or t_0 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 expenditures among continuing goods in that module (for the chained-weighted index, module weights vary by year t , and for the welfare-relevant indices, module weights are fixed at t_1 or t_0 given the assumption that the elasticity of substitution between modules is 1). 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 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). For each module, the welfare-relevant price index based on $t_1 = 2019$ preferences, given price changes between t_0 and t_1 , is

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

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 welfare-relevant price index based on t_0 preferences is

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

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.

Robustness Figure 6 replicates Figure 4 using lower and higher values for the elasticity of substitution across products within modules. 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 3, the bias is small.

Figure 6: Welfare-relevant and chain-weighted price index for continuing products. The welfare-relevant rate is computed assuming that the elasticity of substitution across UPCs in the same module is 2.5 in the left panel and 6.5 in the right panel.

Figure 7 replicates Figure 4 using monthly data. We set $t_1 = \text{June 2019}$ and consider monthly t_0 s rolling back to June 2004. For each t_0 , we further restrict the sample of products to those sold in every month between t_0 and t_1 (and in every month of the year of t_0) and that have a monthly log price change lower than one. The monthly price series are more volatile than the annual ones, but the welfare-relevant numbers are similar to the ones in Figure 4, but the chained-index is much closer to initial tastes than final tastes for longer horizons. This indicates that the second-order approximation is less accurate using higher frequency data, and so the chained measure is not as close to the average of initial and final tastes.

Figure 7: Welfare-relevant and chain-weighted inflation rates for continuing products using monthly data. The welfare-relevant rates are computed assuming that the elasticity of substitution across UPCs in the same module is 4.5 and the elasticity of substitution across modules is one.

Appendix J Non-CES Functional Forms

In this appendix, we generalize Propositions 3 and 8 beyond CES functional forms. To do this, for each producer k with cost function $\mathbf{C}_k$, we define the Allen-Uzawa elasticity of substitution between inputs x and y as

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

where $\epsilon_k(x, y)$ is the elasticity of the demand by producer k for input x with respect to the price p_y of input y , and Ω_{ky} is the expenditure share in cost of input y . Note that the Allen-Uzawa elasticity of substitution is symmetric for any two input pair x and y .

For the household $k = 0$, we use the household's expenditure function in place of the cost function. That is, for the household ($k = 0$), we have $\theta_k(x, y) = \epsilon_{x,y}^H / b_y$, where $\epsilon_{x,y}^H$ is the Hicksian cross-price elasticity and b_y is the budget share on y . The Hicksian cross-price elasticity is, in turn, related to the Marshallian cross-price elasticity by way of Slutsky's equation: $\epsilon_{xy}^H = \epsilon_{xy}^M + \epsilon_x^w b_y$, where ϵ_{xy}^M is the Marshallian cross-price elasticity.

Following Baqaee and Farhi (2019), define the *input-output substitution operator* for pro-

ducer k as

$$\Phi_k(\Psi_{(i)}, \Psi_{(j)}) = - \sum_{1 \leq x, y \leq N+1+F} \Omega_{kx} [\delta_{xy} + \Omega_{ky} (\theta_k(x, y) - 1)] \Psi_{xi} \Psi_{yj}, \quad (59)$$

$$(60)$$

where δ_{xy} is the Kronecker delta.

We can generalize all of our results beyond CES simply by replacing the terms involving covariances with the substitution operator above. 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 the more general case is very similar to the CES case.

That is, Proposition 3 can be generalized to the following.

Proposition 13 (General Approximation of Micro Welfare). *For a consumer with non-homothetic preferences to a second-order approximation, the change in real consumption is*

$$\begin{aligned} \Delta \log Y \approx & \Delta \log I - \mathbb{E}_{b_{t_0}} (\Delta \log p) - \frac{1}{2} \Phi_0 (\Delta \log p, \Delta \log p) \\ & - \frac{1}{2} \text{Cov}_{b_{t_0}} (\Delta \log x, \Delta \log p) - \frac{1}{2} \left[\Delta \log I - \mathbb{E}_{b_{t_0}} (\Delta \log p) \right] \text{Cov}_{b_{t_0}} (\varepsilon, \Delta \log p), \end{aligned}$$

and the change in welfare is

$$\begin{aligned} EV^m \approx & \Delta \log I - \mathbb{E}_{b_{t_0}} (\Delta \log p) - \frac{1}{2} \Phi_0 (\Delta \log p, \Delta \log p) \\ & - \text{Cov}_{b_{t_0}} (\Delta \log x, \Delta \log p) - \left[\Delta \log I - \mathbb{E}_{b_{t_0}} (\Delta \log p) \right] \text{Cov}_{b_{t_0}} (\varepsilon, \Delta \log p). \end{aligned} \quad (61)$$

In the expressions above, $d \log x_t$ is the residual in the Marshallian budget share not explained by income or substitution effects (these are caused by taste shocks). Formally, this is

$$d \log x_{it} = d \log b_{it} - \Phi_0(-d \log p_t, Id_{(:,i)}) - \text{Cov}_{b_t}(\varepsilon_t, I_{(:,i)}) (d \log I_t - \mathbb{E}_{b_t}[d \log p_t]).$$

Proposition 8 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*

$$d \log p_{it} = - \sum_{j \in N} \Psi_{ijt} d \log A_{jt} + \sum_{f \in F} \Psi_{ift}^F d \log \lambda_{ft}. \quad (62)$$

Changes in sales shares for goods and factors are

$$\begin{aligned} \lambda_{it} d \log \lambda_{it} = & \sum_{j \in \{0\} + N} \lambda_{jt} \Phi_j \left(-d \log p_t, \Psi_{(:,t),t} \right) \\ & + \text{Cov}_{\Omega_{(0,:),t}} \left(d \log x_t, \Psi_{(:,i),t} \right) + \text{Cov}_{\Omega_{(0,:),t}} (\varepsilon_t, \Psi_{(:,i),t}) \left(\sum_{k \in N} \lambda_{kt} d \log A_{kt} \right). \end{aligned} \quad (63)$$

Changes in welfare-relevant variables are pinned down by the same set of differential equations above where the second line of (63) is set to zero and the boundary conditions are that $\Omega = \Omega_{t_1}$ and $\Psi = \Psi_{t_1}$.

Appendix K Dynamic Economies

We consider a dynamic multi-sector model with production of consumption goods and investment goods similar to the models that are often used to study structural transformation (Herrendorf et al., 2013). For simplicity, we abstract from growth and restrict our discussion to non-homothetic CES preferences.⁴³

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

$$\mathcal{U}_t = \sum_{s=t}^{\infty} \beta^{s-t} U(C_s), \quad \sum_i \omega_{i0} x_{it} \left(\frac{c_{is}}{C_s^{\xi_i}} \right)^{\frac{\theta_0-1}{\theta_0}} = 1,$$

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 with arbitrary input-output connections

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

where A_{is} is a productivity shifter, l_{is} and k_{is} are labor and capital 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

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

---

⁴³For further discussion of welfare measures in dynamics economies with stable and homothetic preferences, see Licandro et al. (2002), Durán and Licandro (2018), and Duernecker et al. (2021).

where X_s is aggregate investment. Investment goods are produced according to a constant returns technology with arbitrary input-output connections

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

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 15 (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*

$$EV^M = \log \left(\frac{\sum_i p_{it_1} c_{it_1}}{\sum_i p_{it_0} c_{it_0}} \right) + \log \left(\sum_i b_{it_1} \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}. \quad (64)$$

In words, macroeconomic welfare in this dynamic economy is equal to the change in nominal consumption expenditures deflated by the exact-algebra CES price index associated with the t_1 indifference curve, exactly as for the partial equilibrium microeconomic welfare in expression (9), despite the fact that this is a dynamic general equilibrium economy with infinitely many factors.

Proof of Proposition 15. 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 cre-

ated 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

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

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

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 . $\square$