

Online Appendix
Welfare and Output with Income Effects and Taste Shocks
David Baqaee and Ariel Burstein
October 2022

A	Extension to Other Welfare Measures	A1
B	Relation to Konüs Price Indices	A3
C	Comparison of Quality and Taste Changes	A5
D	Non-homothetic CES preferences	A8
	D.1 Derivation of Marshallian budget shares	A8
	D.2 Comparison of welfare and changes in utility index	A9
E	Analytical Examples with Input-Output Connections	A12
F	Additional details on Application I	A14
G	Within-Industry Supply and Demand Shocks	A17
H	Additional Details on Application II	A20
I	Non-CES Functional Forms	A23
J	Dynamic Economies	A25

Appendix A 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

$$v(p_{t_1}, e^{-\phi} I_{t_1}; x_{t_0}) = v(p_{t_0}, I_{t_0}; x_{t_0}). \quad (\text{A1})$$

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} :

$$CV^m = \Delta \log I - \log \left(\sum_i b_{it_0} \left(\frac{p_{it_1}}{p_{it_0}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}. \quad (\text{A2})$$

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

$$\Delta \log CV^m = \Delta \log I - \mathbb{E}_{b_{t_0}}(\Delta \log p) - \frac{1}{2} \sum_{i \in N} \sum_j \Delta \log p_j \frac{\partial b_i}{\partial \log p_j} \Delta \log p_i \quad (\text{A3})$$

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 (A3), we see that up to a second order approximation (without imposing any specific form of preferences),

$$0.5 \times (EV^m + CV^m) \approx \Delta \log Y. \quad (\text{A4})$$

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

$$v(p_{t_1}, I_{t_1}; x_{t_0}) = v(p_{t_0}, e^\phi I_{t_0}; x_{t_0}). \quad (\text{A5})$$

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

$$v(p_{t_1}, e^{-\phi} I_{t_1}; x_{t_1}) = v(p_{t_0}, I_{t_0}; x_{t_1}). \quad (\text{A6})$$

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)

$$CV^M = \sum_{i \in N} \lambda_i \Delta \log A_i + \frac{1}{2} \sum_{i \in N} \sum_{j \in N} \Delta \log A_j \frac{\partial \lambda_i}{\partial \log A_j} \Delta \log A_i, \quad (\text{A7})$$

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_1 . This process does not require knowledge of either the income elasticities ε nor the taste shocks $\Delta \log x$.

Appendix B 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).^{A1}

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 ,

^{A1} CV^m can be calculated by deflating nominal income changes by the Konüs price index corresponding to the *initial* indifference curve.

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.

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 C 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 in the paper becomes the following (for brevity, we assume homothetic CES preferences):

Proposition A1 (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*

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

and the change in welfare is

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

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

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

$$\begin{aligned} EV^m - \Delta \log Y \approx & \underbrace{\mathbb{E}_b(\Delta \log q)}_{\text{average quality}} + \frac{1}{2} \underbrace{(\theta_0 - 1)Var_b(\Delta \log q)}_{\text{dispersion in quality}} + \frac{1}{2} \underbrace{(1 - \theta_0)Cov_b(\Delta \log \tilde{p}, \Delta \log q)}_{\text{covariance of price and quality}} \\ & - \frac{1}{2} \underbrace{Cov_b(\Delta \log \mathbf{x}, \Delta \log \tilde{p})}_{\text{covariance of taste and price}} + \underbrace{Cov_b(\Delta \log \mathbf{x}, \Delta \log q)}_{\text{covariance of taste and quality}}. \end{aligned} \quad (\text{A8})$$

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 (A8) simplifies to

$$EV^m - \Delta \log Y \approx -\frac{1}{2}Cov_b(\Delta \log x, \Delta \log \tilde{p}). \quad (\text{A9})$$

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 (A8) 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 (A8) simplifies to

$$EV^m - \Delta \log Y \approx \mathbb{E}_b(\Delta \log q) + \frac{1}{2}(\theta_0 - 1)Var_b(\Delta \log q) - \frac{1}{2}(\theta_0 - 1)Cov_b(\Delta \log \tilde{p}, \Delta \log q). \quad (\text{A10})$$

Comparing (A9) to (A10) 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.^{A2} 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 (A9) and (A10), the covariance of taste

^{A2}As 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.

shifters and market prices matters, however, in (A10) the sign of the covariance depends on whether the elasticity of substitution is greater than or less than one, whereas in (A9), the sign is always the same.

Appendix D 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.

D.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

$$e(p, u; x) = \left(\sum_{i \in N} x_i p_i^{1-\theta_0} u^{\xi_i} \right)^{\frac{1}{1-\theta_0}}, \quad (\text{A11})$$

with Hicksian demand implicitly defined by

$$c_i = x_i \left(\frac{p_i}{\sum_j p_j c_j} \right)^{-\theta_0} u^{\xi_i}, \quad (\text{A12})$$

and budget shares

$$b_i(p, x, u) \equiv \frac{p_i c_i}{\sum_j p_j c_j} = x_i u^{\xi_i} \left(\frac{p_i}{\sum_j p_j c_j} \right)^{1-\theta_0} = \frac{x_i u^{\xi_i} p_i^{1-\theta_0}}{\sum_{j \in N} x_j u^{\xi_j} p_j^{1-\theta_0}} \quad (\text{A13})$$

Differentiating (A11) and (A13) at any point t ,

$$d \log b_{it} = d \log x_{it} + (1 - \theta_0) (d \log p_{it} - d \log I_t) + \xi_i d \log u_t, \quad (\text{A14})$$

and

$$d \log u_t = \frac{1 - \theta}{\sum_j b_{jt} \xi_j} \left[d \log I_t - \sum_j b_{jt} d \log p_{jt} \right] - \frac{1}{\sum_j b_{jt} \xi_j} \sum_j b_{jt} d \log x_{jt}. \quad (\text{A15})$$

Substituting (A15) into (A14),

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

with demand shifters

$$d \log x_{it} = d \log x_{it} - \frac{\tilde{\xi}_i}{\sum_j b_{jt} \tilde{\xi}_j} \sum_j b_{jt} d \log x_{jt}, \quad (\text{A16})$$

and income elasticities

$$\varepsilon_{it} - 1 = (1 - \theta_0) \left(\frac{\tilde{\xi}_i}{\sum_j b_{jt} \tilde{\xi}_j} - 1 \right). \quad (\text{A17})$$

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.

D.2 Comparison of welfare and changes in utility index

In this appendix, we discuss the difference between changes in welfare as measured by the equivalent variation and changes in the utility index in non-homothetic CES preferences. 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 online Appendix C). 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 (A11). If preferences are homothetic ($\xi_i = \bar{\xi}$ for all i), then $e(p, u) = \left(\sum_i \omega_i p_i^{1-\theta_0}\right)^{\frac{1}{1-\theta_0}} u^{\frac{\bar{\xi}}{1-\theta_0}}$ and we can write

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

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

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

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

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

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)

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

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

The derivatives EV^m and U are in general different. Whereas EV^m is a function of observables, U depends on the cardinalization of the utility function. In particular $\sum_i b_i \bar{\zeta}_i$ 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 D.1. A standard approach in the literature to pin down $\sum_i b_i \bar{\zeta}_i$ is to set one of the $\bar{\zeta}$ to 1.

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

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

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

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

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

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

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

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

Appendix E 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

$$EV^M - \Delta \log Y \approx \frac{1}{2} \left[\Delta \log x' \frac{\partial \lambda^{sum}}{\partial \log x} + \Delta \log V \frac{\partial \lambda^{sum}}{\partial \log u} \right] \Delta \log A, \quad (\text{A18})$$

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 A1. Each of these economies has a roundabout structure. Panel A1a depicts a situation where each producer uses only its own output as an input, Panel A1b a situation where all producers use the same basket of goods (denoted by M) as an intermediate input, and Panel A1c 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 A1: 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 A1a, we get

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

where the covariance is computed across goods $i \in N$ and Ω_{iL} is the labor share for i . Hence, as intermediate inputs become more important, the necessary adjustment becomes larger. This is because, for a given vector of preference shocks, the movement in sales shares is now larger due to the roundabout nature of production.^{A3}

^{A3}Even if all productivity shocks are the same, there may still be an adjustment due to heterogeneity in labor

On the other hand, for Panel A1b, we get^{A4}

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

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.^{A5}

Finally, consider Panel A1c. For clarity, we focus on the case where only producer 1 gets a productivity shock ($\Delta \log A_2 = 0$). In this case, the difference between real GDP and welfare is

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

As the intermediate input share Ω_{21} approaches one, the adjustment goes to zero (since the covariance term goes to zero). Intuitively, as Ω_{21} goes to one, the increase in demand for the first producer from a change in preferences is exactly offset by a reduction in demand from the second producer who buys inputs from the first producer. In this limiting case, changes in consumer preferences have no effect on the overall sales share of the first producer.

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

Appendix F 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 im-

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.

^{A4}For this example, we assume that there are no productivity shocks to the intermediate bundle $\Delta \log A_M = 0$ and we assume that $\Omega_{iM} = 1/N$ for each $i \in N$.

^{A5}As indicated in Footnote A3, if the labor share is heterogeneous across producers, there is an additional adjustment which depends on the covariance between demand shocks and labor shares. If the demand shocks reallocate expenditures towards sectors with high labor shares, then welfare becomes less sensitive to productivity shocks than real GDP.

portant 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 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).^{A6} 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 A1.

The first column in Table A1 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

^{A6}In order to map this nested structure to our baseline model, good 0 is a composite of good 1-3, where good 1 is a composite of primary industries, good 2 is a composite of manufacturing industries, and good 3 is a composite of service industries. Goods 4-65 are the disaggregated industries. Finally, good 66 is the single factor of production.

Table A1: 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%

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 A1 also shows that not all elasticities of substitution are equally important. The results are much more sensitive to changes in the elasticity of substitution across more disaggregated categories, like materials, than aggregated categories, like agriculture, manufacturing, and services.

To see why the results in Table A1 are differentially sensitive to changes in different elasticities of substitution, we use Proposition 9 to obtain the following second-order approximation:

$$\Delta \log TFP^{\text{ev}} \approx \sum_i \lambda_i \Delta \log A_i + \frac{1}{2} \sum_{j \in \{0\} + \mathcal{N}} (\theta_j - 1) \lambda_j \text{Var}_{\Omega_{(j,:)}} \left(\sum_{k \in \mathcal{N}} \Psi_{(:,k)} \Delta \log A_k \right), \quad (\text{A19})$$

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 (A19) 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 A1 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 (A19), setting $\theta_1 = \theta_2 = \theta_3 = 1$ (which is similar to abstracting from heterogeneity within the three broader sectors and heterogeneity within intermediate inputs), then θ_0 is the only parameter that can generate substitution effects in the model. This may help understand why more aggregated models of structural transformation (e.g. Buera et al., 2015 and Alder et al., 2019) require low values of θ_0 to account for the extent of sectoral reallocation in the data.

Appendix G Within-Industry Supply and Demand Shocks

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

Definition A1 (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 A2 (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 A2 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 A2. 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

$$\lambda_i = b_{iI}\lambda_I^c + s_{iI}\lambda_I^f, \quad (\text{A20})$$

$$d\lambda_i = db_{iI}\lambda_I^c + b_{iI}d\lambda_I^c + ds_{iI}\lambda_I^f + s_{iI}d\lambda_I^f, \quad (\text{A21})$$

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

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

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 (A21), 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

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

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

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

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 H 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.^{A7} 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

^{A7}In the period 2004 to 2006 the panel has roughly 40,000 households.

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 quantity.

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

We set $t_1 = 2019$, and $t_0 = 2004, \dots, 2018$. For each t_0 we construct a balanced sample 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) prices 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 A2 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 A3 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 A2: 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 A3: 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 I 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 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 xy and yx .

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 producer k as

$$\Phi_k(\Psi_{(i)}, \Psi_{(j)}) = - \sum_{1 \leq x, y \leq N+1+F} \Omega_{kx} [\delta_{xy} + \Omega_{ky} (\theta_k(x, y) - 1)] \Psi_{xi} \Psi_{yj}, \quad (\text{A22})$$

$$(\text{A23})$$

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 A3 (General Approximation of Micro Welfare). *For a consumer with non-homothetic preferences to a second-order approximation, the change in real consumption is*

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

and the change in welfare is

$$EV^m \approx \Delta \log I - \mathbb{E}_{b_{t_0}} (\Delta \log p) - \frac{1}{2} \Phi_0 (\Delta \log p, \Delta \log p) \quad (\text{A24})$$

$$- \text{Cov}_{b_{t_0}} (\Delta \log \mathbf{x}, \Delta \log p) - \left[\Delta \log I - \mathbb{E}_{b_{t_0}} (\Delta \log p) \right] \text{Cov}_{b_{t_0}} (\varepsilon, \Delta \log p).$$

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

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

Proposition 8 generalizes as follows:

Proposition A4. *At any point in time t , changes in the relevant variables are pinned down by the following system of equations*

$$d \log p_{it} = - \sum_{j \in N} \Psi_{ijt} d \log A_{jt} + \sum_{f \in F} \Psi_{ift}^F d \log \lambda_{ft}. \quad (\text{A25})$$

Changes in sales shares for goods and factors are

$$\lambda_{it} d \log \lambda_{it} = \sum_{j \in \{0\} + N} \lambda_{jt} \Phi_j \left(-d \log p_t, \Psi_{(:,t),t} \right) \quad (\text{A26})$$

$$+ \text{Cov}_{\Omega_{(0,:),t}} \left(d \log \mathbf{x}_t, \Psi_{(:,i),t} \right) + \text{Cov}_{\Omega_{(0,:),t}} (\varepsilon_t, \Psi_{(:,i),t}) \left(\sum_{k \in N} \lambda_{kt} d \log A_{kt} \right).$$

Changes in welfare-relevant variables are pinned down by the same set of differential equations above where the second line of (A26) is set to zero and the boundary conditions are that $\Omega = \Omega_{t_1}$ and $\Psi = \Psi_{t_1}$.

Appendix J 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.^{A8}

^{A8}For 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).

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

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

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

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

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 A5 (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

$$EV^M = \log \left(\frac{\sum_i p_{it_1} c_{it_1}}{\sum_i p_{it_0} c_{it_0}} \right) + \log \left(\sum_i b_{it_1} \left(\frac{p_{it_0}}{p_{it_1}} \right)^{1-\theta_0} \right)^{\frac{1}{1-\theta_0}}. \quad (\text{A27})$$

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 A5. 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 A1 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. \square

Lemma A1. *The steady-state choice of capital (and investment) is the same for any homothetic and stable within-period preferences.*

Proof. Suppose intertemporal welfare is given by

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

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

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

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

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

and the capital accumulation equation is

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

The Lagrangean is

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

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

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

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

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

References

- Alder, S., T. Boppart, and A. Müller (2019). A theory of structural change that can fit the data.
- Baqee, D. R. and E. Farhi (2019). The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem. *Econometrica* 87(4), 1155–1203.
- Buera, F. J., J. P. Kaboski, and R. Rogerson (2015). Skill biased structural change. 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*, 1393–1414.
- 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.
- 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?
- 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.
- Herrendorf, B., R. Rogerson, and A. Valentinyi (2013). Two perspectives on preferences and structural transformation. *American Economic Review* 103(7), 2752–89.
- 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.
- 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.