A UNIFIED APPROACH TO ESTIMATING DEMAND AND WELFARE

Stephen J. Redding
David E. Weinstein

A Unified Approach to Estimating Demand and Welfare
Stephen J. Redding and David E. Weinstein
NBER Working Paper No. 22479
August 2016
JEL No. D11,D12,E01,E31

ABSTRACT

The measurement of price changes, economic welfare, and demand parameters is currently based on three disjoint approaches: macroeconomic models derived from time-invariant utility functions, microeconomic estimation based on time-varying utility (demand) systems, and actual price and real output data constructed using formulas that differ from either approach. The inconsistencies are so deep that the same assumptions that form the foundation of demand-system estimation can be used to prove that standard price indexes are incorrect, and the assumptions underlying standard exact and superlative price indexes invalidate demand-system estimation. In other words, we show that extant micro and macro welfare estimates are biased and inconsistent with each other as well as the data. We develop a unified approach that exactly rationalizes observed micro data on prices and expenditure shares while permitting exact aggregation and meaningful macro comparisons of welfare over time. Using barcode data for the U.S. consumer goods industry, we show that allowing for the entry and exit of products, changing preferences for individual goods, and a value for the elasticity of substitution estimated from the data yields substantially different conclusions for changes in the cost of living from standard index numbers.

Stephen J. Redding
Department of Economics
and Woodrow Wilson School
Princeton University
Fisher Hall
Princeton, NJ 08544
and NBER
reddings@princeton.edu

David E. Weinstein
Columbia University, Department of Economics
420 W. 118th Street
MC 3308
New York, NY 10027
and NBER
dew35@columbia.edu

1 Introduction

The measurement of economic welfare and demand patterns is currently based on three disjoint approaches: macroeconomic methods derived from *time-invariant* utility functions, microeconomic estimation based on *time-varying* utility (demand) systems, and actual price and real output data constructed using formulas that differ from either approach. The inconsistencies are so deep that the same assumptions that form the foundation of demand-system estimation can be used to prove that standard price indexes are incorrect, and the assumptions underlying standard exact and superlative price indexes invalidate demand-system estimation. In other words, we show that extant micro and macro welfare estimates are inconsistent with each other as well as the data.[1]

In order to deal with this problem, our paper presents a new empirical methodology, which we term "the unified approach," that reconciles all major micro, macro, and statistical approaches. Our "unified price index" nests all major price indexes used in welfare or demand system analysis. Thus, how economists and statistical agencies currently measure welfare can be understood in terms of an internally consistent approach that has been altered by ignoring data, moment conditions, and/or imposing particular parameter restrictions. For example, allowing the elasticity of substitution to differ from the Cobb-Douglas assumption of one produces the Sato-Vartia (1976) constant elasticity of substitution (CES) exact price index. Introducing the entry and exit of goods over time generates the Feenstra-CES index (Feenstra (1994). Incorporating demand shocks for each good and estimating the elasticity of substitution using the assumption that these time-varying demand shifts for individual products cancel on average produces the unified index. Other paths are shorter. The Jevons (1865) index—a geometric average of price widely used as an input into many price indexes—is a special case of the unified price index when the elasticity of substitution is infinite. The unified index exactly corresponds to expected utility if consumers have heterogeneous random utility with extreme value distributions (e.g., Logit or Fréchet). Similarly, the Dutot (1738), Carli (1764), Laspeyres (1871) and Paasche (1875) indexes all can be derived from the unified approach by making the appropriate parameter restrictions. Finally, relaxing assumptions necessary to yield the Fisher (1922) and Törnqvist (1936) indexes, yields the broader class of quadratic mean price indexes. The Sato-Vartia index arises naturally in this class, and as we just discussed, yields the unified price index if it is generalized. In other words, many seemingly fundamentally different approaches to welfare measurement—e.g., Laspeyres and Cobb-Douglas indexes—are actually linked together via the unified approach.

The first key insight of the unified approach is that any demand system errors (e.g., taste shocks) must show up in the utility and unit expenditure functions, and therefore the price index. However, all extant exact and superlative indexes (such as the Sato-Vartia, Fisher and the Törnqvist) are derived under the assumption that the demand parameter for each good is *time invariant*. Researchers make this assumption because it is a *sufficient* condition to guarantee the existence of a constant aggregate utility function. Unfortunately, the standard assumption also creates a conundrum. As we show in the paper, if one assumes that demand shocks are time invariant, one can solve for a constant elasticity of substitution *without doing estimation!* Thus,

---

[1]Recent contributions to the measurement of the cost of living and aggregate productivity across countries and over time include Bils and Klenow (2001), Hsieh and Klenow (2009), Jones and Klenow (2016), Feenstra (1994), Neary (2004) and Syverson (2016).

if one believes the assumption underlying all economically motivated price indexes—that demand does not shift—demand system estimation is both wrong (because it assumes demand shifts) and irrelevant (because identification does not require econometrics). Alternatively, if one believes the overwhelming evidence that demand for each good is not constant over time, *i.e.,* demand curves shift, then this violates the assumptions underlying economic approaches to macro price and welfare measurement. In other words, macro and micro approaches are based on contradictory assumptions: either one can believe the assumption of constant demand underlying exact price indexes, which means that demand-system estimates are incorrect, or one can believe micro evidence that demand curves shift, which means that existing price and real output measures are incorrect.

The solution to bridging the micro-macro divide requires our second key insight, which is to show that the assumption of time-invariant preferences for each good is neither the correct nor a necessary condition to make consistent comparisons of welfare over time when there are demand shocks for individual goods. To be able to make such consistent welfare comparisons, one must obtain the same change in the cost of living between a pair of time periods, whether one uses today's preferences for both periods, yesterday's preferences for both periods, or the preferences for each period (so that all three comparisons are consistent with one another). This necessary condition is (trivially) satisfied when preferences for each good are time invariant, but it is also satisfied whenever the utility function is *money metric*, in the sense that the change in the cost of living depends solely on changes in prices and expenditure shares (and does not directly depend on changes in preferences). We provide sufficient conditions for the utility function to be money metric, which require that the demand shocks for the individual goods cancel out and hence do not directly affect the change in the cost of living.

These conditions enable us to write down our "unified price index," which is exact for the CES utility function in the presence of mean-zero, time-varying demand shocks for each good as well as when the set of goods is changing. Moreover, in contrast to many conventional index numbers, our index also has the advantage that it is robust to mean-zero log additively separable measurement error in prices and expenditure shares. Finally, by comparing the Sato-Vartia CES index with ours, we identify a new source of "consumer valuation" bias that arises whenever one measures prices under the assumption that demand never shifts and applies such an index to data in which demand curves actually do move. This bias will be positive whenever demand shifts are positively correlated with expenditure shifts. For example, if positive demand shifts are associated with price and expenditure increases, a conventional price index will tend to overstate changes in the cost of living because it will weight the price increases more heavily than the decreases and fail to take into account the fact that these price increases are partially offset in utility terms by consumers getting more utility per unit from the newly preferred goods.

One of the most surprising results from incorporating demand shocks into the utility function is that we provide a new way to identify the demand parameters. Traditional approaches rely on estimating demand and supply shifts. When the identifying assumptions underlying these approaches are satisfied, they yield consistent estimates of the elasticity of substitution that can be incorporated into our unified price index, but they do not make full use of all of the moment conditions implied by the CES preference structure. In

particular, we show that when there are demand shocks for each good a price index will typically imply that the unit expenditure function for consumers with the initial or final sets of preferences will depend on the precise set of preference shocks. In other words, given the same prices and income consumers in two time periods would report different utility levels. In such circumstances, one cannot write down a money-metric utility function, and standard welfare analysis becomes problematic. To overcome this problem, we introduce a novel estimation technique that makes use of information contained not just in the demand system, but also the unit expenditure function. Surprisingly, this permits identification without specifying the supply side.

The intuition for identification arises from counting equations and unknowns in a simple setup with continuous and differentiable prices and expenditure shares. If we think about a dataset containing price and share changes for k goods, we have k unknowns (one unknown price elasticity and $k-1$ unknown values for each of the preference shocks given a normalization). However, we also have a system of k independent equations ($k-1$ independent demand equations and one equation for the change in the unit expenditure function). Therefore, the system is exactly identified. In other words, given data on prices and expenditure shares and the assumption of a money-metric utility function, one does not need to estimate demand parameters; one just solves for them.

The problem is more complex when there are discrete changes, because price and expenditure share derivatives become finite differences, but the same basic intuition applies. With discrete changes, we show that there are three equivalent ways of writing the change in the unit expenditure function: one using the expenditure shares of consumers in the initial period and the second using the expenditure shares of consumers in the final period, and a third (our unified price index) which uses preferences in both periods. While the UPI will always generate a money-metric utility function for any value of the elasticity of substitution, the other two indexes will produce deviations from money-metric utility that depend on the precise value of the elasticity of substitution used in their computation. We use this fact to develop a "reverse-weighting" estimator that identifies the elasticity of substitution by minimizing the deviations of the first two unit expenditure functions from a money-metric utility function. For small demand shocks, this reverse-weighting estimator consistently estimates the true elasticity of substitution and the preference parameter for each good and time period irrespective of the size of price shocks and the correlation between preference and price shocks, and always yields a unique money-metric utility function. More generally, we show that this reverse-weighting estimator provides a first-order approximation to the data, which becomes exact as demand and price shocks become small.

We focus on the CES functional form, because there is little doubt that this is the preferred approach to modeling product variety across international trade, economic geography and macroeconomics. We also address a number of potential shortcomings of this approach. Our CES price index is not superlative, because it does not approximate any continuous and differentiable utility function. But superlative indexes like Fisher (used in the personal consumption expenditure index) and Törnqvist are closely related to CES indexes, because they arise from quadratic mean utility functions, and can be written as similar functions of price and expenditure share data. Indeed, we find that if we impose similar parameter restrictions on our unified price index (no demand shocks or variety changes), the differences in measured price changes between our index

and superlative indexes in the data are trivial. This result establishes that, empirically, the key differences between the unified and superlative indexes stem from assumptions about the existence of demand shocks or new goods, not functional forms.

A second potential concern is that agents may not be homogeneous. Our unified index features symmetry and homotheticity and exhibits an independence of irrelevant alternatives (IIA) property (the relative expenditure of any two varieties only depends on the characteristics of those varieties and not on the characteristics of other varieties within a market). Building price indexes when this assumption is violated has proven to be a vexing issue for economists. For example, Deaton (1998) writes, "it is unclear that a quality-corrected cost-of-living index in a world with many heterogeneous agents is an operational concept." More recently, Chevalier and Kashyap (2014) have investigated differences in inflation rates in models with consumer heterogeneity. In order to address this concern, we show, as an extension, how to break these features by allowing for heterogeneous consumers with different elasticities of substitution and demand for each good, as in Berry, Levinsohn and Pakes (1995) and McFadden and Train (2000). In this extension, the elasticity of substitution for a given good can vary across markets depending on the composition of heterogeneous types (breaking symmetry), the relative demand for two goods can depend on what other goods are supplied to the market (when it affects the expenditure shares of the heterogeneous types); and differences in the elasticity of substitution and demand parameters across the heterogeneous types allow for non-homotheticities across types. This extension thus unifies the heterogeneous consumer and price index literatures.

Our paper is related to a number of strands of existing research. First, we build on a long line of existing research on price indexes. Price measurement in most national and international agencies is based on the "statistical approach" to price indexes developed by Dutot (1738), Carli (1764), and Jevons (1865). The methodologies developed in these papers form the foundation of 98 percent of all consumer price indexes generated by government statistical agencies (Stoevska 2008). We show how sampling techniques convert these indexes into Laspeyres (1871), Paasche (1875), and Cobb-Douglas price indexes.[2] These indexes as well as the Fisher (1922) and Törnqvist (1936) price indexes are either nested or closely related to our unified price index.

However, the path to the unified price index need not start with the actual price indexes used by statistical agencies. Following Konüs (1924), economic theory has largely rejected the "statistical approach" to price measurement in favor of the "economic approach," which asserts that all price indexes should be derived from consumer theory and correspond to the unit expenditure function. The subsequent economic approach to price measurement, including Diewert (1976), Sato (1976), Vartia (1976), Lau (1979), Feenstra (1994), Moulton (1996), Balk (1999), Caves, Christensen and Diewert (1982), Neary (2004) and Feenstra and Reinsdorf (2010), has focused on exact and superlative index numbers that feature time-invariant demand parameters. Our unified price index also arises naturally when following this economic approach. We show how to relax the assumption of time-invariant demand for each good while preserving a money-metric utility function. Thus, although there has been an international rift in the approach to measuring the cost of living—with the

---

[2]The "Cobb-Douglas" functional form was first used by Wicksell (1898) and the price index was discovered by Konyus (Konüs) and Byushgens (1926). Cobb and Douglas (1928) applied it to U.S. data. For a review of the origins of index numbers, see Chance (1966).

U.S. Department of Labor accepting the economic approach to price measurement and U.K. statistical agencies explicitly rejecting it (Triplett 2001)—we show these debates can be reduced to asking what restrictions should be placed on the unified approach.

It is an interesting feature of the literature that even path-breaking economists who have taught us how to measure time-varying demand parameters often assume these away in the same work when they measure price indexes and welfare. For example, Deaton and Muellbauer (1980) provide extensive discussions of time-varying demand in the estimation of the demand system. However, when they use unit expenditure functions that are standard in the price index literature in order to show how to measure welfare changes, there is no discussion of the fact that these were derived (elsewhere) based on a time-invariant formulation of demand. Similarly, Feenstra (1994), identifies CES parameters based on the heteroskedasticity of demand shocks, and explicitly points out the inconsistency between the demand system estimation and the CES price index, but does not resolve it.

Our study is also related to a more recent, voluminous literature in macroeconomics, trade and economic geography that has used CES preferences. This literature includes, among many others, Anderson and van Wincoop (2003), Antràs (2003), Arkolakis, Costinot and Rodriguez-Clare (2012), Armington (1969), Bernard, Redding and Schott (2007, 2011), Blanchard and Kiyotaki (1987), Broda and Weinstein (2006, 2010), Dixit and Stiglitz (1977), Eaton and Kortum (2002), Feenstra (1994), Helpman, Melitz and Yeaple (2004), Hsieh and Klenow (2009), Krugman (1980, 1991), Krugman and Venables (1995) and Melitz (2003). Increasingly, researchers in international trade and development are turning to bar-code data in order to measure the impact of globalization on welfare. Prominent examples of this include Handbury (2013), Atkin and Donaldson (2015), and Atkin, Faber, and Gonzalez-Navarro (2015), and Fally and Faber (2016). Our contribution relative to this literature is to derive an exact price index that allows for changes in variety and demand for each good, while preserving the property of a money-metric utility function.

Our work is also related to research in macroeconomics aimed at measuring the cost of living, real output, and quality change. Shapiro and Wilcox (1996) sought to back out the elasticity of substitution in the CES index by equating it to a superlative index. Whereas that superlative index number assumed time-invariant demand for each good, we explicitly allow for time-varying demand for each good, and derive the appropriate index number in such a case. Bils and Klenow (2001) quantify quality growth in U.S. prices. We show how to incorporate changes in quality (or subjective taste) for each good into a unified framework for computing changes in the aggregate cost of living over time and estimating the elasticity of substitution.

Finally, our analysis connects with the broader literature on demand systems estimation, including McFadden (1974), Deaton and Muellbauer (1980), Anderson, de Palma and Thisse (1992), Berry (1994), Berry, Levinsohn and Pakes (1995), McFadden and Train (2000), Sheu (2014), and Thisse and Ushchev (2016). A related literature examines the implications of new goods for welfare, including Feenstra (1994), Bresnahan and Gordon (1996), Hausman (1996), Nevo (2003), Broda and Weinstein (2006, 2010) and Petrin (2002). In contrast to these literatures, our method emphasizes the intimate relationship between price indexes and demand systems. We provide an approach that exactly rationalizes the observed data on prices and expenditure for individual goods as an equilibrium of the model, while also preserving a money-metric utility function, and

6

hence permitting meaningful comparisons of welfare over time.

The remainder of the paper is structured as follows. Section 2 develops our theoretical framework and derives our unified price index. Section 3 examines the relationships between this unified price index and the standard price indexes used by economists and statistical agencies. Section 4 incorporates heterogeneous groups of consumers with different substitution parameters. Section 5 shows how our unified approach can be used to estimate the elasticity of substitution. Section 6 uses detailed barcode data for the U.S. consumer goods sector to illustrate our approach and demonstrate its quantitative relevance for measuring changes in the aggregate cost of living. Section 7 concludes.

2   The Unified Price Index

We begin by considering a CES utility function with time-varying demand parameters for each good and write down the price index and demand system that are compatible with it when the set of goods is changing over time. We show how the price index and demand system can be combined to derive our unified price index. For expositional clarity, we develop our approach in the simplest possible setting with a representative consumer, but we relax this assumption in a later section.[3] Although we initially treat the elasticity of substitution as known and solve for the demand parameters for all goods and time periods, we show in later sections how our unified approach can be used to estimate both the elasticity of substitution and the demand parameters.

2.1   Preferences and Demand

Utility ($\mathbb{U}_t$) is defined over the consumption (C_{kt}) of each good k at time t:

$$\mathbb{U}_t = \left[\sum_{k \in \Omega_t} \left(\varphi_{kt} C_{kt} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 1, \qquad \varphi_{kt} > 0, \tag{1}$$

where σ is the elasticity of substitution across goods; φ_{kt} is the preference ("demand") parameter for good k at time t; and the set of goods supplied at time t is denoted by Ω_t.[4] Although we allow demand parameters for individual goods (φ_{kt}) to change over time, we continue to assume a constant elasticity of substitution (σ) over time, as is required for money-metric utility. The corresponding unit expenditure function ($\mathbb{P}_t$) is defined over the price (P_{kt}) of each good k at time t:

$$\mathbb{P}_t = \left[\sum_{k \in \Omega_t} \left(\frac{P_{kt}}{\varphi_{kt}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \tag{2}$$

Applying Shephard's Lemma to this unit expenditure function, we obtain the demand system in which the expenditure share ($S_{\ell t}$) for each good ℓ and time period t is:

$$S_{\ell t} \equiv \frac{P_{\ell t} C_{\ell t}}{\sum_{k \in \Omega_t} P_{kt} C_{kt}} = \frac{\left(P_{\ell t} / \varphi_{\ell t} \right)^{1-\sigma}}{\sum_{k \in \Omega_t} \left(P_{kt} / \varphi_{kt} \right)^{1-\sigma}}, \qquad \ell \in \Omega_t. \tag{3}$$

---

[3]For simplicity, we also assume a single CES tier of utility, but our approach generalizes immediately to a nested CES structure, as shown in Section A.1 of the web appendix.

[4]We focus on CES utility as in Dixit and Stiglitz (1977) and abstract from the generalizations of the love of variety properties of CES in Benassy (1996) and Behrens et al. (2014).

We allow the demand parameters (φ_{kt}) to vary across goods and over time so as to exactly rationalize the observed expenditure shares (S_{kt}) as an equilibrium of the model given the observed prices (P_{kt}) and the elasticity of substitution (σ). These demand parameters (φ_{kt}) are therefore structural residuals that ensure the model explains the observed data. Our unified approach exploits the key insight of duality that these parameters in the demand system are intimately related to those in the unit expenditure function. Assuming time-invariant parameters for each good in the utility function (as in all exact and superlative index numbers) while at the same time assuming time-varying parameters for each good in the demand system (as in all empirical demand systems estimation) is inconsistent with the principles of duality. Instead our unified approach allows the demand parameters for each good to change over time (so that model exactly rationalizes the observed data on prices and expenditure shares) while at the same time preserving a money-metric utility function (so as to make comparisons of aggregate welfare over time).

Another important feature of our framework is that we allow for the entry and exit of goods over time, as observed in the data. In particular, we partition the set of goods in period t (Ω_t) into those "common" to t and $t-1$ ($\Omega_{t,t-1}$) and those added between $t-1$ and t (I_t^+), where $\Omega_t = \Omega_{t,t-1} \cup I_t^+$. Similarly, we partition the set of goods in period $t-1$ (Ω_{t-1}) into those common to t and $t-1$ ($\Omega_{t,t-1}$) and those dropped between $t-1$ and t (I_t^-), where $\Omega_{t-1} = \Omega_{t,t-1} \cup I_{t-1}^-$. We denote the number of goods in period t by $N_t = |\Omega_t|$ and the number of common goods by $N_{t,t-1} = |\Omega_{t,t-1}|$. We assume that $\varphi_{kt} = 0$ for a good k before it enters and after it exits, which rationalizes the observed entry and exit of goods over time.

2.2 Changes in the Cost of Living

We now combine the unit expenditure function (2) and demand system (3) to derive our unified price index, taking into account the entry and exit of goods and changes in demand for each good. We start by expressing the change in the cost of living from $t-1$ to t as the ratio between the unit expenditure functions (2) in the two periods:

$$\Phi_{t-1,t} = \frac{\mathbb{P}_t}{\mathbb{P}_{t-1}} = \left[\frac{\sum_{k \in \Omega_t} \left(P_{kt}/\varphi_{kt} \right)^{1-\sigma}}{\sum_{k \in \Omega_{t-1}} \left(P_{kt-1}/\varphi_{kt-1} \right)^{1-\sigma}} \right]^{\frac{1}{1-\sigma}}. \tag{4}$$

The fact that the set of goods is changing means that the set of goods in the denominator is not the same as that in the numerator. Feenstra (1994) showed that one way around this problem is to express the price index in terms of price index for "common goods" (i.e., goods available in both time periods) and a variety-adjustment term. Summing equation (3) over the set of commonly available goods, we can express expenditure on all common goods as a share of total expenditure in periods t and $t-1$ respectively as:

$$\lambda_{t,t-1} \equiv \frac{\sum_{k \in \Omega_{t,t-1}} \left(P_{kt}/\varphi_{kt} \right)^{1-\sigma}}{\sum_{k \in \Omega_t} \left(P_{kt}/\varphi_{kt} \right)^{1-\sigma}}, \qquad \lambda_{t-1,t} \equiv \frac{\sum_{k \in \Omega_{t,t-1}} \left(P_{kt-1}/\varphi_{kt-1} \right)^{1-\sigma}}{\sum_{k \in \Omega_{t-1}} \left(P_{kt-1}/\varphi_{kt-1} \right)^{1-\sigma}}, \tag{5}$$

where $\lambda_{t,t-1}$ is equal to the total sales of continuing goods in period t divided by the sales of all goods available in time t evaluated at *current* prices. Its maximum value is one if no goods enter in period t and will *fall* as the share of *new* goods *rises*. Similarly, $\lambda_{t-1,t}$ is equal to total sales of continuing goods as share of total sales of all goods in the *past* period evaluated at $t-1$ prices. It will equal one if no goods cease being sold and will *fall* as the share of *exiting* goods *rises*.

Multiplying the numerator and denominator of the fraction inside the square parentheses in (4) by the summation $\sum_{k \in \Omega_{t,t-1}} (P_{kt}/\varphi_{kt})^{1-\sigma}$ over common goods at time t, and using the definition of $\lambda_{t,t-1}$ in (5), we obtain:

$$\Phi_{t-1,t} = \left[\frac{1}{\lambda_{t,t-1}} \frac{\sum_{k \in \Omega_{t,t-1}} (P_{kt}/\varphi_{kt})^{1-\sigma}}{\sum_{k \in \Omega_{t-1}} (P_{kt-1}/\varphi_{kt-1})^{1-\sigma}} \right]^{\frac{1}{1-\sigma}}.$$

Multiplying the numerator and denominator by the summation $\sum_{k \in \Omega_{t,t-1}} (P_{kt-1}/\varphi_{kt-1})^{1-\sigma}$ over common goods at time $t-1$, and using the definition of $\lambda_{t-1,t}$ in (5), we obtain the exact CES price index:

$$\Phi_{t-1,t} = \left[\frac{\lambda_{t-1,t}}{\lambda_{t,t-1}} \frac{\sum_{k \in \Omega_{t,t-1}} (P_{kt}/\varphi_{kt})^{1-\sigma}}{\sum_{k \in \Omega_{t,t-1}} (P_{kt-1}/\varphi_{kt-1})^{1-\sigma}} \right]^{\frac{1}{1-\sigma}} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)^{\frac{1}{\sigma-1}} \frac{\mathbb{P}_t^*}{\mathbb{P}_{t-1}^*}, \tag{6}$$

where we use an asterisk to denote the value of a variable for the common set of goods (*i.e.*, goods available in periods t and $t-1$), such that $\mathbb{P}_t^*$ is the unit expenditure function defined over *common* goods:

$$\mathbb{P}_t^* \equiv \left[\sum_{k \in \Omega_{t,t-1}} \left(\frac{P_{kt}}{\varphi_{kt}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \tag{7}$$

The common goods price index ($\mathbb{P}_t^*/\mathbb{P}_{t-1}^*$) is the change in the cost of living if the set of goods is not changing, and it will prove to be a useful building block in our unified price index. The term multiplying it in equation (6) is the "variety-adjustment" term ($(\lambda_{t,t-1}/\lambda_{t-1,t})^{1/(\sigma-1)}$). This term adjusts the common goods price index for entering and exiting goods. If new goods are more numerous than exiting goods or have lower prices relative to demand (lower (P_{kt}/φ_{kt})), then $\lambda_{t,t-1}/\lambda_{t-1,t} < 1$, and the price index ($\Phi_{t-1,t}$) will fall due to an increase in variety or the entering varieties having higher demand than the exiting varieties.

To complete the derivation of our unified price index, we use the CES demand system (3), which implies that the share of each common good in expenditure on all common goods ($S_{\ell t}^*$) is:

$$S_{\ell t}^* \equiv \frac{P_{\ell t} C_{\ell t}}{\sum_{k \in \Omega_{t,t-1}} P_{kt} C_{kt}} = \frac{(P_{\ell t}/\varphi_{\ell t})^{1-\sigma}}{\sum_{k \in \Omega_{t,t-1}} (P_{kt}/\varphi_{kt})^{1-\sigma}}, \qquad \ell \in \Omega_{t,t-1}. \tag{8}$$

Rearranging terms, we obtain the following useful relationship for the common goods unit expenditure function:

$$(\mathbb{P}_t^*)^{1-\sigma} = \sum_{k \in \Omega_{t,t-1}} (P_{kt}/\varphi_{kt})^{1-\sigma} = \frac{1}{S_{\ell t}^*} (P_{\ell t}/\varphi_{\ell t})^{1-\sigma}, \qquad \ell \in \Omega_{t,t-1}. \tag{9}$$

If we take logs of both sides of equation (9), difference over time, sum across all $\ell \in \Omega_{t,t-1}$, and divide both sides by the number of common goods, we find that the log change in the common goods price index can be written as:

$$\ln \left(\frac{\mathbb{P}_t^*}{\mathbb{P}_{t-1}^*} \right) = \ln \left(\frac{\tilde{P}_t^*}{\tilde{P}_{t-1}^*} \right) + \frac{1}{\sigma-1} \ln \left(\frac{\tilde{S}_t^*}{\tilde{S}_{t-1}^*} \right) - \ln \left(\frac{\tilde{\varphi}_t^*}{\tilde{\varphi}_{t-1}^*} \right), \tag{10}$$

where a tilde over a variable denotes a geometric average and the asterisk indicates that the geometric average is taken for the set of common goods, such that $\tilde{x}_t^* = \left(\prod_{k \in \Omega_{t,t-1}} x_{kt} \right)^{1/N_{t,t-1}}$ and $\tilde{x}_{t-1}^* = \left(\prod_{k \in \Omega_{t,t-1}} x_{kt-1} \right)^{1/N_{t,t-1}}$ for the variables x_{kt} and x_{kt-1}.

Our objective in this paper is to allow for demand shifts for individual goods while still being able to make consistent comparisons of welfare over time. To be able to make such consistent welfare comparisons between a pair of time periods, one must obtain the same change in the cost of living whether one uses today's preferences for both periods, yesterday's preferences for both periods, of the preferences of each period (so that all three comparisons are consistent with one another). This requirement for consistent welfare comparisons implies that the change in the cost of living depends solely on changes in prices and expenditure shares (and not directly on changes in preferences). From equation (10), this condition requires:

$$\ln\left(\frac{\tilde{\varphi}_t^*}{\tilde{\varphi}_{t-1}^*}\right) = 0. \tag{11}$$

This condition marks the main theoretical difference between all extant economically motivated index numbers and our approach. Standard approaches also derive a money-metric utility function by imposing a much stronger assumption: namely, that there are *no* demand shifts, i.e., $\ln\left(\varphi_{kt}/\varphi_{k,t-1}\right) = 0$ for all $k \in \Omega_{t,t-1}$. It immediately follows from this assumption that equation (11) holds. While this standard assumption is obviously a sufficient condition to obtain a money-metric utility function, it is not a necessary one. Indeed, the math only requires that something much weaker hold: that the geometric average of the demand shifts is zero. Thus, the crucial difference between our approach and the conventional one is that the unified price index does not require that there be no demand shifts, only that they cancel *on average*. Crucially, when condition (11) is satisfied, a given set of price and expenditure shares in the two periods will correspond to the same change in the cost of living in (10) whether one uses the period $t-1$ preferences for both periods (because $\ln\left(\tilde{\varphi}_{t-1}^*/\tilde{\varphi}_{t-1}^*\right) = 0$); the period t preferences for both periods (because $\ln\left(\tilde{\varphi}_t^*/\tilde{\varphi}_t^*\right) = 0$) or the preferences for each period (because $\ln\left(\tilde{\varphi}_t^*/\tilde{\varphi}_{t-1}^*\right) = 0$). Therefore, whenever equation (11) holds, the utility function is money metric.

Oddly enough, while assuming that equation (11) holds may seem like a startling departure from the standard time-invariant taste assumption commonly used to build macro price indexes, in applied microeconomics, the assumption is banal. In demand-systems estimation, equation (11) is nothing more than the assumption that the log demand shocks are mean zero: $\frac{1}{N_{t,t-1}}\sum_k \Delta\ln\left(\varphi_{kt}\right) = \ln\left(\tilde{\varphi}_t^*/\tilde{\varphi}_{t-1}^*\right) = 0$. This property suggests another advantage of our approach. We unify both macro and micro approaches to welfare measurement by making the *same* assumption of time-varying, mean-zero stochastic shocks in all equations.

We can impose this condition through a simple normalization. The demand system is homogeneous of degree zero in the demand shifters (φ_{kt}), and hence these preference parameters can be measured up to a normalization. We choose units such that the geometric mean of demand for common goods is equal to one ($\tilde{\varphi}_t^* = \left(\prod_{k\in\Omega_{t,t-1}}\varphi_{kt}\right)^{1/N_{t,t-1}} = 1$), which guarantees that (11) is satisfied.[5] Using this normalization and the expenditure share (3), we can solve explicitly for demand for each good k and time period t in terms of

---

[5]An advantage of this normalization is that it does not depend on the characteristics of the common goods, such as their expenditure shares, which can change endogenously over time. Feenstra and Reinsdorf (2007) assume that demand for each good is stochastic and use a normalization for the demand parameters based on expenditure shares to derive standard errors for index numbers. In Section A.2 of the web appendix, we show that our approach can be generalized to allow for a Hicks-neutral demand shifter that raises or lowers the demand for all goods. This generalization leaves our estimation of the elasticity of substitution (σ) unchanged, because that estimation equates three equivalent expressions for the change in the cost of living, and the change in the Hicks-neutral demand shifter cancels from these three equivalent expressions.

observed prices (P_{kt}) and expenditure shares (S_{kt}) and the elasticity of substitution (σ):

$$\varphi_{kt} = \frac{P_{kt}}{\tilde{P}_t^*} \left(\frac{S_{kt}}{\tilde{S}_t^*} \right)^{\frac{1}{\sigma-1}}. \tag{12}$$

Substituting our normalization (11) and the expression for the common goods price index (10) into the overall CES price index (6) yields our main proposition:

Proposition 1. *The "unified price index" (UPI)—which is exact for the CES preference structure in the presence of changes in the set of goods, demand-shocks that do not directly affect utility, and discrete changes in prices and expenditure shares—is given by*

$$\Phi_{t-1,t}^U = \underbrace{\left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)^{\frac{1}{\sigma-1}}}_{\text{Variety Adjustment}} \underbrace{\left[\frac{\tilde{P}_t^*}{\tilde{P}_{t-1}^*} \left(\frac{\tilde{S}_t^*}{\tilde{S}_{t-1}^*} \right)^{\frac{1}{\sigma-1}} \right]}_{\text{Common-Goods Price Index } \Phi_{t-t,t}^{CG}}. \tag{13}$$

Proof. The proposition follows directly from substituting equations (10) and (11) into (6). □

The UPI also has important similarities with other price indexes that enables us to nest many existing approaches in this framework. For example, as in Feenstra (1994), the unified price index (UPI) expresses the change in the cost of living as a function of a variety-adjustment term and a common-goods component of the unified price index (CG-UPI). The variety adjustment term (namely $(\lambda_{t,t-1}/\lambda_{t-1,t})^{1/(\sigma-1)}$ in equation (13)) captures changes in the unit expenditure function due to product turnover, changes in the number of varieties, and new goods. The CG-UPI (denoted by $\Phi_{t-t,t}^{CG}$ in equation (13)) measures how changes in prices, demand-shifts, and product substitution for common goods affects a consumer's unit expenditure function. It is comprised of two terms. The first term ($\tilde{P}_t^*/\tilde{P}_{t-1}^*$) is none other than the geometric average of price relatives that serves as the basis for lower level of the U.S. Consumer Price Index (also known as the "Jevons" index). Indeed, in the special case in which varieties are perfect substitutes ($\sigma \to \infty$), the UPI collapses to the Jevons index, since both $(\lambda_{t,t-1}/\lambda_{t-1,t})^{1/(\sigma-1)}$ and $(\tilde{S}_t^*/\tilde{S}_{t-1}^*)^{1/(\sigma-1)}$ converge to one as $\sigma \to \infty$.

The last term ($(\tilde{S}_t^*/\tilde{S}_{t-1}^*)^{1/(\sigma-1)}$) is novel and captures heterogeneity in expenditure shares across common goods. This term moves with the ratio of the geometric mean of common goods expenditure shares in the two periods. Critically, as the market shares of common goods in a time period become more uneven, the geometric average will fall. Thus, this term implies that the cost of living will fall if expenditure shares become more dispersed. The intuition for this result can be obtained by considering a simple example. Imagine that there are just two goods in every period and that the price of both goods is the same and unchanging across time. In this example, the variety-adjustment and price terms are one, and we can focus on demand shocks. Now suppose that consumers initially prefer the first good to the second, which means that the first good constitutes a larger share of expenditure. Consider how utility would move if consumers faced a mean-zero demand shock that shifted the preference parameter for the first good up by 1 percent and the preference parameter for the second good down by 1 percent. This would cause the geometric average of the shares to fall because the dispersion in the shares would rise. Importantly, utility would also rise (and the cost of

11

living would fall) because the consumer would benefit more from a positive demand shift for a good that constitutes a large share of expenditure than an equal negative shift for a good that constitutes a small share of expenditure. Thus, demand shifts that *raise* the dispersion in expenditures *lower* the price index because consumers benefit more from positive taste shifts for goods that constitute big shares of expenditures. More generally, when both prices and demand are changing, this term captures the tendency for P_{kt}/φ_{kt} to fall by more for goods with large market shares.[6]

The UPI in (13) has a number of desirable economic and statistical properties. First, this price index and each of its components are time reversible for any value of σ, thereby permitting consistent comparisons of welfare going forwards and backwards in time. Second, given a value for the elasticity of substitution, the common goods price index is unaffected by mean-zero log additive measurement error in either prices or expenditure shares, because such measurement error leaves the geometric means of prices and expenditure shares unchanged. In contrast, most existing price indexes are non-linear functions of observed expenditure shares and are directly affected by such measurement error. Third, the unified price index depends in a simple and transparent way on the elasticity of substitution. Variation in this elasticity leaves the terms in common goods prices unchanged ($\tilde{P}_t^*/\tilde{P}_{t-1}^*$) and affects the variety adjustment ($\left(\lambda_{t,t-1}/\lambda_{t-1,t}\right)^{1/(\sigma-1)}$) and heterogeneity terms ($\left(\tilde{S}_t^*/\tilde{S}_{t-1}^*\right)^{1/(\sigma-1)}$) depending on the extent to which these two expenditure share ratios are greater than or less than one. Finally, the relative magnitude of these variety and heterogeneity corrections in logs is independent of the value of the elasticity of substitution, and depends solely on the relative values of expenditure share moments in the data ($\ln\left(\left(\lambda_{t,t-1}/\lambda_{t-1,t}\right)/\ln\left(\tilde{S}_t^*/\tilde{S}_{t-1}^*\right)\right)$).

3   Relation to Existing Price Indexes

In this section, we compare our unified price index with all of the main economic and statistical price indexes used in the existing theoretical and empirical literature on price measurement. We first discuss the relationship between our index and other indexes for the CES demand system. We next show that all other conventional price indexes are special cases of the unified price index that either impose particular parameter restrictions (on the elasticity of substitution), abstract from the entry and exit of goods, and/or neglect changes in demand for each good.

3.1   Relation to Existing Exact CES Price Indexes

The formula for the UPI differs from the CES price index in Feenstra (1994) because we do not use the Sato (1976) and Vartia (1976) formula for the common goods price index. The formula for the Feenstra index is

---

[6]Our unified price index (13) differs from the expression for the CES price index in Hottman et al. (2016), which did not distinguish entering and exiting goods from common goods (omitting $\left(\lambda_{t,t-1}/\lambda_{t-1,t}\right)^{1/(\sigma-1)}$) and captured the dispersion of sales across common goods in different way (using a different term from $\left(\tilde{S}_t^*/\tilde{S}_{t-1}^*\right)^{1/(\sigma-1)}$).

given by:[7]

$$\frac{\mathbb{P}_t^*}{\mathbb{P}_{t-1}^*} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}}\right)^{\frac{1}{\sigma-1}} \Phi_{t-1,t}^{SV}, \qquad \Phi_{t-1,t}^{SV} \equiv \prod_{k\in\Omega_{t,t-1}} \left(\frac{P_{kt}}{P_{kt-1}}\right)^{\omega_{kt}^*}, \qquad \omega_{kt}^* \equiv \frac{\frac{S_{kt}^*-S_{kt-1}^*}{\ln S_{kt}^*-\ln S_{kt-1}^*}}{\sum_{\ell\in\Omega_{t,t-1}} \frac{S_{\ell t}^*-S_{\ell t-1}^*}{\ln S_{\ell t}^*-\ln S_{\ell t-1}^*}}. \qquad (14)$$

Both indexes require the estimation of σ, but our approach resolves a tension that Feenstra (1994) observed was inherent in his use of the Sato-Vartia formula. The Sato-Vartia index ($\Phi_{t-1,t}^{SV}$) used for $\mathbb{P}_t^*/\mathbb{P}_{t-1}^*$ assumes that demand is constant over time for each good ($\varphi_{kt} = \varphi_{kt-1} = \bar{\varphi}_k$ for all $k \in \Omega_{t,t-1}$ and t), whereas the estimation of σ assumes that demand for goods changes over time ($\varphi_{kt} \neq \varphi_{kt-1}$ for some k and t).

This tension is more pernicious than it might appear because the assumption of time invariant demand is a *crucial* assumption for the derivation of the Sato-Vartia index, and the index cannot be derived if one assumes mean zero log demand shocks. Under the assumption of constant demand for each common good ($\varphi_{kt} = \varphi_{kt-1} = \bar{\varphi}_k$ for all $k \in \Omega_{t,t-1}$), we show in the proposition below that there is no need to estimate σ, because it can be recovered from observed prices and expenditure shares using the weights from the Sato-Vartia price index. Furthermore, the model is overidentified when demand is constant for each common good, with the result that there exists an *infinite* number of approaches to measuring σ. If demand is indeed constant for each common good ($\varphi_{kt} = \varphi_{kt-1} = \bar{\varphi}_k$ for all $k \in \Omega_{t,t-1}$), each of these approaches returns exactly the same value for σ. However, if demand for goods changes over time ($\varphi_{kt} \neq \varphi_{kt-1}$ for some $k \in \Omega_{t,t-1}$), and a researcher falsely assumes constant demand for each good, we show that each of these approaches returns a different value for σ in every time period. Even making the additional assumption that on average the change in demand for goods is zero for common goods does not eliminate the problem. These approaches produce a different value for σ unless demand is constant for every common good.

Proposition 2. *(a) Under the assumption that demand is constant for each common good ($\varphi_{kt} = \varphi_{kt-1} = \bar{\varphi}_k$ for all $k \in \Omega_{t,t-1}$ and t), the elasticity of substitution (σ) is uniquely identified from observed changes in prices and expenditure shares with no estimation. Furthermore, there exists a continuum of approaches to measuring σ, each of which weights prices and expenditure shares with different non-negative weights that sum to one, but returns the same value for σ.*
(b) If demand for common goods changes over time ($\varphi_{kt} \neq \varphi_{kt-1}$ for some $k \in \Omega_{t,t-1}$ and t), but a researcher falsely assumes that demand for each common good is constant, each of these alternative approaches returns a different value for σ, depending on which non-negative weights are used.

Proof. See Section A.3 of the web appendix. □

This proposition makes clear the link between the common-goods component of the unified price index and the standard Sato-Vartia CES price index. If there are no demand shifts, the two indexes are identical. In the presence of non-zero demand shifts, the CG-UPI exactly replicates the observed data on expenditure shares and prices as an equilibrium of the model based on the assumption of a constant elasticity of substitution (σ) and time-varying demand (φ_{kt}). In contrast, the Sato-Vartia index assumes time-invariant demand for each

---

[7]As shown in Banerjee (1983), the Sato-Vartia weights (ω_{kt}^*) are only one of a broader class of weights that can be used to construct the exact common-goods CES price index with constant demand for each common good ($\varphi_{kt} = \bar{\varphi}_k$).

good, which implies that the model does not exactly replicate the observed data on expenditure shares and prices if there are non-zero demand shifts. The elasticity of substitution implied by the Sato-Vartia index will vary with these demand shifts, which makes the Sato-Vartia price index depend on demand parameters and therefore incompatible with a money-metric utility function. The implicit elasticity of substitution in the Sato-Vartia CES price index is is not only time varying (a property we will explore in Section 6.2), but also will differ based on what arbitrary subset of common goods are included in the index and how one weights them. Therefore, if there are demand shifts, standard CES price indexes imply that the elasticity of substitution is not constant within a time period or across them, rendering the utility function time varying and traditional welfare analysis problematic. By contrast, a key advantage of the UPI is that it results in a money-metric utility function even in the presence of these shocks.

This problem also biases any attempt to measure aggregate price changes using a Sato-Vartia formula in the presence of demand shifts as the following proposition demonstrates.

Proposition 3. *In the presence of non-zero demand shocks for some good (i.e., $\ln\left(\varphi_{kt}/\varphi_{kt-1}\right) \neq 0$ for some $k \in \Omega_{t,t-1}$), the Sato-Vartia price index ($\Phi_{t-1,t}^{SV}$) differs from the exact common goods CES price index. The Sato-Vartia price index ($\Phi_{t-1,t}^{SV}$) equals the unified price index (13) plus a demand shock bias term.*

$$\ln \Phi_{t-1,t}^{SV} = \ln \Phi_{t-1,t}^{CG} + \underbrace{\left[\sum_{k\in\Omega_{t,t-1}} \omega_{kt}^* \ln\left(\frac{\varphi_{kt}}{\varphi_{kt-1}}\right)\right]}_{demand\ shock\ bias}, \tag{15}$$

$$where \qquad \varphi_{kt} = \frac{P_{kt}}{\tilde{P}_t^*}\left(\frac{S_{kt}}{\tilde{S}_t^*}\right)^{\frac{1}{\sigma-1}}, \qquad \omega_{kt}^* \equiv \frac{\frac{S_{kt}^*-S_{kt-1}^*}{\ln S_{kt}^*-\ln S_{kt-1}^*}}{\sum_{\ell\in\Omega_{t,t-1}}\frac{S_{\ell t}^*-S_{\ell t-1}^*}{\ln S_{\ell t}^*-\ln S_{\ell t-1}^*}}, \qquad \sum_{k\in\Omega_{t,t-1}} \omega_{kt}^* = 1. \tag{16}$$

Proof. See Section A.4 of the web appendix. □

In order for the Sato-Vartia price index to be unbiased, we require demand shocks ($\varphi_{kt}/\varphi_{kt-1}$) to be uncorrelated with the Sato-Vartia weights (ω_{kt}^*) in the demand shock bias term. However, the Sato-Vartia weights are endogenous and depend on the demand parameter (φ_{kt}). As shown in the proposition below, a positive demand shock for a good mechanically increases the Sato-Vartia weight for that good and reduces the Sato-Vartia weight for all other goods. Other things equal, this mechanical relationship introduces a positive correlation between demand shocks ($\varphi_{kt}/\varphi_{kt-1}$) and the Sato-Vartia weights (ω_{kt}^*), which implies that the Sato-Vartia price index ($\Phi_{t-1,t}^{SV}$) is upward biased.

Proposition 4. *A positive demand shock for a good k (i.e., $\ln\left(\varphi_{kt}/\varphi_{kt-1}\right) > 0$ for some $k \in \Omega_{t,t-1}$) increases the Sato-Vartia weight for that good (ω_{kt}^*) and reduces the Sato-Vartia weight for all other goods $\ell \neq k$ ($\omega_{\ell t}^*$).*

Proof. See Section A.5 of the web appendix. □

Therefore, in the presence of demand shocks, the Sato-Vartia index is not only a noisy measure of the change in the cost of living but is also upward biased, and hence overstates the increase in the cost of living over time. The intuition for why conventional indexes like the Sato-Vartia suffer from this "consumer valuation bias" in the presence of mean-zero demand shocks is simple. Suppose the price of no good changes between $t-1$ and t. All conventional indexes will report a price change of zero. However, if there are *any* demand shocks, consumers with period t preferences will adjust their expenditure shares so that they increase consumption of the goods that they like more in period t and reduce consumption of the goods they like less. However, if no price has changed, they still can consume their original bundle of goods, so they must be better off in period t. More generally, even if prices and demand shifts are positively correlated, the bias will arise as long as demand shifts are associated with higher expenditure shares. If demand shifts in favor of a good and the price of that good rises, a conventional index will tend to overstate the price increase because it implicitly assumes that the failure of the expenditure share to fall for the newly expensive good is due to a low elasticity of substitution and not to a demand shift. Put concretely, if a consumer initially consumes equal amounts of Coke and Pepsi but then starts to like Pepsi more, any relative price increase of Pepsi must be offset by the fact that the consumer is now getting more utility per unit from Pepsi consumption. Thus, the UPI will report a lower change in the cost of living than an index that assumed there was no change in preferences. Our UPI incorporates these implications of changes in relative preferences for goods, while preserving the property that a money-metric utility function exists.[8]

In conclusion, Propositions 2-4 show that there are two major differences between our index (13) and the Feenstra index. First, if one assumes that demand for each good is time invariant when it is in fact time varying, the Sato-Vartia formula arbitrarily implies one of an infinite set of elasticities that are consistent with the CES functional form, and none of these need be consistent with the elasticity identified using econometric techniques. Thus, our index eliminates the inconsistency that Feenstra (1994) identified as arising from imposing no demand shocks when computing the price change for the common goods component of the CES price index while also assuming these shocks to be time varying when estimating σ for the variety correction term $((\lambda_{t,t-1}/\lambda_{t-1,t})^{1/(\sigma-1)})$. Second, we show that the assumption of time-invariant demand in the construction of price indexes introduces an upward "consumer valuation bias" because of the counterfactual assumption that consumers will not shift expenditures towards goods they prefer.

3.2 Relation to Conventional Price Indexes

The unified price index that we have developed is exact for the CES functional form and expresses changes in the cost of living solely in terms of prices and expenditure shares. However, there are two other equivalent expressions for the change in the cost of living in terms of prices, expenditure shares and demands for each good. These equivalent expressions arise from forward and backward differences of the unit expenditure

---

[8]Our analysis focuses on CES preferences, because these yield a tractable specification for controlling for the entry and exit of goods over time and estimating the elasticity of substitution between goods (see Section 5). In Section A.14 of the web appendix, we show that the same bias from neglecting changes in demand for each good arises in the translog functional form. In the presence of time-varying demand for each good, the Törnqvist index differs from the exact translog price index and is upward biased. Section A.15 of the web appendix shows that continuous time index numbers, such as the Divisia index, also make the assumption of constant demand for each good.

function and we now make them explicit in order to relate our approach to other conventional price indexes and to later show how our approach can be used to estimate the elasticity of substitution between goods.

The forward difference of the unit expenditure function evaluates the *increase* in the price index from $t-1$ to t using the expenditure shares of consumers in period $t-1$. Using equations (5), (6), (7) and (8), this forward difference can be written in terms of the change in variety ($\lambda_{t,t-1}/\lambda_{t-1,t}$), the initial share of each common good in expenditure on all common goods (S^*_{kt-1}), and changes in prices (P_{kt}/P_{kt-1}) and demand ($\varphi_{kt}/\varphi_{kt-1}$) for all common goods:

$$
\Phi^F_{t-1,t} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}}\right)^{\frac{1}{\sigma-1}} \frac{\mathbb{P}^*_t}{\mathbb{P}^*_{t-1}} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}}\right)^{\frac{1}{\sigma-1}} \left[\sum_{k\in\Omega_{t,t-1}} S^*_{kt-1} \left(\frac{P_{kt}/\varphi_{kt}}{P_{kt-1}/\varphi_{kt-1}}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}, \tag{17}
$$

as shown in Section A.6 of the web appendix. The backward difference of the unit expenditure function uses the expenditure shares of consumers period t to evaluate the *decrease* in the price index from t to $t-1$. Using equations (5), (6), (7) and (8), this backward difference can be written in an analogous form as:

$$
\Phi^B_{t,t-1} = \left(\frac{\lambda_{t-1,t}}{\lambda_{t,t-1}}\right)^{\frac{1}{\sigma-1}} \frac{\mathbb{P}^*_{t-1}}{\mathbb{P}^*_t} = \left(\frac{\lambda_{t-1,t}}{\lambda_{t,t-1}}\right)^{\frac{1}{\sigma-1}} \left[\sum_{k\in\Omega_{t,t-1}} S^*_{kt} \left(\frac{P_{kt-1}/\varphi_{kt-1}}{P_{kt}/\varphi_{kt}}\right)^{1-\sigma}\right]^{\frac{1}{1-\sigma}}, \tag{18}
$$

where the algebra is again relegated to Section A.6 of the web appendix.[9]

The only variable not in common to the forward and backward differences is the expenditure share (S^*_{kt-1} versus S^*_{kt}). When evaluating the change in the cost of living going forward in time, we use the period $t-1$ expenditure shares, whereas when doing the same going backward in time, we use the period t expenditure shares. The terms in square brackets in (13), (17) and (18) correspond to three equivalent ways of expressing the change in the cost of living for common goods.

In general, the forward and backward differences are not money metric, in the sense that the change in the cost of living does not only depend on changes in prices and expenditure shares, but also depends on changes in the demand for each good. In Section 5 below, we provide conditions under which these forward and backward differences are also money metric, and show how these conditions can be used to estimate the elasticity of substitution between goods. Before doing so, we use the equivalence between the unified price index and these forward and backward differences to show that all conventional price indexes correspond to special cases of our unified price index that impose particular parameter restrictions, abstract from changes in demand for each good, and/or abstract from the entry and exit of goods over time.

According to an International Labor Organization (ILO) survey of 68 countries around the world, the Dutot (1738) index is still the most prominent one for measuring price changes (Stoevska (2008)).[10] This index is the ratio of a simple average of prices in two periods:

---

[9]The forward and backward differences in equations (17) and (18) are related to the comparisons of welfare using initial and final preferences considered in Fisher and Shell (1972). A key difference is that our expressions (17) and (18) include the change in demand for each good ($\varphi_{kt}/\varphi_{kt-1}$), and hence are exactly equal to the unified price index (13), rather than providing bounds for it.

[10]41 percent of countries use this index although historically its popularity was much higher. For example, all U.S. inflation data prior to 1999 is based on this index, and Belgian, German, and Japanese data continues to be based on it. The ILO report can be accessed here: http://www.ilo.org/public/english/bureau/stat/download/cpi/survey.pdf

$$\Phi_{t-1,t}^D \equiv \frac{\frac{1}{N_{t,t-1}}\sum_{k\in\Omega_{t,t-1}}P_{kt}}{\frac{1}{N_{t,t-1}}\sum_{k\in\Omega_{t,t-1}}P_{kt-1}} = \sum_{k\in\Omega_{t,t-1}}\frac{P_{k,t-1}}{\sum_{k\in\Omega_{t,t-1}}P_{kt-1}}\left(\frac{P_{kt}}{P_{k,t-1}}\right) \tag{19}$$

As the above formula shows, this index is simply a price-weighted average change in prices, which does not have a clear rationale in terms of economic theory.

A price-weighted average of price changes is a sufficiently problematic way of measuring changes in the cost of living that most statistical agencies do not just compute unweighted averages of prices in two periods, but select their sample of price quotes based on the largest selling products in the first period. If we think that the probability that a statistical agency picks a product for inclusion in its sample of prices is based on its purchase frequency ($C_{\ell,t-1}/\sum_{k\in\Omega_{t-1,t}}C_{k,t-1}$), then the Dutot index, as it is typically implemented, becomes the more familiar Laspeyres index:

$$\Phi_{t-1,t}^L \equiv \frac{\sum_{k\in\Omega_{t,t-1}}C_{k,t-1}P_{kt}}{\sum_{k\in\Omega_{t,t-1}}C_{k,t-1}P_{kt-1}} = \sum_{k\in\Omega_{t,t-1}}\frac{C_{k,t-1}P_{k,t-1}}{\sum_{k\in\Omega_{t,t-1}}C_{k,t-1}P_{kt-1}}\left(\frac{P_{kt}}{P_{k,t-1}}\right) = \sum_{k\in\Omega_{t,t-1}}S_{kt-1}^*\frac{P_{kt}}{P_{kt-1}}. \tag{20}$$

Written this way, it is clear that the Laspeyres index is a special case of our CES price index (17) in which the utility gain of new goods is exactly offset by the loss from disappearing goods ($\lambda_{t,t-1}/\lambda_{t-1,t} = 1$), the elasticity of substitution equals zero and demand for each good is constant ($\varphi_{kt}/\varphi_{kt-1} = 1$).

The Carli index, used by 19 percent of countries, is another popular index that can be thought of as a variant of the Laspeyres index. The formula for the Carli index is

$$\Phi_{t-1,t}^C \equiv \sum_{k\in\Omega_{t,t-1}}\frac{1}{N_{t,t-1}}\left(\frac{P_{kt}}{P_{k,t-1}}\right) \tag{21}$$

This index is identical to the Laspeyres if all goods have equal expenditure shares. However, as with the Dutot, it is important to remember that statistical agencies are more likely to select a good for inclusion in the sample with a past high sales share ($S_{k,t-1}^*$) for inclusion. In this case, the Carli index also collapses back to the Laspeyres formula.

Similarly, the Paasche index is closely related to the Laspeyres index with the only difference that it weights price changes from $t-1$ to t by their expenditure shares in the *end* period t:

$$\Phi_{t-1,t}^P = \frac{\sum_{k\in\Omega_{t,t-1}}P_{kt}C_{kt}}{\sum_{k\in\Omega_{t,t-1}}P_{kt-1}C_{kt}} = \left[\sum_{k\in\Omega_{t,t-1}}S_{kt}^*\left(\frac{P_{kt}}{P_{kt-1}}\right)^{-1}\right]^{-1}. \tag{22}$$

We can also think of the Paasche index as is a special case of the CES price index (18) in which we apply the same parameter restrictions to derive the Laspeyres index.[11]

Finally, the Jevons index, which forms the basis of the lower level of the U.S. Consumer Price Index, is the second-most popular index currently in use, with 37 percent of countries building their measures of changes in the cost of living based on it.[12] The index is constructed by taking an unweighted geometric mean of price

---

[11]To derive (22) from (18), we use $\Phi_{t-1,t} = 1/\Phi_{t,t-1}$, assume $\lambda_{t-1,t}/\lambda_{t,t-1} = 1$ and $\varphi_{kt}/\varphi_{kt-1} = 1$ for all k, and set $\sigma = 0$.

[12]The percentages do not sum to 100 because 3 percent of sample respondents used other formulas.

changes from $t-1$ to t:

$$\Phi^J_{t-1,t} = \prod_{k \in \Omega_{t,t-1}} \left(\frac{P_{kt}}{P_{kt-1}} \right)^{\frac{1}{N_{t,t-1}}} = \frac{\tilde{P}^*_t}{\tilde{P}^*_{t-1}}. \tag{23}$$

As we discussed earlier, this formula is just a special case of the unified price index (13) in the limit as $\sigma \to \infty$. It is also related to the unified price index through another route. Statistical agencies typically choose products based on their historic sales shares. In this case the Jevons index becomes:

$$\Phi^{CD}_{t-1,t} = \prod_{k \in \Omega_{t,t-1}} \left(\frac{P_{kt}}{P_{kt-1}} \right)^{S^*_{k,t-1}}, \tag{24}$$

which Konyus (Konüs) and Byushgens (1926) proved was exact for the Cobb-Douglas (1928) functional form. This price index is a special case of the CES price index when the elasticity of substitution equals one, demand for each good is constant, and there are no changes in variety.

Existing measures of changes in the cost of living are therefore special cases of the unified approach developed in this paper, and biases can be thought of in terms of parameter restrictions on the unified price index. For example, "substitution bias" arises from building a price index using the wrong elasticity of substitution (σ). Most studies of consumer behavior suggest that this elasticity is greater than one, but in Laspeyres and Paasche indexes it arises because this elasticity is assumed to be zero. The recent move to the Jevons index by many countries reduced the substitution bias by changing the elasticity in the unified price index to infinity or, if one reinterprets the Jevons index as a Cobb-Douglas index, an elasticity of one. Our index corrects for this shortcoming in previous indexes by letting the data determine the correct elasticity.

"Variety" or "New Goods Bias" arises from the assumption that $\lambda_{t,t-1}/\lambda_{t-1,t} = 1$, which means that the utility gain from new goods is exactly offset by the loss from disappearing goods.[13] The fact that we tend to think that price per unit quality of new goods exceeds that of disappearing goods—one gets more utility from paying \$1,000 for a computer today than ten years ago—implies that this assumption is wrong because $\lambda_{t,t-1}/\lambda_{t-1,t} < 1$. In contrast, our index explicitly incorporates new and disappearing goods into the measurement of changes in the cost of living.

The third "consumer valuation" bias is novel and arises because of the assumption that consumer demand for each good is constant over time ($\varphi_{kt}/\varphi_{kt-1} = 1$). Mechanically, this arises whenever a price index specifies that prices should be deflated by a demand parameter that is time varying (as here where the unit expenditure function depends on P_{kt}/φ_{kt}). In this sense, it is isomorphic to the well-known substitution bias that plagues fixed-weight indexes like the Laspeyres. Analogously, the consumer valuation bias arises whenever one fixes the utility parameter associated with a good because it assumes consumers will not change expenditure patterns when their tastes change.

Interestingly, the two remaining "superlative" price indexes (Fisher and Törnqvist) are also closely related to the CES. Taking the geometric mean of the forward and backward differences of the CES price index (17) and (18), which are equal to the unified price index (13), we obtain the following quadratic mean of order

---

[13]The new goods bias is typically stated in terms of an index not allowing for new goods, but this is not technically correct. The absence of new goods would correspond to $\lambda_{t,t-1} = 1$. While it is true that if there are no new or exiting goods, we will have $\lambda_{t,t-1} = \lambda_{t-1,t} = 1$, the validity of common goods price indexes depends on a slightly weaker assumption: $\lambda_{t,t-1}/\lambda_{t-1,t} = 1$.

$2(1 - \sigma)$ price index (Diewert 1976):

$$\Phi_{t,t-1} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)^{\frac{1}{\sigma - 1}} \left[\frac{\sum_{k \in \Omega_{t,t-1}} S^*_{kt-1} \left(\frac{P_{kt}/\varphi_{kt}}{P_{kt-1}/\varphi_{kt-1}} \right)^{1-\sigma}}{\sum_{k \in \Omega_{t,t-1}} S^*_{kt} \left(\frac{P_{kt}/\varphi_{kt}}{P_{kt-1}/\varphi_{kt-1}} \right)^{-(1-\sigma)}} \right]^{\frac{1}{2(1-\sigma)}}, \tag{25}$$

The Fisher index is the geometric mean of the Laspeyres (20) and Paasche (22) price indexes, and corresponds to the special case of (25) in which $\sigma = 0$, the utility gain from new goods is exactly offset by the loss from disappearing goods ($\lambda_{t,t-1}/\lambda_{t-1,t} = 1$), and demand for each good is constant ($\varphi_{kt}/\varphi_{kt-1} = 1$):

$$\Phi^F_{t-1,t} = \left(\Phi^L_{t-1,t} \Phi^P_{t-1,t} \right)^{1/2}. \tag{26}$$

Closely related to the Fisher index is the Törnqvist index, which corresponds to the limiting case of (25) in which $\sigma \to 1$, the utility gain from new goods is exactly offset by the loss from disappearing goods ($\lambda_{t,t-1}/\lambda_{t-1,t} = 1$), and demand for each good is constant ($\varphi_{kt}/\varphi_{kt-1} = 1$):

$$\Phi^T_{t-1,t} = \prod_{k \in \Omega_{t,t-1}} \left(\frac{P_{kt}}{P_{kt-1}} \right)^{\frac{1}{2}\left(S^*_{kt-1} + S^*_{kt}\right)}. \tag{27}$$

Another way of looking at the Törnqvist index is to realize that it is just a geometric average of Cobb-Douglas price indexes defined in equation (24) evaluated at times $t - 1$ and t.

The Fisher and Törnqvist price indexes are exact in the sense that they hold for flexible functional forms: quadratic mean of order-r preferences and the translog expenditure function respectively (Diewert 1976). These price indexes are also superlative in the sense that they provide a local second-order approximation to any continuous and differentiable expenditure function. However, we have shown that both indexes are closely related to the CES price index, and are in fact special cases of the geometric mean of two of our equivalent expressions for the CES price index (25) for a particular value of the elasticity of substitution. Therefore, the CES, Fisher and Törnqvist price indexes for common goods are all closely related functions of the same underlying price and expenditure data. Empirically, we show below that the differences between these three indexes are trivially small for a given set of common goods under the assumption of no changes in demand for each good. Importantly, the exact and superlative properties of the Fisher and Törnqvist indexes are derived under the assumption of no entry and exit of goods and no changes in demand for each good. A key advantage of our unified price index (13) relative to these other two price indexes is that it explicitly takes into account both product turnover and changes in consumer valuations of each good, which we show below to be central features of micro data on prices and expenditure shares.[14]

---

[14]In Section A.14 of the web appendix, we derive the Törnqvist index from the translog expenditure function and show that the consumer valuation bias from changing demand for each good is also present for this expenditure function. Additionally, we derive the generalization of the Törnqvist index to incorporate changes in demand for each good that is analogous to our generalization of the Sato-Vartia price index for the CES expenditure function.

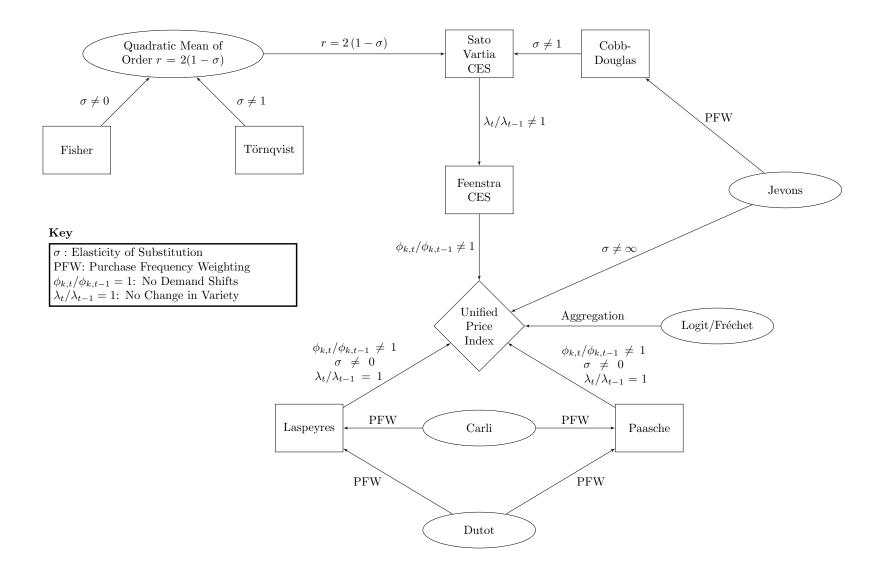Figure 1: Relation Between Existing Indexes and the UPI

Figure 1 summarizes how all major indexes are related to the unified index. Most indexes (such as the Dutot, Carli, Laspeyres, Paasche, Jevons, Cobb-Douglas, Sato-Vartia-CES, Feenstra-CES) are simply special cases of our index. One can think of the standard approach to index numbers, therefore, as versions of the unified approach in which researchers make different parameter restrictions, ignore certain parts of the data (e.g., new goods), ignore certain implications of the model (e.g., changes in tastes in the demand system also enter into the unit expenditure function), and fail to sample based on purchase frequencies. Existing exact CES price indexes assume no demand shocks, and superlative indexes are simply different weighted averages of the same building blocks as those of the unified index under the assumption of no change in the set of goods or consumer tastes. The relaxation of all of these assumptions and restrictions results in the unified approach.

4   The UPI with Heterogeneous Consumers

There are a number of objections that have been raised to using the CES setup. The easiest to dismiss is the one arising from the fact that if consumers had CES preferences they would demand all goods, while in reality we observe consumers that typically have a preferred variety as considered in the discrete choice literature following McFadden (1974). This objection is not really substantive as Anderson, de Palma, and Thisse (1992) showed that the CES preferences of the representative consumer are identical to the aggregate behavior of all consumers in a random utility model in which heterogeneous consumers only demand their preferred good.

A second potential objection is that CES imposes strong assumptions in the form of symmetric substitution elasticities, homotheticity, and the independence of irrelevant alternatives (IIRA). This IIRA property implies that the relative sales of any two varieties depends only on their relative characteristics and not on the characteristics of other varieties supplied to the market. Relaxing these assumptions was one of the key motivations for the random coefficients model with a continuum of unobserved types in Berry, Levinsohn and Pakes (1995). In this section, we also extend the random utility model to allow for multiple types of consumers with different substitution and preference parameters. This extension relaxes the assumptions of symmetry, homotheticity and IIRA using a discrete number of types as in the mixed logit model (see for example McFadden and Train 2000).

In particular, we partition consumers into different types indexed by $r \in \{1, \ldots, R\}$. The utility of an individual i of type r who consumes C_{ik}^r units of product k is:

$$\mathbb{U}_i^r = z_{ik}^r \varphi_k^r C_{ik}^r, \tag{28}$$

where φ_k^r captures type-r consumers' common tastes for product k; z_{ik}^r captures idiosyncratic consumer tastes for each product; and we have omitted the time subscript t on each variable to simplify notation. Each consumer i of type r, therefore, chooses C_{ik}^r units of good k to maximize utility. Since the consumer only consumes their preferred good, their budget constraint implies:

$$C_{ik}^r = \frac{E_i^r}{P_k^r}, \tag{29}$$

where E_i^r is the consumer's expenditure, and P_k^r is the price of the good, which we allow to potentially vary across types r. Using this result, utility (28) can be re-written in the indirect form as:

$$\mathbb{U}_i^r = z_{ik}^r \left(\varphi_k^r / P_k^r \right) E_i^r. \tag{30}$$

These idiosyncratic tastes are assumed to have a Fréchet (Type-II Extreme Value) distribution:

$$G(z) = e^{-z^{-\theta^r}}, \tag{31}$$

where we allow the shape parameter determining the dispersion of idiosyncratic tastes (θ^r) to vary across types. We normalize the scale parameter of the Fréchet distribution to one, because it affects consumer expenditure shares isomorphically to the consumer tastes parameter φ_k^r.[15] Using the monotonic relationship between idiosyncratic tastes and utility, we have:

$$z_{ik}^r = \frac{\mathbb{U}_i^r}{\left(\varphi_k^r / P_k^r \right) E_i^r}.$$

Therefore, the distribution of utility from product k for individual i is:

$$G_{ik}^r (\mathbb{U}^r) = \exp \left(- \left(\frac{\mathbb{U}_i^r P_k^r}{\varphi_k^r E_i^r} \right)^{-\theta^r} \right). \tag{32}$$

From this distribution of utility (32), the probability that an individual i of type r chooses product k is the same across all individuals of that type and equal to:

$$S_{ik}^r = S_k^r = \frac{\left(P_k^r / \varphi_k^r \right)^{-\theta^r}}{\sum_{\ell=1}^N \left(P_\ell^r / \varphi_\ell^r \right)^{-\theta^r}}, \tag{33}$$

as shown in Section A.7 of the web appendix. The probability (33) also equals the share of product k in the expenditure of consumers of type r (S_k^r), since all consumers of the same type are assumed to have the same expenditure: $E_i^r = E^r$. The expected utility of consumer i of type r is:

$$\mathbb{E}[\mathbb{U}^r] = \gamma^r \left[\sum_{k=1}^N (E_i^r)^{\theta^r} (P_k^r / \varphi_k^r)^{-\theta^r} \right]^{\frac{1}{\theta^r}}, \qquad \gamma^r = \Gamma \left(\frac{\theta^r - 1}{\theta^r} \right), \tag{34}$$

where $\Gamma(\cdot)$ is the Gamma function, as also shown in Section A.7 of the web appendix. This expected utility can be re-written as:

$$\mathbb{E}[\mathbb{U}^r] = \frac{E_i^r}{\mathbb{P}^r}, \tag{35}$$

where $\mathbb{P}^r$ is the unit expenditure function for consumers of type r:

$$\mathbb{P}^r = (\gamma^r)^{-1} \left[\sum_{k=1}^N (P_k^r / \varphi_k^r)^{-\theta^r} \right]^{-\frac{1}{\theta^r}}. \tag{36}$$

Total expenditure on a product k across all consumers i of type r is:

$$E_k^r = \sum_i E_{ik}^r = \sum_i S_k^r E_i^r = S_k^r E^r,$$

---

[15] Although we assume a Fréchet (Type-II Extreme Value) distribution for idiosyncratic tastes, because the derivations are more direct, analogous results hold in a closely-related specification with a logit (Type-I Extreme Value) distribution.

which can be re-written as:

$$E_k^r = (\gamma^r)^{\theta^r} (P_k^r / \varphi_k^r)^{-\theta^r} (\mathbb{P}^r)^{\theta^r} E^r, \tag{37}$$

where $\mathbb{P}^r$ is again the unit expenditure function (36) for consumer of type r.

The key point to realize is that if we change notation and define $\theta^r = \sigma^r - 1$ and assume that there is only one type (r) of consumers, equations (33) and (36) become identical to the demand system and unit expenditure function that we derived in the CES case (up to a normalization or choice of units in which to measure φ_{kt}^r to absorb the constant $(\gamma^r)^{-1}$). Thus, the CES demand system and its "love-of-variety" property can be thought of as a means of aggregating "ideal-type" consumers who only consume one of each type of variety.

Proposition 5. *Given data on prices and on expenditure of consumers of each type r, the mixed random utility model defined by the indirect utility function (28) and Type-II Extreme Value distributed idiosyncratic tastes (31) with shape parameter θ^r is isomorphic to a constant elasticity of substitution (CES) model in which consumers of different types r have different demand parameters (φ_k^r) and elasticities of substitution (σ^r). This mixed random utility model implies a demand system (33) and unit expenditure function (36) for consumers of a given type r that are isomorphic (up to a normalization or choice of units for φ_k^r) to those in a mixed CES model with multiple consumer types, where $\theta^r = \sigma^r - 1$.*

Proof. The proposition follows immediately from the demand system (33) and unit expenditure function (36), substituting $\theta^r = \sigma^r - 1$. □

Given data on prices (P_k^r) and expenditures (E_k^r) for consumers of different observable types r (e.g., consumers from different regions or income quantiles), our CES-based methodology can be used to estimate the elasticity of substitution ($\sigma^r = 1 + \theta^r$) and preference parameter (φ_k^r) for each type. This model with multiple types of consumers generates much richer predictions for cross-price elasticities than the CES model with a single consumer type. Considering data on multiple markets with different shares of each consumer type, these cross-price elasticities will vary across markets depending on the shares of each consumer type in overall expenditure and the share of each product in total expenditure for each consumer type. Furthermore, the IIRA property will no longer necessarily hold across these different markets. The relative sales of two products across markets will depend not only on the characteristics of those products, but also on the shares of each consumer type in overall expenditure and the share of each product in total expenditure for each consumer type (which depends on the characteristics of other products). Finally, partitioning consumer types by income, the variation in the substitution and preference parameters across types allows for non-homotheticities in preferences across consumer types.

In this specification with multiple types of consumers, our unified price index now provides the exact price index for each type of consumers that allows for the entry and exit of goods over time, changes in demand for each good over time (where demand for each good and time period can now differ across consumer types) and imperfect substitutability between goods (where the degree of substitutability between goods can vary across consumer types).

Proposition 6. *The "unified price index" (UPI) for consumer type r—which is exact for the mixed random utility model defined by the indirect utility function (28) and Type-II Extreme Value distributed idiosyncratic tastes (31)—is given by*

$$\Phi^{Ur}_{t-1,t} = \underbrace{\left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}}\right)^{\frac{1}{\theta^r}}}_{\text{Variety Adjustment}} \underbrace{\left[\frac{\tilde{P}^*_t}{\tilde{P}^*_{t-1}}\left(\frac{\tilde{S}^*_t}{\tilde{S}^*_{t-1}}\right)^{\frac{1}{\theta^r}}\right]}_{\text{Common-Goods Price Index } \Phi^{CGr}_{t-t,t}}. \tag{38}$$

Proof. The proposition follows from combining the expenditure share (33) and unit expenditure function (36) for each consumer type r, following the same line of argument as for the CES specification with a representative consumer in Section 2.2. □

Our price index therefore has the same functional form but a slightly different interpretation in a random utility model. While it is not valid for any individual consumer, who has idiosyncratic tastes, our index tells us the average movement in the unit expenditure function for consumers of type r. Therefore introducing heterogeneous types of consumers enables us to relax the assumptions of symmetry, homotheticity and IIRA that are inherent in the representative consumer CES specification, while at the same time preserving our ability to compute an exact price index for each type of consumer, which incorporates changes in variety, changes in demand for each good and imperfect substitutability.

5 Estimation of the Elasticity of Substitution

Our unified price index does not contain any demand shift terms and therefore can always be used to deflate expenditures and create a money-metric utility function, but the precise value of the UPI depends on the elasticity parameter. In principle, there are a number of different ways of estimating the elasticity of substitution, including demand system estimation with instrumental variables (as in Berry 1994) and the approach based on double-differenced heteroskedastic demand and supply shocks (introduced by Feenstra 1994). While our unified price index is compatible with any of these approaches, we now show that minimizing deviations from a money-metric utility function itself provides a method of identifying the elasticity of substitution. We provide conditions under which this approach yields consistent estimates of the true elasticity of substitution. We show how this approach provides a metric for quantifying the magnitude of the departure from a money-metric utility function when using an alternative estimate of the elasticity of substitution, such as from an instrumental variables or double-differenced heteroskedastic demand and supply shocks approach.

5.1 The Reverse-Weighting Estimator

We begin by rewriting our forward and backward differences of the CES price index in terms of aggregate demand shifters that summarize the effect of changes in demand for each good on aggregate utility. Using these forward and backward differences ((17) and (18) respectively), the common good expenditure share (8),

and the unified price index (13), we obtain the following system of three equivalent expressions for the change in the cost of living from period $t-1$ to t:

$$\frac{\mathbb{P}_t^*}{\mathbb{P}_{t-1}^*} = \Theta_{t-1,t}^F \left[\sum_{k \in \Omega_{t,t-1}} S_{kt-1}^* \left(\frac{P_{kt}}{P_{kt-1}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \tag{39}$$

$$\frac{\mathbb{P}_t^*}{\mathbb{P}_{t-1}^*} = \left(\Theta_{t,t-1}^B \right)^{-1} \left[\sum_{k \in \Omega_{t,t-1}} S_{kt}^* \left(\frac{P_{kt}}{P_{kt-1}} \right)^{-(1-\sigma)} \right]^{-\frac{1}{1-\sigma}}, \tag{40}$$

$$\frac{\mathbb{P}_t^*}{\mathbb{P}_{t-1}^*} = \frac{\tilde{P}_t^*}{\tilde{P}_{t-1}^*} \left(\frac{\tilde{S}_t^*}{\tilde{S}_{t-1}^*} \right)^{\frac{1}{\sigma-1}}, \tag{41}$$

where the forward and backward aggregate demand shifters ($\Theta_{t-1,t}^F$ and $\Theta_{t,t-1}^B$) can be written as:

$$\Theta_{t-1,t}^F \equiv \left[\sum_{k \in \Omega_{t,t-1}} S_{kt}^* \left(\frac{\varphi_{kt-1}}{\varphi_{kt}} \right)^{\sigma-1} \right]^{\frac{1}{\sigma-1}} \quad \text{and} \quad \Theta_{t,t-1}^B \equiv \left[\sum_{k \in \Omega_{t,t-1}} S_{kt-1}^* \left(\frac{\varphi_{kt}}{\varphi_{kt-1}} \right)^{\sigma-1} \right]^{\frac{1}{\sigma-1}}, \tag{42}$$

as shown in Section A.8 of the web appendix.

The forward and backward aggregate demand shifters in (42) have an intuitive interpretation. Each aggregate demand shifter is an expenditure-share weighted average of the changes in demand for each good, where the expenditure-share weights are either the initial or the final-period expenditure shares. These aggregate demand shifters summarize the impact of demand shocks for each good on overall aggregate utility for the forward and backward differences of the price index. They represent the departures from a money-metric utility function that can potentially arise if consumers have different relative preferences for goods in periods $t-1$ and t.

While all three formulations of the price index are equivalent, only the UPI is consistent with a money-metric utility function because it alone contains no aggregate demand-shift terms. Each of these aggregate demand shifters are functions of the elasticity of substitution, which raises the question of whether there is an elasticity of substitution that renders all three of our formulas for the CES price index consistent with the same money-metric utility function. This property would require the existence of an elasticity parameter such that the following condition holds:

$$\Theta_{t-1,t}^F = \left(\Theta_{t,t-1}^B \right)^{-1} = 1. \tag{43}$$

When this condition is satisfied, the change in the cost of living evaluated using period $t-1$ preferences for both periods (which corresponds to the term in square brackets in (39)) is equal to the change in the cost of living evaluated using period t preferences for both periods (which corresponds to the term in square brackets in (40)), which in turn is equal to the change in the cost of living using the preferences for each period (equation (41)).

We now show how this condition can be used to construct a generalized method of moments (GMM) estimator of the elasticity of substitution (σ). Combining the three equivalent expressions (39)-(41), we obtain the following moment function for each pair of time periods $t-1$ and t:

$$m_t(\sigma) = \begin{pmatrix} m_t^1(\sigma) \\ m_t^2(\sigma) \end{pmatrix} = \begin{pmatrix} \frac{1}{1-\sigma}\ln\left[\sum_{k\in\Omega_{t,t-1}} S_{kt-1}^*\left(\frac{P_{kt}}{P_{kt-1}}\right)^{1-\sigma}\right] - \ln\left[\frac{\tilde{P}_t^*}{\tilde{P}_{t-1}^*}\left(\frac{\tilde{S}_t^*}{\tilde{S}_{t-1}^*}\right)^{\frac{1}{\sigma-1}}\right] \\ -\frac{1}{1-\sigma}\ln\left[\sum_{k\in\Omega_{t,t-1}} S_{kt}^*\left(\frac{P_{kt}}{P_{kt-1}}\right)^{-(1-\sigma)}\right] - \ln\left[\frac{\tilde{P}_t^*}{\tilde{P}_{t-1}^*}\left(\frac{\tilde{S}_t^*}{\tilde{S}_{t-1}^*}\right)^{\frac{1}{\sigma-1}}\right] \end{pmatrix} = \begin{pmatrix} -\ln\left(\Theta_{t-1,t}^F\right) \\ \ln\left(\Theta_{t,t-1}^B\right) \end{pmatrix}. \quad (44)$$

The first of these moment conditions quantifies the deviation of the forward difference from a money-metric utility function, while the second quantifies the deviation of the backward difference from a money-metric utility function. If the forward and backward differences are both consistent with the same money-metric utility function, the aggregate demand shifters will both be one, which requires the following sample moment condition to hold:

$$M(\sigma) = \frac{1}{T}\sum_{t=1}^T m_t(\sigma) = 0. \quad (45)$$

The GMM estimator, $\hat{\sigma}^{RW}$, solves:

$$\hat{\sigma}^{RW} = \arg\min\left\{M\left(\sigma^{RW}\right)' \times \mathbb{I} \times M\left(\sigma^{RW}\right)\right\}, \quad (46)$$

where we weight the two moments for the forward and backward difference equally by using the identity matrix ($\mathbb{I}$) for the weighting matrix.

We term the estimate of the elasticity of substitution that we obtain from this GMM procedure the "reverse-weighting" estimate ($\hat{\sigma}^{RW}$), because it involves equating expressions for the change in the cost of living using both initial and final expenditure share weights. Our use of the identity matrix as the weighting matrix ensures that this estimator minimizes the sum of squared deviations of the log of the forward and backward aggregate demand shifters ($\left(-\ln\left(\Theta_{t-1,t}^F\right)\right)^2 + \left(\ln\left(\Theta_{t,t-1}^B\right)\right)^2$) from zero. Therefore the reverse-weighting estimator minimizes the squared deviations from money-metric utility given our assumption of CES preferences and the observed data on prices and expenditure shares. As the GMM estimator is overidentified—we have one parameter (σ) to identify and two equations to satisfy—the sum of squared deviations of the aggregate demand shifters in general will be different from zero. However, we can obtain a sense of how well we do by quantifying the departures of the forward and backward differences from money-metric utility at the reverse-weighting estimate ($\hat{\sigma}^{RW}$) and compare this magnitude to the departures from money-metric utility for alternative values of the elasticity of substitution.[16] Having estimated the elasticity of substitution ($\hat{\sigma}^{RW}$), we can recover the demand parameter for each good k and period t ($\hat{\varphi}_{kt}^{RW}$) using the CES expenditure share as in (12).

We now provide conditions under which the forward and backward differences are consistent with a money-metric utility function ($\Theta_{t-1,t}^F = \left(\Theta_{t,t-1}^B\right)^{-1} = 1$), and the reverse-weighting estimator consistently estimates the true elasticity of substitution.

Proposition 7. *If the elasticity of substitution is not equal to one ($\sigma \neq 1$) and there is sufficient variation in expenditure shares that:*

$$\sum_{k\in\Omega_{t,t-1}} S_{kt-1}^*\ln\left(\frac{P_{kt}}{P_{kt-1}}\right) \neq \sum_{k\in\Omega_{t,t-1}} S_{kt}^*\ln\left(\frac{P_{kt}}{P_{kt-1}}\right) \neq \sum_{k\in\Omega_{t,t-1}} \frac{1}{N_{t,t-1}}\ln\left(\frac{P_{kt}}{P_{kt-1}}\right),$$

---

[16]When prices and expenditure shares are continuous and differentiable, the forward and backward differences of the CES price index are equivalent, and the elasticity of substitution is exactly identified, as shown in Section A.9 of the web appendix. This specification is the limiting case of the discrete changes considered in the main paper as changes in prices and expenditure shares and the interval between time periods become small, as shown in Section A.10 of the web appendix.

then as changes in demand become small $(((\varphi_{kt}/\varphi_{kt-1}) - 1) \to 0)$, the forward and backward differences are both consistent with a money-metric utility function $(\Theta^F_{t-1,t} \xrightarrow{p} 1/\Theta^B_{t,t-1} \xrightarrow{p} 1)$, and the reverse-weighting estimator consistently estimates the elasticity of substitution $(\hat{\sigma}^{RW} \xrightarrow{p} \sigma)$ and demand $(\hat{\varphi}^{RW}_{kt} \xrightarrow{p} \varphi_{kt})$ for each good k in each time period t.

Proof. See Section A.11 of the web appendix. □

Proposition 7 makes clear that the elasticity of substitution (σ) is identified from variation in expenditure-share-weighted average price changes. In the knife-edge case of Cobb-Douglas preferences $(\sigma = 1)$, expenditure shares are independent of prices, so that there is no variation to identify σ. In another knife-edge case in which all goods have equal expenditure shares in both time periods $(S^*_{kt-1} = S^*_{kt} = 1/N_{t,t-1})$, there is again no variation to identify σ. Outside these two knife-edge cases, the three CES expressions for the price index in equations (39)-(41) have different constant slopes with respect to σ. These three expressions therefore exhibit a single-crossing property that identifies the unique elasticity of substitution σ. Having identified this unique elasticity of substitution, demand for each good and time period (φ_{kt}) can be uniquely determined using the expenditure share (3), as in equation (12).

Proposition 7 holds irrespective of the size and correlation of price changes for each good k and period t. As the demand shocks for each good become small $(((\varphi_{kt}/\varphi_{kt-1}) - 1) \to 0)$, the forward and backward aggregate demand shifters in equation (42) converge to one $(\Theta^F_{t-1,t} \xrightarrow{p} 1$ and $\Theta^B_{t,t-1} \xrightarrow{p} 1)$, and the assumption that all three expression for the price index are consistent with the same money-metric utility function is satisfied $(\Theta^F_{t-1,t} = \left(\Theta^B_{t,t-1}\right)^{-1} = 1)$. In this case, the forward and backward differences of the price index reduce to the expenditure-share-weighted average of the price changes, and hence take a money metric form. More generally, we now show that the assumption that the forward and backward differences are consistent with money-metric utility is satisfied up to a first-order approximation.

Proposition 8. *To a first-order approximation, the forward and backward differences are are consistent with a money-metric utility function $(\Theta^F_{t-1,t} \approx 1/\Theta^B_{t,t-1} \approx 1)$.*

Proof. See Section A.12 of the web appendix. □

Proposition 8 holds for small changes in demand and prices, regardless of the covariance of these demand and price changes across goods, because these covariance terms are second-order small. The intuition is that the demand shocks for each good that enter the forward aggregate demand shifter $(\Theta^F_{t-1,t})$ are the inverse of those that enter the backward aggregate demand shifter $(\Theta^B_{t,t-1})$. Therefore, for small changes, an increase in the forward demand shifter necessarily implies a decrease in the backward demand shifter and *vice versa*. An implication of this result is that the reverse-weighting estimator can be interpreted as a model-consistent way of recovering the elasticity of substitution from the observed data on prices and expenditure shares that holds up to a first-order approximation. Existing exact and superlative price indexes are derived for small price changes under the assumption of no demand shifts. Proposition 8 maintains the assumption of small price changes while generalizing the analysis to allow for small demand changes. We can also numerically consider cases of large changes in prices and demand. In Section A.13, we use a Monte Carlo to show that the

reverse-weighting estimator provides a good approximation to the model's true parameters, even for large changes and a relatively small number of common goods. The intuition for why this works is that for larger changes in prices and demand, the reverse-weighting estimator provides a log-linear approximation to the data.

Proposition 8 also clarifies the relationship between our unified approach, the existing macro approach based on price indexes, and the existing micro approach based on demand system estimation. The existing macro literature on exact and superlative index numbers assumes that demand for each good is time invariant to ensure a money-metric utility function. But this assumption is strongly rejected by observed data on prices and expenditure shares for any plausible functional form for demand. The existing micro literature on demand systems estimation allows for a time-varying error term for each good to explain the observed data on prices and expenditure shares. But this time-varying error term in general violates the assumption of a money-metric utility function, which precludes comparisons of aggregate welfare over time. Our unified approach makes explicit the tension between explaining the observed data on prices and expenditure shares and preserving the property of a money-metric utility function. Our reverse-weighting estimator minimizes the departure from a money-metric utility function conditional on explaining the observed data on prices and expenditure shares. This estimator provides an approximation to the true underlying preference structure that becomes exact for small changes in prices and demand.

Propositions 7 and 8 are extremely useful in understanding a number of index number puzzles. For example, we showed in Section 3.2 that if one assumes that the elasticity of substitution is zero and that there are no demand shifts, one can derive both the Laspeyres and the Paasche indexes from the CES utility function. However, it is well-known that these price indexes generate different inflation rates when applied to real data. Similarly, the Fisher index differs from these even though it is also exact for the CES under the same assumptions. How can it be that all these price indexes can be derived from the same utility function with the same parameter restrictions and yet yield different measures of the change in the cost of living?

The answer is simple in the UPI framework. We know that there are three equivalent expressions of the CES price index given by the forward and backward differences and the UPI (equations (17), (18), and (13)). Moreover, we have shown that while the UPI is money-metric for *any* value of σ, the other two will only be equal to the same money-metric utility function for a *particular* elasticity. If a researcher computes the forward and backward differences using any other elasticity—e.g., the elasticity of zero implicit in the Laspeyres, Paasche, and Fisher—these indices yield different values for the change in the cost of living, and we have an apparent "paradox". Seen through the lens of our model, the reason CES-based indexes differ is that researchers are imposing the wrong elasticity on the data. Our approach shows that with the right elasticity, the generalized forms of the Laspeyres and Paasche (the forward and backward difference) can yield identical measures of price changes that are consistent with the same money-metric utility function. Thus, all index number formulas that can be derived from our generalized CES functional form would yield the same measure of the change in the cost of living if one used the reverse-weighting estimate of the elasticity instead of arbitrarily choosing an elasticity. In this sense, we have flipped index number theory on its head. The project is not to find a theoretical price index (based only on prices and expenditure shares) that is consistent

with any elasticity; the project is to find an elasticity that renders all theoretically-motivated price indexes identical and only based on prices and expenditure shares.

More generally, our reverse-weighting estimator provides a metric for computing departures from money-metric utility when using alternative estimates of the elasticity of substitution from elsewhere. These other estimates can be substituted into the definitions of the forward and backward aggregate demand shifters in equation (42) and used to compute the GMM objective (46) for the sum of squared departures of the log aggregate demand shifters from zero. We compare the value of these departures from money-metric utility $(\Theta_{t-1,t}^F = \left(\Theta_{t,t-1}^B\right)^{-1} = 1)$ for our reverse-weighting parameter estimate (σ^{RW}) and for alternative values of the elasticity of substitution (σ).

5.2 Robustness

We also report a number of robustness checks on our assumption of money-metric utility. First, our GMM estimator (46) is overidentified with two moment conditions (44) to identify the one elasticity of substitution (σ). We therefore also consider exactly identified specifications, in which we use only one of the two moment conditions (either only the forward moment condition $(m_t^1(\sigma))$ or only the backward moment condition $(m_t^2(\sigma))$. The reverse-weighting specification estimates σ by minimizing the sum of squared deviations of $m_t^1(\sigma) = -\ln\Theta_{t-1,t}^F$ and $m_t^2(\sigma) = \ln\Theta_{t,t-1}^B$, i.e., $\left(\ln\Theta_{t-1,t}^F\right)^2 + \left(\ln\Theta_{t,t-1}^B\right)^2$. In contrast, the forward specification estimates σ by setting $m_t^1(\sigma) = -\ln\Theta_{t-1,t}^F$ exactly equal to zero, while the backward specification estimates σ by setting $m_t^2(\sigma) = \ln\Theta_{t,t-1}^B$ exactly equal to zero. To the extent that the assumption of a money-metric utility function $(\Theta_{t-1,t}^F = \left(\Theta_{t,t-1}^B\right)^{-1} = 1)$ is not satisfied in the data, we would expect these three estimators to differ from one another. Therefore comparing the results from these three estimators is an important specification check on our identifying assumption.

Second, we assume the elasticity of substitution to be the same for each pair of years t and $t-1$. Therefore, a further specification check is to compare the results from the reverse-weighting estimator, pooling all years and estimating a single elasticity of substitution with the results from implementing this estimator separately for each pair of years and estimating a separate elasticity of substitution for each pair of years. Third, the size and correlation of demand and price shocks is likely to depend on the time interval over which prices and expenditure shares are differenced. A further specification check on the sensitivity of our estimates to the size and correlation of demand and price shocks is therefore to compare the results from implementing the reverse-weighting estimator for different time intervals $\triangle$ between years t and $t - \triangle$. As a final check on the sensitivity of our estimates of changes in the cost of living over time, we undertake a grid search over a range of possible values for the elasticity of substitution (σ), and compute the unified price index (13) and the implied departures from money-metric utility for each of these values for σ.

6 Results

In this section, we implement our unified price index empirically and compare the results to those using conventional price indexes. We first discuss the barcode data used in our empirical implementation. We next

estimate the elasticity of substitution for each of the product groups in our data using the reverse-weighting estimator. Finally, we compute the unified price index for each product group and in the aggregate and report the results of comparisons with existing exact and superlative price indexes (e.g., Fisher, Törnqvist) and with standard statistical price indexes (e.g., Laspeyres).

6.1 Data

We estimate the model using bar-code data from the Nielsen HomeScan database, which contains price and purchase quantity data for millions of bar codes bought between 2004 and 2014. A major advantage of bar-code data over other types of price and quantity data is that product quality does not vary within a bar code, because any change in observable product characteristics results in the introduction of a new barcode. Barcodes are inexpensive to purchase and manufacturers are discouraged from reusing them because reusing the same bar code for different goods or using several bar codes for the same product can create problems for store inventory systems that inform stores about how much of each product is available. Thus, bar codes are typically unique product identifiers and changes in physical attributes manifest themselves through the creation (and destruction) of bar-coded goods, not changes in the characteristics of existing bar-coded goods. This property means that shifts in demand for bar-coded goods cannot be driven by changes in the physical quality of the good, which makes these data ideal for identifying demand shift parameters, φ_{kt}.

The bar-code dataset we use is from Nielsen.[17] The data is based on a sample of approximately 50,000 households each year who scan in the price and quantity of every bar-coded good they buy each week. Nielson adjusts the data for sampling errors (response rates that are higher or lower for different demographic groups) and enables us to compute national total value and quantity purchased of each bar-coded good. The set of goods represents close to the universe of bar-coded goods available in grocery, mass-merchandise, and drug stores, representing around a third of all goods categories included in the CPI.

Nielsen organizes goods into product groups, which are based on where goods appear in stores. We dropped "variable weight" product groups which contain products whose quality may vary (e.g., fresh foods) and focus on the one hundred product groups constituting "packaged goods." The largest of these are carbonated beverages, pet food, paper products, bread and baked goods, and tobacco. Quantities do not vary for bar codes and are typically defined to be volume, weight, area, length, or counts (e.g., fluid ounces for Carbonated Beverages). We also adjust for multipacks, so we compare the price per battery, not the price per battery pack.

In choosing the time frequency with which to use the barcode data, we face a trade-off. On the one hand, as we work with higher frequency data, we are closer to observing actual prices paid for bar-codes as opposed to averages of prices. Thus high-frequency data has the advantage of allowing for a substantial amount of heterogeneity in price and consumption data. On the other hand, the downside is that the assumption that the total quantity purchased equals the total quantity consumed breaks down in very high-frequency data (e.g., daily or weekly) because households do not consume every item on the same day or even week they purchase

---

[17]Our results are calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business. Information on availability and access to the data is available at http://research.chicagobooth.edu/nielsen

it. Thus, the choice of data frequency requires a tradeoff between choosing a sufficiently high frequency that keeps us from averaging out most of the price variation, and a low enough frequency that enables us to be reasonably confident that purchase and consumption quantities are close.[18]

In order to deal with these issues, we worked with two different data frequencies: quarterly and annual that both produced very similar results. We collapse the household and time-dimensions in the data to construct a quarterly or annual samples of total value sold, total quantity sold, and average price. When using the quarterly data, four-quarter differences are computed by comparing values for the fourth quarter of each year relative to the fourth quarter of the previous year, and cumulative changes are computed by compounding these four-quarter differences.

6.2 Estimates of the Elasticity of Substitution

Figure 2 shows the distribution of our estimated elasticities of substitution for each product group at the four-quarter frequency. The mean and median elasticity of substitution is 4.4 and 4.5, respectively, which is substantially larger than the implicit elasticity of 0 in Laspeyres indexes and 1 in the Cobb-Douglas index. In other words, estimated rates of product substitution based on statistical indexes are likely to dramatically understate the degree of substitution by consumers. In terms of magnitudes, these elasticities do not differ that much from other studies. For example, Hottman, Redding, and Weinstein (2016), using the same data (but a different model nesting structure and estimation methodology) found that the elasticity of substitution had a median value of 3.9 across firms and 6.9 within firms. Our estimate, which pools within and across firms, falls in-between these two values. Figure 3 displays the estimated coefficients for each product group (blue solid line) as well as 95 percent point confidence intervals (dashed red lines).[19] As shown in the figure, the elasticities are precisely estimated, and all are significantly larger than one.

Our reverse-weighting estimator is overidentified, because it is based on both the forward and backward differences of the price index. A reasonable question to ask is how different would our estimates be if we used only the moment condition arising from the forward difference or the backward difference, where under our identifying assumption of money-metric utility ($\Theta_{t-1,t}^F = \left(\Theta_{t,t-1}^B \right)^{-1} = 1$) all three estimators should return the same elasticity of substitution. We denote the elasticity obtained from the reverse-weighting estimator in equation (44) by σ_g^{RW}; we indicate the estimated elasticity using only the forward difference (the first row of the moment vector) by σ_g^F; and we represent the estimated elasticity using only the backward difference (the second row of the moment vector) by σ_g^B. All three specifications yield extremely similar estimates of the elasticity of substitution. The standard deviation of $\sigma_g^F / \sigma_g^{RW}$ is 0.04, while that of $\sigma_g^B / \sigma_g^{RW}$ is 0.05. Therefore, we obtain similar estimates for the elasticity of substitution using all three estimators, consistent with our identifying assumption of money-metric utility.

We can examine this more directly by considering how close our objective function approaches zero,

---

[18]Even so, HomeScan data can sometimes contain coding errors. To mitigate this concern, we dropped purchases by households that reported paying more than three times or less than one third the median price for a good in a quarter or who reported buying twenty-five or more times the median quantity purchased by households buying at least one unit the good. We also winsorized the data by dropping observations whose percentage change in price or market share were in the top or bottom 1 percent.

[19]We compute the confidence intervals from 50 bootstrap replications. Each bootstrap replication for a given product group resamples the observed data on the prices and expenditure shares of goods k in periods t within that product group.
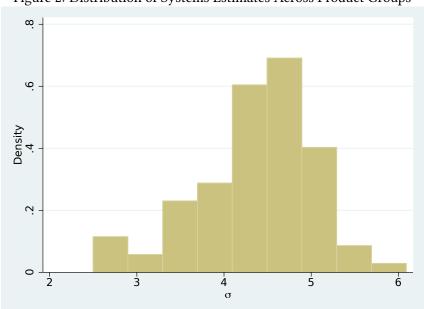
Figure 2: Distribution of Systems Estimates Across Product Groups

Note: Calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business.
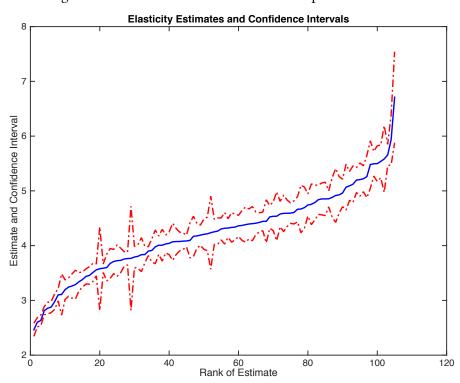
Figure 3: Estimated Elasticities and Bootstrap Standard Errors

Note: Calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business.
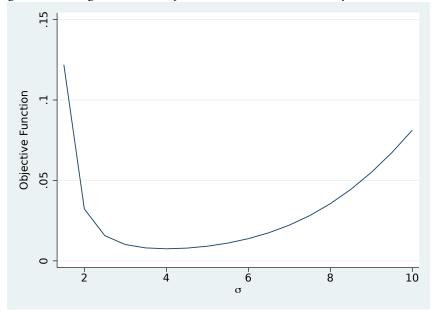
Figure 4: Average Value of Objective Function vs. Elasticity of Substitution

Note: Calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business.

which tells us how close we come to rendering the forward and backward differences consistent with money-metric utility. The identifying assumption of money-metric utility $(-\ln \Theta^F_{t-1,t} = \ln \Theta^B_{t,t-1} = 0)$ corresponds to the case in which our objective function $\left(\left(-\ln \Theta^F_{t-1,t}\right)^2 + \left(\ln \Theta^B_{t,t-1}\right)^2\right)$ is equal to zero. To examine how closely this assumption holds in the data and hence how close we are to a money-metric utility function, we construct an "average" demand shifter, $\overline{\Theta}_t$, that satisfies $\left(-\ln \Theta^F_{t-1,t}\right)^2 + \left(\ln \Theta^B_{t,t-1}\right)^2 = 2\left(\ln \overline{\Theta}_t\right)^2$. In other words, $\left|\ln \overline{\Theta}_t\right|$ is an "average" demand shifter in the sense that if the absolute value of both aggregate demand shifters equaled it, we would we would replicate the actual deviation from a money-metric utility function found in the data. When we do this for all product groups we find that the average value for $\left|\ln \overline{\Theta}_t\right|$ is close to zero (around 0.05), providing support for our identifying assumption. Using alternative values for the elasticity of substitution can generate substantial departures from money-metric utility. In Figure 4, we plot the value of the objective function across alternative values for the elasticity of substitution. Elasticities less than two or greater than eight tend to produce demand shifters that are sixty or more times larger than we obtain from the reverse-weighting estimator.

The fact that the aggregate demand shifters are much further from one when one uses an elasticity of substitution that deviates significantly from the reverse-weighting estimate indicates that standard indexes (e.g., Laspeyres, Paasche or Cobb-Douglas) imply substantial departures from a money-metric utility function. The fluctuations in utility are hidden in these approaches because the elasticity parameter associated with each index is often not made explicit and the violation in money-metric utility is only apparent when one realizes that different index numbers yield different measures of welfare (and therefore different money-metric utility). However, given that the UPI nests these approaches, we know that there must be violations of money-metric utility when researchers use an index whose implicit elasticity of substitution deviates substantially

from σ^{RW}. If one uses an elasticity that differs from the reverse-weighting one, the data will not be consistent with the assumption that a price index based on a forward or backward difference only depends on prices and expenditure shares, resulting in different values for cost-of-living changes depending on whether one uses initial expenditure shares, final shares, or both initial and final expenditure shares. These differences are exactly what $\Theta^F_{t-1,t}$ and $\Theta^B_{t,t-1}$ measure.

A second way of seeing the problem of existing approaches is to impose the assumption of no demand shocks on the data and directly back out what this implies about the utility function. We can do this easily in the CES case. As we showed in Proposition 2, we can solve for the elasticity of substitution using the Sato-Vartia formula in the case of no demand shocks (see equation (A.11) in the web appendix). If demand shocks are small, we would expect this utility parameter to be stable as well. In order to compute how demand shocks affect the implied elasticity of substitution, we denote the implied Sato-Vartia elasticity of substitution for each period by σ^{SV}_{gt} for every four-quarter difference and product-group. We expect these estimates to vary by product group, so we are interested in the dispersion of these estimates relative to the product group mean, or $\left(\sigma^{SV}_{gt} - \frac{1}{T} \sum_t \sigma^{SV}_{gt} \right)$, where T is the number of periods. In the absence of demand shocks, we should expect this number to be close to zero.

Table 1: Distribution of Elasticities for Each Year and Product Group

	Mean	Standard Deviation	10th Percentile	25th Percentile	50th Percentile	75th Percentile	90th Percentile
Sato-Vartia Elasticity	17.34	324.83	-53.92	-15.43	-1.27	12.69	35.69
Reverse-Weighting Elasticity	3.92	1.08	-1.41	-0.63	0.11	0.67	1.26

Note: The mean elasticity is $\frac{1}{GT} \sum_{t,g} \sigma_{gt}$, and the standard deviation is the average across all product groups, g, of the standard deviation of $\left(\sigma_{gt} - \frac{1}{T} \sum_t \sigma_{gt} \right)$, where G is the number of product groups and T is the number of time periods. Percentiles correspond to the distribution of $\left(\sigma_{gt} - \frac{1}{T} \sum_t \sigma_{gt} \right)$. Calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business. The mean is the average of all elasticities of substitution at the product-group level computed using equation (A.11).

Table 1 reports the mean of $\frac{1}{T} \sum_t \sigma^{SV}_{gt}$ in the first column and moments of the distribution of $\left(\sigma^{SV}_{gt} - \frac{1}{T} \sum_t \sigma^{SV}_{gt} \right)$ in the remaining columns. The mean value is 17.3 with a standard deviation of 324.8. Clearly, the implicit elasticities are quite volatile, and while there are some influential outliers, the volatility of the estimates permeates the distribution. Half of all observations are outside the range of -15.4 below the median implied elasticity in a product group to 12.7 above it. This enormous variation in the implied values of the elasticity of substitution, which spans all reasonable and many unreasonable values, means that the assumption of no demand shifts that underlies the Sato-Vartia formula is a deeply flawed way of thinking about consumer behavior. If one believes the underlying assumption of the exact price index—that demand for each good is constant over time—then one must also believe that the substitution parameter between goods in the utility function varies substantially over time. However, if the substitution parameter between goods varies so dramatically across pairs of periods, it is difficult to give any economic interpretation for what the price index is

measuring.

Having established that assuming no demand shifts results in absurd estimates of the elasticity of substitution, we now show that our method resolves this problem. Our estimates so far pooled pairs of time periods and estimated a single elasticity of substitution (by assuming $\sigma_g^{RW} = \sigma_{gt}^{RW}$). However, theoretically, it could be the case that the elasticity of substitution is also time varying. Thus, one might wonder whether the imposition of the assumption of a common elasticity inherent in the UPI also does violence to the data. In order to see if this is is the case, we estimated σ_{gt}^{RW} for every product group and year and report the the distribution of $\left(\sigma_{gt}^{RW} - \frac{1}{T} \sum_t \sigma_{gt}^{RW} \right)$ in Table 1. These estimates are much more tightly distributed around the product-group mean estimate than the time-invariant demand elasticities. The mean and median estimate has the reasonable value of 3.9, very close to the mean value of 4.4 for σ_g^{RW}, and almost all of the annual estimates deviate from the median value for the product group by less than one. In other words, the conventional approach of assuming no demand shocks not only cannot replicate the observed expenditure shares and prices as an equilibrium of the model but also implies wildly-varying elasticities of substitution. In contrast, our unified approach exactly rationalizes the observed data on expenditure shares and prices as an equilibrium of the model for a stable elasticity of substitution. Seen in this light, the data indicates that the unified approach is the only coherent means of reconciling demand data with welfare analysis.

6.3   Comparison with Conventional Index Numbers

We have already argued that our framework nests many existing methods of measuring price changes and welfare. This nesting makes it possible to step-by-step show how important each assumption is in measuring price changes. In each case, we construct price indexes for changes in the cost of living for every product group in our sample. With 10 time periods and 87 product groups, we have a sample of 870 price changes.

The Fisher and Törnqvist indexes are not strictly nested in our setup but are slightly different averages of the same building blocks.[20] The first question we need to address is how much it matters whether one uses a superlative index or a CES index. To the extent these differences are large, one might worry that adopting a CES utility function as opposed to a quadratic mean utility function (e.g., Fisher) or translog expenditure system (e.g., Törnqvist) is driving our results. While these different indexes need not be identical in theory, they are extremely similar in practice.

Figure 5 presents histograms of every four-quarter price change in our data at the product group level for each price index. We express each change in the cost of living as a difference from the superlative Fisher index, so a value of zero implies that the price index coincides with the Fisher index. The most noticeable feature of the graph is that all of the economic indexes yield almost exactly the same changes in the cost of living on average. The Törnqvist and Sato-Vartia CES typically record an average change in the cost of living that is identical to the Fisher index up to less than one decimal place. Moreover, there is very little dispersion in these price indexes. As one can see in Table 2, the standard deviation of the difference with the Fisher is only 0.1 percentage points per year. We also can replicate this same pattern in our Monte Carlo

---

[20]All of these indexes weight price relatives using past and current expenditure shares. For example, the Törnqvist weights the log price changes by an *arithmetic* average of past and current shares while the Sato-Vartia CES index weights them by a *logarithmic* average of the two shares.
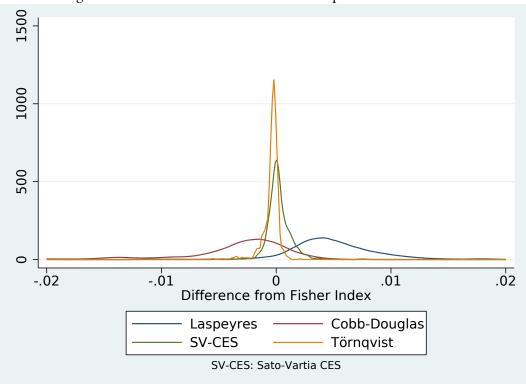
Figure 5: Each Index Differenced from the Superlative Fisher Index

SV-CES: Sato-Vartia CES

Note: Calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business.

exercise, which demonstrates that the result is not simply a feature of using bar-code data. Since the Sato-Vartia CES index is identical to the unified price index under the assumption that there are no new goods and no demand shifts for any good, we can safely say that our adoption of the CES functional form instead of a superlative index matters little for understanding changes in the cost of living. Whatever differences we find in subsequent sections must come from relaxing assumptions about the existence of demand shifts for each good or changes in the set of goods.

The fact that the CES functional form results in changes in the cost of living that are virtually identical to those of superlative indexes does not mean that *any* choice of price index yields similar results. As one can see in Figure 1, two commonly used indexes—the Cobb-Douglas and Laspeyres—are special cases of the CES in which the elasticity of substitution is one or zero, respectively. As one can see from Figure 5 and Table 2, imposing an elasticity of zero or one on the CES functional form instead of using the Sato-Vartia formula to allow the data to dictate the implied elasticity can result in very different measures of cost-of-living changes. For example, imposing an elasticity of zero (*i.e.,* Laspeyres) overstates changes in the cost of living relative to the CES, because it implicitly assumes consumer expenditure patterns do not change when prices change.

6.4   The Unified Price Index

The unified price index differs from the Sato-Vartia because it relaxes two assumptions. First, it allows for demand shifts for each good; and second, it allows for the set of goods to change over time. As we showed

Table 2: Comparisons of Alternative Measures of the Change in the Cost of Living

	Four Quarter	Annual	Cumulative, 2004-2014
Fisher Mean	1.8	1.6	19.6
Standard Deviation of Fisher	4.6	3.6	4.6
5th Percentile of Fisher	-3.1	-3.3	-3.1
50th Percentile of Fisher	1.0	0.6	1.0
95th Percentile of Fisher	9.4	6.6	9.4
Törnqvist Mean	1.8	1.6	19.2
Standard Deviation of Difference from Fisher	0.1	0.1	0.1
5th Percentile of Difference from Fisher	-0.2	-0.2	-0.2
50th Percentile of Difference from Fisher	-0.0	-0.0	-0.0
95th Percentile of Difference from Fisher	0.0	0.0	0.0
Sato-Vartia CES Mean	1.8	1.7	19.9
Standard Deviation of Difference from Fisher	0.1	0.1	0.1
5th Percentile of Difference from Fisher	-0.1	-0.0	-0.1
50th Percentile of Difference from Fisher	0.0	0.0	0.0
95th Percentile of Difference from Fisher	0.2	0.2	0.2
Cobb-Douglas Mean	1.5	1.3	15.8
Standard Deviation of Difference from Fisher	0.7	0.9	0.7
5th Percentile of Difference from Fisher	-1.6	-2.2	-1.6
50th Percentile of Difference from Fisher	-0.2	-0.2	-0.2
95th Percentile of Difference from Fisher	0.2	0.1	0.2
Laspeyres Mean	2.3	1.9	24.8
Standard Deviation of Difference from Fisher	0.4	0.3	0.4
5th Percentile of Difference from Fisher	-0.0	-0.1	-0.0
50th Percentile of Difference from Fisher	0.5	0.3	0.5
95th Percentile of Difference from Fisher	1.1	0.9	1.1
CG-UPI Mean	-1.0	-2.0	-9.6
Standard Deviation of Difference from Fisher	3.4	3.6	3.4
5th Percentile of Difference from Fisher	-8.9	-11.7	-8.9
50th Percentile of Difference from Fisher	-2.7	-3.7	-2.7
95th Percentile of Difference from Fisher	2.0	0.4	2.0
UPI Mean	-5.9	-3.7	-45.4
Standard Deviation of Difference from Fisher	6.8	5.9	6.8
5th Percentile of Difference from Fisher	-21.7	-16.5	-21.7
50th Percentile of Difference from Fisher	-6.8	-5.2	-6.8
95th Percentile of Difference from Fisher	-1.1	-0.8	-1.1

Note: The reported mean for each index is the initial period expenditure share weighted average of the index across product groups and over time. The standard deviation and percentiles for all price indexes except for the Fisher index are computed based on the difference between that index and the Fisher index. In the case of the Fisher index, the standard deviation and percentiles correspond to the actual variation of the index across product groups and over time. Calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business.

in Proposition 5, relaxing the first assumption gives rise to the common-goods component of the unified price index, which we know will lie below the Sato-Vartia index as long as demand shifts are positively correlated with expenditure shifts. The question we now address is how important is this bias. In addition to comparing the Sato-Vartia and Fisher indexes, Table 2 reports an analogous comparison for the common-goods component of the unified price index (the term in square brackets in equation (13)). While the average cost-of-living change for the Sato-Vartia index is 1.8 percent per year, the CG-UPI averages only -1.0 per year. This large difference indicates the importance of the consumer valuation bias in equation (15). Thus, assuming no demand shifts for any good not only results in an unstable elasticity parameter and a failure to be able to match the data, it also results in a substantial consumer valuation bias arising from the counterfactual assumption that consumers do not substitute towards goods that they like more.

Relaxing the second assumption regarding changes in the set of goods moves us from the CG-UPI to the UPI (see equation (13)). The variety-adjustment term, which was first estimated in Feenstra (1994), combines the elasticity of substitution, which tells us how much consumers value varieties, with the rates of product creation and destruction. Figure 6 presents a histogram of the $\lambda_{t,t-1}/\lambda_{t-1,t}$ ratios that drive the variety bias. Importantly, the fact that these ratios are less than one indicates not just product turnover, but substantial product upgrading. If bar codes were just turning over without upgrading, the prices and market shares of exiting bar codes would match those of new products resulting in a $\lambda_{t,t-1}/\lambda_{t-1,t}$ ratio of one. The fact that these ratios are less than one indicates that new goods tend to have lower price relative-to-demand ratios (P_{kt}/φ_{kt}) than disappearing ones. In other words, there is product upgrading. As one can see in Table 2, the relatively rapid rate of new good creation results in the mean and median unified-index price increase across all product groups and times being 4.9 percentage points lower than the common-goods component of the unified index.

We can see these differences at the aggregate level in Figure 7, which plots the expenditure-share-weighted average of the changes in the cost of living across product groups for each of the different index numbers over time, again using the initial period expenditure share weights. Not surprisingly, the Fisher, Törnqvist and Sato-Vartia result in almost identical changes in the cost of living that are bounded by the Paasche and Laspeyres indexes. This similarity is driven by the fact that they all assume no demand shifts for any good and therefore imply a time-varying utility function.

As one can see in Figure 1 and equation (13), there are two equivalent ways of moving from the Sato-Vartia index to the unified price index: one can first relax the assumption of no demand shifts (yielding the CG-UPI) and then make the variety adjustment, or one can first relax the assumption of a constant set of varieties (yielding the Feenstra-CES index) and then relax the assumption of no demand shifts. Thus the distance between the Sato-Vartia index and the CG-UPI tells us the importance of the consumer valuation bias and the distance between the Sato-Vartia and the Feenstra-CPI indicates the value of the variety-adjustment. Both biases suggest that standard indexes overstate cost-of-living changes and both biases are of roughly equal magnitude in many years. As discussed above, the relative magnitude of the variety and heterogeneity corrections in logs is independent of the value of the elasticity of substitution and depends only on the relative expenditure share moments ($\ln\left((\lambda_{t,t-1}/\lambda_{t-1,t})/\ln\left(\tilde{S}_t^*/\tilde{S}_{t-1}^*\right)\right)$). Therefore this finding that the variety and
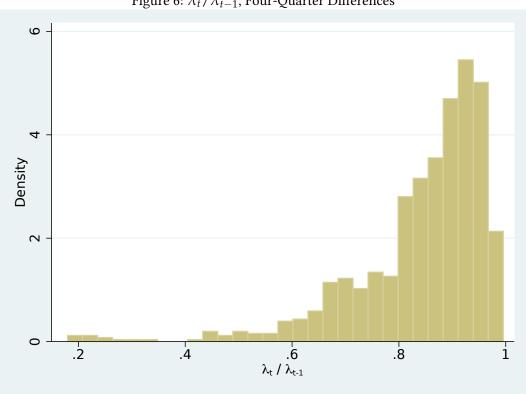
Figure 6: λ_t/λ_{t-1}, Four-Quarter Differences

Note: Calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business.

consumer valuation biases are of roughly equal magnitude to one another is robust to the assumption of alternative values for the elasticities of substitution from our reverse-weighting estimates.

A final question one might ask is whether it is safe to assume that the consumer valuation bias is constant, in which case it might be safe to use a standard index in a "difference-in-differences" approach. Interestingly, the data suggests that the consumer valuation bias does fluctuate. We can see this by computing the correlation between the the various indexes. While the correlation between the Feenstra-CES index and the Sato-Vartia index is 0.94, the correlation between the CG-UPI and the Sato-Vartia is only 0.73. Thus, while the variety bias is fairly stable the consumer valuation bias fluctuates more.[21] This fluctuation in the consumer valuation bias suggests that one should be cautious about interpreting the bias as a constant.

Taken together, these results show that allowing for demand shifts results in substantially different measures of price changes and welfare. Standard price indexes implicitly assume an elasticity parameter that, if used in data work, is equivalent to assuming substantial departures from money-metric utility. Moreover, by assuming that demand shifts do not cause expenditure shares to rise, standard price indexes tend to overstate changes in the cost of living.

---

[21]One can also come to a similar conclusion by computing the correlations between the Feenstra-CES, CG-UPI, and UPI. The correlation between the Feenstra-CES and the UPI is 0.78, and the correlation between the CG-UPI and the UPI is 0.86. Since the Feenstra-CES differs from the UPI only in the assumption about the existence of demand shifts, while the CG-UPI and UPI differ in the assumption about variety changes, the lower correlation between the Feenstra-CES and the UPI also indicates that consumer valuation bias fluctuates more than the variety adjustment.
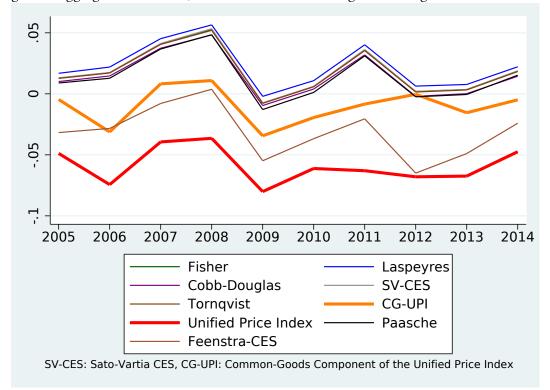
Figure 7: Aggregate Price Index, Calculated as a Share-Weighted Average Price Growth Rate

SV-CES: Sato-Vartia CES, CG-UPI: Common-Goods Component of the Unified Price Index

Note: Calculated based on data from The Nielsen Company (US), LLC and provided by the Marketing Data Center at The University of Chicago Booth School of Business.

7 Conclusions

Economics is broadly divided into macro and micro approaches that offer starkly different methods for evaluating welfare. Macro approaches are based on deflating nominal variables with price indexes that are derived using the assumption that there are no shifts in demand for any good. By contrast, the notion that demand curves shift is a central idea in microeconomic theory and demand-system estimation. This yields a deep inconsistency between the two approaches. If the assumptions underlying economically motivated price indexes are to be believed, there are no demand shifts for any good, key utility parameters can be backed out of the data, and there is no need for econometrics. Microeconomists reject this notion because the approach fails on micro data—expenditure shares are not perfectly explained by prices—there is a time-varying error term. Unfortunately, the existence of non-zero demand shifts undermines the assumptions of standard macro price indexes, leaving us in the uncomfortable state of either being able to consistently estimate key utility parameters but not knowing how to use those parameters to build a micro-founded aggregate price index or having to use price indexes based on assumptions that fail at a micro level and are inconsistent with a money-metric utility function.

We make two main contributions to this problem. First, we develop a unified price index that is consistent with time-varying demand shifts for each good at the micro level and yields a money-metric utility function at the macro level. This price index is time reversible and exact for the CES functional form even in the presence of the entry and exit of goods. We show how this index nests all existing major price indexes. Indeed, existing

price indexes can be thought of as arising from the imposition of parameter restrictions on the unified index. Thus, we bridge the divide between the micro and macro approaches.

Our index also enables us to identify a novel form of bias that arises from the assumption of time-invariant demand in existing price indexes. Consumer valuation bias arises whenever expenditure shares respond to demand shifts. Since conventional indexes assume that expenditure shares are only affected by price changes, they will be biased whenever expenditure share changes are correlated with demand shifts. We show, for example, if demand shifts cause expenditure shares and prices to rise, a conventional index will overstate cost-of-living changes because it will not adjust for the fact that some of the price increase is offset by the higher utility per unit associated with the demand shift.

Our second main contribution is to develop a novel way of estimating the elasticity of substitution. Extant approaches focus on identification from supply and demand systems. However, we show that one can also identify this parameter by combining information from the demand system and unit expenditure function. The intuition stems from the fact that in a totally differentiated CES demand system with k goods, one obtains $k - 1$ independent demand equations and k independent parameters: one for each of the $k - 1$ demand shifters and one for the elasticity of substitution. A key insight of our approach is that the unit expenditure function adds an additional equation to the system that can be exploited to produce an equal number of equations and unknowns, resulting in identification. With discrete changes the system is overidentified, but the basic intuition remains the same. One of the desirable properties of our "reverse-weighting" estimator is that it minimizes departures from money-metric utility for all of three of our equivalent representations of the CES utility function, making it particularly attractive for welfare analysis.

Finally, we use bar-code data to examine the properties of our unified price index and reverse-weighting estimator. We find that we obtain reasonable elasticity estimates in the sense that they are similar to those identified using other methodologies on the same data. Moreover, they are quite stable: year-by-year estimates of the elasticity do not differ much from the average over the full sample period, validating our assumption of a common utility function with time-varying demand parameters for each good. Lastly, the consumer valuation biases in existing indexes appear to be quite substantial, suggesting that allowing for demand shifts is an economically important force in understanding price and real income changes.

In conclusion, we provide a unified approach to demand and welfare estimation that reconciles micro and macro approaches, is easy to implement, and matters for understanding welfare.

References

[1] Anderson, James and Eric van Wincoop (2003) "Gravity with Gravitas: a Solution to the Border Puzzle," *American Economic Review*, 93(1), 170-192.

[2] Anderson, Simon, André de Palma and Jacques-François Thisse (1992) *Discrete Choice Theory of Product Differentiation*, MIT Press.

[3] Antràs Pol (2003) "Firms, Contracts, and Trade Structure," *Quarterly Journal of Economics*, 118(4), 1375-1418.

[4] Arkolakis, Costas and Arnaud Costinot and Andrès Rodriguez-Clare (2012) "New Trade Models, Same Old Gains," *American Economic Review*, 102(1), 94-130.

[5] Armington, Paul S. (1969) "A Theory of Demand for Products Distinguished by Place of Production," IMF Staff Papers, 16, 159-178.

[6] Atkin, David, and Dave Donaldson (2015) "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," *NBER Working Paper*, 21439.

[7] Atkin, David, Benjamin Faber, and Marco Gonzales-Navarro (2016) "Retail Globalization and Household Welfare: Evidence from Mexico," *Journal of Political Economy,* forthcoming.

[8] Balk, Bert M. (1999) "On Curing the CPI's Substitution and New Goods Bias," Department of Statistical Methods, Statistics Netherlands, mimeograph.

[9] Banerjee, K. S. (1983) "On the Existence of Infinitely Many Ideal Log-Change Index Numbers Associated with the CES Preference Ordering," *Statistische Hefte / Statistical Papers*, 24, 141-8.

[10] Behrens, Kristian, Yoshitsugu Kanemoto and Yasusada Murata (2014) "New Trade Models, Elusive Welfare Gains," *CEPR Discussion Paper*, 10255.

[11] Benassy, Jean-Pascal (1996) "Taste for Variety and Optimum Production Patterns in Monopolistic Competition," *Economics Letters*, 52, 41-47.

[12] Bernard, Andrew, Stephen Redding and Peter Schott (2007) "Comparative Advantage and Heterogeneous Firms," *Review of Economic Studies*, 74(1), 31-66.

[13] Bernard, Andrew, Stephen Redding and Peter Schott (2011) "Multi-product Firms and Trade Liberalization," *Quarterly Journal of Economics*, 126(3), 1271-1318.

[14] Berry, Steven (1994) "Estimating Discrete-Choice Models of Product Differentiation," *RAND Journal of Economics*, 25(2), 242-262.

[15] Berry, Steven, James Levinsohn and Ariel Pakes (1995) "Automobile Prices in Market Equilibrium," *Econometrica*, 63(4), 841-890.

[16] Bils, Mark and Klenow, Peter (2001) "Quantifying Quality Growth," *American Economic Review*, 91, 1006-1030.

[17] Blanchard, Olivier Jean and Nobuhiro Kiyotaki (1987) "Monopolistic Competition and the Effects of Aggregate Demand," *American Economic Review*, 77(4), 647-666.

[18] Bresnahan, Timothy F. and Gordon, Robert J. (1996) *The Economics of New Goods*, Chicago: University of Chicago Press.

[19] Broda, Christian and David E. Weinstein (2006) "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2), 541-86.

[20] Broda, Christian and David E. Weinstein (2010) "Product Creation and Destruction: Evidence and Price Implications," *American Economic Review*, 100, 691–723.

[21] Carli, G. (1764) "Del Valore e Della Proporzione de Metalli Monetati," reprinted in (ed) G. G. Destefanis, *Scrittori Classici Italiani di Economia Politica*, Milan, 297-336.

[22] Caves, Douglas W., Laurits R. Christensen and W. Erwin Diewert (1982) "The Economic Theory of Index Numbers and the Measurement of Input, Output and Productivity," *Econometrica*, 50(6), 1393-1414.

[23] Chance, W. A. (1966) "A Note on the Origins of Index Numbers," *Review of Economics and Statistics*, 48(1), 108-110.

[24] Chevalier, Judith A. and Anil Kasyap (2014) "Best Prices: Price Discrimination and Consumer Valuation," NBER Working Paper #20768.

[25] Cobb, Charles W. and Paul H. Douglas (1928) "A Theory of Production," *American Economic Review*, Papers and Proceedings, 18(1), 139-165.

[26] Deaton, Angus and John Muellbauer (1980) *Economics and Consumer Behavior*, Cambridge: Cambridge University Press.

[27] Deaton, Angus (1998) "Getting Prices Right: What Should be Done?" *Journal of Economic Perspectives*, 12(1), 37-46.

[28] Diewert, W. E. (1976) "Exact and Superlative Index Numbers," *Journal of Econometrics*, 4, 115-145.

[29] Dixit, Avinash and Joseph Stiglitz (1977) "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 297-308.

[30] Dutot, Nicolas (1738) *Réflexions Politiques sur les Finances et le Commerce*, Paris: E. Droz.

[31] Eaton, Jonathan and Kortum, Samuel (2002) "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741-1779.

[32] Fally, Thibault and Benjamin Faber (2016) "Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data," mimeo

[33] Feenstra, Robert C. (1994) "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84(1), 157-177.

[34] Feenstra, Robert C. and Marshall B. Reinsdorf (2007) "Should Exact Index Numbers Have Standard Errors? Theory and Application to Asian Growth," in (eds) Ernst R. Berndt and Charles R. Hulten, *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, Chicago: University of Chicago Press.

[35] Feenstra, Robert C. and Marshall B. Reinsdorf (2010) "An Exact Price Index for the Almost Ideal Demand System," *Economics Letters*, 66(2) 159-162.

[36] Fisher, Irving (1922) *The Making of Index Numbers*, Boston: Houghton Mifflin.

[37] Fisher, Franklin M. and Karl Shell (1972) *The Economic Theory of Price Indices: Two Essays on the Effects of Taste, Quality, and Technological Change*, New York: Academic Press.

[38] Hausman, Jerry A. (1996) "Valuation of New Goods under Perfect and Imperfect Competition," Chapter 5 in (eds) Timothy F. Bresnahan and Robert J. Gordon, *The Economics of New Goods*, Chicago: Chicago University Press, 209-247.

[39] Handbury, Jessie (2013) "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," Wharton mimeo.

[40] Helpman Elhanan, Marc J. Melitz and Stephen R. Yeaple (2004) "Export Versus FDI with Heterogeneous Firms," *American Economic Review*, 94(1), 300-316.

[41] Hottman, Colin, Stephen Redding and David Weinstein (2016) "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, forthcoming.

[42] Hsieh, Chang-Tai and Peter Klenow (2009) "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403-1448.

[43] Jevons, William Stanley (1863) *A Serious Fall in the Value of Gold*, Edward Stanford.

[44] Jones, Charles I. and Peter J. Klenow (2016) "Beyond GDP? Welfare Across Countries and Time," *American Economic Review*, forthcoming.

[45] Konus, A. (1924) "The Problem of the True Index of the Cost of Living," translated in *Econometrica* (1939), 7, 10-29.

[46] Konyus, A. A. and S. S. Byushgens (1926) "Probleme Pokupatelnoi Cili Deneg," *Voprosi Konyunkturi*, II(1), 151-172, [English title: Conus, A.A. and S.S. Buscheguennce, 1926, On the Problem of the Purchasing Power of Money, The Problems of Economic Conditions (supplement to the Economic Bulletin of the Conjuncture Institute) 2, 151-172].

[47] Krugman, Paul R. (1980) "Scale Economies, Product Differentiation and the Pattern of Trade," *American Economic Review*, 70, 950-959.

[48] Krugman, Paul R. (1991) "Increasing Returns and Economic Geography," *Journal of Political Economy*, 99(3), 483-499.

[49] Krugman, Paul and Anthony J. Venables (1995) "Globalization and the Inequality of Nations," *Quarterly Journal of Economics*, 857-880.

[50] Laspeyres, Etienne (1871) "Die Berechnung einer mittleren Warenpreissteigerung" in *Jahrbücher für Nationalökonomie und Statistik*.

[51] Lau, L. J. (1979) "On Exact Index Numbers," *Review of Economics and Statistics*, 61, 73-82.

[52] McFadden, Daniel (1974) "Conditional Logit Analysis of Qualitative Choice Behavior," in (ed) Paul Zaremba, *Frontiers in Econometrics*, New York: Academic Press.

[53] McFadden, Daniel and Kenneth Train (2000) "Mixed MNL Models of Discrete Response," *Journal of Applied Econometrics*, 15, 447-470.

[54] Melitz, Marc J. (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695-725.

[55] Moulton, B. R. (1996) "Constant Elasticity Cost-of-Living Index in Share- Relative Form," Washington DC: U. S. Bureau of Labor Statistics, mimeograph.

[56] Neary, J. Peter (2004) "Rationalizing the Penn World Table: True Multilateral Indices for International Comparisons of Real Income," *American Economic Review*, 94(5), 1411-1428.

[57] Nevo, Aviv (2003) "New Products, Quality Changes, and Welfare Measures," *Review of Economics and Statistics*, 85(2), 266–275.

[58] Paasche, Hermann (1875) *Die Geldentwertung zu Halle in den letzten Decennien dieses Jahrhunderts*, Plötz, Halle.

[59] Petrin, Amil (2002) "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110 (4), 705-729.

[60] Sato, Kazuo (1976) "The Ideal Log-Change Index Number," *Review of Economics and Statistics*, 58(2), 223-228.

[61] Shapiro, M. D. and D. W. Wilcox, (1997) "Alternative Strategies for Aggregating Prices in the CPI," *Federal Reserve Bank of St. Louis Review*, 79(3), 113-125.

[62] Sheu, Gloria (2014) "Price, Quality, and Variety: Measuring the Gains from Trade in Differentiated Products," *AEJ: Applied Economics*, 6(4), 66-89.

[63] Stoevska, Valentina (2008) *Consumer Price Index Manual: Theory and Practice*, Report on the 2007 Worldwide Survey on the Usefulness of the CPI Manual, Geneva: International Labor Organization (ILO).

[64] Syverson, Chad (2016) "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown," *NBER Working Papers*, 21974.

[65] Thisse, Jacques-Francois and Philip Ushchev (2016) "When Can a Demand System be Described by a Multinomial Logit with Income Effect?" mimeo.

[66] Törnqvist, Leo (1936) "The Bank of Finland Consumption Price Index," *Bank of Finland Monthly Bulletin*, 10, 1-8.

[67] Triplett, Jack E. (2001) "Should the Cost-of-living Index Provide the Conceptual Framework for a Consumer Price Index?" *Economic Journal*, 111(472), 311-334.

[68] Vartia, Yrjo (1976) "Ideal Log-Change Index Numbers," *Scandinavian Journal of Statistics*, 3(3), 121-126.

[69] Wicksell, K. (1898) *Interest and Prices*, English translation, New York 1936: Sentry Press.