

NBER WORKING PAPER SERIES

MEASURING AGGREGATE PRICE INDEXES WITH DEMAND SHOCKS:
THEORY AND EVIDENCE FOR CES PREFERENCES

Stephen J. Redding
David E. Weinstein

Working Paper 22479
<http://www.nber.org/papers/w22479>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2016, Revised May 21018

Previously circulated as “A Unified Approach to Estimating Demand and Welfare.” We are especially grateful to the editor, five anonymous referees, Dave Donaldson, Pablo Fajgelbaum, Rob Feenstra, Pete Klenow, and Nick Bloom for helpful comments. We are also grateful to many other colleagues and conference and seminar participants for their helpful comments. We also thank Mark Greenan, Mathis Maehlum, Anders Nielsen and Dyanne Vaught for excellent research assistance. All 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. Responsibility for the results, opinions, and any errors lies with the authors alone. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Stephen J. Redding and David E. Weinstein. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Measuring Aggregate Price Indexes with Demand Shocks: Theory and Evidence for CES Preferences

Stephen J. Redding and David E. Weinstein

NBER Working Paper No. 22479

August 2016, Revised May 21018

JEL No. D11,D12,E01,E31

ABSTRACT

We develop a new approach to measuring the cost of living for constant elasticity of substitution (CES) preferences. Our approach allows for demand shocks for individual goods (to rationalize micro data) while preserving a money-metric expenditure function (to compare the cost of living over time). We develop a new “reverse-weighting” estimator of the elasticity of substitution between goods and provide upper and lower bounds to the true parameter value. We show that abstracting from demand shocks introduces a “consumer-valuation bias,” which is analogous to the well-known “substitution bias,” and results in a substantial overestimate of the increase in the cost of living over time.

Stephen J. Redding
Department of Economics
and Woodrow Wilson School
Princeton University
14 College Road
Princeton, NJ 08540
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

Measuring price aggregates is central to international trade and macroeconomics, which depend on being able to distinguish real and nominal income.¹ One of the most influential preferences considered in these literatures is constant elasticity of substitution (CES) preferences. Existing exact CES price indexes assume constant demand (taste or preference) parameters for each good, whereas existing demand-system estimation typically requires time-varying demand parameters for each good to rationalize the observed data. In contrast, we develop a unified approach to the CES price index and demand system, which incorporates demand shocks for individual goods (to rationalize micro data), while preserving a money-metric utility function (to compare the cost of living over time). We develop a new “reverse-weighting” (RW) estimator of the elasticity of substitution and provide upper and lower bounds to the true parameter value. Although we focus on CES preferences because of their prominence in international trade and macroeconomics, we also consider a number of extensions and generalizations of our approach, including non-homothetic CES (indirectly additive), nested CES, mixed CES, logit and translog preferences. We show that abstracting from demand shocks for individual goods introduces a “consumer-valuation bias,” which is analogous to the well-known “substitution bias” from neglecting the response of expenditure shares to price movements, and results in a substantial overestimate of the increase in the cost of living over time.

Our starting point is the CES expenditure function, which determines the cost of obtaining a given level of utility as a function of income, prices, and a “demand parameter” for each good (consumer tastes). We show that we can always define these demand parameters such that they enter the expenditure function inversely with prices. As a result, the consumer’s cost of living depends on demand-adjusted prices, but only unadjusted prices are observed in the data. To overcome this challenge, we invert the CES demand system to substitute for the unobserved time-varying demand parameters in terms of observed prices and expenditure shares. We use this insight to derive a new exact CES price index for “common” (surviving goods) with time-varying demand shocks that generalizes the existing Sato (1976) and Vartia (1976) exact price index. We combine our new common goods price index with the variety correction term from Feenstra (1994) to obtain a new exact CES price index for the overall change in the cost of living that incorporates the entry and exit of goods.

Although we allow for demand shocks for individual goods, we provide conditions on the stochastic process for demand under which the mean of these log-demand shocks across common goods converges to zero as the number of these common goods becomes large. This property allows us to incorporate the demand shocks for individual goods (so that our model is consistent with the observed price and expenditure share data), while ensuring that these demand shocks average out across common goods (so that the change in the cost of living is money-metric in the sense that it depends only on prices and expenditure shares). Our approach is valid under the same set of assumptions as the existing Sato-Vartia price index (no demand shocks for each common good), but it also valid under a much weaker set of assumptions (demand shocks for individual common goods that average out across these goods). To reflect the fact that our approach treats the demand system and the price index in a unified way, we term our overall exact price index incorporating

¹Recent 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).

entry/exit the CES unified price index (CUPI).

We show that abstracting from demand shocks for individual goods introduces a substantial bias into conventional price indexes that we term the “consumer-valuation bias.” This bias is related to the well-known “substitution bias,” which arises when goods are substitutes in consumer preferences, but a researcher measures changes in the cost of living using a Laspeyres price index that implies no substitution across goods. This Laspeyres index overestimates the increase in the cost of living over time, because it does not take into account that when the price of a good rises, the consumer can substitute towards other goods. Our consumer-valuation bias is related, because the demand parameter for each good enters the expenditure function inversely to its price. Hence, if consumer preferences allow substitution towards goods for which demand has increased, but a researcher uses a price index that rules out such substitution, this price index again overestimates the increase in the cost of living. The researcher fails to take into account that an increase in the relative demand for a good is analogous to a reduction in its relative price and induces consumers to substitute towards that good. Empirically, we find this consumer-valuation bias to be substantial, equal to more than one percentage point per annum, and around the same magnitude as the bias that would arise from failing to account for the impact of the entry and exit of goods on the cost of living.

Our results for the unified price index and the consumer valuation bias hold regardless of how the elasticity of substitution between goods is estimated. But we also develop a new estimator for this parameter, which we term the “reverse-weighting estimator,” because it uses both initial-period and final-period expenditure share weights. The conventional approach to estimating this elasticity involves demand-system estimation using instruments for demand and supply shocks. However, developing valid instruments for demand and supply shocks can be challenging in the settings with large numbers of sectors considered in international trade and macroeconomics. In contrast, our estimator combines the demand system with the unit expenditure function and uses the identifying assumption of money-metric utility: the change in the cost of living depends solely on prices and expenditure shares. Using this assumption, our baseline “reverse-weighting” estimator minimizes difference between the implied change in the cost of living using the tastes of the initial or final period. We show that this estimator is consistent if price and demand shocks are small or if they are orthogonal for each good and independently distributed across goods. Our “generalized reverse-weighting” estimator minimizes the difference between the change in the cost of living using initial or final period tastes, after controlling for the component of demand shocks that is correlated with price shocks. We show that that this estimator is consistent if demand and price shocks are correlated for each good but are independently distributed across goods. Both estimators belong to the class of M-Estimators, and we show that they perform well in finite samples using Monte Carlos. Finally, we use our inversion of the demand system to provide upper and lower bounds to the elasticity of substitution that hold regardless of the correlation between demand and price shocks.

Our paper is related to several strands of existing research. First, we contribute to the “economic approach” to price measurement following Konüs (1924), in which price indexes are derived from consumer theory through the expenditure function. This long line of research includes Diewert (1976, 2004), Lau (1979), Feenstra (1994), Moulton (1996), Balk (1999), Caves, Christensen and Diewert (1982), Neary (2004), Feenstra

and Reinsdorf (2007, 2010), Białek (2017), and Diewert and Feenstra (2017). As discussed above, Sato (1976) and Vartia (1976) introduced an exact CES price index for common goods assuming time-invariant demand for each common good, while Feenstra (1994) generalized this price index to incorporate the entry and exit of goods over time. Our contribution relative to this research is to allow for time-varying demand shocks for each common good (to rationalize the micro data) while retaining a money-metric unit expenditure function (to compare the cost of living over time).

Our study is also related to the 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). We show that our methodology also holds for the closely-related logit model, and hence our work also connects with the large body of applied research using the logit model, as synthesized in Anderson, de Palma and Thisse (1992) and Train (2009). Increasingly, researchers in international trade and development are turning to barcode 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 all of these studies is to derive an exact price index that allows for both changes in demand for individual common goods and entry and exit, while preserving the property of a money-metric utility function.

More generally, our work relates to a conceptual debate in applied microeconometrics about whether to incorporate time-varying error terms from the demand system into the utility function, as examined in Nevo (2003). We distinguish two main interpretations of demand shocks that can be taken when estimating demand systems and evaluating welfare. First, our interpretation is that the observed choices of consumers given the observed prices reflect their true preferences. Under this interpretation, a demand shock that affects expenditure shares should also show up in price indexes and welfare, because it reflects the true preferences of consumers. Second, an alternative interpretation is that these observed choices do not reflect the true preferences of consumers, because of random shocks that lead realized choices to diverge from expected choices, or because of measurement error or specification error. Under this alternative interpretation, a demand shock that affects expenditure shares need not show up in price indexes and welfare, because it does not reflect the true or average preferences of consumers.

Although both approaches are valid, there are several advantages of our interpretation. First, our view is consistent with classical revealed preference arguments, in the sense that we use observed choices to infer true preferences. Second, our model exactly rationalizes the observed data without requiring the inclusion of an extraneous error term, which ensures that we can interpret the data in a self-contained way within the model. Third, in our view, the distinction between realized and expected choices is most natural for individual consumers. However, the observed choices in our data are aggregations across thousands of consumers, where one should expect the law of large numbers to apply, such that these aggregations are informative about

true preferences. Indeed, there is an isomorphism between the CES preferences of a representative agent and the aggregation of the idiosyncratic preferences of individual consumers with extreme value distributed preferences, as shown in Anderson, de Palma and Thisse (1992) and our web appendix.² Fourth, our use of barcode scanner data substantially reduces the scope for measurement error, and the log linear functional form of our common goods price index ensures that it is robust to mean-zero measurement error in log prices and expenditure shares. Finally, although specification error remains a potential concern, any model is necessarily an abstraction, and will require a time-varying error term to fit the data. Our approach provides a systematic treatment of demand system estimation and price index measurement in the presence of such a time-varying error term. While our baseline specification focuses on CES preferences, we show that our main insight generalizes to other functional forms below.

Finally, our work connects with 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.

The remainder of the paper is structured as follows. Section 2 introduces our baseline specification for CES preferences, including our new exact price index and reverse-weighting estimator. Section 3 develops a number of extensions and generalizations, including non-homothetic CES (indirectly additive), nested CES, mixed CES, logit and translog preferences. Section 4 introduces the detailed barcode data for the U.S. consumer goods sector used in our empirical analysis. Section 5 presents our main empirical results and demonstrates the quantitative relevance of allowing demand shocks for individual goods for measuring changes in the aggregate cost of living. Section 6 concludes. A web appendix collects together technical derivations, the proofs of propositions, additional information about the data, and supplementary empirical results.

2 Demand and Price Indexes with CES Preferences

In this section, we focus on our baseline specification of CES preferences with a single nest (e.g. an economy consisting a single sector composed of many varieties of goods). We first derive our new CES unified price index (CUPI), before next developing our new reverse-weighting (RW) estimator of the elasticity of substitution. In a later section, we extend the analysis to accommodate multiple CES nests and allow for more flexible functional forms (see Section 3 below).

²As the observed choices in our data are aggregations across thousands of individual consumers, our demand shocks correspond to shifts in a common component of tastes for all consumers, or shifts in a parameter determining the average realization of idiosyncratic tastes for a good in the extreme value distribution.

2.1 Preferences and Demand

Under the assumption of homothetic CES preferences, the unit expenditure function (P_t) depends on the price (p_{kt}) and demand parameter (φ_{kt}) for each good k at time t :

$$P_t = \left[\sum_{k \in \Omega_t} \left(\frac{p_{kt}}{\varphi_{kt}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad \sigma > 1, \quad (1)$$

where σ is the constant elasticity of substitution between goods and Ω_t is the set of goods supplied at time t .³ We can always define the demand parameters for each good (φ_{kt}) so that they enter the unit expenditure function (1) inversely to price, such that the consumer cares about demand-adjusted prices (p_{kt}/φ_{kt}). We assume that the log demand parameter for each good k in each time period t ($\ln \varphi_{kt}$) has a time-invariant component ($\ln \varphi_k$) and a time-varying component ($\ln \theta_{kt}$):

$$\ln \varphi_{kt} = \ln \varphi_k + \ln \theta_{kt}, \quad \ln \theta_{kt} \sim F(\mu_\theta, \chi_\theta^2), \quad (2)$$

The time-invariant term ($\ln \varphi_k$) captures differences in average levels of expenditure across goods (some goods are always more popular than others in all time periods). We assume that the idiosyncratic term ($\ln \theta_{kt}$) is independently and identically distributed across goods and over time with finite mean (μ_θ) and variance (χ_θ^2). We also assume that this idiosyncratic shock to consumer tastes ($\ln \theta_{kt}$) is realized after goods have been supplied to the market, which implies that there is no selection of the goods supplied based on this idiosyncratic shock. But we allow the time-invariant component of consumer tastes ($\ln \varphi_k$) to be observed before goods are supplied to the market, which allows for selection into the market based on this time-invariant component. The overall demand parameter ($\ln \varphi_{kt}$), including both the time-invariant and time-varying components, captures the true preferences of the consumer for each good in each time period.

Using our assumptions that the idiosyncratic shock ($\ln \theta_{kt}$) is independently and identically distributed across goods and realized after goods have been supplied to the market, the weak law of large numbers implies that the mean idiosyncratic shock converges in probability to its population mean of μ_θ as the number of goods becomes large. We make these assumptions on the log demand parameters ($\ln \varphi_{kt}$) so as to ensure that the level of the demand parameter (φ_{kt}) is positive for each good and time period. As CES preferences imply a marginal utility of consumption that is unbounded at zero consumption, all goods supplied to the market with positive values of φ_{kt} are consumed in equilibrium. Therefore, to ensure positive utility given the entry and exit of goods with positive φ_{kt} , our model requires $\sigma > 1$. Note that the conventional case in which the demand parameter for each good is constant over time ($\ln \varphi_{kt} = \ln \varphi_k$ for all k and all t) is simply a special case of our specification in which the distribution $F(\mu_\theta, \chi_\theta^2)$ is degenerate at μ_θ with $\chi_\theta^2 = 0$.

Applying Shephard's Lemma to the unit expenditure function (1), we obtain the demand system in which the expenditure share (s_{kt}) for each good is:

$$s_{kt} \equiv \frac{p_{kt} c_{kt}}{\sum_{\ell} p_{\ell t} c_{\ell t}} = \frac{(p_{kt}/\varphi_{kt})^{1-\sigma}}{\sum_{\ell \in \Omega_t} (p_{\ell t}/\varphi_{\ell t})^{1-\sigma}} = \frac{(p_{kt}/\varphi_{kt})^{1-\sigma}}{P_t^{1-\sigma}}, \quad k \in \Omega_t, \quad (3)$$

³We focus on CES preferences 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).

where c_{kt} denotes consumption of good k at time t .

Two well-known properties of these CES preferences are the independence of irrelevant alternatives (IIIRA) and the symmetry of substitution effects. The first of these properties implies that the relative expenditure share of any two goods depends solely on the relative price and demand parameter of those goods and not on the characteristics of any other goods: $s_{kt}/s_{\ell t} = [(p_{kt}/\varphi_{kt}) / (p_{\ell t}/\varphi_{\ell t})]^{1-\sigma}$. The second of these properties implies that the elasticity of expenditure on any one good ($x_{kt} = p_{kt}c_{kt}$) with respect to a change in the price of another good depends solely on the expenditure share of that other good: $(\partial x_{kt}/\partial p_{\ell t})(p_{\ell t}/x_{kt}) = (\sigma - 1)s_{\ell t}$. We relax these two assumptions in Section 3, where we consider both mixed CES preferences with heterogeneous consumers and translog preferences.

Taking logarithms of the CES expenditure share (3), and using equation (2), we obtain the following demand system for log expenditure shares as a function of log prices:

$$\ln s_{kt} = (\sigma - 1) \ln P_t + (1 - \sigma) \ln p_{kt} + (\sigma - 1) \ln \varphi_{kt}, \quad \varphi_{kt} = \ln \varphi_k + \ln \theta_{kt}. \quad (4)$$

Two points are worthy of mention at this point. First, in general, the demand parameter ($\ln \varphi_{kt}$) can be correlated with prices ($\ln p_{kt}$). Second, the idiosyncratic component of demand ($\ln \theta_{kt}$) is a structural residual that ensures that the model exactly fits the observed log expenditure shares ($\ln s_{kt}$) for each good k and time period t given the observed log prices ($\ln p_{kt}$). Additionally, since the demand system (4) is derived from the unit expenditure function (1), any time-varying demand parameter ($\ln \theta_{kt}$) in the demand system also in general appears in the unit expenditure function. As discussed above, there are a number of potential alternative interpretations to consumer tastes for this time-varying error term in the demand system (4), including changes in product quality, measurement error and specification error. Under some of these interpretations, the time-varying error term in the demand system need not necessarily appear in the unit expenditure function. However, our use of barcode data in our empirical application rules out changes in product quality, because firms have strong incentives of inventory and stock control not to use the same barcode for products with different observable characteristics. Therefore, any change in product characteristics leads to the introduction of a new barcode, and is reflected in the entry and exit of barcodes instead of changes in quality within surviving barcodes. Similarly, our use of barcode data alleviates concerns about measurement error. Although specification error remains a possibility, any model is necessarily an abstraction and will require a time-varying error term to fit the data. We show below that our main insight generalizes to other preference structures, including flexible functional forms such as translog.

2.2 Entry and Exit

An important advantage of CES preferences is that they yield a tractable variety correction term for the pervasive entry and exit observed in micro data, as shown in Feenstra (1994). To implement this variety correction, 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 that enter 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 that exit between $t - 1$ and t (I_{t-1}^-), 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 $N_{t,t-1} = |\Omega_{t,t-1}|$.

Using this notation, the change in the cost of living between periods $t-1$ and t ($\Phi_{t-1,t}$) can be expressed in terms of the change in the share of expenditure on common goods ($\lambda_{t,t-1}/\lambda_{t-1,t}$) and the change in the cost of living for these common goods (P_t^*/P_{t-1}^*):

$$\Phi_{t-1,t} = \frac{P_t}{P_{t-1}} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)^{\frac{1}{\sigma-1}} \frac{P_t^*}{P_{t-1}^*}, \quad (5)$$

where the derivations are reported in Section A.2 of the web appendix. The terms $\lambda_{t,t-1}$ and $\lambda_{t-1,t}$ capture expenditure on common goods as a share of total expenditure in periods t and $t-1$ respectively. We use an asterisk to denote the value of a variable for the common set of goods, such that P_t^* and P_{t-1}^* are the unit expenditure functions for *common* goods:

$$P_t^* \equiv \left[\sum_{k \in \Omega_{t,t-1}} \left(\frac{p_{kt}}{\varphi_{kt}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \quad (6)$$

The ratio of the common goods price index in the two periods (P_t^*/P_{t-1}^*) in equation (5) equals the change in the cost of living if the set of goods is not changing. The term multiplying this ratio in equation (5) is the “variety-adjustment” term ($(\lambda_{t,t-1}/\lambda_{t-1,t})^{1/(\sigma-1)}$), which captures the impact on the cost of living of the entry and exit of goods. If new goods are more numerous than exiting goods or have lower demand-adjusted prices (i.e., lower p_{kt}/φ_{kt}), then $\lambda_{t,t-1}/\lambda_{t-1,t} < 1$, and the cost of living will fall due to an increase in variety or the entering varieties being more appealing given their cost than the exiting varieties.

We can also define the share of individual common good $k \in \Omega_{t,t-1}$ in common goods expenditure (s_{kt}^*):

$$s_{kt}^* \equiv \frac{p_{kt}c_{kt}}{\sum_{\ell \in \Omega_{t,t-1}} p_{\ell t}c_{\ell t}} = \frac{(p_{kt}/\varphi_{kt})^{1-\sigma}}{\sum_{\ell \in \Omega_{t,t-1}} (p_{\ell t}/\varphi_{\ell t})^{1-\sigma}} = \frac{(p_{kt}/\varphi_{kt})^{1-\sigma}}{(P_t^*)^{1-\sigma}}, \quad k \in \Omega_{t,t-1}, \quad (7)$$

which takes the same form as the share of each good in total expenditure in equation (3), except that the summation in the denominator is only over common goods.

2.3 Exact CES Price index for Common Goods

We now derive an exact price index for the change in the cost of living for common goods ($\ln \Phi_{t-1,t}^* = \ln (P_t^*/P_{t-1}^*)$) that allows for demand shocks for individual goods and can be expressed in a money-metric form in terms of observed prices (p_{kt} , p_{kt-1}), expenditure shares (s_{kt}^* , s_{kt-1}^*), and the elasticity of substitution between goods (σ).

Using the common goods unit expenditure function (6) and the expenditure share (7), we obtain the following exact price index for the change in the cost of living for common goods:

$$\ln \Phi_{t-1,t}^* = \sum_{k \in \Omega_{t,t-1}} \omega_{kt}^* \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \sum_{k \in \Omega_{t,t-1}} \omega_{kt}^* \ln \left(\frac{\varphi_{kt}}{\varphi_{kt-1}} \right), \quad (8)$$

where the weights ω_{kt}^* are the logarithmic mean of common goods expenditure shares (s_{kt}^*) in periods t and

$t - 1$ and sum to one,

$$\omega_{kt}^* \equiv \frac{\frac{s_{kt}^* - s_{kt-1}^*}{\ln s_{kt}^* - \ln s_{kt-1}^*}}{\sum_{\ell \in \Omega_{t,t-1}} \frac{s_{\ell t}^* - s_{\ell t-1}^*}{\ln s_{\ell t}^* - \ln s_{\ell t-1}^*}}, \quad (9)$$

and the derivation is reported in Section A.3 of the web appendix.

This exact price index for common goods in equation (8) is a generalization of the Sato-Vartia price index (Sato 1976 and Vartia 1976), which corresponds to the special case in which demand is assumed to be time invariant for each individual common good ($\varphi_{kt}/\varphi_{kt-1} = 1$ for all $k \in \Omega_{t,t-1}$):

$$\ln \Phi_{t-1,t}^{*SV} = \sum_{k \in \Omega_{t,t-1}} \omega_{kt}^* \ln \left(\frac{p_{kt}}{p_{kt-1}} \right). \quad (10)$$

The challenge in implementing the exact price index (8) empirically is that it depends on the log change in demand-adjusted prices ($\ln((p_{kt}/\varphi_{kt})/(p_{kt-1}/\varphi_{kt-1}))$), whereas only unadjusted prices are observed in the data ($\ln(p_{kt}/p_{kt-1})$). We overcome this challenge by inverting the CES demand system to express the unobserved time-varying demand parameter (φ_{kt}) in terms of observed prices (p_{kt}) and common goods expenditure shares (s_{kt}^*). This demand system inversion uses the fact that the CES demand system satisfies the conditions for “connected substitutes” in Berry, Gandhi and Haile (2013). These conditions rule out the possibility that some goods are substitutes while others are complements. Taking logarithms in the common goods expenditure share (7), differencing over time, and then differencing from the mean across common goods within each time period, we obtain the following closed-form expression for the log change in the demand shifter for each common good:

$$\ln \left(\frac{\varphi_{kt}/\tilde{\varphi}_t}{\varphi_{kt-1}/\tilde{\varphi}_{t-1}} \right) = \ln \left(\frac{p_{kt}/\tilde{p}_t}{p_{kt-1}/\tilde{p}_{t-1}} \right) + \frac{1}{\sigma - 1} \ln \left(\frac{s_{kt}^*/\tilde{s}_t^*}{s_{kt-1}^*/\tilde{s}_{t-1}^*} \right), \quad (11)$$

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

We now use our assumptions on the stochastic process for demand shocks for each good in equation (2). In particular, applying the weak law of large numbers across common goods, as the number of common goods becomes large ($N_{t,t-1} \rightarrow \infty$), the mean of the time-varying component of demand ($\ln \theta_{kt}$) converges towards its population mean of zero, which in turn implies that the mean demand shock converges to zero:

$$\lim_{N_{t,t-1} \rightarrow \infty} \ln \left(\frac{\tilde{\varphi}_t}{\tilde{\varphi}_{t-1}} \right) = \lim_{N_{t,t-1} \rightarrow \infty} \ln \left(\frac{\tilde{\theta}_t}{\tilde{\theta}_{t-1}} \right) = \lim_{N_{t,t-1} \rightarrow \infty} \frac{1}{N_{t,t-1}} \sum_{k=1}^{N_{t,t-1}} \ln \left(\frac{\theta_{kt}}{\theta_{kt-1}} \right) = 0. \quad (12)$$

Using this result, and substituting our closed-form solution for the demand shocks (11) into equation (8), we obtain our exact CES common-goods price index (CCG):

$$\ln \Phi_{t-1,t}^{*CCG} = \frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) + \frac{1}{\sigma - 1} \frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \ln \left(\frac{s_{kt}^*}{s_{kt-1}^*} \right). \quad (13)$$

The CCG summarizes the effect of changes in demand-adjusted prices on a consumer’s cost of living. It is comprised of two terms. The first term is none other than the average of log price changes that serves

as the basis for lower level of the U.S. Consumer Price Index (the log of the “Jevons” index). Indeed, in the special case in which varieties are perfect substitutes ($\sigma \rightarrow \infty$), the CCG collapses to this Jevons index, since the second term in equation (13) converges to zero as $\sigma \rightarrow \infty$. This second term is novel and captures heterogeneity in expenditure shares across common goods. This term moves with the average of the log expenditure shares in the two periods. Critically, as the market shares of common goods in a time period become more uneven, the mean of the log expenditure shares will fall (since the log function is concave). Therefore, this term implies that the cost of living will fall if expenditure shares become more dispersed. Intuitively, when varieties are substitutes ($\sigma > 1$), consumers value dispersion in demand-adjusted prices across varieties, because they can substitute consumption towards the varieties with lower demand-adjusted prices. If demand-adjusted prices (p_{kt}/φ_{kt}) are constant for all common goods, the change in the cost of living is necessarily zero. The reason is that the average log change in prices equals the average log change in demand-adjusted prices ($\ln(\tilde{p}_t/\tilde{p}_{t-1}) = \ln((\tilde{p}_t/\tilde{\varphi}_t)/(\tilde{p}_{t-1}/\tilde{\varphi}_{t-1})) = 0$ since $\ln(\tilde{\varphi}_t/\tilde{\varphi}_{t-1}) = 0$) and expenditure shares cannot change if demand-adjusted prices are constant ($\ln(\tilde{s}_t/\tilde{s}_{t-1}) = 0$). However, even if unadjusted prices (p_{kt}) are constant for all common goods, the cost of living can change, because consumer welfare depends on demand-adjusted rather than unadjusted prices.

We use the CCG instead of the Sato-Vartia price index for the log change in the cost of living for common goods ($\ln \Phi_{t-1,t}^* = \ln(P_t^*/P_{t-1}^*)$). Substituting the CCG (13) into equation (5), we obtain our overall CES unified price index (CUPI).

Proposition 1. *The “CES unified price index” (CUPI), which is exact for CES preferences in the presence of changes in the set of goods, demand shocks for individual goods, and discrete changes in prices and expenditure shares for each good, is given by:*

$$\ln \Phi_{t-1,t}^{CUPI} = \underbrace{\frac{1}{\sigma-1} \ln \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)}_{\text{Variety Adjustment}} + \underbrace{\frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) + \frac{1}{\sigma-1} \frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \ln \left(\frac{s_{kt}^*}{s_{kt-1}^*} \right)}_{\text{CES Common-Goods Price Index (CCG)}}. \quad (14)$$

Proof. The proposition follows directly from substituting the CCG (13) into equation (5). $\square$

In deriving the CUPI using equation (11), we use the log linearity of the common-goods expenditure share (7), which reflects the independence of irrelevant alternatives (IIA) property of CES preferences discussed above. Since this IIA property is specific to CES and the closely-related logit preferences discussed below, our unified price index is only exact for CES or logit preferences. However, in Section 3 below, we show that our main insight of using the demand system to substitute for the unobserved demand parameter generalizes to other invertible demand systems. Empirically, we find in Section 5 that the differences between our exact CES price index and those for other flexible functional forms are small relative to the differences that arise from the failure to control for demand shocks for individual goods.⁴

From equations (11)-(14), the CUPI uses unweighted means across common goods, but the log linearity of the common-goods expenditure share (7) implies that one could instead construct an alternative price index using weighted means. However, our argument based on the weak law of large numbers and our reverse-weighting estimator below require that these weights are orthogonal to demand shocks, which restricts the

⁴Our unified price index (14) 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 and captured the dispersion of sales across common goods using a different term.

class of admissible weights. As a robustness test, in Section 5.5 below, we compute price indexes based on random weights that sum to one and are orthogonal to demand shocks by construction. Given our large number of products within each sector, we find a high correlation between the CUPPI based on an unweighted mean across common goods and these other price indexes based on random weights.

A potential alternative approach to deriving a money-metric price index is to assume that the demand shock for a single benchmark common good is equal to zero. This alternative is closely related to the assumption of an outside good in empirical research. The problem is finding such a benchmark good for which one can be certain that relative demand did not change over time. Furthermore, it is not obvious how one could test this assumption for the benchmark good without prior knowledge of the parameters of the demand system. In contrast, our assumption that the demand shocks are mean zero across common goods emerges naturally from primitive assumptions on the stochastic properties of demand shocks.

More broadly, the CUPPI has other attractive economic and statistical properties. First, it is exact under the same assumptions as the Sato-Vartia price index (no demand shocks for each common good), but it is also valid under a much weaker set of assumptions (demand shocks for individual common goods that average out across common goods). Second, this price index is “time reversible” for any value of σ , thereby permitting consistent comparisons of the cost of living going forwards and backwards in time. In other words, given any set of product turnover, price changes, and demand shifts between $t - 1$ and t , the percent change in prices between $t - 1$ and t is the inverse of the change between t and $t - 1$. Third, the CUPPI depends in a simple and transparent way on the elasticity of substitution. Variation in this elasticity leaves the terms in common goods prices ($\frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \ln \left(\frac{p_{kt}}{p_{kt-1}} \right)$) unchanged and affects the variety adjustment ($\frac{1}{\sigma-1} \ln \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)$) and heterogeneity terms ($\frac{1}{\sigma-1} \frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \ln \left(\frac{s_{kt}^*}{s_{kt-1}^*} \right)$) depending on the extent to which these two expenditure share ratios are greater than or less than one. Indeed, the relative size 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.

Finally, as our exact common-goods price index (CCG) depends on the mean of the log prices and expenditure shares of common goods, it is invariant to mean-zero, log-additive measurement error in these prices and expenditure shares. In contrast, the Sato-Vartia price index in equation (10) involves a non-linear transformation of the common goods expenditure shares (s_{kt}^* and s_{kt-1}^* through ω_{kt}^*) and hence is directly affected by such mean-zero measurement error. One remaining concern about such measurement error is that the CCG includes an unweighted mean of the log expenditure shares across common goods. Therefore, it could be affected by measurement error for goods with small expenditure shares. To address this concern, we use the property of CES preferences that the price index for all common goods can be rewritten as equal to the price index for a subset of common goods times the share of expenditure on this subset in all expenditure on common goods. We implement this robustness check in Section 5.5 below using the subset of common goods with above-median expenditure shares. We show that we find similar values for the change in the overall cost of living as in our baseline specification above.

2.4 Relation to Existing Price Indexes

We now compare our CUPÍ to existing price indexes and examine the implications of allowing individual goods to experience demand shocks for the measurement of the cost of living. The existing CES exact price index combines the Feenstra (1994) variety correction with the Sato-Vartia price index:

$$\underbrace{\ln \Phi_{t-1,t}^{FE}}_{\text{Feenstra Price Index}} = \underbrace{\frac{1}{\sigma-1} \ln \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)}_{\text{Variety Adjustment}} + \underbrace{\ln \Phi_{t-1,t}^{*SV}}_{\text{Sato-Vartia Price Index}}. \quad (15)$$

Comparing equations (14) and (15), the CUPÍ uses the same variety correlation term as Feenstra (1994), but uses the CCG from equation (13) instead of the Sato-Vartia price index. 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 ($\ln(\Phi_{t-1,t}^{*SV})$) assumes that demand is constant over time for each common good ($\ln(\varphi_{kt}) = \ln(\varphi_{kt-1}) = \ln(\varphi_k)$ for all $k \in \Omega_{t,t-1}$ and t), whereas the estimation of σ assumes the existence of demand shocks ($\ln(\varphi_{kt}/\varphi_{kt-1}) = \ln(\theta_{kt}/\theta_{kt-1}) \neq 0$ 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 in the derivation of the Sato-Vartia index.

In the presence of non-zero demand shocks for some common good $k \in \Omega_{t,t-1}$, we show in Section A.3 of the web appendix that the true exact CES common goods price index ($\Phi_{t-1,t}^{*CCG}$) equals the Sato-Vartia price index ($\Phi_{t-1,t}^{*SV}$) minus an additional term that we refer to as the consumer-valuation bias:

$$\ln \Phi_{t-1,t}^{*CCG} = \ln \Phi_{t-1,t}^{*SV} - \underbrace{\left[\sum_{k \in \Omega_{t,t-1}} \omega_{kt}^* \ln \left(\frac{\theta_{kt}}{\theta_{kt-1}} \right) \right]}_{\text{consumer valuation bias}}, \quad (16)$$

where from now onwards we use the fact that the change in the overall demand parameter for each good ($\varphi_{kt}/\varphi_{kt-1}$) is determined by the change in the time-varying component of this parameter ($\theta_{kt}/\theta_{kt-1}$).

This consumer-valuation bias affects existing exact CES price indexes for both common goods and the overall cost of living, because both use the Sato-Vartia price index. This consumer-valuation bias arises because the Sato-Vartia index in equation (10) is based on the expenditure-share weighted average of observed price changes, whereas the true exact CES price index for common goods in equation (8) depends on the expenditure-share weighted average of *demand-adjusted* price changes. Therefore, the Sato-Vartia index is only unbiased if the demand shocks ($\ln(\theta_{kt}/\theta_{kt-1})$) are orthogonal to the expenditure-share weights (ω_{kt}^*); it is upward-biased if they are positively correlated with these weights; and it is downward-biased if they are negatively correlated with these weights. In principle, either a positive or negative correlation between the demand shocks ($\ln(\theta_{kt}/\theta_{kt-1})$) and the expenditure-share weights (ω_{kt}^*) is possible, depending on the underlying correlation between demand and price shocks. However, there is a mechanical force for a positive correlation, because the expenditure-share weights themselves are functions of the demand shocks. In particular, a positive demand shock for a good mechanically increases the expenditure-share weight for that

good and reduces the expenditure-share weight for all other goods:

$$\frac{d\omega_{kt}^*}{d\theta_{kt}} \frac{\theta_{kt}}{\omega_{kt}^*} > 0, \quad \frac{d\omega_{\ell t}^*}{d\theta_{kt}} \frac{\theta_{kt}}{\omega_{\ell t}^*} < 0, \quad \forall \ell \neq k, \quad (17)$$

as shown in Section A.4 of the web appendix.

The intuition for this consumer-valuation bias is as follows. An increase in demand for a good ($\theta_{kt}/\theta_{kt-1} > 1$ and hence $\varphi_{kt}/\varphi_{kt-1} > 1$) is analogous to a reduction in price for that good ($p_{kt}/p_{kt-1} < 1$) because consumer preferences depend on demand-adjusted prices (p_{kt}/φ_{kt}). Other things equal, consumers substitute towards goods that experience relative increases in demand, which raises consumer welfare relative to the value it would take if expenditure shares were left unchanged in the face of these different demand-adjusted prices. A price index that rules out such changes in demand by assumption cannot capture this increase in welfare, thereby giving rise to the consumer-valuation bias. This bias is analogous to the well-known “substitution bias,” in which a Laspeyres index that allows for no substitution across goods overstates the increase in the cost of living because it does not take into account that consumers can substitute towards goods that experience reductions in relative prices.

A final metric for the tension inherent in the Sato-Vartia price index’s assumption of time-invariant demand for each common good is to note that under this assumption the elasticity of substitution can be recovered from the observed data on prices and expenditure shares with no estimation (from differencing over time in equation (4) under the assumption that $\varphi_{kt} = \varphi_{kt-1} = \varphi_k$). Indeed, the model is overidentified, with an infinite number of approaches to measuring the elasticity of substitution, each of which uses different weights for each common good, as shown in Section A.5 of the web appendix. If demand for *all* common goods is indeed constant (including no changes in tastes, quality, measurement error or specification error), all of these approaches will recover the same elasticity of substitution. However, if demand for *some* common good changes over time, but a researcher falsely assumes time-invariant demand for all common goods, these alternative approaches will return different values for the elasticity of substitution, depending on which weights are used. We use this metric below to provide evidence on the empirical validity of the assumption of time-invariant demand for all common goods.

2.5 Estimating the Elasticity of Substitution

Our CES unified price index (CUPI) in equation (14) and the existing CES exact price index in equation (15) both require an estimate of the elasticity of substitution (σ). In Subsection 2.5.1, we review existing approaches based on demand systems estimation. In Subsection 2.5.2, we introduce our new “reverse-weighting” (RW) estimator. We show that this estimator is consistent (i) as demand shocks become small or (ii) as the number of common goods becomes large and demand shocks are uncorrelated with price shocks for each good and independently and identically distributed across goods. In Subsection 2.5.3, we extend our analysis to develop a “generalized-reverse-weighting” (GRW) estimator, which retains the assumption that demand shocks are independently and identically distributed across goods, but allows the demand and price shocks for a given good to be correlated with one another. In Subsection 2.5.2, we show how use our inversion of the CES demand system can be used to provide upper and lower bounds for the true elasticity of substitution that

hold regardless of the correlation between demand and price shocks.

2.5.1 Demand Systems Estimation

Using the common-goods expenditure share in equation (7), dividing by its geometric mean across common goods, taking logarithms, and differencing between a pair of periods $t - 1$ and t , we obtain the following structural demand system:

$$\Delta \ln \bar{s}_{kt}^* = (1 - \sigma) \Delta \ln \bar{p}_{kt} - (1 - \sigma) \Delta \ln \theta_{kt}, \quad (18)$$

where Δ denotes the time-difference operator such that $\Delta \ln \bar{p}_{kt} = \ln (\bar{p}_{kt} / \bar{p}_{kt-1})$; a bar above a variable indicates that it is normalized by its geometric mean across common goods such that $\ln (\bar{p}_{kt}) = \ln (p_{kt} / \bar{p}_t)$; and the time-invariant component of demand (φ_k) has differenced out between the two time periods to leave only the change in the time-varying component of demand ($\Delta \ln \theta_{kt}$); and we have used our result that demand shocks average out across common goods ($\ln (\tilde{\theta}_t / \tilde{\theta}_{t-1}) = 0$). This structural demand system (18) has the following reduced-form representation:

$$\Delta \ln \bar{s}_{kt}^* = \beta_0 + \beta_1 \Delta \ln \bar{p}_{kt} + u_{kt}, \quad (19)$$

which is analogous to equation (4) but uses common goods expenditure shares ($\bar{s}_{kt}^*$) and differences both over time and relative to the geometric mean; the constant (β_0) in equation (19) is necessarily equal to zero, because of the normalization of common goods expenditure shares and prices by their geometric mean.

The main challenge in estimating the reduced-form demand system (19) is that shocks to prices ($\Delta \ln \bar{p}_{kt}$) can be correlated with shocks to demand (u_{kt}), giving rise to conventional omitted variable bias. The standard approach to this problem is to specify a supply-side such as:

$$\Delta \ln \bar{p}_{kt} = \gamma_0 + \gamma_1 \Delta \ln \bar{s}_{kt}^* + \gamma_2 z_{kt} + v_{kt}, \quad (20)$$

and to search for instruments (z_{kt}) that are both powerfully correlated with the log change in prices ($\text{cov} (\Delta \ln \bar{p}_{kt}, z_{kt}) \neq 0$) and have no direct effect on expenditure shares ($\text{cov} (u_{kt}, z_{kt}) = 0$). However, finding valid instruments in the settings with many industries considered in international trade and macroeconomics can be challenging.

The main alternative approach to estimating the CES demand system is that of Feenstra (1994). This alternative estimator uses the second differences of the demand system (19) and supply-system (20) across goods and over time, and makes the identifying assumption that these double-differenced demand and supply shocks are orthogonal to one another and heteroskedastic.⁵ The orthogonality assumption defines a rectangular hyperbola for each good in the space of the demand and supply elasticities. The heteroskedasticity assumption implies that these rectangular hyperbolas for different goods do not lie on top of another. Therefore, their intersection separately identifies the demand and supply elasticities.

⁵See Section A.8 of the web appendix for further details.

2.5.2 Reverse-Weighting (RW) Estimator

We now develop an alternative “reverse-weighting” (RW) estimator of the elasticity of substitution, which combines the demand system with the unit expenditure function, and uses the identifying assumption that changes in the cost of living are money metric. Although this approach uses different identifying assumptions from Feenstra (1994), we find that in practice our RW estimates in this section and our GRW estimates in the next section do not differ greatly from those using the Feenstra (1994) estimator, and all three sets of estimates lie within the upper and lower bounds derived below.

We begin by augmenting the UPI in equation (14) with two other equivalent expressions for the change in the cost of living that arise from taking forward and backward differences of the CES unit expenditure function. These forward and backward differences were first introduced for the case without demand shocks by Lloyd (1975) and Moulton (1996) and are sometimes referred to as “Lloyd-Moulton” indexes.⁶ The forward difference 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) and (7), 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 ($\theta_{kt}/\theta_{kt-1}$) for all common goods:

$$\Phi_{t-1,t}^F = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)^{\frac{1}{\sigma-1}} \frac{P_t^*}{P_{t-1}^*} = \left(\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}} \right)^{\frac{1}{\sigma-1}} \left[\sum_{k \in \Omega_{t,t-1}} s_{kt-1}^* \left(\frac{p_{kt}/\theta_{kt}}{p_{kt-1}/\theta_{kt-1}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad (21)$$

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

$$\Phi_{t,t-1}^B = \left(\frac{\lambda_{t-1,t}}{\lambda_{t,t-1}} \right)^{\frac{1}{\sigma-1}} \frac{P_{t-1}^*}{P_t^*} = \left(\frac{\lambda_{t-1,t}}{\lambda_{t,t-1}} \right)^{\frac{1}{\sigma-1}} \left[\sum_{k \in \Omega_{t,t-1}} s_{kt}^* \left(\frac{p_{kt-1}/\theta_{kt-1}}{p_{kt}/\theta_{kt}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}, \quad (22)$$

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

We use these forward and backward differences to develop our RW estimator, but they also can be used to relate our CUPI to most existing economic and statistical price indexes. As shown in Section A.7 of the web appendix, under our assumption of CES preferences, the CUPI coincides with many of these existing price indexes (including Laspeyres, Paasche, Fisher and Törnqvist indexes) for specific parameter values and particular assumptions about entry and exit and changes in demand for common goods. For example, under the assumption of no changes in the set of products ($\lambda_{t-1,t}/\lambda_{t,t-1} = 1$), no demand shifts ($\theta_{kt}/\theta_{kt-1} = 1$), and no substitution across goods ($\sigma = 0$), the forward difference (21) collapses to the Laspeyres index and the backward difference (22) collapses to the inverse of the Paasche index. Nevertheless, there are of course other ways of rationalizing these existing price indexes using alternative functional form assumptions (e.g. the Laspeyres index is exact for Leontief preferences).⁷

⁶See, for example, Diewert (2004) and Bialek (2015).

⁷All of these existing economic and statistical price indexes assume no demand shocks for individual goods, as do continuous time index numbers such as the Divisia index, as also shown in Section A.7 of the web appendix.

Equating our three expressions for the change in the cost of living in equations (14), (21) and (22), the common variety correction term $((\lambda_{t,t-1}/\lambda_{t-1,t})^{\frac{1}{\sigma-1}})$ cancels. Re-arranging terms, we obtain the following two key equalities between equivalent ways of writing the change in the cost of living for common goods:

$$\Theta_{t-1,t}^F \left[\sum_{k \in \Omega_{t,t-1}} s_{kt-1}^* \left(\frac{p_{kt}}{p_{kt-1}} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} = \frac{\tilde{p}_t}{\tilde{p}_{t-1}} \left(\frac{\tilde{s}_t^*}{\tilde{s}_{t-1}^*} \right)^{\frac{1}{\sigma-1}}, \quad (23)$$

$$\left(\Theta_{t,t-1}^B \right)^{-1} \left[\sum_{k \in \Omega_{t,t-1}} s_{kt}^* \left(\frac{p_{kt}}{p_{kt-1}} \right)^{-(1-\sigma)} \right]^{-\frac{1}{1-\sigma}} = \frac{\tilde{p}_t}{\tilde{p}_{t-1}} \left(\frac{\tilde{s}_t^*}{\tilde{s}_{t-1}^*} \right)^{\frac{1}{\sigma-1}}, \quad (24)$$

where we have used our result that $\tilde{\varphi}_t/\tilde{\varphi}_{t-1} = 1$ and $\Theta_{t-1,t}^F$ and $\Theta_{t,t-1}^B$ are forward and backward aggregate demand shifters that are defined respectively as:

$$\begin{aligned} \Theta_{t-1,t}^F &\equiv \left[\frac{\sum_{k \in \Omega_{t,t-1}} s_{kt-1}^* \left(\frac{p_{kt}}{p_{kt-1}} \right)^{1-\sigma} \left(\frac{\theta_{kt}}{\theta_{kt-1}} \right)^{\sigma-1}}{\sum_{k \in \Omega_{t,t-1}} s_{kt-1}^* \left(\frac{p_{kt}}{p_{kt-1}} \right)^{1-\sigma}} \right]^{\frac{1}{1-\sigma}} = \left[\sum_{k \in \Omega_{t,t-1}} s_{kt}^* \left(\frac{\theta_{kt-1}}{\theta_{kt}} \right)^{\sigma-1} \right]^{\frac{1}{\sigma-1}}, \\ \Theta_{t,t-1}^B &\equiv \left[\frac{\sum_{k \in \Omega_{t,t-1}} s_{kt}^* \left(\frac{p_{kt-1}}{p_{kt}} \right)^{1-\sigma} \left(\frac{\theta_{kt-1}}{\theta_{kt}} \right)^{\sigma-1}}{\sum_{k \in \Omega_{t,t-1}} s_{kt}^* \left(\frac{p_{kt-1}}{p_{kt}} \right)^{1-\sigma}} \right]^{\frac{1}{1-\sigma}} = \left[\sum_{k \in \Omega_{t,t-1}} s_{kt-1}^* \left(\frac{\theta_{kt}}{\theta_{kt-1}} \right)^{\sigma-1} \right]^{\frac{1}{\sigma-1}}, \end{aligned} \quad (25)$$

and all derivations for this section are reported in Section A.9 of the web appendix.

The terms in square parentheses that multiply $\Theta_{t-1,t}^F$ and $\left(\Theta_{t,t-1}^B \right)^{-1}$ in equations (23) and (24) are the Lloyd-Moulton indexes without demand shocks discussed above. In standard approaches, the assumption of time-invariant demand for each good implies that these two Lloyd-Moulton indexes should be equal, and the reason why in practice they do not take the same value in the data is left unexplained (see, for example, Shapiro and Wilcox 1996). In contrast, our setup makes clear that if one allows for product-level demand shocks, the various CES price indexes derived under the assumption of time-invariant demand need not take the same value, and the differences between them contain information about the unobserved pattern of demand shocks. The aggregate demand shifters $(\Theta_{t-1,t}^F$ and $\Theta_{t,t-1}^B)$ in equation (25) capture differences between movements in demand-adjusted prices (p_{kt}/φ_{kt}) and unadjusted prices (p_{kt}) . If demand and price shocks are positively correlated, the impact of price increases in raising the cost of living is offset on average by increases in the utility obtained from each unit of the good. In contrast, if demand and price shocks are negatively correlated, the impact of price increases in raising the cost of living is magnified by reductions in the utility obtained from each unit of the good.

Comparing our different expressions for the change in the cost of living on the right and left-hand sides of equations (23) and (24), the CUPPI on the right-hand side is always money-metric for any elasticity of substitution (σ), because of our result that $\tilde{\varphi}_t/\tilde{\varphi}_{t-1} = 1$. In contrast, the forward and backward differences of the unit expenditure function on the left-hand side are not in general money-metric, because they are directly affected by the demand shocks for each common good $(\theta_{kt}/\theta_{kt-1})$ through the aggregate demand shifters $(\Theta_{t-1,t}^F$ and $\Theta_{t,t-1}^B)$. These aggregate effects depend on the correlation between the demand shocks for each common good and initial and end-period expenditure shares, as also shown in equation (25). If the

goods that experience increases in demand account for large shares of expenditure, while the goods that experience falls in demand account for small shares of expenditure, this will reduce the cost of living in the forward and backward differences of the unit expenditure function, and vice versa. This observation raises three related questions: (i) are there any circumstances under which all three expressions for the cost of living are money metric (which requires $\Theta_{t-1,t}^F = \left(\Theta_{t,t-1}^B\right)^{-1} = 1$), (ii) could this identifying assumption be used to construct an estimator of the elasticity of substitution (σ), and (iii) when would this estimator be consistent?

Before turning to these questions, we note that equations (23) and (24) imply that the following relationship between the forward and backward aggregate demand shifters must hold for any elasticity of substitution and any combination of demand and price shocks:

$$\left[\sum_{k \in \Omega_{t,t-1}} s_{kt-1}^* \left(\frac{p_{kt}}{p_{kt-1}} \right)^{1-\sigma} \right]^{-\frac{1}{1-\sigma}} \left[\sum_{k \in \Omega_{t,t-1}} s_{kt}^* \left(\frac{p_{kt}}{p_{kt-1}} \right)^{-(1-\sigma)} \right]^{-\frac{1}{1-\sigma}} = \Theta_{t-1,t}^F \Theta_{t,t-1}^B = \bar{\Theta}. \quad (26)$$

Using this relationship in equations (23) and (24), we find that if the equality between the forward difference of the unit expenditure function and the UPI in equation (23) is satisfied, the equality between the backward difference of the unit expenditure function and the UPI in equation (24) also must be satisfied, and vice versa. Therefore, there is a single value of the elasticity of substitution (σ) that satisfies both of these equations for any constellation of demand and price shocks.

Our reverse-weighting (RW) estimator estimates the elasticity of substitution (σ) using the identifying assumption that all three expressions for the change in the cost of living are money metric:

$$\Theta_{t-1,t}^F = \left(\Theta_{t,t-1}^B\right)^{-1} = 1, \quad (27)$$

which requires that the weighted sum of the demand shocks going forward in time ($(\varphi_{kt}/\varphi_{kt-1})^{\sigma-1}$) using initial-period expenditure shares (s_{kt-1}^*) and the weighted sum of the demand shocks going backward in time ($(\varphi_{kt-1}/\varphi_{kt})^{\sigma-1}$) using final-period expenditure shares (s_{kt}^*) are both equal to one in equation (25).

Using this identifying assumption of money-metric utility in equations (23) and (24), we obtain the following sample moment conditions:

$$M^{RW}(\sigma, \mathbf{X}) = \begin{pmatrix} \frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \left[\frac{1}{1-\sigma} \ln \left(\sum_{\ell \in \Omega_{t,t-1}} s_{\ell t-1}^* \left(\frac{p_{\ell t}}{p_{\ell t-1}} \right)^{1-\sigma} \right) - \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \frac{1}{\sigma-1} \ln \left(\frac{s_{kt}^*}{s_{kt-1}^*} \right) \right] \\ \frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \left[-\frac{1}{1-\sigma} \ln \left(\sum_{\ell \in \Omega_{t,t-1}} s_{\ell t}^* \left(\frac{p_{\ell t}}{p_{\ell t-1}} \right)^{-(1-\sigma)} \right) - \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \frac{1}{\sigma-1} \ln \left(\frac{s_{kt}^*}{s_{kt-1}^*} \right) \right] \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad (28)$$

where $\mathbf{X}$ is the matrix formed by the observed data on prices and expenditure shares for each good k for periods $t-1$ and t . The reverse weighting estimator ($\hat{\sigma}^{RW}$) solves:

$$\hat{\sigma}^{RW} = \arg \min \left\{ M^{RW}(\sigma, \mathbf{X})' \times \mathbb{I} \times M^{RW}(\sigma, \mathbf{X}) \right\}, \quad (29)$$

where $\mathbb{I}$ is the identity matrix.⁸

This RW estimator is overidentified with two moment conditions to estimate one parameter. If the identifying assumption of money-metric utility in equation (27) is satisfied, equations (23) and (24) imply that

⁸In Section A.10 of the web appendix, we show that the reverse-weighting estimator in equations (28) and (29) generalizes to allow for a Hicks-neutral shifter of tastes that is common to all goods because, like the variety correction term, this Hicks-neutral shifter cancels from equations (23)-(24).

both moment conditions in equation (28) are simultaneously satisfied at the same value for the elasticity of substitution. As we use the identity matrix ($\mathbb{I}$) as the weighting matrix in equation (29), the RW estimator weights the two moment conditions equally, and hence minimizes the sum of squared deviations of the log aggregate demand shifters from zero ($\left\{ \ln \Theta_{t-1,t}^F \right\}^2 + \left\{ \ln \Theta_{t,t-1}^B \right\}^2$).

Our identifying assumption of money-metric utility in equation (27) has an attractive economic interpretation. The money-metric forward difference (the left-hand side of equation (23) with $\Theta_{t-1,t}^F = 1$) corresponds to the change in the cost of living evaluated using period $t - 1$ tastes. Similarly, the money-metric backward difference (the left-hand side of equation (24) with $\left(\Theta_{t,t-1}^B \right)^{-1} = 1$) corresponds to the change in the cost of living evaluated using period t tastes. This property implies that the RW estimator ($\hat{\sigma}^{RW}$) minimizes the sum of squared deviations between (i) the change in the cost of living evaluated using the unified price index and tastes in each time period (inverting the demand system to substitute for these unobserved tastes using prices and expenditure share in each period), (ii) the change in the cost of living evaluated using period $t - 1$ tastes, and (iii) the change in the cost of living using period t tastes, as shown in Section A.11 of the web appendix. This property relates to the results of Fisher and Shell (1972), which uses the tastes of the initial or final period to bound the change in the cost of living. Here, we show that the elasticity of substitution itself can be chosen to minimize the difference between the implied change in the cost of living using initial or final-period tastes.

In addition to having an intuitive economic interpretation, we now provide conditions under which our identifying assumption of money metric utility in equation (27) is satisfied, and the RW estimator consistently estimates the true elasticity of substitution (σ^D), where we use the superscript D to indicate the true parameter value. First, we show that the RW estimator consistently estimates the true elasticity of substitution (σ^D) as the demand shocks for each good become small.

Proposition 2. *As changes in demand become small ($(\theta_{kt}/\theta_{kt-1}) \rightarrow 1$), the reverse-weighting (RW) estimator consistently estimates the true elasticity of substitution ($\hat{\sigma}^{RW} \xrightarrow{P} \sigma^D$).*

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

As the demand shocks for each good become small ($(\theta_{kt}/\theta_{kt-1}) \rightarrow 1$), the weighted sum of these demand shocks using either initial or final-period expenditure shares converges to one, and hence the forward and backward aggregate demand shifters in equation (25) converge to one ($\Theta_{t-1,t}^F \xrightarrow{P} 1$ and $\Theta_{t,t-1}^B \xrightarrow{P} 1$). Therefore, as the demand shocks for each good become small, the identifying assumption of money-metric utility is satisfied, and the RW estimator consistently estimates the elasticity of substitution ($\hat{\sigma}^{RW} \xrightarrow{P} \sigma^D$). In Section A.20 of the web appendix, we report Monte Carlos, in which we show that the mean RW estimate across Monte Carlo replications lies close to the true parameter even in finite samples, and the standard deviation of the RW estimate across these replications falls with the standard deviation of the demand shocks. In Section A.13 of the web appendix, we also show that the identifying assumption of money-metric utility in equation (27) is satisfied up to a first-order approximation, which implies that the RW estimator can be interpreted as providing a first-order approximation to the data.

Second, we show that the RW estimator consistently estimates the true elasticity of substitution (σ^D) as the number of common goods becomes large ($N_{t,t-1} \rightarrow \infty$) if demand shocks are uncorrelated with price shocks for each good and independently and identically distributed across goods.

Proposition 3. *Assume that demand shocks are uncorrelated with price shocks for a given good and are independently and identically distributed across goods, such that $(\theta_{kt}/\theta_{kt-1}) \sim \text{i.i.d}(1, \psi_\theta^2)$ for $(\theta_{kt}/\theta_{kt-1}) \in (0, \infty)$. As the number of common goods becomes large ($N_{t,t-1} \rightarrow \infty$), the reverse-weighting (RW) estimator consistently estimates the elasticity of substitution ($\hat{\sigma}^{RW} \xrightarrow{P} \sigma^D$).*

Proof. See Section A.14 of the web appendix. □

If demand shocks are orthogonal to price shocks and independently and identically distributed across goods, the demand shock going forward in time ($(\varphi_{kt}/\varphi_{kt-1})^{\sigma-1}$) is uncorrelated with the initial-period expenditure share (s_{kt-1}^*), and the demand shock going backwards in time ($(\varphi_{kt-1}/\varphi_{kt})^{\sigma-1}$) is uncorrelated with the final-period expenditure share (s_{kt}^*). In the proof of Proposition 3, we show that this property implies that the forward and the inverse of the backward aggregate demand shifters in equation (25) converge to a common probability limit ($\text{plim} [\Theta_{t-1,t}^F] = \text{plim} [\left(\Theta_{t,t-1}^B\right)^{-1}]$). If the expected value of the demand shock for each good is equal to one ($\mathbb{E}(\theta_{kt}/\theta_{kt-1}) = 1$), we also show that this common probability limit is equal to one. Therefore, the identifying assumption of money-metric utility is satisfied asymptotically, and the RW estimator again consistently estimates the elasticity of substitution ($\hat{\sigma}^{RW} \xrightarrow{P} \sigma^D$). In Section A.20 of the web appendix, we report Monte Carlos, in which we show that the mean RW estimate is close to the true parameter value if the assumptions in Proposition 3 are satisfied, even in finite samples with a relatively small number of common goods.

In Section A.15 of the web appendix, we use the consistency results from Propositions 2 and 3 together with the fact that the RW estimator belongs to the class of M-estimators (Newey and McFadden 1994 and Wooldridge 2002) to show that the RW estimates are asymptotically normal.

2.5.3 Generalized-Reverse-Weighting Estimator (GRW)

We now develop our “generalized-reverse-weighting” (GRW) estimator, which retains the assumption of independence across goods, but allows demand and price shocks for any given good to be correlated with one another. In particular, we assume that demand shocks can be partitioned into a component that is correlated prices and an idiosyncratic component that is uncorrelated with prices for each good and independently and identically distributed across goods:

$$\ln\left(\frac{\theta_{kt}}{\theta_{kt-1}}\right) = \gamma \ln\left(\frac{p_{kt}/\tilde{p}_t}{p_{kt-1}/\tilde{p}_{t-1}}\right) + \ln\left(\frac{\epsilon_{kt}}{\epsilon_{kt-1}}\right); \quad (30)$$

where γ is the projection coefficient of demand shocks on price shocks, i.e.,

$$\gamma = \rho \frac{\chi_\theta}{\chi_p}, \quad \left(\frac{\epsilon_{kt}}{\epsilon_{kt-1}}\right) \perp \left(\frac{p_{kt}/\tilde{p}_t}{p_{kt-1}/\tilde{p}_{t-1}}\right); \quad (31)$$

where χ_θ is the standard deviation of demand shocks ($\ln(\theta_{kt}/\theta_{kt-1})$); χ_p is the standard deviation of price shocks ($\ln((p_{kt}/\tilde{p}_t)/(p_{kt-1}/\tilde{p}_{t-1}))$); ρ is the correlation of demand and price shocks; and all derivations for this section are reported in Section A.16 of the web appendix.

In that section of the web appendix, we show that this specification in equation (30) can be derived from the assumption that demand and price shocks are joint log-normally distributed with a general variance-covariance matrix. Therefore, our GRW estimator imposes additional structure on the underlying distributions of demand and price shocks. So far, we have used our assumption that demand shocks are independently and identically distributed across goods to derive the property that mean log-demand shocks are zero. Now, we require that the stochastic processes for demand and price shocks have a joint log-normal distribution. Under this distributional assumption, the standard deviations (χ_p, χ_θ) and correlation (ρ), and hence the projection coefficient (γ), are structural parameters in addition to the elasticity of substitution (σ).

Our RW moment conditions in equation (28) are defined over price shocks and initial- and end-period expenditure shares and do not directly include demand shocks. Under the conditions specified in Propositions 2 and 3 above, these end-period expenditure shares (s_{kt}^*) are not systematically influenced by these demand shocks, either because these demand shocks are close to one ($\theta_{kt}/\theta_{kt-1} \rightarrow 1$), or because demand and price shocks for a given good are uncorrelated with one another ($\text{cov}[(\theta_{kt}/\theta_{kt-1}), (p_{kt}/p_{kt-1})] = 0$). However, if demand shocks are large ($\theta_{kt}/\theta_{kt-1} \neq 1$) and correlated with price shocks for a given good ($\text{cov}[(\theta_{kt}/\theta_{kt-1}), (p_{kt}/p_{kt-1})] \neq 0$), the true direct effect of these price shocks on the end-period expenditure shares through the elasticity of substitution (σ) is obscured by their correlation with demand shocks.

We now characterize the direction of the bias in the RW estimator with large and correlated demand and price shocks, before developing our GRW estimator below that allows for this correlation.

Proposition 4. *If demand and price shocks are positively correlated ($\gamma > 0$), the RW estimator is asymptotically downward biased ($\text{plim}(\hat{\sigma}^{RW}) < \sigma$), whereas if demand and price shocks are negatively correlated ($\gamma < 0$), the RW estimator is asymptotically upward biased ($\text{plim}(\hat{\sigma}^{RW}) > \sigma$).*

Proof. See Section A.17 of the web appendix. □

Intuitively, if demand and price shocks are positively correlated ($\gamma > 0$), the impact of an increase in price for a good is offset on average by an increase in demand for that good, which reduces the responsiveness of expenditure shares to changes in prices. If a researcher assumes no demand shocks or that demand shocks are uncorrelated with price shocks, she will conclude that this unresponsiveness of expenditure shares to price changes reflects inelastic demand, whereas in fact it reflects the positive correlation between demand and price shocks. Conversely, if demand and price shocks are negatively correlated ($\gamma < 0$), she will conclude that the responsiveness of expenditure shares to price shocks is explained by elastic demand, whereas in fact it is explained by the negative correlation between demand and price shocks. In the Monte Carlos in Section A.20 of the web appendix, we illustrate this result in finite samples, where the mean RW estimate lies above the true parameter value when demand and price shocks are negatively correlated and below the true parameter value when demand and price shocks are positively correlated.

To allow for correlated demand and price shocks, we now use the structure of the model to generate counterfactual end-period expenditure shares (S_{kt}^*) that are purged of the component of demand shocks that is correlated with price shocks. Note that if we could observe these demand shocks ($\theta_{kt}/\theta_{kt-1}$), we could remove their entire effect on end-period expenditure shares (s_{kt}^*) by using the CES demand system to construct counterfactual end-period expenditure shares (S_{kt}^*) as follows:

$$S_{kt}^* = \frac{(\theta_{kt}/\theta_{kt-1})^{-(\sigma-1)} s_{kt}^*}{\sum_{\ell \in \Omega_{t,t-1}} (\theta_{\ell t}/\theta_{\ell t-1})^{-(\sigma-1)} s_{\ell t}^*}, \quad (32)$$

where we use an upper case S_{kt}^* to denote counterfactual common goods expenditure shares.

Although we do not observe these demand shocks ($\theta_{kt}/\theta_{kt-1}$), we can use equation (30) to remove the component of demand shocks that is correlated with price shocks. In particular, we first combine equations (30) and (31) with CES demand (7) to obtain the following closed-form solution for the projection coefficient (γ) in terms of the elasticity of substitution (σ) and observed moments:

$$\gamma = \gamma(\sigma) = \frac{1}{\sigma - 1} \left[\sigma - 1 + \frac{\chi_{ps}}{\chi_p^2} \right], \quad (33)$$

where χ_{ps} is the covariance of price and sales shocks. In Section A.16 of the web appendix, we show that we can also solve in closed-form for the standard deviation of demand shocks ($\chi_\theta(\sigma)$) and the correlation between demand and price shocks ($\rho(\sigma)$) as a function of the elasticity of substitution (σ) and observed moments, where we require $\chi_\theta(\sigma) \geq 0$ and $|\rho(\sigma)| \leq 1$.

We next compute our counterfactual end-period expenditure shares ($S_{kt}^*(\sigma)$) purged of the component of demand that is correlated with price shocks as follows:

$$S_{kt}^*(\sigma) = \frac{(p_{kt}/p_{kt-1})^{-\gamma(\sigma-1)} s_{kt}^*}{\sum_{\ell \in \Omega_{t,t-1}} (p_{\ell t}/p_{\ell t-1})^{-\gamma(\sigma-1)} s_{\ell t}^*}, \quad (34)$$

where these counterfactual shares ($S_{kt}^*(\sigma)$) are expressed solely as a function of σ and observed data, because γ is a function of σ and observed data from equation (33) immediately above.

Finally, we use these counterfactual end-period expenditure shares ($S_{kt}^*(\sigma)$) to construct the following two moment conditions for our Generalized-Reverse-Weighting (GRW) estimator:

$$M^{GRW}(\sigma, \mathbf{X}) = \begin{pmatrix} \frac{1}{N_{t,t-1}} \sum_{k \in N_{t,t-1}} \left[\frac{1}{1-\sigma} \ln \left(\sum_{\ell \in \Omega_{t,t-1}} s_{\ell t-1}^* \left(\frac{p_{\ell t}}{p_{\ell t-1}} \right)^{1-\sigma} \right) - \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \frac{1}{\sigma-1} \ln \left(\frac{S_{kt}^*(\sigma)}{s_{kt-1}^*} \right) \right] \\ \frac{1}{N_{t,t-1}} \sum_{k \in N_{t,t-1}} \left[-\frac{1}{1-\sigma} \ln \left(\sum_{\ell \in \Omega_{t,t-1}} S_{\ell t}^*(\sigma) \left(\frac{p_{\ell t}}{p_{\ell t-1}} \right)^{-(1-\sigma)} \left(\frac{\theta_{\ell t}}{\theta_{\ell t-1}} \right)^{-(\sigma-1)} \right) - \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \frac{1}{\sigma-1} \ln \left(\frac{S_{kt}^*(\sigma)}{s_{kt-1}^*} \right) \right] \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}. \quad (35)$$

Our GRW estimator estimates the elasticity of substitution (σ) by solving:

$$\hat{\sigma}^{GRW} = \arg \min \left\{ M^{GRW}(\sigma, \mathbf{X})' \times \mathbb{I} \times M^{GRW}(\sigma, \mathbf{X}) \right\}, \quad (36)$$

where $\mathbb{I}$ is again the identity matrix.

The moment conditions for the GRW and RW estimators in equations (28) and (35) take the same form, except that the GRW moment conditions use the counterfactual final-period expenditure shares ($S_{kt}^*(\sigma)$) instead of the actual final-period expenditure shares (s_{kt}^*). Therefore, the RW estimator makes the identifying assumption of money-metric utility ($\Theta_{t-1,t}^F = \left(\Theta_{t,t-1}^B \right)^{-1} = 1$) using the *observed* expenditure shares for the

initial and final periods (s_{kt-1}^*, s_{kt}^*), whereas the GRW estimator makes the identifying assumption of money-metric utility after controlling for the correlation between demand and price shocks. Intuitively, as discussed in the previous subsection, the RW estimator chooses the elasticity of substitution to minimize the difference in the implied change in the cost of living using the tastes of the initial or final period. In contrast, the GRW estimator chooses the elasticity of substitution to minimize the difference between these two measures of the cost of living after controlling for the correlation between demand and price shocks.

We now show that the GRW estimator consistently estimates the elasticity of substitution (σ) regardless of the correlation between demand and price shocks for a given good as number of common goods becomes large ($N_{t,t-1} \rightarrow \infty$).

Proposition 5. *Assume that demand shocks for each good ($\theta_{kt}/\theta_{kt-1} \in (0, \infty)$) can be partitioned into a component that is correlated with price shocks and an orthogonal component ($(\epsilon_{kt}/\epsilon_{kt-1}) \sim i.i.d (1, \psi_\epsilon^2)$ for $\epsilon_{kt}/\epsilon_{kt-1} \in (0, \infty)$), and are independently and identically distributed across goods. As the number of common goods becomes large ($N_{t,t-1} \rightarrow \infty$), the Generalized-Reverse-Weighting (GRW) estimator consistently estimates the elasticity of substitution ($\hat{\sigma}^{GRW} \xrightarrow{P} \sigma^D$).*

Proof. See Section A.18 of the web appendix. □

In the Monte Carlos in Section A.20 of the web appendix, we show that the mean GRW estimator lies close to the true parameter value regardless of the correlation between demand and price shocks, although it is less precisely estimated with larger standard errors than the RW estimator. As for the RW estimator in Subsection 2.5.2 above, the GRW estimator belongs to the class of M-estimators (Newey and McFadden 1994 and Wooldridge 2002), as shown in Section A.15 of the web appendix. Therefore, the GRW estimator inherits the same asymptotic normality properties as discussed for the RW estimator above.

2.5.4 Bounding the Elasticity of Substitution

As a check on our RW and GRW estimates, we now use our inversion of the CES demand system and assumption of joint log normality to provide upper and lower bounds for the elasticity of substitution (σ) regardless of the correlation between demand and price shocks. From the CES demand system (11), we have the following expressions for the covariance between log price shocks ($\ln((p_{kt}/\tilde{p}_t) / (p_{kt-1}/\tilde{p}_{t-1}))$) and log sales shocks ($\ln((s_{kt}/\tilde{s}_t) / (s_{kt-1}/\tilde{s}_{t-1}))$) and the variance of log sales shocks:

$$\chi_{ps} = (1 - \sigma) \left[\chi_p^2 - \chi_{p\theta} \right], \quad (37)$$

$$\chi_s^2 = (1 - \sigma)^2 \left[\chi_p^2 + \chi_\theta^2 - 2\chi_{p\theta} \right], \quad (38)$$

where the definitions of the variance and covariance terms (χ_{ps} , χ_p^2 , $\chi_{p\theta}$, χ_s^2 , χ_θ^2) and all other results for this section are reported in Section A.19 of the web appendix. Under our assumption of joint log normality, the correlation between price and demand shocks ($\chi_{p\theta}$) and the variance of demand shocks (χ_θ^2) are both parameters.

Using equation (37) to substitute for $\chi_{p\theta}$ in equation (38), we obtain the following relationship that implicitly defines the elasticity of substitution (σ) as a function of the observed moments (χ_{ps} , χ_p^2 , χ_s^2) for each assumed value for the variance of demand shocks (χ_θ^2):

$$\chi_\theta^2 = \frac{\chi_s^2}{(1 - \sigma)^2} + \chi_p^2 + \frac{2}{\sigma - 1} \chi_{ps}. \quad (39)$$

As discussed above, our model requires $\sigma > 1$ to ensure positive utility given the entry and exit of goods with positive demand ($\varphi_{kt} > 0$). Therefore, our lower bound for the elasticity of substitution is one ($\underline{\sigma} = 1$). In Section A.19 of the web appendix, we show that a necessary and sufficient condition for the elasticity of substitution (σ) implied by equation (39) to be monotonically decreasing in the assumed variance of demand shocks (χ_θ^2) is that the variance of demand shocks exceeds the variance of price shocks ($\chi_\theta^2 > \chi_p^2$). Evaluating equation (39) using our sample moments for (χ_s^2 , χ_p^2 , χ_{ps}) for each of our product groups, we find that the implied value of σ is indeed monotonically decreasing in the assumed value of χ_θ^2 , which implies that this necessary and sufficient condition is satisfied in our data. Using these properties that the variance of demand shocks exceeds the variance of price shocks ($\chi_\theta^2 > \chi_p^2$) and the implied elasticity of substitution is decreasing in the assumed variance of demand shocks (χ_θ^2), our upper bound for the elasticity of substitution ($\bar{\sigma}$) is obtained by solving equation (39) for the lowest possible value for the variance of demand shocks ($\chi_\theta^2 = \chi_p^2$).

We thus obtain set identification for the elasticity of substitution (σ) using the CES demand system and joint log normality, given the sample moments for the variances and covariance of price and sales shocks (χ_p^2 , χ_s^2 , χ_{ps}). Using our assumption that demand and price shocks are independent across goods, as the number of common goods becomes large ($N_{t,t-1} \rightarrow \infty$), the sample moments (χ_{ps} , χ_p^2 , χ_s^2) converge to their population counterparts. Therefore, the true elasticity of substitution (σ) necessarily lies within this identified set as the number of common goods becomes large.

Proposition 6. *Assume that demand and price shocks can be correlated with one another for each good but are independently and identically distributed across goods. As the number of common goods becomes large ($N_{t,t-1} \rightarrow \infty$), equation (39) identifies the set of possible values for elasticity of substitution $\sigma \in (1, \bar{\sigma})$ consistent with the observed data on prices and expenditure shares (p_{kt} , s_{kt}^*) under our assumptions of CES demand and joint log normality. As the number of common goods becomes large ($N_{t,t-1} \rightarrow \infty$), the true elasticity of substitution (σ) necessarily lies within this identified set.*

Proof. See Section A.19 of the web appendix. □

Intuitively, the observed variances and covariance of price and sales shocks (χ_{ps} , χ_p^2 , χ_s^2) contain enough information for set identification under the structure imposed by CES demand and joint log normality. This set identification is closely related to the use of forward and reverse regression to bound the extent of measurement error in Klepper and Leamer (1984) and the partial identification through inequality constraints in Leamer (1981). In contrast to the RW and GRW estimators, this approach does not use the CES unit expenditure function, and hence provides a useful specification check on these other estimators (and any other

estimate for the elasticity of substitution). We find in our empirical results below that the consumer-valuation bias is larger for lower values of the elasticity of substitution, because the ability to substitute towards goods for which demand has increased has a greater impact on utility the less substitutable are goods. Therefore, what matters for obtaining a lower bound on this consumer-valuation bias is the upper bound for the elasticity of substitution, as discussed further below.

3 Extensions and Generalizations

In this section, we consider a number of extensions and generalizations of our approach, including non-homothetic CES (indirectly additive), nested CES, mixed CES, logit, and translog preferences. We show that our main insight that the demand system can be inverted to express unobserved demand shocks for individual goods in terms of observed prices and expenditure shares generalizes to each of these specifications. Therefore, in each case, we can use this demand system inversion to derive a money-metric expression for the change of cost of living, and existing price indexes that assume time-invariant demand for each good are subject to a consumer-valuation bias.

3.1 Non-homothetic CES

We begin by generalizing our approach to allow for non-homotheticities, as recently emphasized in Fajgelbaum and Khandelwal (2016). We consider the non-separable class of CES functions in Sato (1975), which satisfy implicit additivity in Hanoch (1975), as recently used in the macroeconomics literature in Comin, Lashkari and Mestieri (2015). We suppose that we observe data on households indexed by $h \in \{1, \dots, H\}$ that differ in income and total expenditure (E_t^h). The non-homothetic CES consumption index for household h (C_t^h) is defined by the following implicit function:

$$\sum_{k \in \Omega_t} \left(\frac{\varphi_{kt}^h c_{kt}^h}{(C_t^h)^{(\epsilon_k - \sigma)/(1 - \sigma)}} \right)^{\frac{\sigma - 1}{\sigma}} = 1, \quad (40)$$

where c_{kt}^h denotes household h 's consumption of good k at time t ; φ_{kt}^h is household h 's demand parameter for good k at time t , which evolves according to equation (2); σ is the constant elasticity of substitution between varieties; ϵ_k is the constant elasticity of consumption of good k with respect to the consumption index (C_t^h) that allows for non-homotheticity. Assuming that goods are substitutes ($\sigma > 1$), we require $\epsilon_k < \sigma$ for the consumption index (40) to be globally monotonically increasing and quasi-concave, and hence to correspond to a well-defined utility function. Our baseline homothetic CES specification from Section 2 above corresponds to the special case of equation (40) in which $\epsilon_k = 1$ for all $k \in \Omega$.

Solving the household's expenditure minimization problem, we obtain the following expressions for the price index (P_t^h) dual to the consumption index (C_t^h) and the expenditure share for an individual good k (s_{kt}^h):

$$P_t^h = \left[\sum_{k \in \Omega_t} \left(p_{kt} / \varphi_{kt}^h \right)^{1 - \sigma} (C_t^h)^{\epsilon_k - 1} \right]^{\frac{1}{1 - \sigma}}, \quad (41)$$

$$s_{kt}^h = \frac{(p_{kt}/\varphi_{kt}^h)^{1-\sigma} (C_t^h)^{\epsilon_k-1}}{\sum_{\ell \in \Omega_t} (p_{\ell t}^h/\varphi_{\ell t}^h)^{1-\sigma} (C_t^h)^{\epsilon_\ell-1}} = \frac{(p_{kt}/\varphi_{kt}^h)^{1-\sigma} (E_t^h/P_t^h)^{\epsilon_k-1}}{(P_t^h)^{1-\sigma}}, \quad (42)$$

where we assume that all households h face the same price for a given good (p_{kt}) and the derivation for all results in this section is reported in Section A.21 of the web appendix.

As for the homothetic case above, the price index (41) depends on demand-adjusted prices (p_{kt}/φ_{kt}^h) rather than observed prices (p_{kt}). One challenge relative to the homothetic CES case is that the overall CES price index (P_t^h) enters the numerator of the expenditure share in equation (42). To overcome this challenge, we work with the share of each good in *overall* expenditure (s_{kt}^h) rather than the common goods expenditure share (s_{kt}^{h*} in our earlier notation), but we still take averages across the common goods, because only those common goods are supplied in both time periods. In particular, taking logarithms of the overall expenditure share in equation (42) for an individual common good, differencing between the two time periods, averaging across the common goods, and exponentiating, we obtain the following generalization of our CES unified price index to the non-homothetic case for each household h :

$$\frac{P_t^h}{P_{t-1}^h} = \left(\frac{\tilde{p}_t}{\tilde{p}_{t-1}} \right)^{\frac{1}{1+\vartheta}} \left(\frac{s_t^h}{s_{t-1}^h} \right)^{\frac{1}{(\sigma-1)(1+\vartheta)}} \left(\frac{E_t^h}{E_{t-1}^h} \right)^{\frac{\vartheta}{1+\vartheta}}, \quad (43)$$

$$\vartheta \equiv \frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \frac{\epsilon_k - 1}{1 - \sigma},$$

where the tilde above a variable denotes a mean across common goods; we have used our result that the mean log-demand shock across common goods is zero ($\ln(\tilde{\varphi}_t/\tilde{\varphi}_{t-1}) = 0$); the derived parameter ϑ captures the average across the common goods of the elasticity of expenditure with respect to the consumption index (ϵ_k) relative to the elasticity of substitution (σ); and the change in the household's cost of living (P_t^h/P_{t-1}^h) now depends directly on the change in income (and hence total expenditure) for parameter values for which preferences are non-homothetic ($\epsilon_k \neq 1$ for some k and hence $\vartheta \neq 0$).

Therefore, our unified approach to the demand system and the price index can be extended to accommodate non-homotheticity. In Section A.21 of the web appendix, we show that our reverse-weighting estimation procedure also can be generalized to estimate both the elasticity of substitution between goods (σ) and the elasticity of consumption of each good with respect to the consumption index (ϵ_k).

3.2 Nested CES

In our baseline specification in Section 2, we focus on a single CES tier of utility, which can be interpreted as a single sector consisting of many goods. In this section, we show that our analysis generalizes to a nested CES specification with multiple tiers of utility, by adding an additional upper tier of utility that is defined across sectors. For simplicity, we return to our baseline specification of homothetic CES. In particular, we assume that the aggregate unit expenditure function is defined across sectors $g \in \Omega^G$ as follows:

$$P_t = \left[\sum_{g \in \Omega^G} \left(\frac{P_{gt}^G}{\varphi_{gt}^G} \right)^{1-\sigma^G} \right]^{\frac{1}{1-\sigma^G}}, \quad \sigma^G > 1, \quad (44)$$

where σ^G is the elasticity of substitution across sectors; P_{gt}^G is the unit expenditure function for each sector, which is defined as in equation (1) across barcodes within that sector; φ_{gt}^G is the demand parameter for each sector; we assume for simplicity that the set of sectors is constant over time and denote the number of elements in this set by $N^G = |\Omega^G|$; and the derivations for this section of the paper are reported in Section A.22 of the web appendix.

All of the results for entry and exit and the exact CES price index with time-varying demand shocks from Section 2 continue to hold for this nested demand structure. We assume that demand for each sector (φ_{gt}^G) and demand for each good within each sector (φ_{kt}^K) take the same form as in equation (2). Therefore, as the number of sectors and the number of goods within each sector become large, the means of these log-demand shocks converge to zero. Furthermore, our CES common goods price index (CCG) involves taking the mean of logged variables, where the mean is a linear operator. Hence, we can apply this operator recursively across tiers of utility, and the change in the aggregate cost of living remains log linear:

$$\begin{aligned} \ln\left(\frac{P_t}{P_{t-1}}\right) &= \frac{1}{N^G} \sum_{g \in \Omega^G} \frac{1}{N_{gt,t-1}^K} \sum_{k \in \Omega_{gt,t-1}^K} \ln\left(\frac{p_{kt}^K}{p_{kt-1}^K}\right) + \frac{1}{N^G} \sum_{g \in \Omega^G} \frac{1}{\sigma_g^K - 1} \frac{1}{N_{gt,t-1}^K} \sum_{k \in \Omega_{gt,t-1}^K} \ln\left(\frac{s_{gkt}^{K*}}{s_{gkt-1}^{K*}}\right) \\ &\quad + \frac{1}{N^G} \sum_{g \in \Omega^G} \frac{1}{\sigma_g^K - 1} \ln\left(\frac{\lambda_{gt,t-1}^K}{\lambda_{gt-1,t}^K}\right) + \frac{1}{\sigma^G - 1} \frac{1}{N^G} \sum_{g \in \Omega^G} \ln\left(\frac{s_{gt}^G}{s_{gt-1}^G}\right), \end{aligned} \quad (45)$$

where we have used the result that mean log-demand shocks are zero across both sectors

(($(1/N^G) \sum_{g \in \Omega^G} \ln(\varphi_{gt}^G / \varphi_{gt-1}^G) = 0$) and across common goods within each sector

(($(1/N_{gt,t-1}^K) \sum_{k \in \Omega_{gt,t-1}^K} \ln(\varphi_{kt}^K / \varphi_{kt-1}^K) = 0$); $N_{gt,t-1}^K$ is the number of common goods for each sector g ; s_{gkt}^{K*} is the share of an individual common good k in expenditure on sector g at time t ; $(1/(\sigma_g^K - 1)) \ln(\lambda_{gt,t-1}^K / \lambda_{gt-1,t}^K)$ is the variety correction term for the entry and exit of goods within sector g ; and s_{gt}^G is the share of sector g in aggregate expenditure at time t).

Although, for simplicity, we focus on two tiers of utility here, this procedure can be extended from the highest tier of utility all the way down to the lowest. In general, our RW estimator can be applied recursively to each of these tiers of utility. However, conventional measures of the overall cost of living often aggregate categories using expenditure-share weights. Therefore, we assume that the upper tier of utility across sectors is Cobb-Douglas ($\sigma^G = 1$), and use our RW estimator to estimate the elasticity of substitution across barcodes within sectors (σ_k^K), which yields an estimated elasticity for each sector.

3.3 Mixed CES

The non-homothetic specification in Section 3.1 assumes that the only source of heterogeneity across consumers is differences in income and that all consumers have the same elasticity of substitution (σ). In this section, we introduce a mixed CES specification that allows both the elasticity of substitution and the demand parameters to differ across groups and does not restrict the ways in which these parameters differ. This mixed CES specification relaxes both the independence of irrelevant alternatives and symmetric substitution assumptions of our baseline CES specification.⁹ We consider a setting in which there are multiple groups of

⁹This mixed CES specification is used, for example, in Adao, Costinot and Rodriguez-Clare (2017) and is different from but related to the random coefficients model of Berry, Levinsohn and Pakes (1995).

heterogeneous consumers indexed by $h \in \{1, \dots, H\}$, which differ in terms of both their demand parameter for each good (φ_{kt}^h) and their substitution elasticities between goods (σ^h). In particular, we assume that the unit expenditure function (P_t^h) and expenditure share (s_{kt}^h) for a household from group h are given by:

$$P_t^h = \left[\sum_{k \in \Omega_t} \left(\frac{p_{kt}}{\varphi_{kt}^h} \right)^{1-\sigma^h} \right]^{\frac{1}{1-\sigma^h}}, \quad (46)$$

$$s_{kt}^h = \frac{(p_{kt}/\varphi_{kt}^h)^{1-\sigma^h}}{\sum_{\ell \in \Omega_t} (p_{\ell t}/\varphi_{\ell t}^h)^{1-\sigma^h}} = \frac{(p_{kt}/\varphi_{kt}^h)^{1-\sigma^h}}{(P_t^h)^{1-\sigma^h}}, \quad (47)$$

where s_{kt}^h is a share of product k in the expenditure of group h at time t ; we assume for simplicity that all groups face the same prices (p_{kt}); we suppose that demand for each each good for each group (φ_{kt}^h) takes the same form as in equation (2) above; we assume that the set of products available (Ω_t) is the same for all groups; but we allow for the possibility that some groups do not consume some products, which we interpret as corresponding to the limiting case in which the demand parameter converges to zero for that group and product ($\lim \varphi_{kt}^h \rightarrow 0$); and the derivation for all results in this section is reported in Section A.23 of the web appendix.

This specification relaxes the independence of irrelevant alternatives (IIIRA) of CES, because the differences in preferences across groups imply that the relative expenditure shares of two goods in two different markets depend on the relative size of the groups in those markets. This specification also relaxes the symmetric cross-substitution properties of CES, because the elasticity of expenditure on one variety with respect to a change in the price of another variety in two different markets also depends on group composition:

$$\frac{\partial x_{kt}}{\partial p_{\ell t}} \frac{p_{\ell t}}{x_{kt}} = \frac{1}{s_{kt}} \sum_{h=1}^H f_t^h (\sigma^h - 1) s_{kt}^h s_{\ell t}^h, \quad (48)$$

where s_{kt} is the share of product k in total expenditure; s_{kt}^h is the share of product k in total expenditure for group h ; and f_t^h is the share of group h in total expenditure.

All of our results from our baseline specification in Section 2 now hold for each group of consumers separately. Therefore, we can apply these results to calculate the change in the cost of living for each group separately.¹⁰ Following the same analysis as in Section 2.3, the exact CES unified price index for each group, allowing for entry and exit and demand shocks, takes the same form as in equation (14):

$$\ln \Phi_{t-1,t}^{hCUIPI} = \underbrace{\frac{1}{\sigma^h - 1} \ln \left(\frac{\lambda_{t,t-1}^h}{\lambda_{t-1,t}^h} \right)}_{\text{Variety Adjustment}} + \underbrace{\frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) + \frac{1}{\sigma^h - 1} \frac{1}{N_{t,t-1}} \sum_{k \in \Omega_{t,t-1}} \ln \left(\frac{s_{kt}^{h*}}{s_{kt-1}^{h*}} \right)}_{\text{Common-Goods Price Index}}, \quad (49)$$

where $(1/(\sigma^h - 1)) \ln (\lambda_{t,t-1}^h / \lambda_{t-1,t}^h)$ is the variety correction term for the entry and exit of goods for group h ; s_{kt}^{h*} is the share of an individual common good k in all expenditure on common goods for group h ; and we have used our result that demand shocks are mean zero in logs.

¹⁰In order to aggregate across groups, we would need to impose additional assumptions in the form of a social welfare function that specifies how to weight the preferences of each group.

We can use our RW estimator to estimate the elasticities of substitution (σ_g^h) for each group separately using the data on prices and expenditure shares for that group. In our empirical analysis in Section 5 below, we report such a robustness test for high- and low-income households, and compare both the estimated elasticities of substitution (σ^h) and changes in the cost of living for each group (P_t^h / P_{t-1}^h).

3.4 Logit

A well-known result in the discrete choice literature is that CES preferences can be derived as the aggregation of the choices of individual consumers with extreme-value-distributed idiosyncratic preferences, as shown in Anderson de Palma and Thisse (1992) and Train (2009). In this section, we briefly use this result to show that our unified price index and RW estimator for CES preferences also can be applied for logit preferences, as widely used in applied microeconomic research. Following McFadden (1974), we suppose that the utility of an individual consumer i who consumes c_{ikt} units of product k at time t is given by:

$$U_{it} = \ln \varphi_{kt} + \ln c_{ikt} + z_{ikt}, \quad (50)$$

where φ_{kt} captures the component of consumer tastes for each product that is common across consumers; z_{ikt} captures idiosyncratic consumer tastes for each product that are drawn from an independent Type-I Extreme Value distribution, $G(z) = e^{-e^{(-z/\nu + \kappa)}}$, where ν is the shape parameter of the extreme value distribution and $\kappa \approx 0.577$ is the Euler-Mascheroni constant.

Each consumer has the same expenditure E_t and chooses their preferred product given the observed realizations for idiosyncratic tastes. Using the properties of the extreme value distribution, we show in Section A.24 of the web appendix, that the expenditure share for each product and expected utility take exactly the same form as in our baseline CES specification in Section 2 of the paper, where $1/\nu = \sigma - 1$. Therefore, all our results for the unified price index and RW estimator can be applied for the logit model. Additionally, in the same way that our baseline CES specification can be generalized to accommodate mixed CES (as in Section 3.3 above), the baseline logit model in this section can be generalized to accommodate a mixed logit specification, as in McFadden and Train (2000).

3.5 Translog

In this final generalization, we show that our approach also holds for the flexible functional form of translog preferences. We focus for simplicity on a homothetic translog specification, which provides an arbitrary close local approximation to any continuous and twice-differentiable homothetic expenditure function.¹¹ In particular, we consider a translog unit expenditure function defined over the price (p_{kt}) and demand parameter (φ_{kt}) for a constant set of goods $k \in \Omega$ with number of elements $N = |\Omega|$:

$$\ln P_t = \ln \alpha_0 + \sum_{k \in \Omega} \alpha_k \ln \left(\frac{p_{kt}}{\varphi_{kt}} \right) + \frac{1}{2} \sum_{k \in \Omega} \sum_{\ell \in \Omega} \beta_{k\ell} \ln \left(\frac{p_{kt}}{\varphi_{kt}} \right) \ln \left(\frac{p_{\ell t}}{\varphi_{\ell t}} \right), \quad (51)$$

¹¹In the same way that our baseline homothetic CES specification can be extended to non-homothetic CES in Section 3.1 above, so this baseline homothetic translog specification can be extended to non-homothetic translog.

where the parameters β_{kl} control substitution patterns between goods; symmetry between goods requires $\beta_{kl} = \beta_{lk}$; symmetry and homotheticity together imply $\sum_{k \in \Omega} \alpha_k = 1$ and $\sum_{k \in \Omega} \beta_{kl} = \sum_{l \in \Omega} \beta_{lk} = 0$.

As for our baseline specification of homothetic CES preferences in Section 2 of the paper, the exact price index for these translog preferences depends on demand-adjusted prices (p_{kt} / φ_{kt}):

$$\ln \Phi_{t-1,t}^{TR} = \ln \left(\frac{P_t}{P_{t-1}} \right) = \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) \ln \left(\frac{\varphi_{kt}}{\varphi_{kt-1}} \right), \quad (52)$$

where the derivation of all results in this section is reported in Section A.25 of the web appendix; the weights are the arithmetic means of expenditure shares in the two periods ($(1/2) (s_{kt} + s_{kt-1})$).

In the same way that our unified price index for CES is a generalization of the Sato-Vartia price index to allow for demand shocks for each good, so the translog exact price index ($\ln \Phi_{t-1,t}^{TR}$) in equation (52) is a generalization of the Törnqvist index ($\ln \Phi_{t-1,t}^{TO}$), which corresponds to the special case of equation (52) in which demand is assumed to be constant for all goods ($(\varphi_{kt} / \varphi_{kt-1}) = 1$ for all $k \in \Omega$):

$$\ln \Phi_{t-1,t}^{TO} = \ln \left(\frac{P_t}{P_{t-1}} \right) = \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) \ln \left(\frac{p_{kt}}{p_{kt-1}} \right). \quad (53)$$

Comparing equations (52) and (53), the exact translog price index with time-varying demand shocks ($\ln \Phi_{t-1,t}^{TR}$) differs from the conventional Törnqvist index that assumes time-invariant demand ($\ln \Phi_{t-1,t}^{TO}$) by an additional term that we again refer to as the consumer-valuation bias. Comparing equation (52) for translog with equation (16) for CES, this consumer-valuation bias takes a similar form as for CES, except that the demand shock for each good is weighted by the arithmetic mean of expenditure shares in the two time periods instead of the logarithmic mean of these expenditure shares. The source of the consumer evaluation bias is again that the true exact price index ($\ln \Phi_{t-1,t}^{TR}$) depends on demand-adjusted price changes, whereas the Törnqvist index ($\ln \Phi_{t-1,t}^{TO}$) is based on observed price changes. Therefore, the Törnqvist index does not take into account that an increase in demand for a good is analogous to a reduction in its price. In response to such a fall in the demand-adjusted price for a good, consumers can obtain a higher level of welfare by substituting towards that good and away from other goods. A price index that rules out such demand shocks by assumption cannot capture this substitution in response to changes in demand, in the same way that a Laspeyres index cannot capture substitution in response to changes in price.

As for our CES specification in Section 2, the challenge in implementing the exact price index (52) empirically is that demand-adjusted prices (p_{kt} / φ_{kt}) are not directly observed in the data. Again we overcome this challenge by inverting the demand system to solve for the demand parameters (φ_{kt}) as a function of the observed prices and expenditure shares (p_{kt}, s_{kt}). Applying Shephard's Lemma to the unit expenditure function, and differencing over time, we obtain the following expression for the change in the expenditure share for each product:

$$\Delta s_{kt} = \sum_{\ell \in \Omega} \beta_{k\ell} [\Delta \ln (p_{\ell t}) - \Delta \ln (\theta_{\ell t})], \quad (54)$$

where demand for each good is specified as in equation (2) in the paper with $\varphi_{kt} = \varphi_k \theta_{kt}$. We assume that each good's expenditure share is decreasing in its own demand-adjusted price ($\beta_{kk} < 0$), and increasing in the

demand-adjusted price of other goods ($\beta_{k\ell} > 0$ for $\ell \neq k$), which ensures that this demand system satisfies the “connected substitutes” conditions from Berry, Gandhi and Haile (2013).

We solve for the unobserved demand shocks ($\Delta \ln(\theta_{\ell t})$) by inverting the system of expenditure shares in equation (54), as shown in Section A.25 of the web appendix. The demand system (54) consists of a system of equations for the change in the expenditure shares (Δs_{kt}) of the N goods that is linear in the change in the log price ($\Delta \ln p_{kt}$) and log demand parameter ($\Delta \ln \theta_{kt}$) for each good. These changes in expenditure shares must sum to zero across goods, because the expenditure shares sum to one. Furthermore, under our assumptions of symmetry and homotheticity, the rows and columns of the symmetric matrix formed by the coefficients $\{\beta_{k\ell}\}$ for all pairs of goods must each sum to zero. Therefore, without loss of generality, we can omit the equation for one good. We can nevertheless recover the demand shock for all goods (including the omitted one) using the property that the demand shocks average out across goods ($(1/N) \sum_{k \in N} \Delta \ln \theta_{kt} = 0$), as shown in the web appendix. We thus obtain the unobserved demand shock for each good in terms of observed prices and expenditure shares: $\Delta \ln \varphi_{kt} = \Delta \ln \theta_{kt} = S^{-1}(\Delta \mathbf{s}_t, \Delta \ln \mathbf{p}_t, \{\beta_{kl}\})$.

Substituting for these unobserved demand shocks in equation (52), we obtain the following exact money-metric price index in terms of prices and expenditure shares:

$$\ln \Phi_{t-1,t}^{TR} = \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) S^{-1}(\Delta \mathbf{s}_t, \Delta \ln \mathbf{p}_t, \{\beta_{kl}\}), \quad (55)$$

which corresponds to the analogous common goods price index for translog preferences as our CES common goods price index ($\ln \Phi_{t-1,t}^{CCG}$) in equation (13) above.

Therefore, our main insight that the demand system can be unified with the unit expenditure function to construct a price index that allows for time-varying demand shocks for individual goods and yet remains money metric is not specific to CES, but also holds for the flexible functional form of translog preferences. Furthermore, the consumer-valuation bias is again present for this flexible functional form, because a price index that rules out demand shocks by assumption cannot capture the potential for consumers to increase welfare by substituting towards goods for which increases in demand reduce demand-adjusted prices.

4 Data

Our data source is the Nielsen HomeScan database,¹² which contains sales and purchase quantity data for millions of barcodes bought between 2004 and 2014. Nielsen collects its barcode data by providing handheld scanners to on average 55,000 households a year to scan each good purchased that has a barcode.¹³ Prices are either downloaded from the store in which the good was purchased or hand entered, and the household records any deals used that may affect the price. Barcode data have a number of advantages for the purpose of our analysis. First, product quality does not vary within a barcode, because any change in observable product characteristics results in the introduction of a new barcode. Barcodes are inexpensive to purchase

¹²Our 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. Further information on availability and access to the data is available at <http://research.chicagobooth.edu/nielsen>

¹³The data for 2004 through 2006 come from a sample of 40,000 households, and the data for 2007 through 2014 come from a sample of 60,000 households.

and manufacturers are discouraged from assigning the same barcode to more than one product, because it can create problems for store inventory systems that inform stores about how much of each product is available. Thus, barcodes are typically unique product identifiers and changes in physical attributes (such as product quality) manifest themselves through the creation (and destruction) of barcoded goods, not changes in the characteristics of existing barcoded goods. Thus, a barcode is the closest thing we have empirically to the theoretical concept of a good.

In the raw Nielsen data, some households with particular demographic characteristics are more likely to be sampled by design. In order to construct national or regional expenditure shares and purchase quantities that represent the populations in these regions, Nielsen provides sampling weights that enable us to reweight the data so that the average expenditures and prices are representative of the actual demography in each region rather than the Nielsen sample. Therefore, these households represent a demographically-balanced sample of households in 42 cities in the United States. The set of goods included represents close to the universe of barcoded goods available in grocery, mass-merchandise, and drug stores, representing around a third of all goods categories included in the CPI. For our baseline CES specification, we collapse the household dimension in the data and collapse the weekly purchase frequency to construct a national quarterly database by barcode on the total value sold, total quantity sold, and average price. In a robustness test for our mixed CES specification, we construct national datasets on total value sold, total quantity sold, and average price for high- and low-income households separately. We define low-income households as those with incomes below the median income bracket in our Nielsen data (\$50-59,000 in all but three years) and classify the remaining households as high-income.

Nielsen organizes goods into product groups, which are based on where goods appear in stores. We dropped “magnet data,” which corresponds to products that do not use standard barcodes (e.g., non-branded fruits, vegetables, meats, and in-store baked goods), but kept barcoded goods within these product groups (e.g., Perdue Chicken Breasts, Dole Baby Spinach, etc.). The 5 largest of our 104 product groups are carbonated beverages, pet food, paper products, bread and baked goods, and tobacco. We report a full list of the product groups and summary statistics for each product group in the web appendix. Output units are common within a product group: typically volume, weight, area, length, or counts. Importantly, we deflate by the number of units in the barcode, so prices are expressed in price per unit (e.g., price per ounce). When the units are in counts, we also deflate by the number of goods in a multipack, so for instance, we would measure price per battery for batteries sold in multipacks. Although about two thirds of these barcoded items correspond to food items, the data also contain significant amounts of information about nonfood items like medications, housewares, detergents, and electronics.

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 barcodes 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

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.¹⁴

We resolve this trade-off using a quarterly frequency in our baseline specification (though we find very similar results in a robustness test using an annual frequency). Four-quarter differences were then computed by comparing values for the fourth quarter of each year relative to the fourth quarter of the previous year. In Table 1, we report summary statistics for our sample. For each economic variable of interest, we first compute the average of that variable across years for a given product group, before reporting the mean and standard deviation of this time-average characteristics across product groups, as well as percentiles of its distribution across product groups. As shown in the first row, the median number of price and quantity observations (“Sector Sample Size”) is 44,552, with the sectors in the fifth percentile of observations only having just short of 8,269 data points and those in the 95th percentile having over 142,198 observations. The median number of barcodes per product group is just over 11,000, with 95 percent of these product groups having more than 1,700 unique products, and the largest five percent of them encompassing over 45,000 unique products.

We find substantial entry and exit of products, with the typical life of a barcoded good being only three to four years. On average, 31 percent of all products in a given year exit the sample in the following year, while 32 percent of products sold in a year were not available in the previous year. In comparison, the net growth in the number of barcodes is on average 3 percent across all product groups. These averages mask substantial heterogeneity in innovation rates across product groups, with the average life of a cottage cheese product equal to 5.9 years, whereas the average life of an electronics product is only 1.7 years. High rates of product turnover are reflected in shares of common goods in total expenditure ($\lambda_{t,t-1}$ and $\lambda_{t-1,t}$) of less than one, although these again vary substantially across product groups from a low of 0.34 to a high of 0.99. Consistent with entering products being more numerous or more attractive to consumers than exiting products, we find that common products account for a larger share of expenditure in $t - 1$ than in t (a value of $\lambda_{t,t-1} / \lambda_{t-1,t}$ of less than one). We also report means and standard deviations for the log change in prices ($\Delta \ln p_{kt}$) and expenditure shares ($\Delta \ln s_{kt}$), where these expenditure shares are defined as a share of expenditure within each product group. As apparent from the table, we find that expenditure shares are substantially more variable than prices, which in our model is explained by a combination of elastic demand and demand shocks.

5 Empirical Results

We now present our main empirical results. In Section 5.1, we report our estimates of the elasticity of substitution across barcodes within each of the product groups in our data. In Section 5.2, we use these estimated elasticities of substitution to invert the demand system, and provide evidence on the properties of the resulting demand parameters. In Section 5.3, we show that exact CES price indexes yield similar measures of the change of the cost of living to superlative price indexes under the same assumption of time-invariant

¹⁴Even so, HomeScan data can sometimes contain entry 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 of the good. We also winsorized the data by dropping observations whose percentage change in price or value were in the top or bottom one percent.

Table 1: Product Group Descriptive Statistics

	N	Mean	Sd	Min	P5	P25	P50	P75	P95	Max
Sector Sample Size	104	59,971	50,672	1,999	8,269	24,294	44,552	86,784	142,198	253,668
Number of UPCs	104	15,683	14,852	751	1,706	5,188	11,201	21,711	45,310	79,576
Mean No. Years UPC is in Market	104	3.50	0.97	1.63	2.07	2.91	3.37	4.21	5.33	5.88
Mean $\lambda_{t,t-1}$	104	0.82	0.12	0.34	0.62	0.78	0.85	0.91	0.96	0.97
Mean $\lambda_{t-1,t}$	104	0.91	0.07	0.57	0.75	0.90	0.94	0.96	0.98	0.99
Mean $\frac{\lambda_{t,t-1}}{\lambda_{t-1,t}}$	104	0.90	0.08	0.53	0.73	0.86	0.91	0.95	0.98	0.99
Percent of UPCs that Enter in a Year	104	31.91	10.44	12.98	16.75	24.03	31.51	37.05	50.90	63.27
Percent of UPCs that Exit in a Year	104	30.99	10.09	13.86	16.90	23.47	30.17	36.87	49.54	62.23
Percent Growth Rate in UPCs	104	3.08	14.29	-5.27	-1.48	0.36	1.26	3.40	5.90	145.26
Mean $\Delta \ln p_{kt}$	104	0.01	0.02	-0.08	-0.03	-0.00	0.01	0.02	0.04	0.06
sd($\Delta \ln p_{kt}$)	104	0.21	0.03	0.11	0.17	0.18	0.20	0.22	0.27	0.32
Mean $\Delta \ln s_{kt}$	104	-0.20	0.11	-0.65	-0.42	-0.24	-0.17	-0.11	-0.08	-0.05
sd($\Delta \ln s_{kt}$)	104	1.40	0.11	1.13	1.21	1.33	1.40	1.47	1.59	1.68

Note: Sample pools all households and aggregates to the national level using sampling weights to construct a nationally-representative quarterly database by barcode (UPC) on the total value sold, total quantity sold, and average price; $\lambda_{t,t-1}$ and $\lambda_{t-1,t}$ are the shares of expenditure on common goods in total expenditure in time t and $t - 1$ respectively as defined in equation (A.57) in the web appendix; N is the number of product groups; we compute statistics for each product group as the average value across time periods; mean, standard deviation (sd ($\cdot$)), maximum, minimum and percentiles p5-p95 are based on the distribution of these time-averaged values across product groups. 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.

demand for each good. In Section 5.4, we implement our new exact CES unified price index that allows for time-varying demand shocks for each good. We show that abstracting from these demand shocks leads to a consumer-valuation bias in existing exact CES price indexes that is around as large as the bias from abstracting from the entry and exit of goods. In Section 5.5, we report a number of robustness tests, including our mixed CES specification that allows for more flexible substitution patterns between goods.

5.1 Estimates of the Elasticity of Substitution

We estimate the elasticity of substitution across barcodes for each product group separately. We begin by stacking the moment conditions for all time periods and estimating a single elasticity of substitution for each product group. In Table 2, we report percentiles of the distribution of these estimates across the product groups for the RW estimator, the GRW estimator, the Feenstra (1994) estimates and our bounds. As shown in Column (1), our RW estimates of the elasticity of substitution range from 3.07 at the 5th percentile to 5.66 at the 95th percentile, with a median elasticity of 4.62. These estimated elasticities imply substantially more substitution between barcodes than implied by an elasticity of zero in conventional Laspeyres price indexes or

an elasticity of one implied by a conventional Jevons Index using expenditure share sampling weights. These differences are not only economically large but also statistically significant. In Figure 1, we show the RW estimates (solid black line) and their 95 percent confidence intervals (gray shading) for each product group.¹⁵ We comfortably reject the null hypothesis of an elasticity of substitution of one or zero at conventional levels of statistical significance for all product groups. Therefore, these estimates suggest that the elasticities implicit in conventional price indexes substantially understate the degree to which consumers can substitute between barcodes, confirming the empirical relevance of the well-known substitution bias.


Table 2: Percentiles of the Distribution of Estimated Elasticities of Substitution (σ) Across Product Groups

Percentile	Reverse Weighting (RW)	Generalized Reverse Weighting (GRW)	Feenstra (1994)	Lower Bound	Upper Bound
Min	2.50	4.51	4.39	1.00	10.51
5th	3.07	5.79	5.11	1.00	11.98
25th	3.92	6.86	5.69	1.00	13.48
50th	4.62	7.51	6.48	1.00	14.52
75th	5.00	8.26	7.25	1.00	16.47
95th	5.66	11.77	8.51	1.00	20.20
Max	6.96	13.07	20.86	1.00	21.49

Note: Percentiles of the distribution of estimated elasticities of substitution across product groups; the reverse-weighting (RW) estimator uses the moment conditions in equation (28); the generalized-reverse-weighting (GRW) estimator uses the moment conditions in equation (35); the Feenstra (1994) estimator uses as moment conditions the orthogonality of double-differenced demand and supply shocks; the lower bound is 1; the upper bound is computed as discussed in Section 2.5.4 above; for each estimator, the moment conditions for each pair of time periods are stacked together 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.

¹⁵We 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 1: RW Estimates and 95 percent Confidence Intervals Across Product Groups



Note: Estimated reverse-weighting (RW) elasticities of substitution (black line) and 95 percent point confidence interval (gray shading) for each product group; product groups are ranked by their estimated RW elasticity; confidence intervals based on 50 bootstrap replications. 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.

As shown in Column (2) of Table 2, our GRW estimates are higher for each product group than our RW estimates, ranging from 5.79 at the 5th percentile to 11.77 at the 95th percentile, with a median elasticity of 7.51. From Propositions 4 and 5, the GRW estimator is consistent regardless of the correlation between demand and price shocks, whereas the RW estimator is downward biased when demand and price shocks are positively correlated and upward biased when they are negatively correlated. Therefore, this pattern of lower RW estimates than GRW estimates is consistent with the idea that demand and price shocks are positively correlated with one another. In Figure 2, we show our GRW estimates (black line) and their 95 percent confidence intervals for each product group (gray shading). Consistent with our earlier Monte Carlo results, we find that the GRW estimates are less precisely estimated than the RW estimates, as reflected in the larger confidence intervals in Figure 2 than in Figure 1.

Figure 2: GRW Estimates and 95 percent Confidence Intervals Across Product Groups



Note: Estimated generalized-reverse-weighting (GRW) elasticities of substitution (black line) and 95 percent point confidence interval (gray shading) for each product group; product groups are ranked by their estimated GRW elasticity; confidence intervals based on 50 bootstrap replications. 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.

As a point of comparison, Column (3) of Table 2 reports percentiles of the distribution of the Feenstra (1994) estimates across product groups. This estimator makes different identifying assumptions (orthogonality and heteroskedasticity of price and demand shocks) from the RW estimator (which drops the heteroskedasticity assumption) and the GRW estimator (which drops the orthogonality and heteroskedasticity assumptions). Therefore, in general, we expect the Feenstra (1994), RW and GRW estimators to differ from one another. However, in practice, we find that the three sets of estimates do not differ greatly from one another, with the Feenstra estimator frequently lying in between our RW and GRW estimators. Therefore, the CES demand system imposes sufficient structure on the data that the estimated elasticities under these different identifying assumptions are not greatly different from one another.

As another check on our parameter estimates, the last two columns of Table 2 report our upper and lower bounds for the elasticity of substitution. As discussed above and shown in Column (4), our lower bound for the elasticity of substitution is one, while as shown in Column (5), our upper bound for the elasticity of substitution is typically around 15. In general, our RW and GRW estimates need not necessarily lie within these bounds in any finite sample, because of sampling variation. Furthermore, our three estimators make somewhat different identifying assumptions, which also could explain differences between them. In particular, the GRW and bounds estimators make the additional assumption of joint log normality, whereas the RW estimator does not. Additionally, the RW and GRW estimators use the CES unit expenditure function, whereas the bounds estimator does not. Nevertheless, in practice, we find that RW, GRW, and Feenstra estimators all lie within the bounds for all product groups. Therefore, we find that the results of our different estimators

corroborate one another. More generally, the bounds identify a relatively compact interval of values for the elasticity of substitution consistent with the observed data and the assumptions of CES demand and joint log normality. In our robustness checks below, we show that we find a substantial consumer-valuation bias for all parameter values within this interval.

Finally, to provide evidence on the importance of allowing for demand shocks for individual goods, we now impose the assumption that the demand parameters are time-invariant for each common good ($\varphi_{kt} = \varphi_{k,t-1} = \varphi_k$ for all $k \in \Omega_{t,t-1}$). As discussed in Section 2.4 above, in this special case of no demand shocks, we can directly solve for the elasticity of substitution for each pair of time periods using the Sato-Vartia formula (see equation (A.84) in Section A.5 of the web appendix). If the assumption of no demand shocks is indeed satisfied, we would expect the resulting estimates of the elasticity of substitution to be stable across time periods. To examine the extent to which this is the case, we compute this Sato-Vartia elasticity of substitution (σ_{gt}^{SV}) for each four-quarter difference and product group. We expect these estimates to vary by product group, so we compute the dispersion of these estimates relative to the product group mean, or $\left(\sigma_{gt}^{SV} - \frac{1}{T} \sum_t \sigma_{gt}^{SV}\right)$, where T is the number of periods. In the absence of demand shocks, we expect this number to be zero.

Table 3: Distribution of Elasticities for Each Year and Product Group

	Elasticity		Deviation from Time Mean					
	Mean	Median	SD	p10	p25	p50	p75	p90
Sato-Vartia	-1.12	-2.55	177.10	-51.72	-16.64	-0.29	12.54	34.74
Reverse-Weighting	4.14	4.16	1.07	-1.42	-0.62	0.08	0.74	1.23
Generalized Reverse-Weighting	11.53	7.91	10.07	-8.02	-4.75	-1.62	1.09	9.60

Note: Elasticities are estimated for each product group and pair of time periods; Sato-Vartia elasticity is estimated using equation (A.84) in Section A.5 of the web appendix; mean is the average of these elasticities across product groups and over time ($\frac{1}{GT} \sum_{t,g} \sigma_{gt}$); standard deviation is the average across product groups of the standard deviation over time in these estimated elasticities normalized by their time mean ($\sigma_{gt} - \frac{1}{T} \sum_t \sigma_{gt}$) within each product group; percentiles are based on the distribution across product groups of the standard deviation over time in these normalized elasticities ($\sigma_{gt} - \frac{1}{T} \sum_t \sigma_{gt}$) within each product group. For the Sato-Vartia elasticity only, we exclude the top and bottom one-percent market share changes within each product group to limit the influence of outliers (including these observations results in an even higher standard deviation for the Sato-Vartia elasticity). 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 the top row of Table 3, we report the mean of $\frac{1}{T} \sum_t \sigma_{gt}^{SV}$ in the first column and moments of the distribution of $\left(\sigma_{gt}^{SV} - \frac{1}{T} \sum_t \sigma_{gt}^{SV}\right)$ in the remaining columns. As apparent from the table, we find substantial volatility in these Sato-Vartia elasticities of substitution. The median elasticity of substitution is -2.55, with a standard deviation of 177, and the mean elasticity is also negative. Over half of the elasticities implied by the Sato-Vartia formula have the wrong sign, and the estimates obtained for different years within the same product group vary wildly: half of them are more than 16.6 below the average value for the product group or 12.5 above it. While there are several possible sources for time-varying demand parameters (including consumer tastes, measurement error and specification error), this pattern of results strongly rejects the Sato-Vartia assumption of no demand shocks for each common good.

We now examine the extent to which the elasticity of substitution is stable over time once one allows for time-varying demand for each common good. Our estimates so far pooled pairs of time periods and estimated a single elasticity of substitution (by assuming $\sigma_{gt} = \sigma_g$). We now estimate separate RW and GRW elasticities of substitution for each product group and time period separately. In the remaining rows of Table 3, we report the mean values of these estimates ($\frac{1}{T} \sum_t \sigma_{gt}^{RW}$ and $\frac{1}{T} \sum_t \sigma_{gt}^{GRW}$) across product groups and time periods and the dispersion of these estimates relative to the mean for each product group ($(\sigma_{gt}^{RW} - \frac{1}{T} \sum_t \sigma_{gt}^{RW})$ and $(\sigma_{gt}^{GRW} - \frac{1}{T} \sum_t \sigma_{gt}^{GRW})$). As apparent from the table, both sets of estimates are much more tightly distributed around the product-group mean estimate than the Sato-Vartia elasticities. The median estimates for the RW and GRW elasticities are close to those commonly found in other studies, ranging from 4.1 for the median RW estimate and 7.9 for the median GRW estimate. As expected, the RW estimates are particularly tightly distributed around the product-group mean with 80 percent of the annual estimates lying between -1.4 and 1.2 units larger than the average estimate, but even the GRW estimates have a standard deviation is seventeen times smaller than that implied by the Sato-Vartia formula.

Taking the results of this section together, our RW, GRW and bounds estimates identify a relatively narrow range of possible values for the elasticity of substitution consistent with the observed data. We strongly reject the Sato-Vartia assumption of time-invariant demand for each common good. In contrast, once we allow for time-varying demand shocks for each common good, we find that the data are consistent with a stable underlying elasticity of substitution for each product group.

5.2 Properties of the Demand Shocks

Using our estimates for the elasticity of substitution (σ) for each product group, we can invert the CES demand system to solve for the time-varying demand parameter ($\ln \varphi_{kt}$) for each product, as in equation (11). In Table 4, we examine the properties of these demand parameters, using our GRW estimated elasticities of substitution that allow demand and price shocks to be correlated. In Panel A, we report correlations in levels and changes between (i) sales and price, (ii) sales and demand, and (iii) price and demand. Consistent with the idea that it is more costly to produce products that are more appealing to consumers, we find a strong positive correlation between the log level of prices and demand, which ranges from 0.78 to close to one. Over time, we also find that log changes in prices and demand are positively correlated, which is consistent with our GRW estimates typically lying above our RW estimates.

Table 4: Correlations of Expenditure Shares, Prices and Demand and Variance Decompositions for Expenditure Shares

Panel A: Correlations								
Percentile	Levels			Changes				
	Sales-Price	Sales-Demand	Demand-Price	Sales-Price	Sales-Demand	Demand-Price		
Min	-0.26	0.04	0.78	0.01	0.53	0.48		
5th	-0.15	0.15	0.84	0.04	0.56	0.62		
25th	-0.05	0.29	0.89	0.06	0.73	0.68		
50th	0.02	0.39	0.94	0.08	0.75	0.71		
75th	0.14	0.47	0.98	0.10	0.79	0.74		
95th	0.28	0.67	0.99	0.13	0.84	0.90		
Max	0.49	0.71	1.00	0.20	0.92	0.90		

Panel B: Variance Decomposition								
Percentile	Levels				Changes			
	Variance Price	Variance Demand	Covariance	Variance Share	Variance Price	Variance Demand	Covariance	Variance Share
Min	1.33	2.64	-2.97	1.00	0.20	1.28	-0.48	1.00
5th	2.05	3.57	-4.62	1.00	0.48	1.61	-1.09	1.00
25th	3.76	4.57	-7.34	1.00	0.70	1.86	-1.55	1.00
50th	7.00	7.89	-13.90	1.00	0.87	1.96	-1.83	1.00
75th	17.06	19.33	-35.39	1.00	1.09	2.31	-2.40	1.00
95th	58.56	61.64	-119.20	1.00	3.62	5.14	-7.76	1.00
Max	107.00	110.63	-216.62	1.00	3.79	5.20	-7.99	1.00

Note: The first three columns of Panel A report percentiles across product groups of pairwise correlations between the levels of log sales ($\log(s_{kt}/\bar{s}_t)$), log price ($\log(p_{kt}/\bar{p}_t)$) and log demand ($\log(\varphi_{kt})$) relative to their geometric means; the second three columns of Panel A report analogous results for changes in these variables; the first three columns of Panel B report percentiles across product groups of decompositions of the variance of the level of log sales shares ($\text{var}[\log(s_{kt}/\bar{s}_t)]$) into the contributions of the variance of log prices ($\text{var}[(1-\sigma)\log(p_{kt}/\bar{p}_t)]$), the variance of log demand ($\text{var}[(\sigma-1)\log(\varphi_{kt})]$) and the covariance ($2\text{cov}[(1-\sigma)\log(p_{kt}/\bar{p}_t), (\sigma-1)\log(\varphi_{kt})]$); the second three columns of Panel B report analogous results for changes in these variables. In Panel B, to ensure that the components of the variance decomposition add up to one across the columns, we require that a given row corresponds to the same product group across these columns. To ensure that this is the case, the percentiles in Panel B correspond to product groups at the relevant percentile of the distribution across product groups of the variance of log prices ($\text{var}[(1-\sigma)\log(p_{kt}/\bar{p}_t)]$) divided by the variance of log sales shares ($\text{var}[\log(s_{kt}/\bar{s}_t)]$).

In the lower panel of the table, we report percentiles of a variance decomposition for the level and change of log sales into the contributions of (i) the variance of log prices, (ii) the variance of log demand, and (iii) the covariance between log prices and demand. We implement this decomposition by using the CES demand system in equation (3) to express log sales ($\log(s_{kt}/\bar{s}_t)$) as a function of log price ($(1-\sigma)\log(p_{kt}/\bar{p}_t)$) and log demand ($(\sigma-1)\log(\varphi_{kt}/\bar{\varphi}_t)$) terms. The log price term has a median variance that is 87 percent of the variance of log sales, which implies that prices are quantitatively relevant in accounting for the observed variation in sales. However, the most striking features of the table are the even larger (in magnitude) contributions from the demand variance and the covariance terms, highlighting the role of the time-varying demand shocks in reconciling the choices of the consumer with the observed data on prices and expenditure shares.

Finally, we examine the time-series properties of the demand shocks by running separate regressions for each product group of the log of the demand parameter for each barcode on a barcode fixed effect. Consistent with there being an important time-invariant component of log demand, we find that the average R^2 across all product groups is 0.87, which suggests that much of the variation in demand is time invariant. When we add

a lag of the log demand parameter to the right-hand side, the average R^2 across product groups rises to 0.89, and the coefficient on the lagged level of the demand parameter has an average value of 0.21 across product groups. This pattern of results implies that about 80 percent of any deviation in demand in one year dissipates in the next, which is consistent with reversion to the mean that in part could be the result of measurement error. As discussed above, an important advantage of our exact CES price index in equation (14) is that it allows for mean-zero measurement error in logs for prices and expenditure shares.

5.3 Comparison with Conventional Index Numbers


We now turn to examine the implications of our results for the measurement of changes in the cost of living. In general, there are three reasons why price indexes can differ: differences in the specification of substitution patterns, differences in the treatment of new goods, and differences in assumptions about demand shocks. In the remainder of this section, we show that exact CES price indexes yield similar measures of the change of the cost of living to superlative price indexes under the same assumption of time-invariant demand for each good. Therefore, the differences between our new CES unified price index and existing price indexes in the next section reflect the treatment of entry and exit and demand shocks for surviving goods rather than alternative assumptions about substitution patterns between goods.

For each product group and time period, we compute four conventional price indexes that are discussed in further details in Section A.7 of the web appendix: (i) the Laspeyres index, which assumes a zero elasticity of substitution and weights goods by initial-period expenditure shares; (ii) the Cobb-Douglas index, which assumes an elasticity of substitution of one; (iii) the Fisher index, which is a superlative index that equals the geometric average of the Laspeyres and Paasche indexes, and is exact for quadratic mean of order- r preferences with time-invariant demand parameters; (iv) the Törnqvist index, which is also superlative and is exact for translog preferences with time-invariant demand parameters; and (v) the Sato-Vartia price index, which is exact for CES preferences with time-invariant demand parameters. All these price indexes are defined for common goods that are supplied in both time-periods and hence abstract from entry and exit between time periods. With ten pairs of time periods and 104 product groups, we have a sample of just over 1,000 price changes across products and over time.

In Figure 3, we display kernel density estimates of the distribution of four-quarter price changes across product groups and over time. We express each of the other price indexes 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 the Törnqvist and Sato-Vartia CES price indexes yield almost exactly the same change in the cost of living as the Fisher index, with a difference between them of less than one tenth of a percentage point per year. In contrast, assuming an elasticity of substitution of zero (the Laspeyres index) or one (the Cobb-Douglas index) can result in measures of cost-of-living changes that vary by around a percentage point.

Since the Sato-Vartia CES index is identical to the CES unified price index under the assumption that there are no new goods and no demand shifts for any good, these results suggest that assuming a CES functional form instead of a flexible functional form (as assumed Fisher and Törnqvist price indexes) has relatively little

Figure 3: Differences in Price Indexes from the Fisher Index



Note: Kernel densities of the distribution across product groups and over time of the difference between price indexes and the Fisher Index. Price indexes are measured as proportional four-quarter changes for each product group for the set of common goods ($(P_{gt}^* - P_{gt-1}^*) / P_{gt-1}^*$). SV-CES is the Sato-Vartia price index (equation (10)). 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.

impact on the measured change in the cost of living under a common set of assumptions of no entry and exit and no demand shocks for common goods.


5.4 The CES Unified Price Index

We now maintain the assumption of CES preferences but allow for the entry and exit of goods and demand shocks for individual common goods. We show that abstracting from these two features of the data introduces a substantial bias into measures of the change in the standard of living. We find that our new consumer-valuation bias is as around as large as the bias from abstracting from the entry and exit of goods and equal to more than a percentage point per year.

We start with the variety adjustment term that captures the impact of entry and exit and was first introduced by Feenstra (1994). This term depends on both the elasticity of substitution (σ_g) and relative expenditure shares on common goods ($\lambda_{gt,t-1} / \lambda_{g,t-1,t}$). It controls both the difference between the Feenstra and Sato-Vartia price indexes in equation (15) and the difference between the CUPI in equation (14) and the CCG in equation (13). In Figure 4, we display a histogram of the relative expenditure shares on common goods ($\lambda_{gt,t-1} / \lambda_{g,t-1,t}$) across product groups and over time. If entering barcodes had similar characteristics to exiting barcodes, the prices and market shares of exiting barcodes would match those of new products, resulting in a $\lambda_{gt,t-1} / \lambda_{g,t-1,t}$ ratio of one. The fact that these ratios are almost always less than one indicates that new goods tend to be more attractive than disappearing ones in terms of having lower demand-adjusted (p_{kt} / φ_{kt}). Moreover, while the occasional λ -ratio in excess of unity indicates that one sometimes observes a negative new-good bias for a particular product group in a given year, these λ -ratios are less than one for every product group over the full set of years (as shown in the web appendix). In other words, there is per-

vasive product upgrading over time. In barcode data, this product upgrading is fully captured in the entry and exit term, because as discussed above any change in the physical characteristics of a good leads to the introduction of a new barcode.

Figure 4: Shares of Common Goods in Expenditure in period t relative to period $t - 1$ ($\lambda_{gt,t-1} / \lambda_{gt-1,t}$), Four-Quarter Differences by Product Group




Note: Histogram of relative expenditure shares on common goods ($\lambda_{gt,t-1} / \lambda_{gt-1,t}$, as defined in equation (A.57) of the web appendix for each product group g) across product groups and over time. Time periods are four-quarter differences. 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.

We now quantify the relative importance of the biases from abstracting from entry and exit and demand shocks for common goods. We compare our CES unified price index (CUPI) that incorporates both of these features of the data to existing price indexes that abstract from one or more of these sources of bias in the measurement of changes in the cost of living. For each product group and time period, we compute alternative measures of changes in the cost of living, and then aggregate across product groups using expenditure-share weights to compute a measure of the change in the aggregate cost of living.

In Figure 5, we plot the resulting measures of the change in the aggregate cost of living using our CUPI and a range of alternative price indexes. It is well-known that conventional indexes—Fisher, Törnqvist and Sato-Vartia (CES)—are bounded by the Paasche and Laspeyres indexes. Thus, we can think of conventional indexes as giving us a band of cost-of-living changes that is determined by assumptions about consumer substitution patterns, under the assumption of no entry and exit and no shifts in demand for any common good. Consistent with our results in the previous section, we find a relatively small gap between the Laspeyres and Paasche price indexes, implying that different assumptions about substitution patterns have a relatively minor impact on the measurement of the cost of living.

Figure 5: Four-Quarter Proportional Changes in the Aggregate Cost of Living ($(P_t - P_{t-1}) / P_{t-1}$)



Note: Proportional change in the aggregate cost of living is computed by weighting the four-quarter proportional change in the cost of living for each of the product groups in our data ($(P_{gt} - P_{gt-1}) / P_{gt-1}$) by their expenditure shares. CCG stands for CES common-good price index (equation (13)). CUPI stands for CES unified price index (equation (14)). The suffixes RW and GRW denote whether the elasticity of substitution was estimated using the reverse-weighting or generalized reverse weighting procedure. Both the CUPI and Feenstra CES correct for the entry and exit of varieties, but the CUPI uses our CES common goods price index (equation (14)), whereas Feenstra-CES uses the Sato-Vartia price index for common goods (equation (15)). 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 bias from abstracting from the entry and exit of goods can be seen in Figure 5 from comparing either the Feenstra and Sato-Vartia price indexes (from equation (15)) or our CUPI and the CCG price indexes (from equations (13) and (14)). We compute this comparison using both our baseline RW estimates and our generalized GRW estimates of the elasticity of substitution between barcodes within each product group. In either case, we find a substantial impact of entry and exit on the measurement of the cost of living, equal to more than one percentage point per year. This bias falls with the estimated value with the elasticity of substitution, because the superior characteristics of entering goods relative to exiting goods result in a larger reduction in the cost of living when goods are less substitutable. Consequently, this bias is somewhat lower using our GRW estimates than using our RW estimates, but remains substantial. Therefore, if one abstracts from the fact that new goods tend to be systematically better than disappearing goods (as measured in the CES demand system by their relative expenditure shares), one systematically overstates the increase in the cost of living over time.

The consumer-valuation bias from neglecting demand shocks for individual common goods can be discerned in Figure 5 from comparing our CUPI and the Feenstra price index (from equations (15), (13) and (14)). Both of these price indexes are exact for CES preferences and allow for the entry and exit of goods. However, the Feenstra price index uses the Sato-Vartia price index for common goods (assuming time-invariant demand), whereas the CUPI uses the CCG price index for common goods (allowing for time-varying demand shocks).¹⁶ As shown in the figure, we find that this bias is around as large as that from abstracting from

¹⁶Another way of seeing the bias from abstracting from demand shocks is to consider the price indexes for common goods and compare the CCG (which allows for demand shocks under our assumption of CES preferences) to the band defined by the Laspeyres and Paasche price indexes (which provide upper and lower bounds for the common goods price index for different functional form

entry and exit, and more than one percentage point per annum.¹⁷ This bias again falls with the elasticity of substitution, because the ability to substitute towards goods for which demand has risen results in a larger reduction in the cost of living when goods are more differentiated. As a result, we find a somewhat smaller bias using our GRW estimates than using our RW estimates, but in both cases it is substantial. This bias arises because consumer welfare depends on demand-adjusted prices, such that increases in demand are directly analogous to falls in price for a good. Therefore, if one abstracts from the fact that consumers can substitute towards goods for which demand has risen, and away from goods for which demand has fallen, one again systematically overstates the increase in the cost of living over time.

Taken together, these results suggest that conventional price indexes are subject to a substantial upward bias that results from abstracting from two first-order features of the data: entry/exit and demand shocks for individual goods.

5.5 Robustness

We now report a number of robustness checks on our baseline CES specification, including the mixed CES specification discussed in Section 3.3, the use of alternative values for the elasticity of substitution, a comparison with official CPI categories, sensitivity to measurement error for goods with small expenditure shares, and the use of alternative weights across common goods to the uniform weights used in the UPI. Across each of these robustness checks, we show that we continue to find a substantial consumer-valuation bias.

5.5.1 High- and Low-Income Households

We begin by implementing the mixed CES specification with heterogeneous groups of consumers discussed in Section 3.3 above. We use low-income and high-income households as our groups, as defined in the data section above. Although the differences in income between these groups are substantial, they are of course smaller than in other settings, such as in settings comparing developed and developing countries. We allow both the elasticity of substitution and the demand parameter for each good to differ between the two groups. Therefore, this specification incorporates non-homotheticities in a more flexible way than the non-homothetic CES specification in Section 3.1 above, which imposed a common elasticity of substitution for all consumers.


In Figure 6, we report our estimates of the elasticities of substitution for high- and low-income households using both the RW and GRW estimators. The black lines correspond to the RW estimates and the gray lines to the GRW ones. As apparent from the figure, we find that for any given estimation procedure, the estimated elasticities for the pooled, high-income, and low-income samples are quite similar, which suggests that high- and low-income households have similar elasticities of substitution. Unsurprisingly, the GRW estimates are somewhat noisier than the RW estimates, but nevertheless the correlations between elasticities for the two groups of households are strong: 0.9 using the RW estimator and 0.6 using the GRW estimator. Therefore, although this mixed CES specification allows for heterogeneity in the elasticity of substitution, we

assumptions about substitution possibilities including CES preferences).

¹⁷The average value from 2005 to 2014 of the Paasche index is 1.1 percent; the Laspeyres, 2.0; the CCG-GRW is -0.2 percent; the CCG-RW is -1.2 percent; the CUPI-GRW is -1.8 percent; the CUPI-RW is 4.5 percent; the Feenstra-CES-RW is -1.8 percent; and the Feenstra-CES-GRW is -0.1 percent.

find similar substitution behavior for the high- and low-income households in our data on barcoded goods.

Figure 6: Estimated Elasticities for High- and Low-Income Households




Note: Estimated elasticities of substitution for each product group. Product groups are ranked by the reverse-weighting (RW) estimated elasticity for our baseline sample (including both high- and low-income households). GRW denotes generalized-reverse-weighting estimator. High- and low-income households are defined as those with incomes above or below the median income in our Nielsen data. 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 Figure 7, we plot the log demand shifters ($\ln \phi_{kt}$) for each group of households against the average of those for the two groups for each barcode. We use our RW estimates and pool observations across product groups, where the demand shifters for each product group are normalized to have a mean in logs of zero. We use a bin scatter with 100 percentiles and also display the regression relationships between the variables. We find a strong positive and statistically significant correlation between the tastes parameters for the two groups of households of 0.89, which is reflected in the figure in both regression lines lying close to the diagonal. Therefore, on average, we find strong agreement between high- and low-income households about which products are more or less appealing.

Another feature of Figure 7 is that the slope for low-income households lies below that for high-income households. This result suggests that high-income households tend to value more appealing barcodes relatively more than low-income households. If average rates of price increase differ between the goods preferred by high- and low-income households, this can induce differences in the inflation rate for the two groups. These differences were the main focus of Jaravel (2017), which showed that the average change in the cost of living for high-income households exceeds that for low-income households by 0.65 percent per year for common goods and by 0.78 percent per year once the entry and exit of goods is taken into account. We find the same pattern of differences in the cost of living between the two groups, as shown in Figure 8 using the RW estimator. On average, the CCG and the CUPI price indexes for low-income households are 0.51 and 0.77 percent per year higher than those for high-income households. Therefore, our index captures the same properties of the data as found in other studies.


We now examine the magnitude of the consumer-valuation bias for the two groups of households. As

Figure 7: Demand Parameters for High- and Low-Income Consumers



Note: Regression lines and bin scatters of the estimated log demand parameter ($\log \phi_{kt}$) for each barcode and time period for high- and low-income households (vertical axis) against the corresponding estimate for our baseline sample including all households (horizontal axis). Both sets of demand parameters are computed using our RW estimates. Log demand parameters have a mean of zero for each product group and time period. Time periods are four-quarter differences. 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 8: Four-Quarter Proportional Changes in the Aggregate Cost of Living ($(P_t - P_{t-1}) / P_{t-1}$), All Households and High- and Low-Income Households



Note: Proportional change in the aggregate cost of living is computed by weighting the four-quarter proportional change in the cost of living for each of the product groups in our data ($(P_{gt} - P_{gt-1}) / P_{gt-1}$) by their expenditure shares. CCG is our CES common goods price index (equation (13)); CUPI is our CES unified price index (equation (14)); both are computed using our reverse-weighting (RW) estimator and our baseline sample of all households. High- and low-income versions of these indexes were computed using only price and expenditure data for households with above and below the median household income respectively. Change in the aggregate cost of living is computed by weighting the change in the cost of living for each of the product groups in our data by their expenditure shares. 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.

evident from Figure 8, most of the variance in annual changes in the cost of living is due to price changes that affect high- and low-income households similarly. The variance in the difference in the cost-of-living between the two groups is less than one fifth as large as the variance in the change in the cost of living measured on average for each year. Although our overall price indexes from the previous section need not always lie in between those for two groups separately, because they are not weighted averages of those for the two groups separately, we find that they do tend to fall in between them in practice. Over the full sample period, the CCG rose by 0.29 percent per year on average, which was between the CCG for low-income households (0.31 percent per year) and that for high-income households (0.21 percent per year). We see a similar pattern for the CUPI, which fell by 2.6 percent on average, which was between the 2.5 and 3.2 percent per year declines for high- and low-income households. Thus, the change in the cost of living that we obtain from our baseline specification on average lies between those of our two groups.

Taken together these results suggest that while we can find evidence of heterogeneity in the demand parameters for individual goods between high- and low-income households, we find similar elasticities of substitution across goods for these two groups, and this heterogeneity in demand parameters essentially shifts the cost of living for each group of households separately up or down around our central estimate.

5.5.2 Other Robustness Checks

We also considered a number of other robustness checks. First, we examine the sensitivity of our consumer-valuation bias to the value of the elasticity of substitution by undertaking a grid search over alternative values for this elasticity. From Table 2 above, our lower bound for the elasticity of substitution is one, and our upper bound ranges from around 10 to 20 across product groups. Therefore, we consider a grid of thirty-eight evenly spaced values for this parameter ranging from 1.5 to 20. For each parameter value on this grid, we first compute our CCG and CUPI for each product group and year. We next compute an overall measure of the cost of living by aggregating across product groups using expenditure-share weights.

We present a plot of these results in Figure A.5 in Section A.27.1 of the web appendix. Interestingly, despite the fact that the Laspeyres and Fisher indexes registered average changes of 2.0 and 1.6 percent per year over this time period, the aggregate change in the cost of living measured by the CUPI is negative for all values of the elasticity below 20. As one should expect, the change in the cost of living captured by the CCG tends to be lower when the elasticity of substitution is small, because demand shifts matter more for welfare if goods are less substitutable. Similarly, smaller values of the elasticity of substitution are associated with greater gains from variety (and therefore a lower cost of living) because a low elasticity means that new varieties are considered more differentiated and hence more valuable to households. The results suggest that both the CCG and the CUPI register substantial price falls over the full time period when the elasticity of substitution is very small (e.g., less than 3), but the differences in average changes in the cost of living measured by the CUPI vary by only 2.4 percentage points per year if we restrict ourselves to the range of median elasticities we found in Table 2 (4.5 to 7.5).

Second, as an illustration of the relevance of our results for official measures of the consumer price index (CPI), we compare conventional price indexes computed using the Nielsen data to official CPI price indexes.

As discussed further in Section A.27.2 of the web appendix, we were able to map 89 of our 104 product groups into CPI categories. We again aggregate across these price sub-indexes for each of the 89 product groups using expenditure-share weights to construct a measure of the overall change in the cost of living. In Figure A.6 of the web appendix, we compare the resulting aggregate price indexes using the Nielsen data and the official CPI sub-indexes.

We find a strong positive and statistically significant correlation of 0.98 between the Laspeyres (based on Nielsen data) and the CPI measures of the change in the overall cost of living. Moreover, the average changes in the cost of living as measured by the Laspeyres index and the CPI are almost identical: 2.30 versus 2.37 percent respectively. The Paasche index (based on Nielsen data) has the same correlation with the CPI, but has an average change that is only 1.5 percent per year. In other words, annual movements in changes in the cost of living as measured by the BLS for this set of goods can be closely approximated by using a Laspeyres index and the Nielsen data, and the difference between the Laspeyres and the Paasche indexes in the Nielsen data is less than one percentage point per year (consistent with the findings of the Boskin Commission in Boskin et al. 1996). In contrast, we find a substantial bias from abstracting from entry/exit and changes in demand for surviving goods, with our CUPi-RW and CUPi-GRW registering average changes in the cost of living that are more than two percentage points below the CPI.

Third, we examined the sensitivity of our results to measurement error in expenditure shares for goods that account for small shares of expenditure. In particular, we implemented the robustness test discussed in Subsection 2.3 above, in which we rewrite the price index for common goods as the product of two terms: the price index for the subset of common goods with above-median expenditure shares and the share of expenditure on this subset in all expenditure on common goods. This robustness specification is less sensitive to measurement error for goods that account for small shares of expenditure, because expenditures on goods with below-median expenditure shares only enter the denominator of the second term, and even then only matter through total expenditure on goods with below-median shares. As reported in Section A.27.3 of the web appendix, the resulting measure of the change in the cost of living tracks closely that in our baseline specification above.

Fourth, we explored the sensitivity of our results to the use of alternative weights for common goods from the uniform weights used in the CUPi. Both our weak law of large numbers argument for the demand shocks averaging out across common goods and our RW estimator require that these weights are uncorrelated with demand shocks. Therefore, we experimented with constructing price indexes based on random weights that sum to one and are orthogonal to demand shocks by construction. To do this, we first assigned each good a random number based on a draw from a uniform distribution. We converted these numbers into weights that summed to one by dividing them by the sum of all of these numbers in each year and product-group combination and then computed the average change in the CCG over all years using these random weights. We repeated this exercise 100 times. The median of these randomly weighted CCGs differed by less than 0.002 percentage points from the CCG and the 5th and 95th percentiles differed by only 0.1 percentage points from the CCG. Thus, we find little difference in CUPis computed based on unweighted averages of products and those based on randomly weighted averages of products.

6 Conclusions

Measuring price aggregates is central to international trade and macroeconomics, which depend critically on being able to distinguish real and nominal income. We make three main contributions to aggregate price measurement with the constant elasticity of substitution (CES) preferences that dominate these fields. First, we develop a new exact price index for CES preferences that treats the demand system and unit expenditure function in a unified way. We incorporate demand shocks for individual goods (so that our model is consistent with the observed price and expenditure share data), while ensuring that these demand shocks average out across common goods (so that the change in the cost of living is money-metric in the sense that it depends only on prices and expenditure shares). The key insight behind our approach is to invert the demand system to express the unobserved demand shocks in terms of observed changes in prices and expenditure shares. Our approach is valid under the same set of assumptions as the existing Sato-Vartia price index (no demand shocks for each common good), but it is also valid under a much weaker set of assumptions (demand shocks for individual common goods that average out across these goods).

Second, we show that abstracting from demand shocks for individual goods introduces a substantial bias into conventional price indexes that we term the “consumer-valuation bias.” This bias is analogous to the well-known substitution bias, and arises because consumer welfare depends on demand-adjusted prices, whereas only unadjusted prices are observed in the data. If consumer preferences allow substitution towards goods for which demand has increased, but a researcher assumes a price index that rules out such substitution, this price index overestimates the increase in the cost of living. The researcher fails to take into account that an increase in the relative demand for a good is analogous to a reduction in its relative price and consumers can obtain higher utility by substituting towards that good and away from other goods. Empirically, we find this consumer-valuation bias to be substantial, equal to more than one percentage point per annum, and around the same magnitude as the bias from abstracting from the entry and exit of goods.

These first two contributions hold regardless of how the elasticity of substitution between goods is estimated. Our third main contribution is to develop a new estimator of this parameter that uses the identifying assumption of money-metric utility. Our baseline “reverse-weighting” estimator minimizes the difference between the implied change in the cost of living using the tastes of the initial or final period. Our “generalized reverse-weighting” estimator minimizes the difference between the change in the cost of living using initial or final period tastes, after controlling for the component of demand shocks that is correlated with price shocks. We provide conditions under which these estimators are consistent and show that they perform well in finite samples using Monte Carlo. Finally, we use our inversion of the demand system to provide upper and lower bounds to the elasticity of substitution that hold regardless of the correlation between demand and price shocks. We show that there is a relatively narrow range of possible values for the elasticity of substitution that is consistent with the observed data and the assumption of a CES demand system.

Although we focus for most of our analysis on CES preferences because of their dominance in international trade and macroeconomics, we also consider a number of extensions and generalizations of our approach, including non-homothetic CES (indirectly additive), nested CES, mixed CES, logit and translog

preferences. As long as preferences satisfy the property of connected substitutes, the demand system can be inverted to express the unobserved demand shocks in terms of observed changes in prices and expenditure shares. As long as consumers can substitute towards good for which demand has risen, conventional price indexes that abstract from such demand shocks are subject to the consumer-valuation bias. In a robustness check using mixed CES, which relaxes the independence of irrelevant alternatives and symmetric substitution assumptions of CES, we show that this consumer-valuation bias remains substantial and is again around as large as the bias from abstracting from entry and exit.

References

- [1] Adao, Rodrigo, Arnaud Costinot and Dave Donaldson (2017) “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade,” *American Economic Review*, 107(3), 633-689.
- [2] Anderson, James and Eric van Wincoop (2003) “Gravity with Gravitas: a Solution to the Border Puzzle,” *American Economic Review*, 93(1), 170-192.
- [3] Anderson, Simon, André de Palma and Jacques-François Thisse (1992) *Discrete Choice Theory of Product Differentiation*, MIT Press.
- [4] Antràs Pol (2003) “Firms, Contracts, and Trade Structure,” *Quarterly Journal of Economics*, 118(4), 1375-1418.
- [5] Arkolakis, Costas and Arnaud Costinot and Andrés Rodríguez-Clare (2012) “New Trade Models, Same Old Gains,” *American Economic Review*, 102(1), 94-130.
- [6] Armington, Paul S. (1969) “A Theory of Demand for Products Distinguished by Place of Production,” IMF Staff Papers, 16, 159-178.
- [7] Atkin, David, and Dave Donaldson (2015) “Who’s Getting Globalized? The Size and Implications of Intra-national Trade Costs,” *NBER Working Paper*, 21439.
- [8] Atkin, David, Benjamin Faber, and Marco Gonzales-Navarro (2016) “Retail Globalization and Household Welfare: Evidence from Mexico,” *Journal of Political Economy*, forthcoming.
- [9] Balk, Bert M. (1999) “On Curing the CPI’s Substitution and New Goods Bias,” Department of Statistical Methods, Statistics Netherlands, mimeograph.
- [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, James Levinsohn and Ariel Pakes (1995) “Automobile Prices in Market Equilibrium,” *Econometrica*, 60(4), 889-917.
- [15] Berry, Steven, Amit Gandhi, and Philip Hale (2013) “Connected Substitutes and Invertibility of Demand,” *Econometrica*, 81(5), 2087-2111.

- [16] Białek, Jacek (2017) “Approximation of the Fisher Price Index by Using the Lloyd–Moulton Index: Simulation Study,” *Communications in Statistics - Simulation and Computation*, 46(5), 3588-3598.
- [17] Bils, Mark and Klenow, Peter (2001) “Quantifying Quality Growth,” *American Economic Review*, 91, 1006-1030.
- [18] Blanchard, Olivier Jean and Nobuhiro Kiyotaki (1987) “Monopolistic Competition and the Effects of Aggregate Demand,” *American Economic Review*, 77(4), 647-666.
- [19] Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches and Dale W. Jorgenson (1996) “Toward a More Accurate Measure of the Cost of Living,” Final Report to the Senate Finance Committee. Washington, DC: Senate Finance Committee.
- [20] Broda, Christian and David E. Weinstein (2006) “Globalization and the Gains from Variety,” *Quarterly Journal of Economics*, 121(2), 541-86.
- [21] Broda, Christian and David E. Weinstein (2010) “Product Creation and Destruction: Evidence and Price Implications,” *American Economic Review*, 100, 691–723.
- [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] Comin, Diego A., Danial Lashkari and Martí Mestieri (2015) “Structural Change with Long-run Income and Price Effects,” *NBER Working Paper*, 21595.
- [24] Diewert, W. E. (1976) “Exact and Superlative Index Numbers,” *Journal of Econometrics*, 4, 115-145.
- [25] Diewert, Erwin (2004) *Consumer Price Index Manual: Theory and Practice*, Washington DC: International Monetary Fund.
- [26] Diewert, Erwin and Robert Feenstra (2017) “Estimating the Benefits and Costs of New and Disappearing Products,” University of California at Davis, mimeograph.
- [27] Dixit, Avinash and Joseph Stiglitz (1977) “Monopolistic Competition and Optimum Product Diversity,” *American Economic Review*, 297-308.
- [28] Eaton, Jonathan and Kortum, Samuel (2002) “Technology, Geography, and Trade,” *Econometrica*, 70(5), 1741-1779.
- [29] Fajgelbaum, Pablo D. and Amit K. Khandelwal (2016) “Measuring the Unequal Gains from Trade,” *Quarterly Journal of Economics*, 1113-1180.
- [30] Fally, Thibault and Benjamin Faber (2016) “Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data,” mimeo

- [31] Feenstra, Robert C. (1994) "New Product Varieties and the Measurement of International Prices," *American Economic Review*, 84(1), 157-177.
- [32] 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.
- [33] Feenstra, Robert C. and Marshall B. Reinsdorf (2010) "An Exact Price Index for the Almost Ideal Demand System," *Economics Letters*, 66(2) 159-162.
- [34] 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.
- [35] Handbury, Jessie (2013) "Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities," Wharton mimeo.
- [36] Hanoch, Giora (1975) "Production and Demand Models with Direct or Indirect Implicit Additivity," *Econometrica*, 43(3), 395-419.
- [37] Helpman Elhanan, Marc J. Melitz and Stephen R. Yeaple (2004) "Export Versus FDI with Heterogeneous Firms," *American Economic Review*, 94(1), 300-316.
- [38] Hottman, Colin, Stephen Redding and David Weinstein (2016) "Quantifying the Sources of Firm Heterogeneity," *Quarterly Journal of Economics*, 131(3), 1291-1364.
- [39] Hsieh, Chang-Tai and Peter Klenow (2009) "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403-1448.
- [40] Jaravel, Xavier (2017) "The Unequal Gains from Product Innovation: Evidence from the US Retail Sector," London School of Economics, mimeograph.
- [41] Jones, Charles I. and Peter J. Klenow (2016) "Beyond GDP? Welfare Across Countries and Time," *American Economic Review*, forthcoming.
- [42] Klepper, Steven and Edward E. Leamer (1984) "Consistent Sets of Estimates for Regressions with Errors in All Variables," *Econometrica*, 52(1), 163-184.
- [43] Konüs, A. (1924) "The Problem of the True Index of the Cost of Living," translated in *Econometrica* (1939), 7, 10-29.
- [44] Krugman, Paul R. (1980) "Scale Economies, Product Differentiation and the Pattern of Trade," *American Economic Review*, 70, 950-959.
- [45] Krugman, Paul R. (1991) "Increasing Returns and Economic Geography," *Journal of Political Economy*, 99(3), 483-499.

- [46] Krugman, Paul and Anthony J. Venables (1995) "Globalization and the Inequality of Nations," *Quarterly Journal of Economics*, 857-880.
- [47] Lau, L. J. (1979) "On Exact Index Numbers," *Review of Economics and Statistics*, 61, 73-82.
- [48] Leamer, Edward E. (1981) "Is it a Demand Curve, or is it a Supply Curve? Partial Identification Through Inequality Constraints," *Review of Economics and Statistics*, 3, 319-327.
- [49] Lloyd, P. J. (1975) "Substitution Effects and Biases in Non-True Price Indices," *American Economic Review*, 65(3), 301-313.
- [50] McFadden, Daniel (1974) "Conditional Logit Analysis of Qualitative Choice Behavior," in (ed) Paul Zarembka, *Frontiers in Econometrics*, New York: Academic Press.
- [51] McFadden, Daniel and Kenneth Train (2000) "Mixed MNL Models of Discrete Response," *Journal of Applied Econometrics*, 15, 447-470.
- [52] Melitz, Marc J. (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695-725.
- [53] Moulton, B. R. (1996) "Constant Elasticity Cost-of-Living Index in Share- Relative Form," Washington DC: U. S. Bureau of Labor Statistics, mimeograph.
- [54] 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.
- [55] Newey, W. K. and D. McFadden (1994) "Large Sample Estimation and Hypothesis Testing," in (eds) R. F. Engle and D. McFadden, *Handbook of Econometrics*, Vol. 4, Amsterdam: North Holland, 2111-2245.
- [56] Nevo, Aviv (2003) "New Products, Quality Changes, and Welfare Measures Computed from Estimated Demand System," *Review of Economics and Statistics*, 85(2), 266-275.
- [57] Sato, R (1975) "The Most General Class of CES Functions," *Econometrica*, 43, 999-1003.
- [58] Sato, Kazuo (1976) "The Ideal Log-Change Index Number," *Review of Economics and Statistics*, 58(2), 223-228.
- [59] 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.
- [60] Syverson, Chad (2016) "Challenges to Mismeasurement Explanations for the U.S. Productivity Slow-down," *NBER Working Papers*, 21974.
- [61] Train, Kenneth E. (2009) *Discrete Choice Methods with Simulation*, Cambridge: Cambridge University Press.
- [62] Vartia, Yrjo (1976) "Ideal Log-Change Index Numbers," *Scandinavian Journal of Statistics*, 3(3), 121-126.

- [63] Wooldridge, Jeffrey M. (2002) *Econometric Analysis of Cross-Section and Panel Data*, Cambridge MA: MIT Press.