

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

Stephen J. Redding[†] David E. Weinstein

October 15, 2019

Abstract

We develop an approach to measuring the cost of living for CES preferences that treats demand shocks as taste shocks that are equivalent to price shocks. In the presence of relative taste shocks, the Sato-Vartia price index is upward biased because an increase in the relative consumer taste for a variety lowers its taste-adjusted price and raises its expenditure share. By failing to allow for this association, the Sato-Vartia index underweights drops in taste-adjusted prices and overweights increases in taste-adjusted prices, leading to what we term a “taste-shock bias.” We show that this bias generalizes to other invertible demand systems.

JEL CLASSIFICATION: D11, D12, E01, E31

KEYWORDS: price index, taste-shock bias, variety bias

*We thank Columbia University and Princeton University for research support. 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 Bamba, Anna Blender, Molly Borden, Mathis Maehlum, Anders Nielsen, Rafael Parente, Dyanne Vaught and Yingjie (Angela) Wu for excellent research assistance. We also want to thank Young and Rubicam for sharing their brand asset value data. Researcher(s) own analyses calculated (or derived) based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researcher(s) and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein. Responsibility for the results, opinions, and any errors lies with the authors alone.

[†]Corresponding author: Department of Economics and WWS, Julis Romo Rabinowitz Building, Princeton, NJ 08544. Tel: 1 609 258 4016. Fax: 1 609 258 6419. Email: reddings@princeton.edu.

I. Introduction

Measuring welfare changes is a fundamental issue in economics, which arises in a number of contexts, not least in the measurement of changes in the cost of living.¹ This problem arises frequently in macroeconomics and international trade, where constant elasticity of substitution (CES) preferences are commonly used to compute the change in the cost of living between a pair of time periods. A challenge in such an analysis is that one's assumed preference structure does not perfectly fit the data in both time periods without the inclusion of a residual, which is necessarily directly related to observed expenditure shares, and can reflect systematic unobserved changes in tastes. In response to this challenge, one can either take the position that all demand (taste) shocks are equivalent to price shocks and compute the change in the cost of living by adjusting the observed price shocks using the time-varying demand residuals. Or one can take the position of holding the taste parameters constant (e.g. at their initial values) and compute the change in the cost of living using the observed price shocks and ignoring changes in the demand residuals.

Though both interpretations are valid, a key building block for the existing exact price index for CES preferences (the Sato-Vartia price index for varieties common to a pair of time periods) is inconsistent with either interpretation.² Under the assumption of time-invariant tastes, this price index uses observed expenditure shares and changes in prices to compute changes in the cost of living. However, the final-period expenditure shares used in the weights for each variety in this price index include time-varying demand residuals, which gives rise to an internal inconsistency. If these demand residuals are not taste shocks, they are correctly excluded from the measured changes in prices, but incorrectly included in the expenditure share weights. If these demand residuals are taste shocks, they are correctly included in the expenditure share weights, but incorrectly excluded from the measured changes in prices. The treatment of these demand residuals becomes particularly important in a world in which consumer tastes can move substantially (with fads, fashions and other fluctuations in tastes) and exhibit systematic patterns in the data.

In this paper, we develop a new exact price index for CES preferences that consistently treats demand shocks as taste shocks that are equivalent to price shocks. We propose computationally feasible measures of these taste shocks and an intuitive correction to existing measures of the change in the cost of living. An exact price index (or equivalently a money-metric utility function) measures the change in the cost of living solely in terms of observed *prices* and *expenditures*. Therefore, the existence of such an exact price index requires that we rule out the possibility of a change in the cost of living when all prices and expenditures remain unchanged. As expenditures depend on *relative* consumer tastes (and not the absolute level of these consumer tastes), an implication is that an exact price index rules out an equiproportional change in consumer tastes. We therefore require a normalization or constant choice of units in which to measure consumer tastes. If there exists good for which a researcher is confident that consumer tastes did not change, the assumption of a constant taste parameter for that good is one such normalization. In the absence of an obvious candidate for such a good

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 (2017).

²The existing exact price index for CES preferences combines the Feenstra (1994) entry and exit correction (which allows for time-varying taste parameters for varieties that are not supplied in both periods) with the Sato (1976) and Vartia (1976) price index for varieties that are common to both periods (which assumes constant taste parameters for each common variety).

in our barcode data, we assume that the geometric mean of the consumer taste parameters is constant. This normalization is consistent with the assumption that the log demand shocks are mean zero in the estimation of the demand system. But as other normalizations for consumer tastes are possible, we examine the robustness of our results to alternative normalizations using the class of generalized means.

Our approach uses the invertibility of the CES demand system to recover unique values for unobserved consumer tastes for each variety (up to this normalization). We use this result to derive an exact price index for the change in the cost of living in terms of only prices and expenditure shares, while allowing for changes in relative consumer tastes across varieties. We show that our estimated taste shifts are not simply measurement or specification errors, because they are strongly related to separate measures of brand asset values from a marketing firm using a completely different methodology: a survey of consumer preferences. A limitation is that we focus for most of our analysis on homothetic CES preferences, because of their prominence in international trade, macroeconomics and economic geography. Nonetheless, we show that our approach can be implemented for other invertible demand systems, including non-homothetic CES (indirectly additive), nested CES, mixed CES, logit, mixed logit, translog and almost ideal demand system (AIDS) preferences.

We show that the inconsistent treatment of the time-varying demand residuals in the Sato-Vartia index introduces a bias that we term the “taste-shock bias.” In particular, through its use of observed rather than taste-adjusted prices, the Sato-Vartia index fails to take account that an increase in taste for a variety is analogous to a fall in its price. Therefore, for varieties experiencing increases in tastes, the measured contribution to the cost of living is above the true contribution. In contrast, for varieties experiencing reductions in tastes, the measured contribution to the cost of living is below the true contribution. This introduces a systematic bias, because an increase in consumer taste for a variety raises its expenditure share and hence its weight in the cost of living. As a result, the errors from ignoring increases in tastes (which reduce the true cost of living below the measured cost of living) are weighted more highly than the errors from ignoring reductions in tastes (which raise the true cost of living above the measured cost of living), thereby giving rise to an upward bias in the Sato-Vartia index. This bias is related to the well-known “quality bias” from failing to take into account quality improvements, because the taste parameter for each variety enters the expenditure function in the same way as quality and inversely to price. However, an important difference with the quality bias is that the taste-shock bias is present even if taste changes are mean zero, and it is not dependent on unmeasured average quality rising. We show that our taste-shock bias is not eliminated by using chain-weighted rather than fixed-weighted price indexes, because it arises from the internal inconsistency of using observed expenditure shares (which include time-varying demand residuals) and observed prices (which do not include these time-varying demand residuals). Empirically, we find this taste-shock bias to be substantial, on average 0.4 percentage points per annum, and sizable relative to the bias from failing to take account of the entry and exit of varieties in measuring the cost of living.

Our approach addresses two notable limitations with the existing Sato-Vartia price index. First, this existing price index rules out by assumption the possibility that a consumer’s relative tastes for any two varieties can change over time, whereas it is intuitively plausible that such movements in relative tastes can occur (e.g. with changes in fashion, societal trends, lifestyle, or product knowledge). Second, while our exact price index allows

for these changes in relative tastes, it is also valid under the Sato-Vartia index's assumption of no changes in consumer tastes for each common variety. Therefore, under this conventional assumption of time-invariant consumer tastes, one should obtain the same change in the cost of living using our exact price index as using the Sato-Vartia index. Contrary to this prediction, we find the Sato-Vartia index is biased upwards by 0.4 percentage points per year relative to our index, which raises the question of what explains these differences. One potential explanation could be departures from the CES functional form. However, we show below that in barcode data the Sato-Vartia index generates similar measured changes in the cost of living as existing superlative indexes that are exact for flexible functional forms, such as the Fisher and Törnqvist indexes, which suggests that CES preferences provide a reasonable approximation to the data. Our approach presents a natural alternative explanation for these differences between our exact price index and the Sato-Vartia index, in terms of changes in relative consumer tastes. We show that the same taste-shock bias from abstracting from changes in relative tastes is present for superlative price indexes such as the Törnqvist index.

Our paper is most closely related to the literature on the estimation of CES price indices, including Feenstra (1994) and Broda and Weinstein (2006, 2010). We point out the tension that arises in this literature from combining a variety adjustment term based on the estimation of a CES demand system and a Sato-Vartia price index for continuing varieties. The issue is that the demand system estimation involves the inclusion of a time-varying demand residual, whereas the Sato-Vartia price index for continuing varieties is only consistent with CES demand if there are no demand shocks. To resolve this tension and be fully consistent with both CES preferences and the demand system estimation, we drop the Sato-Vartia index for continuing varieties, and replace it with our new exact price index that allows for a time-varying demand residual for each continuing good, given our normalization or choice of units in which to measure these demand residuals.

More broadly, our paper is related to research on 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 Fisher and Shell (1972), Lloyd (1975), Diewert (1976, 2004), Sato (1976), Vartia (1976), Lau (1979), Feenstra (1994), Hausman (1997), Moulton (1996), Balk (1999), Caves, Christensen and Diewert (1982), Nevo (2003), Neary (2004), Feenstra and Reinsdorf (2007, 2010), Białek (2017), and Diewert and Feenstra (2017). Our paper is also related to the voluminous literature in macroeconomics, trade and economic geography that has used CES preferences, including Dixit and Stiglitz (1977), Krugman (1980, 1991), Antràs (2003), Melitz (2003), Hsieh and Klenow (2009), and Arkolakis, Costinot and Rodriguez-Clare (2012), among others. We show that our methodology also holds for the closely-related logit model, and hence our work connects with the large body of applied research using this model, as synthesized in Anderson, de Palma and Thisse (1992) and Train (2009). Increasingly, researchers in 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 (2017). Our contribution relative to these studies is to allow for both changes in tastes for each variety and entry and exit, while preserving an exact price index in terms of prices and expenditure shares.

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 tastes for each variety, we explicitly allow for time-varying tastes for each variety, 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 consumer tastes) for each variety 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 II. introduces our new exact price index for CES preferences. Section III. develops a number of extensions and generalizations, including non-homothetic CES (indirectly additive), nested CES, mixed CES, logit, mixed logit, translog and AIDS preferences. Section IV. introduces our barcode data for the U.S. consumer goods sector. Section V. presents our main empirical results for CES preferences and demonstrates the quantitative relevance of allowing for changes in tastes for the measurement of the cost of living. Section VI. introduces our mixed CES extension. Section VII. contains a number of further robustness tests. Section VIII. concludes. An online appendix collects together technical derivations, additional information about the data, and supplementary empirical results.

II. Demand and Price Indexes with CES Preferences

We begin by deriving our new exact price index for CES preferences. To simplify the exposition, we consider a single nest of utility (e.g. an economy consisting of a single sector including many varieties). In Section III., we extend our analysis to accommodate multiple CES nests and more flexible functional forms.

II.A. Preferences

Under the assumption of homothetic CES preferences, the unit expenditure function (P_t) depends on the price (p_{kt}) and consumer taste (φ_{kt}) for each variety 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 varieties; we assume that varieties are substitutes ($\sigma > 1$); and Ω_t is the set of varieties supplied at time t .³ The parameter φ_{kt} captures consumer tastes and allows for both differences in the average level of tastes across varieties (some varieties are always more popular than others in all time periods) and also changes in tastes for individual varieties over time (some varieties become more or less popular relative to others over time).

II.B. Demand System

Applying Shephard's Lemma to this unit expenditure function (1), we obtain the demand system, in which the expenditure share (s_{kt}) for each variety k and time period t is:

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).

$$s_{kt} \equiv \frac{p_{kt}c_{kt}}{\sum_{\ell \in \Omega_t} p_{\ell t}c_{\ell t}} = \frac{(p_{kt}/\varphi_{kt})^{1-\sigma}}{\sum_{\ell \in \Omega_t} (p_{\ell t}/\varphi_{\ell t})^{1-\sigma}} = \frac{(p_{kt}/\varphi_{kt})^{1-\sigma}}{P_t^{1-\sigma}}, \quad k \in \Omega_t, \quad (2)$$

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

Rearranging this expenditure share (2), we obtain the following equivalent expression for the unit expenditure function that must hold for each variety $k \in \Omega_t$:

$$P_t = \frac{p_{kt}}{\varphi_{kt}} s_{kt}^{\frac{1}{\sigma-1}}. \quad (3)$$

To allow for the entry and exit of varieties over time, we define the common set of varieties between a pair of time periods t and $t-1$ (Ω_t^*) as those that are supplied in both periods (such that $\Omega_t^* = \Omega_t \cap \Omega_{t-1}$). Summing expenditures across these common varieties, we obtain the following expression for the aggregate share of common varieties in total expenditure in period t (λ_t):

$$\lambda_t = \frac{\sum_{k \in \Omega_t^*} p_{kt}c_{kt}}{\sum_{k \in \Omega_t} p_{kt}c_{kt}} = \frac{\sum_{k \in \Omega_t^*} (p_{kt}/\varphi_{kt})^{1-\sigma}}{\sum_{k \in \Omega_t} (p_{kt}/\varphi_{kt})^{1-\sigma}}. \quad (4)$$

Using this expression, the share of an individual variety in total expenditure (s_{kt}) in equation (2) can be re-written as its share of expenditure on common varieties (s_{kt}^*) times this aggregate share of common varieties in total expenditure (λ_t):

$$s_{kt} = \lambda_t s_{kt}^* = \lambda_t \frac{(p_{kt}/\varphi_{kt})^{1-\sigma}}{\sum_{\ell \in \Omega_t^*} (p_{\ell t}/\varphi_{\ell t})^{1-\sigma}}, \quad k \in \Omega_t^*, \quad (5)$$

where we use an asterisk to denote the value of a variable for common varieties $k \in \Omega_t^*$.

Two well-known properties of this CES demand system are the independence of irrelevant alternatives (IIA) and the symmetry of substitution effects. The first property implies that the relative expenditure share of any two varieties in equation (5) depends solely on their relative prices and taste parameters and not on the characteristics of any other varieties: $s_{kt}/s_{\ell t} = [(p_{kt}/\varphi_{kt}) / (p_{\ell t}/\varphi_{\ell t})]^{1-\sigma}$. The second property implies that the elasticity of expenditure on any one variety ($x_{kt} = p_{kt}c_{kt}$) with respect to a change in the price of another variety depends solely on the expenditure share of that other variety: $(\partial x_{kt}/\partial p_{\ell t})(p_{\ell t}/x_{kt}) = (\sigma-1)s_{\ell t}$. We relax both these assumptions in Section III. below, where we consider mixed CES preferences with heterogeneous consumers and the flexible functional forms of translog and AIDS preferences.

We treat the time-varying residual in the demand system (5) as a consumer taste shock (φ_{kt}) that also appears in the unit expenditure function (1). We note that there are other possible interpretations, including changes in product quality, measurement error and specification error. Our use of barcode data in our empirical application implies that changes in product quality are unlikely, because firms have strong incentives of inventory and stock control not to use the same barcode for products with different observable characteristics. Therefore, changes in product characteristics lead to the introduction of a new barcode, and are 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 demand residual to fit the data.

While we assume CES preferences in our baseline specification, we show below that our main insight generalizes to other preference structures, including flexible functional forms such as translog. In developing our approach, we highlight that the Sato-Vartia index's assumption of time-invariant tastes is inconsistent with using expenditure-share weights that include time-varying demand residuals, and demonstrate the quantitative relevance of this inconsistency for the measurement of changes in the cost of living.

II.C. Price Index

We now combine the unit expenditure function (3) and the relationship between expenditure shares (5) to measure the change in the cost of living over time (P_t/P_{t-1}). Using these two equations, the change in the cost of living can be written in terms of the change in the price (p_{kt}/p_{kt-1}), tastes ($\varphi_{kt}/\varphi_{kt-1}$) and common variety expenditure share (s_{kt}^*/s_{kt-1}^*) of any individual common variety $k \in \Omega_t^*$ and a variety correction term that controls for entry and exit (λ_t/λ_{t-1}):

$$\frac{P_t}{P_{t-1}} = \frac{p_{kt}/\varphi_{kt}}{p_{kt-1}/\varphi_{kt-1}} \left(\frac{\lambda_t}{\lambda_{t-1}} \frac{s_{kt}^*}{s_{kt-1}^*} \right)^{\frac{1}{\sigma-1}}, \quad k \in \Omega_t^*. \quad (6)$$

We now use this relationship to derive an exact price index that expresses the change in the cost of living solely in terms of observed *prices* and *expenditures*. As expenditures depend on *relative* consumer tastes (and not the absolute level of consumer tastes), the existence of such an exact price index requires that we rule out the possibility of a change in the cost of living, even though all prices and expenditures remain unchanged. We therefore require a normalization or constant choice of units in which to measure consumer tastes. If there exists good for which a researcher is confident that consumer tastes did not change, the assumption of a constant taste parameter for that good is one such normalization. In the absence of an obvious candidate for such a good in our barcode data, we assume that the geometric mean of the consumer taste parameters is constant across common varieties:

$$\tilde{\varphi}_t = \prod_{k \in \Omega_t^*} (\varphi_{kt})^{\frac{1}{N_t^*}} = \prod_{k \in \Omega_{t-1}^*} (\varphi_{kt-1})^{\frac{1}{N_t^*}} = \tilde{\varphi}_{t-1}, \quad (7)$$

where we use a tilde above a variable to denote a geometric mean across common varieties; $N_t^* = |\Omega_t^*|$ denotes the number of these common varieties.

Our assumption (7) is consistent with the assumption that the log demand shocks are mean zero in the estimation of the demand system. We examine the robustness of our results to alternative normalizations using the class of generalized means below. Taking geometric means across common varieties in equation (6), and using our assumption (7), we obtain our exact CES price index:

$$\Phi_t^{CUPI} = \frac{P_t}{P_{t-1}} = \left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{\frac{1}{\sigma-1}} \Phi_t^{*CCV}, \quad (8)$$

$$\Phi_t^{*CCV} = \frac{P_t^*}{P_{t-1}^*} = \frac{\tilde{p}_t}{\tilde{p}_{t-1}} \left(\frac{\tilde{s}_t^*}{\tilde{s}_{t-1}^*} \right)^{\frac{1}{\sigma-1}}. \quad (9)$$

We refer to this exact price index (Φ_t^{CUPI}) as the CES unified price index (CUPI), because our approach treats demand shocks in the same way in both the demand system and the unit expenditure function. This exact

price index has an intuitive interpretation. The first term $((\lambda_t/\lambda_{t-1})^{\frac{1}{\sigma-1}})$ on the right-hand side of equation (8) is the standard Feenstra (1994) variety correction term, which takes account of the entry and exit of varieties. If entering varieties are more attractive than exiting varieties (in the sense of having lower taste-adjusted prices p_{kt}/φ_{kt}), the share of common varieties in total expenditure will be smaller in period t than in period $t-1$ ($\lambda_t/\lambda_{t-1} < 1$), which reduces the cost of living (since $\sigma > 1$).

The second term ($\Phi_t^{*CCG} = P_t^*/P_{t-1}^*$) on the right-hand side of equation (8) is our new CES exact price index for common varieties (CCV), which itself has two components, as shown in equation (9). The first component ($\tilde{p}_t/\tilde{p}_{t-1}$) is the geometric mean of price relatives for common varieties, and is none other than a ‘‘Jevons’’ index, which serves as the basis for the lower level of the U.S. Consumer Price Index. Indeed, in the special case in which varieties are perfect substitutes ($\sigma \rightarrow \infty$), the CCV collapses to this Jevons index, since the exponent on the expenditure share term in equation (9) converges to zero as $\sigma \rightarrow \infty$.

The second component $((\tilde{s}_t^*/\tilde{s}_{t-1}^*)^{1/(\sigma-1)})$ is novel and depends on the geometric mean of relative expenditure shares for common varieties in the two time periods. This second component captures changes in the degree of heterogeneity in taste-adjusted prices across common varieties and moves with the average of the log expenditure shares in the two time periods.⁴ Critically, as the expenditure shares of common varieties become more uneven, the mean of the log expenditure shares falls, because the log function is concave. Therefore, this second term becomes smaller if taste-adjusted prices, and hence expenditure shares, become more dispersed across common varieties. The intuition is that consumers value dispersion in taste-adjusted prices across varieties if these varieties are substitutes ($\sigma > 1$). The reason is that they can substitute away from varieties with high taste-adjusted prices and towards varieties with low taste-adjusted prices.

If all taste-adjusted prices (p_{kt}/φ_{kt}) are constant, the log change in the cost of living for common varieties in equation (9) is necessarily zero. However, even if observed prices (p_{kt}) are constant, the cost of living for common varieties can change with movements in taste-adjusted prices, because the unit expenditure function depends on taste-adjusted rather than observed prices. Nevertheless, our normalization (7) rules out a pure change in consumer tastes, in which tastes for all common varieties are scaled by the same proportion. Therefore, the change in the cost of living depends on movements in prices and *relative* consumer tastes.

II.D. Demand System Inversion

We now show that the CUPI implicitly inverts the CES demand system (5) to substitute out for unobserved changes in tastes ($\varphi_{kt}/\varphi_{kt-1}$) in terms of observed changes in prices (p_{kt}/p_{kt-1}) and common variety expenditure shares (s_{kt}^*/s_{kt-1}^*). We use this result in later sections to generalize our approach to other invertible demand systems, including nested, mixed and non-homothetic CES, logit and mixed logit, translog and AIDS preferences.

In particular, under our baseline assumption of CES preferences, the change in the cost of living for common varieties (P_t^*/P_{t-1}^*) in equation (9) can be written in terms of common variety expenditure shares and taste-adjusted prices as follows:

⁴Our unified price index (8) differs from the expression for the CES price index in Hotman et al. (2016), which did not distinguish entering and exiting varieties from common varieties and captured the dispersion of sales across common varieties using a different term.

$$\ln \left(\frac{P_t^*}{P_{t-1}^*} \right) = \sum_{k \in \Omega_t^*} \omega_{kt}^* \ln \left(\frac{p_{kt} / \varphi_{kt}}{p_{kt-1} / \varphi_{kt-1}} \right), \quad (10)$$

where the weights ω_{kt}^* are the logarithmic mean of common variety expenditure shares (s_{kt}^*) in periods t and $t - 1$ and sum to one:

$$\omega_{kt}^* \equiv \frac{\frac{s_{kt}^* - s_{kt-1}^*}{\ln s_{kt}^* - \ln s_{kt-1}^*}}{\sum_{\ell \in \Omega_t^*} \frac{s_{\ell t}^* - s_{\ell t-1}^*}{\ln s_{\ell t}^* - \ln s_{\ell t-1}^*}}, \quad \sum_{k \in \Omega_t^*} \omega_{kt}^* = 1, \quad (11)$$

and the derivation is reported in Section A.2 of the online appendix.⁵

This expression for the change in the cost of living in equation (10) is a generalization of the Sato-Vartia price index, which corresponds to the special case in which tastes are time invariant for each common variety ($\varphi_{kt} / \varphi_{kt-1} = 1$ for all $k \in \Omega_t^*$). The challenge in implementing equation (10) empirically is that it depends on taste-adjusted prices (p_{kt} / φ_{kt}), whereas only unadjusted prices are observed in the data (p_{kt}). To overcome this challenge, we invert the CES demand system to express the unobserved time-varying taste parameter (φ_{kt}) in terms of observed prices (p_{kt}) and common variety expenditure shares (s_{kt}^*). Dividing the common variety expenditure share (5) by its geometric mean across common varieties, taking logarithms and differencing over time, we obtain the following closed-form expression for the log change in the taste parameter for each common variety $k \in \Omega_t^*$:

$$\ln \left(\frac{\varphi_{kt}}{\varphi_{kt-1}} \right) = \ln \left(\frac{p_{kt} / \tilde{p}_t}{p_{kt-1} / \tilde{p}_{t-1}} \right) + \frac{1}{\sigma - 1} \ln \left(\frac{s_{kt}^* / \tilde{s}_t^*}{s_{kt-1}^* / \tilde{s}_{t-1}^*} \right), \quad (12)$$

where we have used our normalization that the geometric mean of the taste parameters across common varieties is constant: $\tilde{\varphi}_t = \tilde{\varphi}_{t-1}$.

Substituting this closed-form expression for the taste shocks (12) into the change in the cost of living for common varieties in equation (10), we obtain our exact CES common variety price index (CCV) in equation (9), as shown in Section A.2 of the online appendix. This alternative derivation of the CCV highlights the role of the inversion of the demand system in deriving our exact price index. A sufficient condition for the demand system to be invertible in this way is that it satisfies the conditions for “connected substitutes” in Berry, Gandhi and Haile (2013). These conditions rule out the possibility that some varieties are substitutes while others are complements. Our assumptions that demand is CES and varieties are substitutes ($\sigma > 1$) ensure that “connected substitutes” is satisfied.

II.E. Taste-Shock Bias

We now compare our CCV to existing exact CES price indexes. Both our CCV and the Sato-Vartia index use the observed expenditure shares. Our CCV assumes that movements in these expenditure shares reflect changes in both relative prices and relative tastes. Therefore, we adjust the observed price movements for the changes in

⁵To derive this equivalent expression for the exact CES price index for common varieties, take the ratio of the common goods expenditure shares in the two time periods (s_{kt}^* / s_{kt-1}^*), take logarithms, and rearrange terms to obtain: $\left[\ln \left(\frac{P_t^*}{P_{t-1}^*} \right) - \ln \left(\frac{p_{kt} / \varphi_{kt}}{p_{kt-1} / \varphi_{kt-1}} \right) \right] / \ln \left(\frac{s_{kt}^*}{s_{kt-1}^*} \right) = \frac{1}{\sigma - 1}$. Multiply both sides of this equation by $s_{kt}^* - s_{kt-1}^*$, sum across all common varieties, and re-arrange terms to obtain equation (10), as shown in Section A.2 of the online appendix.

relative tastes implied by the demand system when we compute the change in the cost of living. In contrast, the Sato-Vartia index assumes that tastes for each common variety are time-invariant, and interprets the observed movements in the expenditure shares as reflecting *only* changes in relative prices. Hence, the Sato-Vartia index uses the observed prices without making any adjustment for changes in relative tastes in computing the change in the cost of living. We show below that this assumption that relative prices are the sole source of movements in expenditure shares is hard to reconcile with empirical estimates of the demand system. In the remainder of this section, we demonstrate that this assumption introduces a taste-shock bias into the measurement of the cost of living, because the Sato-Vartia index uses initial and final-period expenditure shares (where the final-period shares are affected by demand shocks), while also using observed prices (without adjusting for demand shocks).

From equations (9) and (10), the Sato-Vartia index equals the true exact CES common variety price index (CCV) plus an additional term in consumer taste shocks that we refer to as the taste-shock bias:

$$\ln \Phi_t^{*CCV} = \ln \Phi_t^{*SV} - \underbrace{\left[\sum_{k \in \Omega_t^*} \omega_{kt}^* \ln \left(\frac{\varphi_{kt}}{\varphi_{kt-1}} \right) \right]}_{\text{taste-shock bias}}, \quad (13)$$

where the Sato-Vartia index ($\ln \Phi_t^{*SV}$) is the special case of equation (10) in which $\varphi_{kt}/\varphi_{kt-1} = 1$.

Therefore, the Sato-Vartia index is only unbiased if the taste shocks ($\ln(\varphi_{kt}/\varphi_{kt-1})$) 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 taste shocks ($\ln(\varphi_{kt}/\varphi_{kt-1})$) and the expenditure-share weights (ω_{kt}^*) is possible, depending on the underlying correlation between taste and price shocks. However, there is a mechanical force for a positive correlation, because the expenditure-share weights themselves are functions of the taste shocks. In particular, a positive taste shock for a variety mechanically increases the expenditure-share weight for that variety and reduces the expenditure-share weight for all other varieties:

$$\frac{d\omega_{kt}^*}{d\varphi_{kt}} \frac{\varphi_{kt}}{\omega_{kt}^*} > 0, \quad \frac{d\omega_{\ell t}^*}{d\varphi_{kt}} \frac{\varphi_{kt}}{\omega_{\ell t}^*} < 0, \quad \forall \ell \neq k, \quad (14)$$

as shown in Section A.3 of the online appendix.

The intuition for this taste-shock bias is as follows. The Sato-Vartia index fails to take into account that increases in tastes are like reductions in prices and decrease the cost of living, while reductions in tastes are analogous to increases in prices and increase the cost of living. If the weights placed on varieties in the Sato-Vartia index were uncorrelated with changes in tastes, these measurement errors would average out across varieties. However, other things equal, varieties experiencing increases in tastes have systematically higher weights in the Sato-Vartia index than varieties experiencing reduction in tastes, because an increase in the relative taste for a variety raises its expenditure share and hence its weight in the Sato-Vartia index. Therefore, the errors from ignoring increases in tastes (which reduce the true cost of living below the measured cost of living) are weighted more highly than the errors from ignoring reductions in tastes (which raise the true cost of living above the measured cost of living), thereby introducing an upward bias in the Sato-Vartia index. Even

though our use of barcode data ensures that changes in quality within common varieties are unlikely (because changes in product attributes lead to the introduction of a new barcode), this taste-shock bias is analogous to the well-known “quality bias” from neglecting changes in product quality, because such changes in quality would enter the unit expenditure function in the same way as changes in tastes.

Another metric for the tension inherent in the Sato-Vartia index’s assumption that movements in expenditure shares reflect only changes in relative prices 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. Indeed, the model is overidentified, with an infinite number of approaches to recovering the elasticity of substitution, each of which uses different weights for each common variety, as shown in Section A.4 of the online appendix. If tastes for *all* common varieties are indeed constant (including no changes in preferences, quality, measurement error or specification error), all of these approaches will recover the same elasticity of substitution. However, if tastes for *some* common variety change over time, but a researcher falsely assumes time-invariant tastes for all common varieties, 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 tastes for all common varieties.

From equation (8), the overall change in the cost of living equals the Feenstra (1994) variety correction for entry/exit plus the change in the cost of living for common varieties. Therefore, if the Sato-Vartia index is used to measure the change in the cost of living for common varieties, this translates into a bias in the measurement of the overall cost of living. In contrast, using our CCV to measure the change in the cost of living for common varieties eliminates this bias in the measurement of the overall cost of living.

II.F. Robustness to Alternative Normalizations

Although our normalization (7) is consistent with the assumption that the log demand shocks are mean zero in the estimation of the demand system, it is not the only possible normalization. Therefore, we report robustness tests for a range of alternative normalizations, in which we rule out a pure change in consumer tastes by requiring that a generalized mean of order- r of the taste parameters is constant:

$$\left[\frac{1}{N_t^*} \sum_{k \in \Omega_t^*} \varphi_{kt}^r \right]^{\frac{1}{r}} = \left[\frac{1}{N_t^*} \sum_{k \in \Omega_t^*} \varphi_{kt}^r \right]^{\frac{1}{r}}. \quad (15)$$

Using equations (3) and (5), and taking a generalized mean across common varieties, we obtain the following expression for the change in the cost of living between periods t and $t - 1$:

$$\frac{P_t}{P_{t-1}} = \left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{\frac{1}{\sigma-1}} \frac{P_t^*}{P_{t-1}^*} = \left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{\frac{1}{\sigma-1}} \frac{\left[\frac{1}{N_t^*} \sum_{k \in \Omega_t^*} p_{kt}^r (s_{kt}^*)^{\frac{r}{\sigma-1}} \right]^{\frac{1}{r}}}{\left[\frac{1}{N_t^*} \sum_{k \in \Omega_t^*} p_{kt-1}^r (s_{kt-1}^*)^{\frac{r}{\sigma-1}} \right]^{\frac{1}{r}}}, \quad (16)$$

as shown in Section A.5 of the online appendix. Comparing this expression with our CUPI in equation (8), the variety correction term is unchanged, and the common variety price index takes a similar form as the exact price index in the quadratic mean of order- r expenditure function in Diewert (1976).

As we vary the value of r we place different weights on low versus high values for consumer tastes, with negative values of r placing greater weight on low values for consumer tastes (as $r \rightarrow -\infty$ we obtain the minimum value) and high values of r giving more weight to high values for consumer tastes (as $r \rightarrow \infty$ we obtain the maximum value). Our CUPI corresponds to the limiting case of equation (16) in which $r \rightarrow 0$. We show in Section V.E. that we find quantitatively similar measured changes in the cost of living over time using generalized means ranging from the harmonic mean ($r = -1$), through the geometric mean ($r = 0$) and the arithmetic mean ($r = 1$), to the quadratic mean ($r = 2$).

III. 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, mixed logit, translog and AIDS preferences. We show that our main insight that the demand system can be inverted to express unobserved relative taste shocks for individual varieties 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 an exact price index in terms of only prices and expenditure shares, and existing price indexes that assume time-invariant relative tastes for each variety are subject to a taste-shock bias.

III.A. Non-homothetic CES

We now generalize our approach to allow for non-homotheticities using the non-separable class of CES functions in Sato (1975), Comin, Lashkari and Mestieri (2015), and Matsuyama (2019), which satisfy implicit additivity in Hanoch (1975). Although this specification is more restrictive than the flexible specifications of non-homotheticities in Fajgelbaum and Khandelwal (2016) and Atkin, Faber, Fally and Gonzalez-Navarro (2018), it allows us to show that our approach does not depend on assuming homotheticity, and we analyze the more flexible functional forms of translog and AIDS preferences in later sections below.

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:

$$\sum_{k \in \Omega_t} \left(\frac{\varphi_{kt}^h c_{kt}^h}{(C_t^h)^{(\epsilon_k - \sigma)/(1 - \sigma)}} \right)^{\frac{\sigma - 1}{\sigma}} = 1, \quad (17)$$

where c_{kt}^h denotes household h 's consumption of variety k at time t ; φ_{kt}^h is household h 's taste parameter for variety k at time t ; σ is the constant elasticity of substitution between varieties; ϵ_k is the constant elasticity of consumption of variety k with respect to the consumption index (C_t^h) that allows for non-homotheticity. Assuming that varieties are substitutes ($\sigma > 1$), we require $\epsilon_k < \sigma$ for the consumption index (17) 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 II. above corresponds to the special case of equation (17) in which $\epsilon_k = 1$ for all $k \in \Omega_t$.

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 variety k (s_{kt}^h):

$$P_t^h = \left[\sum_{k \in \Omega_t} (p_{kt} / \varphi_{kt}^h)^{1-\sigma} (C_t^h)^{\epsilon_k - 1} \right]^{\frac{1}{1-\sigma}}, \quad (18)$$

$$s_{kt}^h = \frac{(p_{kt} / \varphi_{kt}^h)^{1-\sigma} (C_t^h)^{\epsilon_k - 1}}{\sum_{\ell \in \Omega_t} (p_{\ell t}^h / \varphi_{\ell t}^h)^{1-\sigma} (C_t^h)^{\epsilon_\ell - 1}} = \frac{(p_{kt} / \varphi_{kt}^h)^{1-\sigma} (E_t^h / P_t^h)^{\epsilon_k - 1}}{(P_t^h)^{1-\sigma}}, \quad (19)$$

where we assume that all households h face the same price for a given variety (p_{kt}) and the derivation for all results in this section is reported in Section A.6 of the online appendix.

As for the homothetic case in the previous section, the price index (18) depends on taste-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 price index (P_t^h) enters the numerator of the expenditure share in equation (19). To overcome this challenge, we work with the share of each variety in *overall* expenditure (s_{kt}^h) rather than the common variety expenditure share (s_{kt}^{h*} in our earlier notation). In particular, rearranging this overall expenditure share (19), and taking ratios between periods t and $t - 1$, we obtain the following expression for the change in the cost of living, which must hold for each common variety that is supplied in both periods:

$$\frac{P_t^h}{P_{t-1}^h} = \frac{p_{kt} / \varphi_{kt}^h}{p_{kt-1} / \varphi_{kt-1}^h} \left(\frac{E_t^h / P_t^h}{E_{t-1}^h / P_{t-1}^h} \right)^{\frac{\epsilon_k - 1}{1-\sigma}} \left(\frac{s_{kt}^h}{s_{kt-1}^h} \right)^{\frac{1}{\sigma-1}}, \quad k \in \Omega_t^*. \quad (20)$$

Using our normalization that the geometric mean of consumer tastes across common varieties for household h is constant, and taking the geometric mean across common varieties in equation (20), we obtain the following generalization of our CES unified price index to the non-homothetic case for each household h :

$$\frac{P_t^h}{P_{t-1}^h} = \left(\frac{\tilde{p}_t}{\tilde{p}_{t-1}} \right)^{\frac{1}{1+\vartheta}} \left(\frac{\tilde{s}_t^h}{\tilde{s}_{t-1}^h} \right)^{\frac{1}{(\sigma-1)(1+\vartheta)}} \left(\frac{E_t^h}{E_{t-1}^h} \right)^{\frac{\vartheta}{1+\vartheta}}, \quad (21)$$

$$\vartheta \equiv \frac{1}{N_t^*} \sum_{k \in \Omega_t^*} \frac{\epsilon_k - 1}{1 - \sigma},$$

where the tilde above a variable denotes a geometric mean across common varieties; the derived parameter ϑ captures the average across the common varieties 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 $\vartheta \neq 0$.

III.B. Nested CES

In our baseline specification in Section II., we focus on a single CES tier of utility, which can be interpreted as a single sector consisting of many varieties. 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. This nested demand structure introduces additional flexibility into the substitution patterns between

varieties, depending on whether those varieties are in the same or different nests. 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:

$$P_t = \left[\sum_{g \in \Omega^G} \left(\frac{P_{gt}^G}{\varphi_{gt}^G} \right)^{1-\sigma^G} \right]^{\frac{1}{1-\sigma^G}}, \quad \sigma^G > 1, \quad (22)$$

where σ^G is the elasticity of substitution across sectors; P_{gt}^G is the unit expenditure function for each sector, which is defined across varieties within each sector; φ_{gt}^G is the taste parameter for each sector; we assume for simplicity that the set of sectors is constant over time and denote the number of sectors by $N^G = |\Omega^G|$; and the derivations for this section of the paper are reported in Section A.7 of the online appendix.

All of the results for entry and exit and the exact CES price index with time-varying taste shocks from Section II. continue to hold for this nested demand structure. We use analogous normalizations for the taste parameters as above: we normalize the geometric mean of sector tastes (φ_{gt}^G) to be constant across sectors and the geometric mean of variety tastes (φ_{kt}^K) to be constant across common varieties within sectors. As the log of the price index for each nest of utility is the mean of the log prices within that nest, and the mean is a linear operator, we can apply this operator recursively across nests of utility. Using this property, the log aggregate price index can be expressed as follows:

$$\begin{aligned} \ln \left(\frac{P_t}{P_{t-1}} \right) &= \frac{1}{N^G} \sum_{g \in \Omega^G} \frac{1}{N_{gt}^{K^*}} \sum_{k \in \Omega_{gt}^{K^*}} \ln \left(\frac{p_{kt}^K}{p_{kt-1}^K} \right) + \frac{1}{N^G} \sum_{g \in \Omega^G} \frac{1}{\sigma_g^K - 1} \frac{1}{N_{gt}^{K^*}} \sum_{k \in \Omega_{gt}^{K^*}} \ln \left(\frac{s_{gkt}^{K^*}}{s_{gkt-1}^{K^*}} \right) \\ &+ \frac{1}{N^G} \sum_{g \in \Omega^G} \frac{1}{\sigma_g^K - 1} \ln \left(\frac{\lambda_{gt}^K}{\lambda_{gt-1}^K} \right) + \frac{1}{\sigma^G - 1} \frac{1}{N^G} \sum_{g \in \Omega^G} \ln \left(\frac{s_{gt}^G}{s_{gt-1}^G} \right), \end{aligned} \quad (23)$$

where we have used our normalizations that $(1/N^G) \sum_{g \in \Omega^G} \ln \left(\varphi_{gt}^G / \varphi_{gt-1}^G \right) = 0$ and $(1/N_{gt}^{K^*}) \sum_{k \in \Omega_{gt}^{K^*}} \ln \left(\varphi_{kt}^K / \varphi_{kt-1}^K \right) = 0$; $N_{gt}^{K^*}$ is the number of common varieties for each sector g ; $s_{gkt}^{K^*}$ is the share of an individual common variety k in expenditure on sector g at time t ; $(1/(\sigma_g^K - 1)) \ln \left(\lambda_{gt}^K / \lambda_{gt-1}^K \right)$ is the variety correction term for the entry and exit of varieties within sector g ; and s_{gt}^G is the share of sector g in aggregate expenditure at time t .

Although, for simplicity, we focus here on two nests of utility, this procedure can be extended for any number of nests of utility, from the highest to the lowest. Conventional measures of the overall cost of living typically aggregate categories using expenditure-share weights. Therefore, we assume in our empirical analysis below that the upper tier of utility across sectors is Cobb-Douglas ($\sigma^G = 1$), and estimate the elasticity of substitution across varieties within sectors (σ_g^K) separately for each sector.

III.C. Mixed CES

The non-homothetic specification in Section III.A. assumes that the only source of heterogeneity across consumers is differences in income and that all consumers have the same elasticity of substitution across varieties. In this section, we introduce a mixed CES specification that allows both the elasticity of substitution and the

taste parameters to vary in an unrestricted way across groups.⁶ In particular, we consider a setting with multiple groups of heterogeneous consumers indexed by $h \in \{1, \dots, H\}$, in which the unit expenditure function (P_t^h) and expenditure share (s_{kt}^h) for a consumer from group h are:

$$P_t^h = \left[\sum_{k \in \Omega_t} \left(\frac{p_{kt}}{\varphi_{kt}^h} \right)^{1-\sigma^h} \right]^{\frac{1}{1-\sigma^h}}, \quad (24)$$

$$s_{kt}^h = \frac{(p_{kt}/\varphi_{kt}^h)^{1-\sigma^h}}{\sum_{\ell \in \Omega_t} (p_{\ell t}/\varphi_{\ell t}^h)^{1-\sigma^h}} = \frac{(p_{kt}/\varphi_{kt}^h)^{1-\sigma^h}}{(P_t^h)^{1-\sigma^h}}, \quad (25)$$

where s_{kt}^h is a share of variety k in the expenditure of group h at time t ; σ^h is the elasticity of substitution across varieties for group h ; φ_{kt}^h denotes consumer tastes for group h ; and the derivation for all results in this section is reported in Section A.8 of the online appendix. We assume for simplicity that all groups face the same prices (p_{kt}) and set of varieties available (Ω_t). Nevertheless, we allow for the possibility that some groups do not consume some varieties, which we interpret as corresponding to the limiting case in which the taste parameter converges to zero for that group and variety ($\lim \varphi_{kt}^h \rightarrow 0$ for some k and h).

This specification relaxes the independence of irrelevant alternatives (IIA) property of CES, because the differences in preferences across groups imply that the relative expenditure shares of two varieties 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:

$$\frac{\partial x_{kt}}{\partial p_{\ell t}} \frac{p_{\ell t}}{x_{kt}} = \frac{1}{s_{kt}} \sum_{h=1}^H f_t^h (\sigma^h - 1) s_{kt}^h s_{\ell t}^h, \quad (26)$$

where s_{kt} is the share of variety k in total expenditure; s_{kt}^h is the share of variety 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 II. now hold for each group of consumers separately.⁷ Following the same analysis as in Section II.D. above, the exact CES unified price index for each group, allowing for entry and exit and taste shocks, takes the same form as in equation (8):

$$\ln \Phi_t^{hCUPI} = \underbrace{\frac{1}{\sigma^h - 1} \ln \left(\frac{\lambda_t^h}{\lambda_{t-1}^h} \right)}_{\text{Variety Adjustment}} + \underbrace{\frac{1}{N_t^*} \sum_{k \in \Omega_t^*} \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) + \frac{1}{\sigma^h - 1} \frac{1}{N_t^*} \sum_{k \in \Omega_t^*} \ln \left(\frac{s_{kt}^{h*}}{s_{kt-1}^{h*}} \right)}_{\text{Common-Variety Price Index}}, \quad (27)$$

where $(1/(\sigma^h - 1)) \ln(\lambda_t^h/\lambda_{t-1}^h)$ is the variety correction term for the entry and exit of varieties for group h ; s_{kt}^{h*} is the share of an individual common variety k in all expenditure on common varieties for group h ; and we have used our normalization that the geometric mean of consumer tastes for each group is constant.

⁶This 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).

⁷In 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.

To implement this mixed CES specification, we estimate the elasticities of substitution (σ_g^h) for each group separately using the data on prices and expenditure shares for that group. In Section VI. 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).

III.D. 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 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_{ik} units of variety k at time t is given by:

$$U_{it} = \ln \varphi_{kt} + \ln c_{ikt} + z_{ikt}, \quad (28)$$

where φ_{kt} captures the component of consumer tastes for each variety that is common across consumers; z_{ikt} captures idiosyncratic consumer tastes for each variety 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.

Consumers are assumed to have the same expenditure level E_t and choose their preferred variety given the realizations for their idiosyncratic tastes. Using the properties of the extreme value distribution, we show in Section A.9 of the online appendix that the expenditure share for each variety and the consumer's expected utility take the same form as in our baseline CES specification in Section II. of the paper, where $1/\nu = \sigma - 1$. Therefore, all our results can be applied for the logit model. Additionally, in the same way that our baseline CES specification can be generalized to mixed CES (as in Section III.C. above), this baseline logit model can be generalized to a mixed logit, as in McFadden and Train (2000).

III.E. Flexible Functional Forms

Finally, we show that our approach also holds for the flexible functional forms of homothetic translog preferences and the non-homothetic almost ideal demand system (AIDS). In this section, we briefly review the homothetic translog case. In Section A.10 of the online appendix, we report the derivations for both the homothetic translog and non-homothetic AIDS specifications.

Homothetic translog preferences provide an arbitrarily close local approximation to any continuous and twice-differentiable homothetic utility function. In particular, we consider the following unit expenditure function that is defined over the price (p_{kt}) and taste parameter (φ_{kt}) for a constant set of varieties $k \in \Omega$ with number of elements $N = |\Omega|$:

$$\ln P_t = \ln \alpha_0 + \sum_{k \in \Omega} \alpha_k \ln \left(\frac{p_{kt}}{\varphi_{kt}} \right) + \frac{1}{2} \sum_{k \in \Omega} \sum_{\ell \in \Omega} \beta_{k\ell} \ln \left(\frac{p_{kt}}{\varphi_{kt}} \right) \ln \left(\frac{p_{\ell t}}{\varphi_{\ell t}} \right), \quad (29)$$

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

As for CES preferences in equation (10), the change in the cost of living can be written as an expenditure-share-weighted average of the change in taste-adjusted prices (p_{kt}/φ_{kt}) for each variety:

$$\ln \Phi_t^{TR} = \ln \left(\frac{P_t}{P_{t-1}} \right) = \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) \ln \left(\frac{\varphi_{kt}}{\varphi_{kt-1}} \right), \quad (30)$$

where the weights for translog are the arithmetic mean of expenditure shares in the two time periods ($(1/2)(s_{kt} + s_{kt-1})$) instead of the logarithmic mean for CES (equation (11)).

This expression for the change in the cost of living (30) is a generalization of the Törnqvist index ($\ln \Phi_t^{TO}$), which corresponds to the special case of equation (30) in which taste is assumed to be constant for all varieties ($(\varphi_{kt}/\varphi_{kt-1}) = 1$ for all $k \in \Omega$). Therefore, the Törnqvist index for translog is subject to a similar taste-shock bias as the Sato-Vartia index for CES, except that the taste shock for each variety is weighted by the arithmetic mean of expenditure shares in the two time periods instead of the logarithmic mean. The source of this bias is again the failure to take into account that an increase in taste for a variety is analogous to a fall in its price, which induces a systematic overstatement of the increase in the cost of living, because consumers substitute towards varieties that become more desirable. Once again, we have the result that a positive taste shift lowers the taste-adjusted price for a variety and raises its expenditure share, while a negative one has the reverse effect. Since the Törnqvist index also does not take the association between taste shifts and expenditure shares into account, it underweights reductions in relative taste-adjusted prices and overweights increases in taste-adjusted prices, just like the Sato-Vartia index.

Again, we overcome the challenge that consumer tastes are not observed in the data by inverting the demand system to solve for tastes (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 variety:

$$\Delta s_{kt} = \sum_{\ell \in \Omega} \beta_{k\ell} [\Delta \ln (p_{\ell t}) - \Delta \ln (\varphi_{\ell t})]. \quad (31)$$

We assume that each variety's expenditure share is decreasing in its own taste-adjusted price ($\beta_{kk} < 0$), and increasing in the taste-adjusted price of other varieties ($\beta_{k\ell} > 0$ 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 taste shocks ($\Delta \ln (\varphi_{\ell t})$) by inverting the system of expenditure shares in equation (31), as shown in Section A.10 of the online appendix. The demand system (31) consists of a system of equations for the change in the expenditure shares (Δs_{kt}) of the N varieties that is linear in the change in the log price ($\Delta \ln p_{kt}$) and log taste parameter ($\Delta \ln \varphi_{kt}$) for each variety. These changes in expenditure shares must sum to zero across varieties, 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 β for all pairs of varieties must each sum to zero. Therefore, without loss of generality, we can omit the equation for one variety. We can nevertheless recover the taste shock for all varieties (including the omitted one) using our normalization that the geometric mean of tastes is constant (which implies $(1/N) \sum_{k \in \Omega} \Delta \ln \varphi_{kt} = 0$), as

shown in Section A.10 of the online appendix. We thus obtain the unobserved taste shock for each variety in terms of observed prices and expenditure shares:

$$\Delta \ln \varphi_{kt} = S_{kt}^{-1} (\Delta \mathbf{s}_t, \Delta \ln \mathbf{p}_t, \boldsymbol{\beta}). \quad (32)$$

Substituting for these unobserved taste shocks in equation (30), we obtain the following exact price index in terms of prices and expenditure shares:

$$\ln \Phi_t^{TR} = \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) \ln \left(\frac{p_{kt}}{p_{kt-1}} \right) - \sum_{k \in \Omega} \frac{1}{2} (s_{kt} + s_{kt-1}) S_{kt}^{-1} (\Delta \mathbf{s}_t, \Delta \ln \mathbf{p}_t, \boldsymbol{\beta}), \quad (33)$$

which corresponds to the analogous common variety price index for translog preferences as our CCV for CES preferences in equation (9) above.

Therefore, our main insight that the demand system can be unified with the unit expenditure function to construct an exact price index that allows for time-varying taste shocks for individual varieties is not specific to CES, but also holds for flexible functional forms. Furthermore, the taste-shock bias is again present, because a conventional price index that assumes time-variant tastes interprets all movements in expenditure shares as reflecting changes in prices, and hence does not take into account that these movements in expenditure shares are also influenced by the time-varying demand residual.

IV. Data

Our data source is the Nielsen Homescan Consumer Panel,⁸ 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.⁹ Prices are either downloaded from the store in which the barcode was purchased or hand entered, and the household records any deals used that may affect the price. Measuring varieties using barcodes has 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 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 barcodes, not changes in the characteristics of existing barcodes.

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 re-weight

⁸Our 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>

⁹The 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.

the data so that the average expenditures and prices are representative of the actual demography in each region rather than the Nielsen sample. We use these weights to construct a demographically-balanced sample of households in 42 cities in the United States. The set of goods included corresponds to 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 both the household dimension in the data and 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 barcoded goods (varieties) into product groups (sectors), which are based on where they 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 sectors (e.g., Perdue Chicken Breasts, Dole Baby Spinach, etc.). The 5 largest of our 104 sectors in order of expenditure are pet food, carbonated beverages, paper products, bread and baked goods, and candy. We report a full list of the sectors and summary statistics for each sector in the online appendix. Output units are common within a sector: 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 barcoded 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.¹⁰ 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.

¹⁰Even 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 barcode 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 barcode. We also winsorized the data by dropping observations whose percentage change in price or value were in the top or bottom one percent.

Our baseline sample of barcodes is an unbalanced panel that includes both barcodes that survive throughout our entire sample period and those that enter or exit at some point during the sample period. When we construct our price indexes, we need to define the subset of varieties (barcodes) that are common across periods, which requires jointly deciding on the number of years (four-quarter differences) over which the common set is defined and when a variety counts as entering and exiting this common set. For the number of years, we consider a range of definitions of the set of common varieties that require a good to be present (a) only in years $t - 1$ and t , (b) for the entire sample period, (c) in years $t - 1$ and t and an intermediate number of years that is less than the full sample period. When we examine the sensitivity of our results across these alternative definitions in Section V.E., we find a stable pattern of results across these alternative definitions of common varieties. We choose as our baseline definition of common varieties the set of varieties present in both years $t - 1$ and t and more than a total of six years.

For determining entry and exit into the set of common varieties, a variety can appear at the beginning or end of the fourth quarter of either years $t - 1$ or t , which affects measured rates of growth over the four-quarter difference. More generally, varieties can experience dramatic increases in sales in the first few quarters as they are rolled out into stores or rapid declines in sales in the last few quarters of their life as stores sell out and deplete their inventories. These features can make it appear as if consumer tastes for a common variety are changing rapidly when in fact they are not. To make sure that these entry and exit events are not driving our results, we also require that a variety must be available for three quarters before the fourth quarter of year $t - 1$ and for three quarters after the fourth quarter of year t to appear in the common set of varieties.

In Table I, we report summary statistics for our baseline sample including common, entering and exiting varieties. For each variable, we first compute the time mean across years for a given sector. We next report in the table the mean and standard deviation of these time means across sectors, as well as percentiles of their distribution across sectors. As shown in the first row, the median number of price and quantity observations (“Sector Sample Size”) is 47,747, with the sectors in the fifth percentile of observations only having just short of 9,033 data points and those in the 95th percentile having over 151,930 observations. The median number of barcodes per sector is just over 11,000, with 95 percent of these sectors having more than 1,700 unique barcodes, and the largest five percent of them encompassing over 45,000 unique barcodes.

We find substantial entry and exit of barcodes, with the typical life of a barcode being only three to four years. On average, 37 percent of all barcodes in a given year exit the sample in the following year, while 38 percent of barcodes 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 sectors. These averages mask substantial heterogeneity in innovation rates across sectors, with the average life of a cottage cheese barcode being 6.4 years, whereas the average life of an electronics barcode is only 1.7 years. High rates of barcode turnover are reflected in shares of common barcodes in total expenditure (λ_t and λ_{t-1}) of less than one, although these again vary substantially across sectors from a low of 0.34 to a high of 0.99. Consistent with entering barcodes being more numerous or more attractive to consumers than exiting barcodes, we find that common barcodes account for a larger share of expenditure in $t - 1$ than in t (a value of λ_t/λ_{t-1} 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 sector. 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 taste shocks and elastic demand.

V. Empirical Results

We now present our main empirical results. In Section V.A., we report our estimates of the elasticity of substitution across barcodes within each of the sectors in our data. In Section V.B., we use these estimated elasticities of substitution to invert the demand system and recover consumer tastes, and provide evidence on the properties of our estimated consumer tastes. In Section V.C., 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 assumptions of time-invariant tastes and no entry/exit. In Section V.D., we implement our new exact CES unified price index that allows for time-varying tastes for each variety. We show that abstracting from these taste shocks leads to a substantial taste-shock bias in existing exact CES price indexes. In Section V.E., we report a joint specification test on our assumptions of CES demand and a constant geometric mean for consumer tastes, and report a robustness test using alternative normalizations for tastes.

V.A. Estimates of the Elasticity of Substitution

We estimate the elasticity of substitution across barcodes for each sector separately using the conventional Feenstra (1994) estimator. We demonstrate the robustness of our results to alternative elasticities of substitution using a grid search over the range of plausible values for the elasticity of substitution in Section VII. below. The Feenstra (1994) estimator uses double-differences of the log common-variety expenditure share (5), where the first difference is across time and the second difference is across barcodes within sectors. The demand elasticity is separately identified from the supply elasticity using the assumption that the double-differenced demand and supply shocks are orthogonal and heteroskedastic, as shown in Section A.11 of the online appendix. We follow Broda and Weinstein (2006) in stacking these moment conditions for each barcode and estimating the elasticity of substitution using the generalized method of moments (GMM).

In Table II, we report percentiles of the distribution of these estimates across the sectors. We find estimated elasticities of substitution that range from 5.11 at the 5th percentile to 8.51 at the 95th percentile, with a median elasticity of 6.48. These estimated elasticities are in line with those estimated in other studies using barcode data and imply substantially more substitution between barcodes than implied by an elasticity of zero in the conventional Laspeyres index 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. We comfortably reject the null hypothesis of an elasticity of substitution of one or zero at conventional levels of statistical significance for all sectors. 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.

Finally, we now highlight the tension between combining a variety adjustment term based on the estimation

of a CES demand system and a Sato-Vartia price index for continuing varieties. In particular, under the Sato-Vartia index’s assumption of no taste shock for any continuing variety, we can directly solve for the elasticity of substitution for each pair of time periods, as shown in equation (A.24) in Section A.4 of the online appendix. If the Sato-Vartia index’s assumption were 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 sector. We expect these estimates to vary by sector, so we compute the dispersion of these estimates relative to the sector mean, or $(\sigma_{gt}^{SV} - \frac{1}{T} \sum_t \sigma_{gt}^{SV})$, where T is the number of periods. In the absence of demand shocks, we expect this number to be zero.

In Table III, we report the mean and median of $\frac{1}{T} \sum_t \sigma_{gt}^{SV}$ in the first two columns and moments of the distribution of $(\sigma_{gt}^{SV} - \frac{1}{T} \sum_t \sigma_{gt}^{SV})$ 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, and the mean elasticity is also negative, with a standard deviation of 196. Therefore, we find strong evidence against the assumption of the Sato-Vartia index that movements in expenditure shares reflect only changes in relative prices. In contrast, once we allow for time-varying tastes using the Feenstra (1994) estimator, we obtain plausible estimated elasticities of substitution in our baseline specification above.

V.B. Properties of the Demand Residuals

Using our estimates of the elasticity of substitution (σ) for each sector, we invert the CES demand system to solve for the time-varying demand residuals ($\ln \varphi_{kt}$) for each variety, as shown in equation (12). As discussed above, there are the two potential approaches that one can take with respect to these time-varying demand residuals: (i) one can interpret them as consumer taste shocks that are equivalent to price shocks and compute the change in the cost of living using “quality-adjusted” prices; (ii) or one can hold the taste parameters constant (e.g. at their initial values) and compute the change in the cost of living using the observed price shocks and ignoring changes in the demand residuals.

In this section, we provide evidence that these time-varying demand residuals have systematic and intuitive properties that are consistent with our treatment of them as consumer taste shocks. In particular, we estimate the following first-order autoregressive process for the log demand residuals for each sector g using our baseline sample:

$$\ln \varphi_{kt} = \rho_g \ln \varphi_{kt-1} + d_{gt} + \epsilon_{kt}, \quad (34)$$

where we pool observations across barcodes and over time for each sector; the autoregressive parameter (ρ_g) captures the degree of serial correlation in the demand residuals over time; our normalization requires that the mean log change in the demand residuals is equal to zero across common varieties within each sector, but we include the year fixed effects (d_{gt}) for each sector to control for common macro shocks; the stochastic error (ϵ_{kt}) captures idiosyncratic shocks to the demand residuals for each barcode; we also consider augmented versions of this specification including age, firm, brand or barcode fixed effects.

In Table IV, we report the results of estimating this regression (34) for each sector separately. In the top panel, we give the estimated coefficient on the lagged dependent variable. In the bottom panel, we list the R^2

of the regression. In each case, we present percentiles of the distribution of the estimates across sectors. In Column (1), we estimate equation (34) including only the lagged dependent variable and year fixed effects. We find a positive and statistically significant coefficient on the lagged dependent variable that is just below one and a high regression R^2 , which reassuringly suggests that our estimates of consumer tastes are persistent over time—in three quarters of the sectors more than 77 percent of the variation in consumer tastes for a barcode in period t can be explained by the consumer tastes for the barcode in period $t - 1$.

In Column (2), we augment this specification with age fixed effects, which are statistically significant at conventional critical values in this and all subsequent specifications. We find an intuitive pattern in which the estimated coefficients on the age fixed effects decline in the first year of a barcode’s existence and then are stable, which is consistent with consumers valuing novelty. Both the coefficient on the lagged dependent variable and the regression R^2 , however, remain close to unaffected. In Column (3), we augment this specification with firm fixed effects. We find that these estimated firm fixed effects are highly statistically significant, which suggests that our estimated consumer tastes are capturing systematic differences in the appeal of the barcodes supplied by different firms (e.g. because of advertising, branding and marketing). Controlling for these persistent characteristics of firms somewhat increases the explanatory power of the regression (as reflected in the regression R^2) and reduces the coefficient on the lagged dependent variable.

In Column (4), we replace the firm fixed effects with brand fixed effects and find a similar pattern of results, both in terms of the coefficient on the lagged dependent variable and the regression R^2 . In Column (5), we replace the brand fixed effects with barcode fixed effects. We find that these barcode fixed effects are also highly statistically significant, which is consistent with consumers valuing some of the barcodes supplied by a firm more than others (e.g. because of superior product design and characteristics). Once we control for these persistent characteristics of barcodes, we find a lower, but still statistically significant, coefficient on the lagged dependent variable and an additional increase in the regression R^2 . This lower estimated coefficient on the lagged dependent variable is in line with the idea that consumers tastes for individual colors, sizes, styles and models of goods can fluctuate from one year to the next with fads and changes in fashion. The increase in the regression R^2 is consistent with the fact that many barcodes are present in the data for only a few years, such that the barcode fixed effects capture much of the variation in consumer tastes (they capture all of the variation for barcodes only present for a single cross-section between years $t - 1$ and t).

To provide additional evidence on the extent to which our estimates capture consumer tastes rather than measurement or specification error, we obtained data from Young and Rubicam (the U.S. subsidiary of the world’s largest marketing firm WPP). Young and Rubicam measure consumer preferences or “brand asset values” (BAVs) by conducting annual surveys of approximately 17,000 U.S. consumers. These BAVs are composed of four basic components: “energized differentiation” measures perceived innovativeness of a product; “relevance” measures whether a product is suitable for consumers given their preferences; “esteem” captures brand prestige; and “knowledge” measures how familiar consumers are with a brand. Since marketing data is reported at the level of the brand rather than the barcode, we estimate the nested CES specification from Section III.B. above with sectors and brands as our nests, as discussed in Section A.14A. of the online appendix. Inverting the demand system, we recover estimates of consumer tastes for each brand and for each barcode

within each brand, where we normalize brand tastes to have a constant geometric mean within each sector and barcode tastes to have a constant geometric mean within each brand.

We begin by regressing the level of our estimates of brand tastes on each of the BAV components, including sector-time fixed effects to control for common macro shocks across brands within each sector. As shown in Section A.14A. of the online appendix, we find that our estimates of brand tastes are positively and significantly correlated with each of the four BAV components. This finding that brands with high estimated demand residuals correspond to brands that consumers rate highly in surveys is consistent with the idea that our estimated demand residuals do indeed capture consumer tastes. As the taste-shock bias in the conventional Sato-Vartia index depends on changes in tastes, we next regress our estimated brand tastes against BAVs in a specification that also includes brand and sector-time fixed effects. The inclusion of the brand fixed effects means that the relationship between our estimates of consumer tastes and the BAV components is identified solely from time-series variation. We find that our estimates of consumer tastes are positively correlated with each of the four BAV components, and the coefficients on relevance, esteem, and knowledge are statistically significant at conventional critical values. Therefore, in both levels and changes, our estimated demand residuals are systematically related to separate measures of brand asset values from consumer surveys, consistent with them capturing consumer tastes.

Taken together, the results of this section are consistent with the view that each barcode has some underlying level of consumer appeal based on its time-invariant physical attributes (captured by the barcode fixed effect), and consumer tastes for the barcode evolve stochastically over time around this underlying level of appeal. In the Sato-Vartia index, these stochastic shocks to tastes are incorporated into the expenditure share weights (implicitly including changes in consumer tastes), but are omitted from the price terms (excluding changes in consumer tastes). In contrast, our CUPPI consistently treats these demand residuals as consumer tastes in both the expenditure-share weights and price terms.

V.C. 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 varieties, and differences in assumptions about taste 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 assumptions of no entry and exit and time-invariant tastes for each common variety. Therefore, the differences between our new CES unified price index and existing price indexes in the next section reflect the treatment of entry/exit and taste shocks for common varieties rather than alternative assumptions about substitution patterns between varieties.

For each sector and time period, we compute five conventional price indexes: (i) the Laspeyres index, which assumes a zero elasticity of substitution and weights varieties by their 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 tastes; (iv) the Törnqvist index, which is also

superlative and is exact for translog preferences with time-invariant tastes; and (v) the Sato-Vartia index, which is exact for CES preferences with time-invariant tastes. All of these price indexes are defined for common varieties. We use our baseline definition of the set of common varieties for the fourth quarters of years $t - 1$ and t , which includes barcodes that appear in those two periods and for more than six years, and that are available for three quarters before the fourth quarter of year $t - 1$ and for three quarters after the fourth quarter of year t . With nine years and 104 sectors, we have a sample of just over 800 price changes across sectors and over time.

In Figure I, we display kernel density estimates of the distribution of four-quarter price changes across sectors 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 indexes yield almost exactly the same change in the cost of living as the Fisher index. The standard deviation of the difference between the Sato-Vartia and Fisher indexes is 0.06 percentage points per year, which compares closely with the corresponding standard deviation of the difference between the Törnqvist and Fisher indexes of 0.04 percentage points per year. In contrast, indexes that assume an elasticity of substitution of zero (the Laspeyres index) or one (the Cobb-Douglas index) have standard deviations from the Fisher index that both equal 0.3 percentage points per year—about five times larger than that for the Sato-Vartia index—and the Laspeyres index has an upward bias of around 0.5 percent per year.

Therefore, these results suggest that assuming a CES functional form instead of a flexible functional form (as for the Fisher and Törnqvist price indexes) has relatively little impact on the measured change in the cost of living under the common set of assumptions of no entry and exit and no taste shocks for common varieties.

V.D. The CES Unified Price Index

We now maintain the assumption of CES preferences but allow for the entry and exit of varieties and taste shocks for individual common varieties. Using our estimated elasticities of elasticity of substitution (σ), we compute our common variety price index (9), the variety correction term, and hence our unified price index (8) for each sector using our baseline definition of the set of common varieties.

We start with the Feenstra (1994) variety adjustment term that captures the impact of entry and exit. This term depends on both the elasticity of substitution (σ_g) and relative expenditure shares on common varieties ($\lambda_{gt}/\lambda_{gt-1}$). It controls both the difference between the Feenstra and Sato-Vartia price indexes and the difference between our CUPI in equation (8) and our CCV in equation (9). In Figure II, we display a histogram of the relative expenditure shares on common varieties ($\lambda_{gt}/\lambda_{gt-1}$) across sectors and over time. If entering varieties had similar characteristics to exiting varieties, the prices and market shares of exiting varieties would match those of new varieties, resulting in a $\lambda_{gt}/\lambda_{gt-1}$ ratio of one. The fact that these ratios are almost always less than one indicates that new varieties tend to be more attractive than disappearing ones. In barcode data, this variety upgrading is fully captured in the entry and exit term, because as discussed above any change in physical attributes of a variety leads to the introduction of a new barcode.

We now quantify the relative magnitude of the biases from abstracting from entry and exit and taste shocks

for common varieties. 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 sector and time period, we compute alternative measures of changes in the cost of living, and then aggregate across sectors using expenditure-share weights to compute a measure of the change in the aggregate cost of living.

In Figure III, we plot the resulting measures of the change in the aggregate cost of living using our CUPI and a range of alternative price indexes. We show the 95 percent bootstrapped confidence intervals for the CUPI using gray shading.¹¹ 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 taste shocks for any common variety. 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.

The bias from abstracting from the entry and exit of varieties can be seen in Figure III by comparing the CUPI and the CCV price indexes (from equations (9) and (8)). We find a substantial impact of entry and exit on the measurement of the cost of living, equal to around one percentage point per year. Therefore, if one abstracts from the fact that new varieties tend to be systematically better than disappearing varieties (as measured in the CES demand system by their relatively greater expenditure shares), one systematically overstates the increase in the cost of living over time.

The taste-shock bias from neglecting taste shocks for individual common varieties can be discerned in Figure III from comparing our CUPI and the Feenstra index. Both of these price indexes are exact for CES preferences and allow for the entry and exit of varieties. However, the Feenstra price index uses the Sato-Vartia index for common varieties (assuming time-invariant tastes for each common variety), whereas the CUPI uses the CCV price index for common varieties (allowing for changes in relative tastes across common varieties). As shown in the figure, we find that this bias is around half as large as that from abstracting from entry and exit (0.4 percentage points per annum).¹² Therefore, the internal inconsistency in the Sato-Vartia index from including time-varying demand residuals in expenditure share weights but not in measured prices has quantitatively relevant effects on the measurement of changes in the cost of living. Conventional price indexes overstate the increase in the cost of living over time, because other things equal varieties experiencing an increase in tastes (for which the change in observed prices is greater than the true change in taste-adjusted prices) receive a higher expenditure-share weight than varieties experiencing a decrease in tastes (for which the change in observed prices is smaller than the true change in taste-adjusted prices).

¹¹We randomly sample barcodes by sector with replacement for 100 bootstrap replications. For each replication, we estimate the elasticities of substitution for each sector, and recompute the CUPI using our baseline definition of common varieties. Using the distribution of values for the CUPI across these 100 replications, we construct the 95 percent confidence intervals.

¹²The average value from 2005 to 2013 of the Paasche index is 1.7 percent; the Laspeyres, 2.6; the CCV is 1.8; the CUPI is 0.9 percent; and the Feenstra-CES is 1.3 percent.

V.E. Specification Checks

Our baseline specification assumes both CES demand and that consumer tastes have a constant geometric mean across common goods. Under these assumptions, we recover a time-varying consumer taste parameter for each good, such that our model exactly rationalizes the observed data on prices and expenditure shares. In this section, we present two specification checks on the reasonableness of these assumptions. In the next section, we provide further evidence on the adequacy of these assumptions by comparing our baseline CES estimates with those from a more flexible mixed CES specification.

Our first specification check uses the IIA property of CES, which implies that the change in the cost of living can be computed either (i) using all common varieties and an entry/exit term or (ii) choosing a subset of common varieties and adjusting the entry/exit term for the omitted common varieties. If preferences are CES and taste shocks average out across varieties such that the geometric means of tastes are constant for both definitions of common varieties, we should obtain the same change in the cost of living from these two different specifications:

$$\frac{P_t}{P_{t-1}} = \left(\frac{\lambda_t}{\lambda_{t-1}} \right)^{\frac{1}{\sigma-1}} \frac{\tilde{p}_t}{\tilde{p}_{t-1}} \left(\frac{\tilde{s}_t^*}{\tilde{s}_{t-1}^*} \right)^{\frac{1}{\sigma-1}} = \left(\frac{\mu_t}{\mu_{t-1}} \right)^{\frac{1}{\sigma-1}} \frac{\tilde{\tilde{p}}_t}{\tilde{\tilde{p}}_{t-1}} \left(\frac{\tilde{\tilde{s}}_t^{**}}{\tilde{\tilde{s}}_{t-1}^{**}} \right)^{\frac{1}{\sigma-1}}, \quad (35)$$

where $\Omega_t^{**} \subset \Omega_t^* \subset \Omega_t$ denotes a subset of common varieties; μ_t and μ_{t-1} are the shares of this subset in total expenditure:

$$\mu_t = \frac{\sum_{k \in \Omega_t^{**}} p_{kt} c_{kt}}{\sum_{k \in \Omega_t} p_{kt} c_{kt}}, \quad \mu_{t-1} = \frac{\sum_{k \in \Omega_{t-1}^{**}} p_{kt-1} c_{kt-1}}{\sum_{k \in \Omega_{t-1}} p_{kt-1} c_{kt-1}}. \quad (36)$$

Building on our earlier notation, we use s_{kt}^{**} to denote the share of an individual common variety in overall expenditure on this subset of common varieties:

$$s_{kt}^{**} = \frac{(p_{kt} / \varphi_{kt})^{1-\sigma}}{\sum_{\ell \in \Omega_t^{**}} (p_{\ell t} / \varphi_{\ell t})^{1-\sigma}}; \quad (37)$$

where we use a double tilde to denote a geometric mean across this subset of common varieties: $\tilde{\tilde{s}}_t^{**} = \left(\prod_{k \in \Omega_t^{**}} s_{kt}^{**} \right)^{\frac{1}{N_t^{**}}}$; $N_t^{**} = |\Omega_t^{**}|$ is the number of varieties in this subset; and the derivations for all results in this section are reported in Section A.12 of the online appendix. In contrast, the two expressions for the change in the cost of living in equation (35) differ in general if (i) preferences are not CES and/or (ii) the geometric means of consumer tastes are not constant for both definitions of common varieties (in which case these two expressions differ by the ratio of the changes in these geometric means).

We implement this joint specification test using alternative definitions of the set of common varieties. We start with the most restrictive definition, in which we require barcodes to be present in all T years of our sample. Using this definition, we compute the change in the cost of living for each four-quarter difference, and chain these four-quarter differences to measure the change in the cost of living over time. We next progressively relax this definition, such that the set of common varieties for the fourth quarters of years $t-1$ and t is defined as the subset of varieties present in those periods and for a total number of x years, where $1 < x \leq T$. For each value of x , we again chain the four-quarter differences to measure the change in the cost of living over time. We

continue this process until we arrive at the most inclusive definition of common varieties, which corresponds to the set of varieties present in the fourth quarters of years $t - 1$ and t ($x = 2$).

In Figure IV, we show the change in the aggregate cost of living for each year in our sample for these alternative definitions of the set of common varieties, where we again aggregate across sectors using expenditure-share weights. For common variety definitions using thresholds (x) of more than six years or above, we find a relatively similar change in the cost of living, which is within or close to the 95 percent confidence interval for our baseline threshold of more than six years. This pattern of results suggests that the joint assumption of CES preferences and a constant geometric mean of tastes across each of these alternative definitions of common varieties provides a reasonable approximation to the data. For common goods definitions using thresholds of less than six years, we find a lower change in the cost of living, which lies outside the 95 percent confidence for our baseline threshold of more than six years. This pattern of results is consistent with these more inclusive definitions of common goods including a disproportionate number of varieties that are only present in the sample for a few years. As changes in the estimated time-varying demand residuals are greatest immediately after entry, the inclusion of these barcodes with strong demand dynamics increases the magnitude of the taste-shock bias, and reduces the CUPI relative to conventional price indexes. Based on the stability of our empirical results using thresholds of more than six years or above, and to be conservative in terms of the magnitude of the taste-shock bias, we used a threshold of more than six years for the definition of common varieties in our baseline specification, as discussed above.

In our second specification check, we use the class of generalized means of order r to compute the change in the cost of living for alternative normalizations for consumer tastes. In Figure V, we show the aggregate change in the cost of living for each year in our sample using our baseline definition of common varieties for different values for r , ranging from a constant harmonic mean ($r = -1$), through a constant geometric mean ($r = 0$) and a constant arithmetic mean ($r = 1$), to a constant quadratic mean ($r = 2$). We find a substantial taste-shock bias across these different values of r , with the CUPI falling further below the Fisher index as we increase r . For example, the average percentage point four-quarter differences between the Fisher Index and the CUPI over the period 2005-2013 are as follows: 0.016 ($r = -1$), 0.013 ($r = 0$), 0.015 ($r = 1$) and 0.014 ($r = 2$). Therefore, our baseline specification of a constant geometric mean ($r = 0$) is again relatively conservative in terms of the magnitude of the taste-shock bias.

As a final check on the sensitivity of our results to alternative normalizations, we recompute the log of the CUPI in equation (8) using the initial expenditure-share-weighted average of the ratio of prices and expenditure shares for each common variety rather than the unweighted average:

$$\ln \left(\frac{P_t}{P_{t-1}} \right) = \frac{1}{\sigma - 1} \ln \left(\frac{\lambda_t}{\lambda_{t-1}} \right) + \sum_{k \in \Omega_t^*} s_{kt-1}^* \left[\ln \left(\frac{p_{kt}}{p_{kt-1}} \right) + \frac{1}{\sigma - 1} \ln \left(\frac{s_{kt}^*}{s_{kt-1}^*} \right) \right].$$

This specification normalizes the initial-expenditure-share weighted average (instead of the unweighted average) of the demand shocks to zero: $\sum_{k \in \Omega_t^*} s_{kt-1}^* \ln(\varphi_{kt}/\varphi_{kt-1}) = 0$. As also shown in Figure V, our exact price index again lies below the Fisher index, with the absolute magnitude of this difference typically larger using the initial-expenditure-weighted mean than in our baseline specification (CUPI, $r = 0$).

Taken together, these specification checks confirm the robustness of our finding of a substantial taste-shock

bias to alternative normalizations for consumer tastes, and suggest that our joint assumption of CES preferences and log taste shocks that average out across common varieties provides a reasonable approximation to the data for sets of common varieties present for a relatively large number of years.

VI. Mixed CES Specification

In this section, we examine the robustness of our results to relaxing the independence of irrelevant alternatives (IIA) and symmetric substitution properties of CES by considering a mixed CES specification with heterogeneous groups of consumers, as discussed in Section III.C. above. We use low-income and high-income households as our two groups, based on those households with above-median and below-median income. Although the income differences between these two groups are substantial (recall that median income is around \$50-59,000), they are of course smaller than in other settings, such as in developed versus developing countries. We allow both the elasticity of substitution and the taste parameter for each variety to differ between these two groups of households. Therefore, this specification incorporates non-homotheticities between these two groups in a more flexible way than the non-homothetic CES specification in Section III.A. above, which imposed a common elasticity of substitution for all consumers.

In Figure VI, we report our estimates of the elasticities of substitution for high- and low-income households using the Feenstra (1994) estimator. As shown in the figure, we find similar estimated elasticities for the pooled, high-income, and low-income samples, with a correlation between the estimated elasticities for the high- and low-income households of 0.80. Therefore, although this mixed CES specification allows in principle for heterogeneity in estimated elasticities of substitution, we find in practice similar substitution behavior for these two groups of households in our data on barcoded goods.

In Figure VII, we show a bin scatter of the log taste parameters ($\ln \varphi_{kt}$) for each group of households against the average of the two groups. We pool observations across sectors, where the log taste parameters for each sector and year are normalized to have a mean of zero. We use a bin scatter with 100 percentiles and also display the regression line. We find a strong positive and statistically significant relationship between the taste parameters for the two groups, with a correlation of 0.96, as reflected in the regression line lying close to the diagonal. Therefore, on average, we find strong agreement between high- and low-income households about which varieties are more or less appealing.

Another feature of Figure VII is that the slope for low-income households is lower than 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 varieties 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 (2019), which showed that the average change in the cost of living for low-income households exceeds that for high-income households by 0.65 percent per year for common varieties and by 0.78 percent per year once the entry and exit of varieties is taken into account. We find the same pattern of differences in the cost of living between the two groups, as shown in Figure VIII using our baseline definition of the set of common varieties. On average, the CCV and the CUPI price indexes

for low-income households are 0.22 and 0.37 percent per year higher than those for high-income households. Therefore, our price indexes capture the same properties of the data as found in other studies.

We now examine the magnitude of the taste-shock bias for the two groups of households. As evident from Figure VIII, 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 around one fifth as large as the variance in the change in the cost of living measured on average for each year. Over the full sample period, the CCV rose by 2.1 percent per year on average, which compares to 2.1 percent for the CCV for low-income households and 1.9 percent for the CCV for high-income households. We see a similar pattern for the CUPI, which rose by 1.4 percent on average, compared to 1.4 and 1.1 percent for low- and high-income households respectively.

Taken together these results suggest that while we can find evidence of heterogeneity in the taste parameters for individual varieties between high- and low-income households, we find similar elasticities of substitution across varieties for these two groups. Furthermore, the heterogeneity in taste parameters does not shift the cost of living for each group of households substantially away from our central estimate.

VII. Further Robustness Checks

In this section, we report a number of further robustness checks. We examine the sensitivity of our measured changes in the cost of living to the Feenstra (1994) estimated elasticities. Next, we illustrate the relevance of our results for official measures of the consumer price index (CPI). Finally, we demonstrate the robustness of our results to the treatment of varieties with smaller expenditure shares for which measurement error could be relatively more important.

To assess the first point, we use a grid search over the parameter space to demonstrate the robustness of our results across the range of plausible values for the elasticity of substitution. In particular, we consider a grid of thirty-eight evenly-spaced values for this elasticity ranging from 1.5 to 20. For each value on the grid, we compute our CCV and CUPI for each sector and year, and aggregate across sectors using expenditure-share weights. In Figure A.3 in Section A.14B. of the online appendix, we compare these changes in the cost of living to the Fisher index. A smaller elasticity of substitution implies that varieties are more differentiated, which increases the absolute magnitude of the variety correction term for entering varieties being more desirable than exiting varieties ($(1/(\sigma - 1)) \ln(\lambda_t/\lambda_{t-1}) < 0$). As a result, the CCV and CUPI fall further below the Fisher index as the elasticity of substitution becomes small. Nevertheless, across the entire range of plausible values for this elasticity, we find a quantitatively relevant taste-shock bias.

Next, we illustrate the relevance of results computed using the Nielsen Homescan data for official measures of the CPI by mapping 89 of our 104 sectors to official CPI categories, as discussed further in Section A.14C. of the online appendix. We again aggregate across sectors using expenditure share weights to construct measures of the aggregate cost of living. As shown in Figure A.4 of the online appendix, we find that conventional price indexes computed using the Homescan data are remarkably successful in replicating properties of official price indexes, with a positive and statistically significant correlation of 0.99 between the Laspeyres (based on the

Homescan data) and the CPI. In other words, annual movements in changes in the cost of living as measured by the BLS for these sectors can be closely approximated by using a Laspeyres index and the Nielsen data. We find that 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 larger bias from abstracting from both entry/exit and taste shocks, with our CUPI more than one percentage point below the CPI. As a further validation of our measures of changes in the cost of living using the Nielsen Homescan data, Appendix A.14E. repeats our analysis using the Nielsen Retail Scanner Dataset as an alternative source of scanner data. In contrast to the Homescan data, these retail scanner data are not nationally representative, because they are based on a non-random set of participant stores, and purchases not made at these stores are omitted. Nonetheless, we show that we find the same pattern of year-to-year changes in the cost of living, with a substantial taste-shock bias.

Finally, we examine the sensitivity of our results to measurement error in expenditure shares for varieties that account for small shares of expenditure using the properties of CES demand discussed in Section V.E. above. In particular, we use the property that the change in the cost of living can be computed either (i) using all common varieties and an entry/exit term or (ii) choosing a subset of common varieties and adjusting the entry/exit term for the omitted common varieties. Using this property, we recompute the CUPI using the subset of our baseline sample of common varieties with above-median expenditure shares. This specification is less sensitive to measurement error for varieties that account for small shares of expenditure, because expenditures on varieties with below-median expenditure shares only enter the change in the cost of living through the aggregate share of expenditure on varieties with above-median expenditure shares. As reported in Section A.14D. of the online appendix, we find a similar change in the aggregate cost of living as in our baseline specification above.

VIII. Conclusions

Measuring price aggregates is central to international trade and macroeconomics, which depend critically on being able to distinguish real and nominal income. In such an analysis, one typically faces the challenge that whatever preferences are assumed do not perfectly fit the data in both time periods without time-varying demand residuals. We show that a key building block for the existing exact price index for CES preferences (the Sato-Vartia index for varieties common to a pair of time periods) implicitly treats these demand residuals in an inconsistent way. On the one hand, this price index assumes time-invariant tastes, and uses observed price changes and expenditure shares to compute changes in the cost of living. On the other hand, the observed final-period expenditure shares used in this price index include the time-varying demand residuals.

In this paper, we develop a new exact price index for CES preferences that consistently treats demand shocks as taste shocks that are equivalent to price shocks. This exact price index expresses the change in the cost of living solely in terms of observed *prices* and *expenditures*. As expenditures depend on *relative* consumer tastes (and not the absolute level of consumer tastes), the existence of such an exact price index requires that we rule out the possibility of a change in the cost of living, even though all prices and expenditures remain

unchanged. To rule out such an equiproportional change in tastes, we assume that the taste parameters have a constant geometric mean across common varieties, which is consistent with the assumption that the log demand shocks are mean zero in the estimation of the demand system. We demonstrate the robustness of our results to alternative normalizations that rule out such a pure change in consumer tastes using the class of generalized means.

Our approach uses the invertibility of the CES demand system to recover unique values for unobserved consumer tastes for each variety (up to our normalization). We use this result to express the change in the cost of living in terms of only prices and expenditure shares, while allowing for changes in *relative* consumer tastes across varieties. We show that the Sato-Vartia index is subject to a substantial taste-shock bias, because it fails to take into account that an increase in taste for a variety is analogous to a fall in its price. This failure leads to a systematic overstatement of the change in the cost of living, because consumers substitute towards varieties that become more desirable. Therefore, other things equal, varieties experiencing an increase in tastes (for which the change in observed prices is greater than the true change in taste-adjusted prices) receive a higher expenditure-share weight than varieties experiencing a decrease in tastes (for which the change in observed prices is smaller than the true change in taste-adjusted prices). In our empirical application using barcode data, we show that the taste-shock bias is around 0.4 percentage points per year, and is sizable relative to the bias from abstracting from the entry and exit of varieties.

Although we focus on CES preferences because of their prominence in international trade, macroeconomics and economic geography, we show that our approach generalizes to other invertible demand systems, including non-homothetic CES (indirectly additive), nested CES, mixed CES, logit, mixed logit, translog and AIDS. In each case, conventional price indexes assume time-invariant tastes and interpret all movements in expenditure shares as the result of changes in prices, but use observed final-period expenditure shares, which are influenced by time-varying demand residuals. Through failing to recognize that an increase in tastes is analogous to a reduction in price, these conventional price indexes are subject to the taste-shock bias. In contrast, our approach of inverting the demand system to express unobserved taste shocks in terms of observed prices and expenditure shares can be used to derive an exact price index that treats these time-varying demand residuals in an internally consistent way.

Princeton University, NBER and CEPR

Columbia University and NBER

References

- [1] Adao, Rodrigo, Arnaud Costinot, and Dave Donaldson, “Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade,” *American Economic Review*, 107 (2017), 633-689.
- [2] Anderson, Simon, André de Palma, and Jacques-François Thisse, *Discrete Choice Theory of Product Differentiation* (Cambridge, MA: MIT Press, 1992).
- [3] Antràs, Pol, “Firms, Contracts, and Trade Structure,” *Quarterly Journal of Economics*, 118 (2003), 1375-1418.
- [4] Arkolakis, Costas, Arnaud Costinot, and Andrés Rodríguez-Clare, “New Trade Models, Same Old Gains,” *American Economic Review*, 102 (2012), 94-130.
- [5] Atkin, David and Dave Donaldson, “Who’s Getting Globalized? The Size and Implications of Intra-national Trade Costs,” NBER Working Paper, 21439, 2015.
- [6] Atkin, David, Benjamin Faber, and Marco Gonzales-Navarro, “Retail Globalization and Household Welfare: Evidence from Mexico,” *Journal of Political Economy*, 126 (2016), 1-73.
- [7] Atkin, David, Ben Faber, Thibault Fally, and Marco Gonzalez-Navarro, “A New Engel on the Gains from Trade,” MIT, mimeograph (2018).
- [8] Balk, Bert M., “On Curing the CPI’s Substitution and New Goods Bias,” Department of Statistical Methods, Statistics Netherlands, mimeograph (1999).
- [9] Behrens, Kristian, Yoshitsugu Kanemoto and Yasusada Murata, “New Trade Models, Elusive Welfare Gains,” CEPR Discussion Paper, 10255, 2014.
- [10] Benassy, Jean-Pascal, “Taste for Variety and Optimum Production Patterns in Monopolistic Competition,” *Economics Letters*, 52 (1996), 41-47.
- [11] Berry, Steven, James Levinsohn, and Ariel Pakes, “Automobile Prices in Market Equilibrium,” *Econometrica*, 60 (1995), 889-917.
- [12] Berry, Steven, Amit Gandhi, and Philip Hale, “Connected Substitutes and Invertibility of Demand,” *Econometrica*, 81 (2013), 2087-2111.
- [13] Białek, Jacek, “Approximation of the Fisher Price Index by Using the Lloyd–Moulton Index: Simulation Study,” *Communications in Statistics - Simulation and Computation*, 46 (2017), 3588-3598.
- [14] Bills, Mark and Klenow, Peter, “Quantifying Quality Growth,” *American Economic Review*, 91 (2001), 1006-1030.
- [15] Boskin, Michael J., Ellen R. Dulberger, Robert J. Gordon, Zvi Griliches, and Dale W. Jorgenson, “Toward a More Accurate Measure of the Cost of Living,” Final Report to the Senate Finance Committee. (Washington, DC: Senate Finance Committee, 1996).
- [16] Broda, Christian and David E. Weinstein, “Globalization and the Gains from Variety,” *Quarterly Journal of Economics*, 121 (2006), 541-86.
- [17] Broda, Christian and David E. Weinstein, “Product Creation and Destruction: Evidence and Price Implications,” *American Economic Review*, 100 (2010), 691–723.
- [18] Caves, Douglas W., Laurits R. Christensen, and W. Erwin Diewert, “The Economic Theory of Index Numbers and the Measurement of Input, Output and Productivity,” *Econometrica*, 50 (1982), 1393-1414.

- [19] Comin, Diego A., Danial Lashkari, and Martí Mestieri, “Structural Change with Long-run Income and Price Effects,” NBER Working Paper, 21595, 2015.
- [20] Diewert, W. E., “Exact and Superlative Index Numbers,” *Journal of Econometrics*, 4 (1976), 115-145.
- [21] Diewert, Erwin, *Consumer Price Index Manual: Theory and Practice*, (Washington DC: International Monetary Fund, 2004).
- [22] Diewert, Erwin and Robert Feenstra, “Estimating the Benefits and Costs of New and Disappearing Products,” University of California at Davis, mimeograph (2017).
- [23] Dixit, Avinash and Joseph Stiglitz, “Monopolistic Competition and Optimum Product Diversity,” *American Economic Review*, 67 (1977), 297-308.
- [24] Fajgelbaum, Pablo D. and Amit K. Khandelwal, “Measuring the Unequal Gains from Trade,” *Quarterly Journal of Economics*, 131 (2016), 1113-1180.
- [25] Fally, Thibault and Benjamin Faber, “Firm Heterogeneity in Consumption Baskets: Evidence from Home and Store Scanner Data,” NBER Working Paper, 23101, 2017.
- [26] Feenstra, Robert C., “New Product Varieties and the Measurement of International Prices,” *American Economic Review*, 84 (1994), 157-177.
- [27] Feenstra, Robert C. and Marshall B. Reinsdorf, “Should Exact Index Numbers Have Standard Errors? Theory and Application to Asian Growth,” in *Hard-to-Measure Goods and Services: Essays in Honor of Zvi Griliches*, Ernst R. Berndt and Charles R. Hulten, ed. (Chicago: University of Chicago Press, 2007).
- [28] Feenstra, Robert C. and Marshall B. Reinsdorf, “An Exact Price Index for the Almost Ideal Demand System,” *Economics Letters*, 66 (2010), 159-162.
- [29] Fisher, Franklin M. and Karl Shell, *The Economic Theory of Price Indices: Two Essays on the Effects of Taste, Quality, and Technological Change*, (New York: Academic Press, 1972).
- [30] Handbury, Jessie, “Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities,” Wharton, mimeograph (2013).
- [31] Hanoch, Giora, “Production and Demand Models with Direct or Indirect Implicit Additivity,” *Econometrica*, 43 (1975), 395-419.
- [32] Hausman, Jerry, “Valuation of New Goods Under Perfect and Imperfect Competition,” in *The Economics of New Goods*, Tim Bresnahan and Robert Gordon, ed., 209-237, (Chicago: University of Chicago Press, 1997) .
- [33] Hottman, Colin, Stephen Redding and David Weinstein, “Quantifying the Sources of Firm Heterogeneity,” *Quarterly Journal of Economics*, 131 (2016), 1291-1364.
- [34] Hsieh, Chang-Tai and Peter Klenow, “Misallocation and Manufacturing TFP in China and India,” *Quarterly Journal of Economics*, 124 (2009), 1403-1448.
- [35] Jaravel, Xavier, “The Unequal Gains from Product Innovation: Evidence from the US Retail Sector,” *Quarterly Journal of Economics*, forthcoming (2019).
- [36] Jones, Charles I. and Peter J. Klenow, “Beyond GDP? Welfare Across Countries and Time,” *American Economic Review*, 106 (2016), 2426-2457.

- [37] Klepper, Steven and Edward E. Leamer, "Consistent Sets of Estimates for Regressions with Errors in All Variables," *Econometrica*, 52 (1984), 163-184.
- [38] Konüs, A., "The Problem of the True Index of the Cost of Living," (1924) translated in *Econometrica*, 7 (1939), 10-29.
- [39] Krugman, Paul R., "Scale Economies, Product Differentiation and the Pattern of Trade," *American Economic Review*, 70 (1980), 950-959.
- [40] Krugman, Paul R., "Increasing Returns and Economic Geography," *Journal of Political Economy*, 99 (1991), 483-499.
- [41] Lau, L. J., "On Exact Index Numbers," *Review of Economics and Statistics*, 61 (1979), 73-82.
- [42] Lloyd, P. J., "Substitution Effects and Biases in Non-True Price Indices," *American Economic Review*, 65 (1975), 301-313.
- [43] Matsuyama, Kiminori, "Engel's Law in the Global Economy: Demand-Induced Patterns of Structural Change, Innovation, and Trade," *Econometrica*, 87 (2019), 497-528.
- [44] McFadden, Daniel, "Conditional Logit Analysis of Qualitative Choice Behavior," in *Frontiers in Econometrics*, Paul Zarembka, ed. (New York: Academic Press, 1974).
- [45] McFadden, Daniel and Kenneth Train, "Mixed MNL Models of Discrete Response," *Journal of Applied Econometrics*, 15 (2000), 447-470.
- [46] Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71 (2003), 1695-725.
- [47] Moulton, B. R., "Constant Elasticity Cost-of-Living Index in Share-Relative Form," U. S. Bureau of Labor Statistics, Washington DC, mimeograph (1996).
- [48] Neary, J. Peter, "Rationalizing the Penn World Table: True Multilateral Indices for International Comparisons of Real Income," *American Economic Review*, 94 (2004), 1411-1428.
- [49] Nevo, Aviv, "New Products, Quality Changes, and Welfare Measures Computed from Estimated Demand System," *Review of Economics and Statistics*, 85 (2003), 266-275.
- [50] Sato, R., "The Most General Class of CES Functions," *Econometrica*, 43 (1975), 999-1003.
- [51] Sato, Kazuo, "The Ideal Log-Change Index Number," *Review of Economics and Statistics*, 58 (1976), 223-228.
- [52] Shapiro, M. D. and D. W. Wilcox, "Alternative Strategies for Aggregating Prices in the CPI," *Federal Reserve Bank of St. Louis Review*, 79 (1997), 113-125.
- [53] Syverson, Chad, "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown," *Journal of Economic Perspectives*, 31 (2017), 165-186.
- [54] Train, Kenneth E., *Discrete Choice Methods with Simulation*, (Cambridge: Cambridge University Press, 2009).
- [55] Vartia, Yrjo, "Ideal Log-Change Index Numbers," *Scandinavian Journal of Statistics*, 3 (1976), 121-126.

TABLE I: Descriptive Statistics

	N	Mean	Sd	Min	P5	P25	P50	P75	P95	Max
Sector Sample Size	104	64,189	54,210	1,999	9,033	26,372	47,747	93,222	151,930	273,286
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	3.80	1.08	1.65	2.09	3.16	3.63	4.57	5.81	6.43
Mean λ_t	104	0.82	0.12	0.34	0.62	0.78	0.85	0.91	0.96	0.97
Mean λ_{t-1}	104	0.91	0.07	0.57	0.75	0.90	0.94	0.96	0.98	0.99
Mean $\frac{\lambda_t}{\lambda_{t-1}}$	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	38.33	9.64	20.89	24.32	30.93	37.80	43.10	55.36	66.60
Percent of UPCs that Exit in a Year	104	37.49	9.39	21.69	24.45	30.43	36.55	42.64	54.15	65.65
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 on the total value sold, total quantity sold, and average price; λ_t and λ_{t-1} are the shares of expenditure on common goods in total expenditure in year t and $t - 1$ respectively (four-quarter difference); N is the number of sectors; we compute statistics for each sector as the average value across years; mean, standard deviation ($\text{sd}(\cdot)$), maximum, minimum and percentiles p5-p95 are based on the distribution of these time-averaged values across sectors. 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.

TABLE II: Percentiles of the Distribution of Estimated Feenstra (1994) Elasticities of Substitution (σ) Across Sectors

Percentile	Estimated Feenstra Elasticity
Min	4.39
5th	5.11
25th	5.69
50th	6.48
75th	7.25
95th	8.51
Max	20.86

Note: Percentiles of the distribution of estimated Feenstra (1994) elasticities of substitution across sectors. 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.

TABLE III: Distribution of Sato-Vartia Elasticities for Each Year and Sector

	Mean	Median	Deviation from Time Mean					
			SD	p10	p25	p50	p75	p90
Sato-Vartia	-1.90	-2.55	196.27	-53.24	-16.64	-0.23	12.63	35.21

Note: Sato-Vartia elasticities are estimated for each sector and pair of time periods using equation (A.24) in Section A.4 of the online appendix; mean is the average of these elasticities across sectors and over time ($\frac{1}{GT} \sum_{t,g} \sigma_{gt}$); standard deviation is the average across sectors 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 sector; percentiles are based on the distribution across sectors of the standard deviation over time in these normalized elasticities ($\sigma_{gt} - \frac{1}{T} \sum_t \sigma_{gt}$) within each sector. We exclude the top and bottom one-percent market share changes within each sector 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.

TABLE IV: Persistence of Estimated Time-Varying Barcode Tastes (φ_{kt})

Percentile	(1)	(2)	(3)	(4)	(5)
Coefficient					
Min	0.81	0.81	0.69	0.51	-0.25
5th	0.83	0.83	0.74	0.70	-0.01
25th	0.89	0.89	0.81	0.79	0.05
50th	0.93	0.93	0.89	0.86	0.11
75th	0.96	0.96	0.94	0.93	0.17
95th	0.98	0.98	0.97	0.97	0.25
Max	0.99	0.99	0.99	0.99	0.50
R²					
Min	0.64	0.64	0.67	0.67	0.83
5th	0.68	0.69	0.70	0.71	0.85
25th	0.77	0.77	0.78	0.79	0.89
50th	0.85	0.85	0.86	0.86	0.93
75th	0.93	0.93	0.93	0.93	0.97
95th	0.96	0.96	0.96	0.96	0.98
Max	0.98	0.98	0.98	0.98	0.99
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes
Age Fixed Effects		Yes	Yes	Yes	Yes
Firm Fixed Effects			Yes		
Brand Fixed Effects				Yes	
Barcode Fixed Effects					Yes

Note: Columns in Table IV report the results of estimating the regression (34) of the log demand residuals for each barcode ($\ln \varphi_{kt}$) on the lagged dependent variable and the controls specified at the bottom of the table. Regression sample for each sector is an unbalanced panel that pools common, entering and exiting barcodes over time. Barcodes are nested within brands which are nested into firms which are in turn nested within sectors (product groups). The top panel reports percentiles of the estimated coefficient on the lagged dependent variable across sectors. The second panel reports percentiles of the estimated R^2 across sectors. 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 I: Differences in Price Indexes from the Fisher Index

Note: Kernel densities of the distribution across sectors and over time of the difference between price indexes and the Fisher index. Price indexes are measured as proportional four-quarter changes for each sector for our baseline sample of common barcodes: $\left(\frac{P_{gt}^* - P_{gt-1}^*}{P_{gt-1}^*}\right)$. SV-CES is the Sato-Vartia price index (the special case of equation (10) with $\varphi_{kt} / \varphi_{kt-1} = 1$ for all $k \in \Omega_t^*$). 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 II: Shares of Common Goods in Expenditure in period t relative to period $t - 1$ ($\lambda_{gt} / \lambda_{gt-1}$),
Four-Quarter Differences by Sector

Note: Histogram of relative expenditure shares on common barcodes across sectors g and over time t ($\lambda_{gt} / \lambda_{gt-1}$) using our baseline sample of common barcodes. 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 III: 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 sectors in our data $((P_{gt} - P_{gt-1}) / P_{gt-1})$ by their expenditure shares. CCV stands for CES common-variety price index (equation (9)). CUPI stands for CES unified price index (equation (8)). The grey band around the CUPI shows its 95 percent bootstrapped confidence interval (100 replications). Both the CUPI and Feenstra CES correct for the entry and exit of varieties, but the CUPI uses our CES common variety price index (equation (9)), whereas Feenstra-CES uses the Sato-Vartia price index for common varieties (the special case of equation (10) in which $\varphi_{kt} / \varphi_{kt-1} = 1$ for all $k \in \Omega_t^*$). 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 IV: CUPI for Alternative Definitions of the Set of Common Varieties

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 sectors in our data $((P_{gt} - P_{gt-1}) / P_{gt-1})$ by their expenditure shares. CUPI stands for CES unified price index (equation (8)). The CUPI is calculated for different definitions of the common set of barcodes. Each definition requires that a barcode is present in years $t - 1$ and t and a different total number of years in the sample. The grey band around the CUPI for our baseline definition of common goods (>6 year) shows its 95 percent bootstrapped confidence interval (100 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.

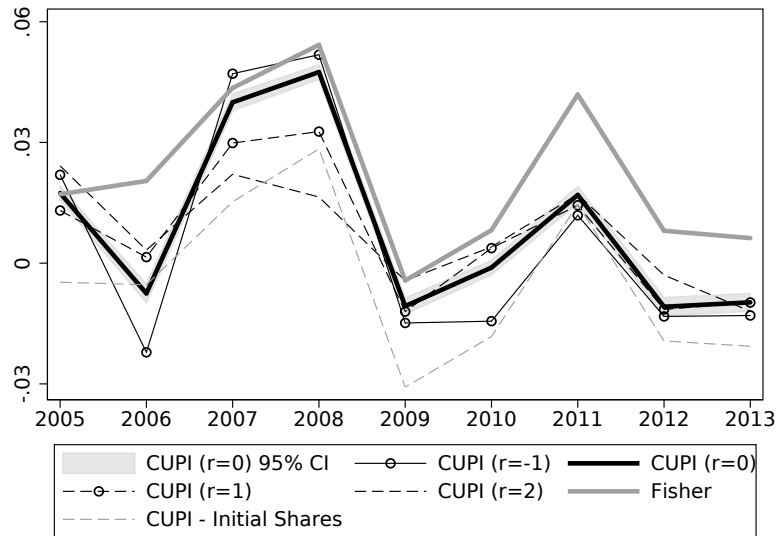


FIGURE V: CUPI and CCV with Different Normalizations

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 sectors in our data ($(P_{gt} - P_{gt-1}) / P_{gt-1}$) by their expenditure shares. CUPI stands for CES unified price index. The grey band around the CUPI shows its 95 percent bootstrapped confidence interval (100 replications). The CUPI is calculated for different normalizations of consumer tastes, in which generalized means of consumers tastes are held constant (see equation (16)), including the harmonic mean ($r = -1$), the geometric mean ($r = 0$), the arithmetic mean ($r = 1$), and the quadratic mean ($r = 2$). We also report a robustness test in which the CUPI is calculated normalizing the initial-expenditure-share weighted mean of the log taste shocks to equal zero. 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 VI: Estimated Elasticities for High- and Low-Income Households

Note: Estimated elasticities of substitution for each sector. Sectors are ranked by the estimated elasticity for our baseline sample (including both high- and low-income households). 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.

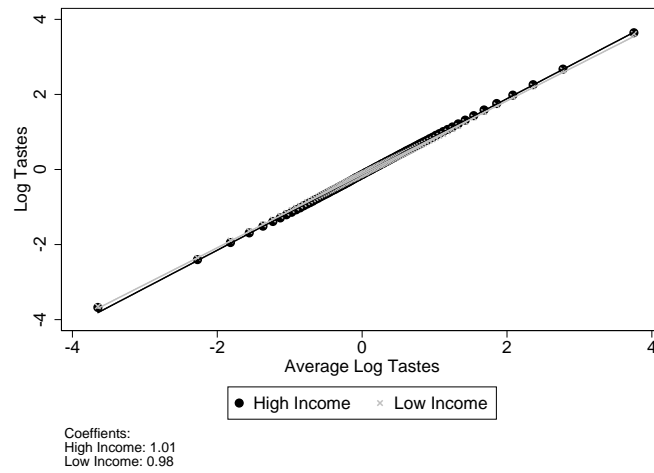


FIGURE VII: Taste Parameters for High- and Low-Income Consumers

Note: Regression lines and bin scatters of the estimated log taste parameter ($\log \varphi_{kt}$) 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). Log taste parameters have a mean of zero for each sector 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 VIII: 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 sectors in our data $((P_{gt} - P_{gt-1}) / P_{gt-1})$ by their expenditure shares. CCV is our CES common variety price index (equation (9)); CUPI is our CES unified price index (equation (8)). The grey band around the CUPI shows its 95 percent bootstrapped confidence interval (100 replications). 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. 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.