

What does consumer heterogeneity mean for measuring changes in the cost of living?

Robert S. Martin*

Division of Price and Index Number Research, Bureau of Labor Statistics
2 Massachusetts Ave, NE, Washington, DC 20212, USA. Email: Martin.Robert@bls.gov

Proceedings of the 2018 Federal Committee on Statistical Methodology (FCSM) Research Conference

Abstract

This paper investigates the assumption that consumer preferences are homothetic. Under this simplification, price indexes like the Chained CPI (C-CPI-U) measure changes in the cost of living for a representative consumer. An active area of research explores departures from homotheticity and their impact on inflation measurement. Recently, Hottman and Monarch (2018) find evidence of nonhomotheticity in demand for imported goods and substantial heterogeneity in import price inflation experienced by different income groups. This paper replicates their model and estimation method using consumer scanner data. While I find this data also rejects homotheticity, it is not clear that it is suitable for consumer price index estimation using Hottman and Monarch's formula because unrestricted parameter estimates fall outside of theoretical bounds.

Keywords: Cost of living index; nonhomothetic preferences; scanner data; heterogeneity

JEL Codes: D12, E31

1 Introduction

The Bureau of Labor Statistics (BLS) aims to measure changes in the cost of living with the Consumer Price Index (CPI). The purpose of this project is to evaluate the assumption that consumer preferences are homothetic, under which price indexes like the Chained CPI (C-CPI-U) measure changes in the cost of living. In consumer theory, homotheticity means that the rate at which a consumer unit is willing to substitute one good for another is independent of its income level (National Research Council, 2002). This implies homogeneity in expenditure patterns across income groups (i.e., consumer units spend the same share of their total expenditures on food and housing regardless of differences in income).¹ On the one hand, homotheticity is a convenient simplification; under this assumption, we can combine prices and aggregate expenditure weights using relatively simple index number formulas that have a powerful interpretation. On the other hand, empirical studies generally reject homotheticity's prediction of constant expenditure shares across income groups (Deaton and Muellbauer, 1980). For instance, the Consumer Expenditure (CE) survey shows higher income deciles are associated with lower shares spent on food and housing (Bureau of Labor Statistics, 2017). Therefore, it is quite sensible to explore the impact of heterogeneous behavior on inflation measurement.

Recently, Hottman and Monarch (2018) (henceforth HM) proposed a model of nonhomothetic consumer preferences and a method for estimating it using detailed transactions data. The model yields cost of living indexes that allow heterogeneity across income groups, a changing basket of goods, and shifts in demand. HM estimate the model using detailed trade data, and find significant heterogeneity in import price inflation across income groups. As the issue of nonhomotheticity is also quite relevant for consumer price index measurement, and so I replicate HM's model and estimation method using household scanner data from Nielsen. This study fits with a broader mission at the BLS to investigate alternative data sources and price index formulas.

*The views expressed herein are those of the author and not necessarily those of the Bureau of Labor Statistics or the U.S. Department of Labor.

¹The CPI-U implicitly assumes expenditure homogeneity across other characteristics as well, such as age, number of consumer units in the group, etc.

There is a substantial literature on the challenges of cost of living measurement. A comprehensive review is beyond the scope of this paper, but I select a few relevant references. First of all, Groshen et al. (2017) describe recent and current efforts at BLS and other agencies to account for issues like quality change and new goods, as well as the budgetary and feasibility constraints that exist within production environments. Moreover, heterogeneity and aggregation across households are among many issues discussed at length in recommendations by the National Research Council’s Panel on Conceptual, Measurement, and Other Statistical Issues in Developing Cost-of-Living Indexes (National Research Council, 2002). It reviews a number of post-WWII studies that find that heterogeneity in consumption patterns does not necessarily translate into significant variation in of inflation rates across income groups. However, modest differences have occurred for periods of time. For example, Cage, Garner, and Ruiz-Castillo (2002) find that differences in housing expenditure shares drove slightly lower price inflation among lower income groups during the 1980s. However, some more recent studies have found more substantial consequences for aggregate measurement. For example, Oulton (2008) finds evidence that in the United Kingdom from 1974-2004, the poor experienced higher inflation than the rich. While their results do not necessarily generalize to other countries, their finding that of a “path-dependence bias” in conventional chained indexes of as large as +0.45% or -0.43% per year (depending on the base year for utility) is striking.²

Hottman and Monarch (2018), whose work motivates this paper, find that higher income deciles in the U.S. experienced lower import price inflation from 1998 to 2014. Unlike many previous studies, which simply aggregate price relatives with group-specific expenditure weights, the HM model also allows for heterogeneous substitution patterns between very similar product varieties. However, they find that heterogeneity across sectoral spending patterns to be more important for driving differences in price inflation. Finally, the HM proposal builds on previous efforts that use constant elasticity of substitution (CES) preferences to model consumer demand. Their model and identification strategy, for example, closely relate to price index methods that incorporate the impact of new goods and shifts in demand.³

My initial results can be summarized as follows. Direct comparisons between my parameter estimates using Homescan and HM’s using trade data are difficult due to substantial differences in data scope and detail. The Homescan data imply higher average price elasticities than HM’s estimates, though they are comparable to Broda and Weinstein (2010), who also use Homescan with a different model. Within the framework proposed by HM, I find the Homescan data also reject the standard homothetic CES model. However, unrestricted parameter estimates do not satisfy theoretical bounds, so I am cautious about the appropriateness of this model for index calculation using this particular dataset.

This paper proceeds as follows. Section 2 reviews HM’s model and estimation procedure, while Section 3 describes the Homescan data and replication exercise. Section 4 presents and discusses the replication results. Finally, Section 5 draws some preliminary conclusions and proposes some next steps.

2 Review of model and estimation procedure

In this section, I sketch the main elements of HM’s partial equilibrium model and reproduce their proposed cost of living index formulas. Consumer preferences, are based on the S-Branch utility tree of Brown and Heien (1972), while firms are assumed to be monopolistically competitive. Using the resulting demand and supply equations, HM use an extension of the Feenstra method to estimate the demand parameters that feed the cost of living index formula.

2.1 Theoretical model

2.1.1 Consumer preferences

In each period t , households (indexed by h) choose quantities of product varieties (indexed by v) which belong to sectors (indexed by s). At the highest level of aggregation, consumer preferences are standard CES. Following HM’s

²As an additional example, Handbury (2013) also uses Homescan data, though a different methodology, to estimate city-specific price indexes that account for nonhomotheticity.

³See Feenstra (1994), Broda and Weinstein (2006), Broda and Weinstein (2010), Hottman, Redding, and Weinstein (2016), and Redding and Weinstein (2016). This identification strategy is not without controversy—see Hausman (1996), for example. Estimators may also suffer from bias when time series are short in length (Soderbery, 2010, 2015).

notation, aggregate utility is given by:

$$V_{ht} = \left[\sum_{s \in S} \left(\varphi_{hst}^{\frac{\sigma-1}{\sigma}} Q_{hst}^{\frac{\sigma-1}{\sigma}} \right) \right]^{\frac{\sigma}{\sigma-1}}, \quad (1)$$

where Q_{hst} is the consumption index for sector s , φ_{hst} is a demand shifter, σ is a substitution elasticity parameter, and S is the set of consumer goods sectors.⁴

The departure from CES comes from the within-sector model, which compared to CES, includes an additional parameter α_v for each variety. The sectoral consumption index is given by:

$$Q_{hst} = \left[\sum_{v \in G_s} \left(\varphi_{vst}^{\frac{\sigma^s-1}{\sigma^s}} (q_{hvt} - \alpha_v)^{\frac{\sigma^s-1}{\sigma^s}} \right) \right]^{\frac{\sigma^s}{\sigma^s-1}}, \quad (2)$$

where q_{hvt} is household h 's quantity consumed of variety v in period t , σ^s is an elasticity parameter for sector s , G^s is the set of varieties in sector s , and φ_{vst} is a demand shifter.⁵ The α_v can be loosely interpreted as subsistence quantities that must be consumed, though they are allowed to be negative. To satisfy regularity conditions for the utility maximization problem, σ , φ_{hst} , σ^s , and φ_{vst} must all be strictly positive, and we must have $k_{hv} < \alpha_v < q_{hvt}$, where $k_{hv} < 0$ is a lower bound on α_v .⁶

The utility maximization problem is solved in two stages. First, for each sector s , q_{hvt} is chosen to maximize the right hand side of Eq. 2 subject to a given sectoral budget allocation Y_{hst} . The price of variety v is denoted by p_{vt} , and the resulting expression also depends on a sectoral price aggregate given by

$$P_{st} = \sum_{v \in G_s} \left(p_{vt}^{1-\sigma^s} \varphi_{vst}^{\sigma^s-1} \right)^{\frac{1}{1-\sigma^s}}. \quad (3)$$

The maximized expression for Q_{hst} is then substituted into Eq. 1 and the result is maximized by choosing the sectoral allocations $\{Y_{hst} : s \in S\}$ subject to total household expenditure Y_{ht} .⁷ Nonhomotheticity is evident when examining the utility-maximizing variety-level expenditure shares (s_{hvt}) and sector-level expenditure shares (S_{hst}), given by Eq.'s 4 and 5.

$$s_{hvt} \equiv \frac{p_{vt} q_{hvt}}{Y_{hst}} = \alpha_v \frac{p_{vt}}{Y_{hst}} + \left(\frac{p_{vt}^{1-\sigma^s} \varphi_{vst}^{\sigma^s-1}}{\sum_{j \in G_s} \left(\frac{p_{jt}}{\varphi_{jst}} \right)^{1-\sigma^s}} \right) \left(\frac{Y_{hst} - \sum_{j \in G_s} \alpha_j p_{jt}}{Y_{hst}} \right) \quad (4)$$

$$S_{hst} = \frac{\sum_{v \in G_s} \alpha_v p_{vt}}{Y_{ht}} + \frac{\varphi_{hst}^{\sigma-1} P_{st}^{1-\sigma}}{\sum_{r \in S} \varphi_{hrt}^{\sigma-1} P_{rt}^{1-\sigma}} \left(\frac{Y_{ht} - \sum_{r \in S} \sum_{j \in G_r} \alpha_j p_{jt}}{Y_{ht}} \right) \quad (5)$$

Examination of Eq. 4 reveals that what HM refer to as ‘‘within-sector’’ nonhomotheticity is governed by α_v . If all α_v terms are zero, then s_{hvt} no longer depends on Y_{hst} and is therefore constant across households. ‘‘Cross-sector’’ nonhomotheticity shows up from both S_{hst} 's dependence on both α_v and the household-specific sectoral demand shifter φ_{hst} . If either all the α_v terms are zero, or φ_{hst} is constant over h , then S_{hst} will also be constant across households.

The cost of living inflation for household h from time t to $t+i$ is defined as the proportional change in its expenditure function, evaluated at a reference utility level V_{hk} . Using the indirect utility function to write V_{hk} in terms

⁴In HM, given the focus on imports, S is defined as the set of tradable consumer goods sectors.

⁵HM note that in principle, α_v and the within-sector demand shifter φ_{vst} are household-level parameters, but they do not allow them to vary by household due to lack of variety specific demand data at the household. However, the aggregate sectoral shifter φ_{hst} is estimated for different income deciles using CE data. I maintain HM's approach to α_v and φ_{vst} , despite having household data, because we observe households purchasing only a small subset of available varieties in any given period.

⁶The lower bound for α_v is defined as $k_{hv} = - \left(\frac{p_{vt}^{-\sigma^s} \varphi_{vst}^{\sigma^s-1}}{\sum_{j \in G_s} \left(\frac{p_{jt}}{\varphi_{jst}} \right)^{1-\sigma^s}} \right) (Y_{hst} - \sum_{j \in G_s} \alpha_j p_{jt})$.

⁷In this framework, total household expenditures equals household income exactly, since neither savings nor transfers are included in the model.

of a reference income level Y_{hk} . The expression is given by Eq. 6, where the impact of within-sector and cross-sector nonhomotheticity is clear.

$$\frac{P_{ht+i}}{P_{ht}} = \frac{\left(\sum_{s \in S} \varphi_{hst+i}^{\sigma-1} P_{st+i}^{1-\sigma}\right)^{\frac{1}{\sigma-1}}}{\left(\sum_{s \in S} \varphi_{hst}^{\sigma-1} P_{st}^{1-\sigma}\right)^{\frac{1}{\sigma-1}}} \left(\frac{Y_{hk} - \sum_{s \in S} \sum_{v \in G_s} \alpha_v p_{vt}}{Y_{hk}} \right) + \left(\frac{\sum_{s \in S} \sum_{v \in G_s} \alpha_v p_{vt+i}}{Y_{hk}} \right) \quad (6)$$

Statistical agencies typically calculate price indexes using weights from aggregate expenditure data, which is related to the notion of a “representative consumer.” In theory, if a statistical agency had enough data and reliable estimates of the model parameters, it might wish to calculate a version of Eq. 6 for each household, from which an average index could be computed that is either democratic (equal weight to each household) or plutocratic (weighted by total household expenditure)(National Research Council, 2002). In practice, however, this would be difficult for many reasons. For example, an individual household purchases a relatively sparse subset of varieties out of the total set available, so even with consumer scanner datasets, calculating a household-level price index is no trivial task.⁸ Moreover, estimation of common model parameters is usually done with aggregate demand data (e.g., see Broda and Weinstein (2010)). The theoretical property that allows the representative consumer formulation is known as “exact linear aggregation.” In the HM model, this condition requires that $\alpha_v = 0$ for non-continuing varieties.⁹ Under the additional restrictions, market demand for variety v is given by

$$q_{vt} = (\alpha_v n_t) + \left(\frac{p_{vt}^{-\sigma_s} \varphi_{vt}^{\sigma_s-1}}{\sum_{j \in G_s} \left(\frac{p_{jt}}{\varphi_{jt}} \right)^{1-\sigma_s}} \right) \left(Y_t - \sum_{j \in G_s} (\alpha_j n_t) p_{jt} \right) \quad (7)$$

where n_t is the number of households in the U.S. at time t . This equation is used in estimating the underlying parameters.

The change in the aggregate cost of living index from time t to $t+i$ is given by:

$$\frac{P_{t+i}}{P_t} = \frac{\left(\sum_{s \in S} \varphi_{st+i}^{\sigma-1} P_{st+i}^{1-\sigma}\right)^{\frac{1}{\sigma-1}}}{\left(\sum_{s \in S} \varphi_{st}^{\sigma-1} P_{st}^{1-\sigma}\right)^{\frac{1}{\sigma-1}}} \left(\frac{Y_k - \sum_{s \in S} \sum_{v \in G_s} (\alpha_v n_t) p_{vt}}{Y_k} \right) + \frac{\sum_{s \in S} \sum_{v \in G_s} (\alpha_v n_{t+i}) p_{vt+i}}{Y_k} \quad (8)$$

where Y_k is the total expenditure level for some reference period.

2.1.2 Firm behavior

To complete the model, the market structure from the firm’s perspective is assumed as monopolistic competition.¹⁰ For their main results, HM treat firms and varieties interchangeably, acting as if all varieties belong to separate firms. Marginal costs are increasing in quantity, given by

$$c_{vt} = \delta_{vt} (1 + \omega_s) q_{vt}^{\omega_s}, \quad (9)$$

where $\omega_s \geq 0$ governs the convexity of the cost function for sector s and $\delta_{vt} > 0$ is a variety-level cost shifter.

The first order condition (FOC) of the firm’s profit maximization problem is

$$p_{vt} = \frac{\varepsilon_{vt}}{\varepsilon_{vt} - 1} c_{vt}, \quad (10)$$

where ε_{vt} is the monopolistically competitive own-price elasticity of demand derived from the market demand equation 7.

$$\varepsilon_{vt} \equiv - \frac{\partial q_{vt} p_{vt}}{\partial p_{vt} q_{vt}} = \left(\frac{q_{vt} - \alpha_v n_t}{q_{vt}} \right) \sigma^s \quad (11)$$

⁸Calculating aggregate chained price indexes using scanner data also presents many challenges. See, for example Ivancic, Diewert, and Fox (2011).

⁹Also required is that $Y_{hst} > \sum_{v \in S} \alpha_v p_{vt}$ and $\alpha_v > k_{hv}$ (the lower limit from earlier). The requirement is that α_v “is such that each household income group buys some amount of each variety in each time period t that has positive sales in the aggregate at that time t .”

¹⁰HM also consider an oligopoly model as a robustness check.

The FOC implies that $\varepsilon_{vt} > 1$. The convexity of demand perceived by firm is

$$\zeta_{vt} \equiv -\frac{q_{vt} \frac{\partial^2 p_{vt}(q_{vt})}{\partial^2 q_{vt}}}{\frac{\partial p_{vt}(q_{vt})}{\partial q_{vt}}} = \frac{-\frac{p_{vt} \frac{\partial^2 q_{vt}(p_{vt})}{\partial^2 p_{vt}}}{\frac{\partial q_{vt}(p_{vt})}{\partial p_{vt}}}}{\varepsilon_{vt}} = \left(\frac{\sigma^s + 1}{\sigma^s} \right) \left(\frac{q_{vt}}{q_{vt} - \alpha_v n_t} \right). \quad (12)$$

The firm's second order condition requires that $\zeta_{vt} < 2$. Combining Eq.'s 10 and 11 gives the firm's pricing equation

$$p_{vt} = \frac{\varepsilon_{vt}}{\varepsilon_{vt} - 1} \delta_{vt} (1 + \omega_s) q_{vt}^{\omega_s} \quad (13)$$

in which the markup over marginal cost can be rewritten as

$$\frac{\varepsilon_{vt}}{\varepsilon_{vt} - 1} = \frac{(q_{vt} - \alpha_v n_t) \sigma^s}{(q_{vt} - \alpha_v n_t) \sigma^s - q_{vt}}. \quad (14)$$

Specification of the firm's pricing equation is important for identifying the demand side parameters, since observed prices and quantities are set simultaneously in a market equilibrium.

2.2 Estimation

HM estimate the model in two stages. First, they estimate the parameters of the variety demand and price functions (i.e. σ^s, ω_s , and α_v for $v \in G_s$, and $s \in S$) at the aggregate market level. Second, they estimate the aggregate substitution elasticity parameter, σ using sectoral expenditures by household income group.

2.2.1 Stage One

The first stage of estimation extends the method of Feenstra (1994), which uses heteroskedasticity across varieties to identify the sector-link demand parameters.¹¹ While the generalized-CES model allows different subsistence quantities α_v for every variety v , as part of their empirical specification, HM restrict these in the following way:

$$\alpha_v = \frac{1}{n_1} \left[\beta_s^j \min_t (q_{vt} | q_{vt} > 0) \right], \text{ for } v \in G_s^j, j \in \{C, E\} \quad (15)$$

where n_1 is the number of households in period 1, G_s^C is the set of "continuing" varieties in sector s that are present in the data (i.e. have positive quantity sold in the aggregate) for the entire sample period, and G_s^E is the set of "non-continuing" varieties that are missing from the data in one or more periods. This specification is notable because it pins down the subsistence quantities in two aspects. First, it requires them to be proportional to minimum quantity demanded observed in the data for variety v , $\min_t (q_{vt} | q_{vt} > 0)$. Second, the coefficient to estimate, β_s^j , is only allowed two possible values. The first aspect seems quite natural. As HM explain, in the case where $\beta_s^C = 1$, for example, the total subsistence quantity for a variety $v \in G_s^C$ sold in period 1 is $\alpha_v n_1$, which is just the minimum quantity of that variety observed. The second aspect, limiting to only two possible parameter values is by necessity, otherwise the model is not identified.

Estimation of $\sigma^s, \omega_s, \beta_s^C$, and β_s^E for each $s \in S$ is complicated by the fact that market price and quantity are determined in equilibrium, which is described by Eq. 7 and Eq. 13. The essence of the Feenstra method for identification is to use the variation in a variety's prices and quantity both over time and with respect to a reference variety in the same sector. Starting from Eq. 7, multiplying both sides by p_{vt} , taking logs, differencing over time and with respect to another variety k , we have:

$$\Delta^{k,t} \ln (p_{vt} q_{vt} - \alpha_v n_t p_{vt}) = (1 - \sigma^s) \Delta^{k,t} \ln p_{vt} + \nu_{vt}, \quad (16)$$

where $\nu_{vt} = (1 - \sigma^s) [\Delta^t \ln \varphi_{kt} - \Delta^t \ln \varphi_{vt}]$. The notation $\Delta^{k,t}$ represents the double difference, and ν_{vt} is the unobserved error. Similarly for the supply side, we start with the supply-side Eq. 13, multiply both sides by $p_{vt}^{\omega_s}$, take logs, and double difference to get:

¹¹Feenstra assumes standard CES preferences and constant demand over time. Broda and Weinstein (2006) and Broda and Weinstein (2010) also use a variation of Feenstra's method.

$$\Delta^{k,t} \ln p_{vt} = \frac{\omega_s}{1 + \omega_s} \Delta^{k,t} \ln (p_{vt} q_{vt}) + \frac{1}{1 + \omega_s} \Delta^{k,t} \ln \left(\frac{\varepsilon_{vt}}{\varepsilon_{vt} - 1} \right) + \kappa_{vt}, \quad (17)$$

where $\kappa_{vt} = \frac{1}{1 + \omega_s} [\Delta^t \ln \delta_{vt} - \Delta^t \ln \delta_{kt}]$ is the unobserved error.

The identifying assumption is that for each variety, the double-differenced demand and supply shocks are orthogonal to each other, implying:

$$G(\beta_s) = \mathbb{E}_t [u_{vt}(\beta_s)] = 0, \quad (18)$$

where $\beta_s = (\sigma^s, \beta_s^C, \beta_s^E, \omega_s)'$, and $u_{vt} = \nu_{vt} \kappa_{vt}$. The t subscript on the expectations operator is to emphasize that the population moment is for the time series generated by variety v .¹² HM argue that double-differencing removes “variation in prices due to markup variation and movements along upward sloping supply curves”, and that “remaining supply shocks take the form of idiosyncratic shifts in the intercept of the variety-level supply curve and are unlikely to be correlated with idiosyncratic shifts in the intercept of the variety-level demand curve.” As a robustness check, they find similar results using only multi-product firms and take differences relative to reference varieties within the same firm, which amounts to removing firm-time fixed effects.

The estimator that makes use of 18 is Generalized Method of Moments (GMM).¹³ For each sector, supposing the estimation sample consists of N_s varieties, then the N_s sample analogs of Eq. 18 are stacked to form the following objective function:

$$\hat{\beta}_s = \operatorname{argmin}_{\beta_s} \{G^*(\beta_s)' W G^*(\beta_s)\}, \quad (19)$$

where $G^*(\beta_s)$ is the stacked sample analogs of Eq. 18 (written in terms of observables), and W is a weighting matrix.¹⁴ HM use a one-step estimator where W , following Broda and Weinstein (2006), is formed to give more weight to varieties if they are either present in the data for more time periods or have larger market shares.

Substituting for ν_{vt} and κ_{vt} and rearranging the expectation in 18 yields:

$$\begin{aligned} & \mathbb{E}_t \left[\left(\Delta^{k,t} \ln p_{vt} \right)^2 \right] \left(\frac{\omega_s}{1 + \omega_s} \mathbb{E} [\Delta^{k,t} \ln (p_{vt} q_{vt}) \Delta^{k,t} \ln p_{vt}] \left(\right. \right. \\ & \quad \left. \left. - \frac{1}{\sigma^s - 1} \mathbb{E} [\Delta^{k,t} \ln p_{vt} \Delta^{k,t} \ln (p_{vt} q_{vt} - \alpha_v n_t p_{vt})] \right) \left(\right. \right. \\ & + \frac{\omega_s}{(1 + \omega_s)(\sigma^s - 1)} \mathbb{E} [\Delta^{k,t} \ln (p_{vt} q_{vt}) \Delta^{k,t} \ln (p_{vt} q_{vt} - \alpha_v n_t p_{vt})] \left(\right. \\ & \quad \left. \left. + \frac{1}{1 + \omega_s} \mathbb{E} \left[\Delta^{k,t} \ln p_{vt} \Delta^{k,t} \ln \left(\frac{\varepsilon_{vt}}{\varepsilon_{vt} - 1} \right) \right] \right) \left(\right. \right. \\ & \left. \left. + \frac{1}{(1 + \omega_s)(\sigma^s - 1)} \mathbb{E} \left[\Delta^{k,t} \ln (p_{vt} q_{vt} - \alpha_v n_t p_{vt}) \Delta^{k,t} \ln \left(\frac{\varepsilon_{vt}}{\varepsilon_{vt} - 1} \right) \right] \right) \right) \end{aligned} \quad (20)$$

The sample analog $G^*(\beta_s)$ is formed from the time averages of the quantities in Eq. 2.2.1, and similar in spirit to Feenstra (1994), the GMM estimator is implemented using weighted nonlinear least squares. Identification comes from sector s having more varieties than parameters, and from variation in the second moment terms that comprise Eq. 2.2.1 across v .

2.2.2 Stage Two

HM’s second stage estimates the aggregate elasticity parameter σ using sectoral price aggregates P_{st} (defined in Eq. 3) and variation in sectoral expenditures across household income deciles (Y_{hst}).¹⁵ The double-differenced sectoral

¹²In Feenstra (1994), the asymptotic arguments hold N_s fixed while $T \rightarrow \infty$.

¹³See Hansen (1982).

¹⁴Note that p_{vt} , q_{vt} , and n_t (and derived quantities) are the observables in the estimating equation. The markup term $\frac{\varepsilon_{vt}}{\varepsilon_{vt} - 1}$ depends on these and the unknown parameters, as shown in Eq. 14.

¹⁵Before calculation of P_{st} , the demand shifters φ_{vt} are recovered under the normalization that $\widetilde{\varphi}_{vt} = \widetilde{\varphi}_v = 1$ where the tilde denotes the geometric mean.

demand equation is

$$\Delta^{k,t} \ln(Y_{hst} - \sum_{v \in G_s} \alpha_v p_{vt}) = (1 - \sigma) \Delta^{k,t} \ln(P_{st}) + \nu_{hst}, \quad (21)$$

where $\nu_{hst} = (\sigma - 1) \Delta^{k,t} \ln \varphi_{hst}$ is the double differenced error. Pooled across household income deciles, sectors, and time, this is estimated via instrumental variables.¹⁶ The instrument for P_{st} is a measure of the “change in dispersion of quality-adjusted variety-level prices within a sector” given by:

$$\Delta^{k,t} \frac{1}{\sigma^s - 1} \ln \left(\frac{1}{N_{st}^v} \sum_{v \in G_{st}} \left(\frac{p_{vt}}{P_{st}} \right)^{1-\sigma^s} \right) \quad (22)$$

Validity of this instrument is argued in Hottman, Redding, and Weinstein (2016).

3 Data and replication

Hottman and Monarch utilize administrative data on import prices and quantities from the Linked-Longitudinal Firm Trade Transaction Database (LFTTD). The issues they raise are also relevant for consumer prices. As previously stated, their methods require linked price and quantity data over time for detailed items. One obstacle to using data collected by the BLS with this method is that price quotes are collected in a separate survey and at a higher frequency than the expenditure data used to calculate weights.

3.1 Homescan

I utilize the Homescan database published by Nielsen, which has been used in previous studies of consumer demand and price index measurement.¹⁷ In 2012, BLS purchased data covering purchases made from 2008 to 2010. The panel consists of roughly 60,000 U.S. households recording roughly 58 million transactions per year. Panelists are instructed to scan the bar codes of every good they purchase. Collected data include quantity, dollar value of the transaction, brand, as well as Nielsen’s codes for department, product group, product module. The dataset also includes information on the retail outlets and demographic information on the sample households. Nielsen includes demographic weights to aggregate the household data to nationally representative sums and averages. To a degree, panelists select into the sample, however, so we must assume their unobservable characteristics are representative of the broader population.

One great advantage of a dataset like Homescan is its disaggregation.¹⁸ Transaction quantities and values are reported at the level of the universal product code (UPC). Industry best practices include using different UPCs for different product varieties. Because of this, Broda and Weinstein (2010) argue it is reasonable to assume that products with the same UPC are homogenous, and that products with different UPCs differ along some dimension that affects consumer willingness to pay. Therefore, one key strength of scanner data is that for the included sectors of the economy, a much greater number of individual products is observed than can be sampled for the CPI. This granular detail is attractive for implementation of the Hottman and Monarch model of consumer preferences of product varieties. On the other hand, Homescan covers only a subset of consumer expenditure categories, mainly consisting of food, personal care, and other general merchandise categories found in grocery and retail chain stores. As noted by Broda and Weinstein (2010), Homescan coverage overlaps significantly with the entry level item (ELI) categories covered by the CPI. However, over 2008 to 2010, the weighted sum of annual expenditures across panelists in Homescan equal about only 7% of total yearly consumer expenditures as measured by the CE.¹⁹ Therefore, caution should be taken when generalizing inferences drawn from Homescan or other similar scanner datasets about consumer behavior.

HM define a sector as the four-digit Harmonized System (HS) code associated with a variety. They define variety as a unique firm-10 digit HS code pair. I treat the product groups defined by Nielsen as the sectors in HM’s model, while I treat UPCs as varieties. Estimating parameters at the product group-level is consistent with Broda and Weinstein

¹⁶Recall the Y_{hst} are choice variables in the consumer’s utility maximization problem, so they are determined in equilibrium along with prices.

¹⁷See, for example Broda and Weinstein (2010) and Hottman, Redding, and Weinstein (2016).

¹⁸See Broda and Weinstein (2010) for a more detailed description of the data.

¹⁹Author’s calculation using Bureau of Labor Statistics (2013).

(2010) and Hottman, Redding, and Weinstein (2016), though these papers also include a nesting below the sector level.²⁰ A key difference is that the data is less aggregated in Homescan—in HM’s data some closely related varieties may be aggregated into the same firm-HS10 pairs. As a consequence, I expect to measure somewhat more elastic substitution patterns in the Homescan data.

3.2 Data cleaning and preparation

While HM use annual trade data, I follow Broda and Weinstein (2010) by aggregating the sales and quantity data to the quarterly level, and then from these calculate unit-value prices for each UPC. As in HM, I winsorize the data by removing observations with double differenced prices or sales values that are either in the top or bottom one percent. Since the Feenstra method relies on sample moments over time, HM also restrict their estimation sample to varieties present in the data for at least six time periods, as a large proportion of varieties are present for only one or two time periods. With three years of Homescan, I cannot make as strict a requirement, but I make a similar one in spirit. Moreover, it is quite common for UPCs in Homescan to enter and drop out of the dataset multiple times. Therefore, I require UPCs to be present for at least five quarter-to-quarter price changes, meaning a variety must be present for at least six pairs of consecutive quarters.

A related challenge concerns sorting UPCs into “continuing” and “non-continuing” varieties for purposes of estimating β_v^C and β_v^E . HM consider a variety to be “continuing” if it is present in the data in every period, grouping all other varieties as “non-continuers.” As their dataset covers the vast majority of U.S. imports, this is likely to be a reliable way of sorting. However, the number of UPCs in Homescan far outweighs the number of transactions per household in the sample each year, meaning the set of UPC’s observed in the data is subject to sampling error. As a consequence, it is difficult to assess whether or not the disappearance and reappearance of UPC in consecutive periods truly represents the exit and re-entry of a good.²¹

HM estimate the aggregate elasticity parameter σ using household income decile level sectoral expenditures Y_{hst} based on CE and Census data. As a preliminary exercise, I avoid issues of mapping UPCs and Nielsen product groups to CE or Census categories by exploiting the demographic variables included in the Homescan dataset. I follow the method described by Handbury (2013) to adjust the categorical income variables in Homescan for characteristics like family size. I then average the sectoral expenditures within the deciles of this adjusted income measure. A future step will be to better estimate group specific expenditures using non-categorical income data from the CE.

3.3 Constraints

Finally, deriving the partial equilibrium for each sector requires regularity of the household’s utility maximization problem and the firm’s profit maximization problem, and this implies several constraints on the parameters and on dependent quantities like the own-price elasticities of demand. While acknowledging that these conditions will not necessarily hold in general, HM ensure they hold in their data either by constraining the parameters directly or by adding penalty terms to Eq. 19. The conditions are: $\beta_s^C < 1$, $\beta_s^E \leq 0$, $\varepsilon_{vt} \geq 1.01$, $\omega_s \geq 0$, $\zeta_{vt} \leq 1.99$, and $k_{hv} < \alpha_v$.²² I estimate Eq. 19 unconstrained, with constraints on β_s^C and β_s^E only, and with constraints and penalties for β_s^C , β_s^E , ω_s , ε_{vt} , and ζ_{vt} .²³

4 Results

This section summarizes my results. Reporting the estimates for each sector is infeasible. To match HM’s reporting, Tables 1 and 2 summarize the estimates with the 10th percentile, median, and 90th percentile across sectors. The IV estimates of the top-level σ are presented along with their 95% confidence intervals and associated OLS estimates

²⁰While Homescan only has 122 product groups versus around 980 sectors in HM’s analysis, their scope is presumably broader considering LFTTD has the universe of imported goods whereas Homescan covers only a small portion of consumer expenditures.

²¹Like HM report for the trade data, continuers in the Homescan data tend to have higher sales than noncontinuers. I also test an alternative, weaker rule whereby a variety is considered “continuing” if it is present in the first and last quarter of data, but found this had little impact on results.

²²These are the conditions stated in HM’s paper. To ensure $\alpha_v < q_{vt}$, it must be that $\beta_s^C < n_1/n_t$, which likely was always the case with HM’s estimates.

²³I did not impose $k_{hv} < \alpha_v$ for computational reasons, but found that it held in 99.64% of observations when β_s^C , β_s^E , ω_s , ε_{vt} , and ζ_{vt} were constrained.

for reference.²⁴ Column 1 contains the results reported by HM using supplier trade data, while columns 2 through 5 present my replications using Homescan, where I have varied the estimator and constraints. As outlined in Section 3, HM's study and this replication are based on different data sources, so I do not expect our estimates to line up precisely. Nevertheless, I hope to make a similar evaluation of within-sector nonhomotheticity. Later in this section, I speculate on other potential reasons for differences between my results and HM's.

4.1 Estimates

Table 1: Summaries of parameter estimates

	HM	Replications with Homescan			
	(1)	(2)	(3)	(4)	(5)
σ^s					
10%	3.06	8.21	10.07	7.85	6.88
Median	4.93	13.79	16.43	12.25	11.73
90%	8.59	21.07	25.93	18.06	17.36
ω_s					
10%	0.16	0.13	0.11	0.12	0.11
Median	0.44	0.22	0.19	0.21	0.17
90%	1.59	0.37	0.29	0.32	0.27
β_s^C					
10%	9.96E-05	0.43	0.46	0.40	0.38
Median	0.33	0.63	0.62	0.58	0.45
90%	0.39	0.82	0.81	0.77	0.46
β_s^E					
10%	-5.97E-05	0.44	0.49	-6.5E-13	-6.0E-13
Median	-2.55E-09	0.66	0.69	0.00	-2.1E-38
90%	-1.08E-10	0.82	0.83	0.00	0.00
σ					
OLS est	0.82	3.54	3.89	3.86	3.91
IV est	2.78	2.59	4.67	4.69	4.93
IV 95% CI	[2.60, 2.97]	[1.89, 3.29]	[3.85, 5.50]	[4.08, 5.31]	[4.31, 5.55]
# sectors	980	113	113	113	113
Stage 1 GMM ^a	One step	One step	Two step	One step	One step
Constraints ^b	Yes	No	No	β_s^j only	Yes ^c

^a GMM estimator type used for first stage parameters. "Two step" uses optimal weighting matrix in second step.

^b Constraints: $\varepsilon_{vt} > 1.01$, $\omega_s \geq 0$, $\zeta_{vt} \leq 1.99$, $\beta_s^C < 1$, $\beta_s^E \leq 0$, $k_{hv} < \alpha_v$.

^c The constraint $k_{hv} < \alpha_v$ was not imposed in estimation, but was satisfied in 99.64% of observations.

Table 1 summarizes the parameter estimates from the estimator described in Eq. 19 implemented for each sector, as well as the IV estimation of Eq. 21. Following HM, I only estimate the model for sectors with at least 30 varieties, leaving 113 product groups in Homescan.²⁵ In all cases, the estimated σ^s (the substitution elasticity parameters) tend to be significantly larger than those reported by HM, with medians ranging from 11.73 to 16.43, compared to HM's median of 4.93.²⁶ As shown in column 3, using optimal GMM increases the reported quantiles of the estimated σ^s somewhat compared to one step GMM, with a median of 16.43, as compared to 13.79 for the one step estimator (column 2).²⁷ Imposing constraints in estimation lowers the distribution of σ^s estimates slightly, with a median of 12.25 when only the β_s^j are constrained and 11.73 with the full complement of constraints. Across all replicatons, the distributions of marginal cost elasticities, ω_s , have significantly lower medians than HM, ranging from 0.17 to 0.22, compared to HM's 0.44. The distributions of ω_s also have a much lower 90th percentile in the HS replications.

²⁴The 95% confidence intervals for columns 2 through 5 are based on heteroskedasticity-robust standard errors of White (1980).

²⁵Similar to HM, I started the stage one numerical optimization from eight points, and of those that converged, I picked the one with the lowest value of the objective function.

²⁶Earlier results from this replication showed even higher σ^s estimates due to a data coding error.

²⁷Feenstra (1994) originally proposed using optimal GMM, which in this case means the second step weights are the inverses of squared residuals, which are found using initial estimates.

The parameters β_s^C and β_s^E are particularly interesting because they govern within-sector non-homotheticity in HM's model. Like HM, I find nonzero estimates of β_s^j and β_s^E , which is evidence against standard CES preferences. The signs and magnitudes of the estimates, however, seem to depend on the constraints imposed. In the unconstrained cases (columns 2 and 3), I find the β_s^j tend to be positive, statistically significant, and higher than HM's results. For instance, the median unconstrained β_s^C estimate is 0.63, versus 0.33 in HM. The median unconstrained β_s^E estimate is similar at 0.66, but this is in stark contrast to HM's estimated quantiles, which are negative and very close to zero. I come closer to replicating HM's estimates when I impose the constraints that $\beta_s^C < 1$ and $\beta_s^E \leq 0$ (column 4). Now the median β_s^C is 0.58, closer to HM, but only slightly lower than when unconstrained. The constraint on β_s^E appears to bind, as now all of my estimates are either negative or zero at Stata's level of precision. Imposing the rest of HM's constraints (column 5), lowers the median β_s^C to 0.45 and compresses the 10th and 90th percentiles to 0.38 and 0.46, respectively.

As HM find, the second stage estimates of σ suggest that assuming Cobb-Douglas form at the upper tier of utility ($\sigma = 1$) is inappropriate. The replicated IV point estimates are all statistically significantly greater than one, ranging from 2.59 to 4.93 depending on the particulars of the stage one estimates. This interval includes HM's IV estimate of σ of 2.78, which is reassuring, though the estimates that use constrained Stage One estimates are a bit higher. This may also be due to scope differences in data.

Table 2: Summaries of other estimated quantities

	HM	Replications with Homescan			
	(1)	(2)	(3)	(4)	(5)
P.E.D. $(\varepsilon_{vt})^d$					
10%	Unrep.	5.95	7.06	5.94	5.72
Median	Unrep.	9.78	11.48	9.49	9.41
90%	Unrep.	13.25	16.86	13.06	14.32
Markups $(\frac{\varepsilon_{vt}}{\varepsilon_{vt}-1})^d$					
10%	1.13	1.10	1.08	1.10	1.08
Median	1.25	1.14	1.12	1.13	1.12
90%	1.48	1.22	1.19	1.22	1.22
Pr($\varepsilon_{vt} < 1$)					
10%	0.00	0.00	0.00	0.00	0.00
Median	0.00	0.00	0.00	0.00	0.00
90%	0.00	0.00	0.00	0.00	0.00
Pr($\zeta_{vt} > 2$)					
10%	0.00	0.06	0.08	0.00	0.00
Median	0.00	0.18	0.17	0.07	0.00
90%	0.00	0.22	0.22	0.12	0.00
# sectors	980	113	113	113	113
Stage 1 GMM ^a	One step	One step	Two step	One step	One step
Constraints ^b	Yes	No	No	β_s^j only	Yes ^c

^a GMM estimator type used for first stage parameters. "Two step" uses optimal weighting matrix in second step.

^b Constraints: $\varepsilon_{vt} > 1.01$, $\omega_s \geq 0$, $\zeta_{vt} \leq 1.99$, $\beta_s^C < 1$, $\beta_s^E \leq 0$, $k_{hv} < \alpha_v$.

^c The constraint $k_{hv} < \alpha_v$ was not imposed in estimation, but was satisfied in 99.64% of observations.

^d Underlying estimates are sales-weighted averages for each sector.

Table 2 reports on other relevant quantities implied by the HM model and parameter estimates. Two related measures are the price elasticity of demand perceived by the firm (PED), ε_{vt} (see Eq. 11), and the profit maximizing markup $\frac{\varepsilon_{vt}}{\varepsilon_{vt}-1}$. For each sector, I take the sales-weighted averages across all varieties and time periods, then report the percentiles listed. I also calculated the proportion of observations in each sector where the estimated PED ε_{vt} was less than one, and the estimated convexity of demand ζ_{vt} was greater than two, each of which is a violation of conditions for profit maximization.

HM do not report estimates of ε_{vt} , likely because their low α_v estimates imply that the average ε_{vt} is distributed

very similarly to the σ^s estimates. In part because I find higher β_s^C and β_s^E , the implied average price elasticities are all lower than their corresponding σ^s , with medians ranging from 9.4 to 11.5 depending on the specification. Even in the closest case, the estimates are still fairly large compared to HM’s σ^s estimates. However, the price elasticities are more comparable to Broda and Weinstein (2010) who use UPC level Homescan data from 1994 and 1999-2003. They find a median elasticity of substitution of 11.5 between UPCs within the same brand module (closely related goods from the same firm).²⁸ A priori, we would expect the elasticities in HM’s model to be slightly lower than those from Broda and Weinstein (2010) because substitutions are also being measured between varieties in different brand modules within the same product group. Not surprisingly given the high estimates of ε_{vt} , the average markups are also lower than those reported by HM, with the estimated median ranging from 1.12 to 1.14. These results are closer to the median markup of 1.16 estimated by Hottman, Redding, and Weinstein (2016) using elasticities of substitution between UPCs produced by the same firm.

Finally, Table 2 evaluates how well the perceived price elasticities of demand (ε_{vt}) and convexities of demand (ζ_{vt}) satisfy conditions implied by profit maximization. The constraints and penalty functions are successful in ensuring that these conditions hold (column 5). Without any constraints, I find only a handful of cases in which ε is less than one. The restriction on the perceived convexity of demand ($\zeta_{vt} < 2$) appears slightly more substantive. When left unrestricted, the restriction fails in the median product group 17% or 18% of the time when the estimates are unrestricted. Restricting just the β_s^j reduces the failure rate by a little more than half. Finally, I did not find a case where either σ^s was less than one or ω_s was less than zero.

4.2 Discussion

In this subsection, I comment on this replication’s findings on homotheticity. I also speculate on some issues that may relate to the differences between these results and HM’s.

4.2.1 The role of constraints

Using their estimates of β_s^C and β_s^E , HM argue that CES “is not a good way of summarizing” the behavior of continuing varieties, but that CES is a “reasonable assumption” for noncontinuers. While the Homescan replication supports their conclusion as it relates to the continuers, the unconstrained results for β_s^E suggest that CES may not fit the noncontinuers either. The similarity of the unrestricted β_s^C and β_s^E estimates could also reflect difficulty in sorting varieties into sets of continuers and noncontinuers, as discussed in Section 3. At face value, the unrestricted, positive estimates of β_s^E imply violations of the condition for utility maximization ($\alpha_v \leq 0$ for noncontinuing varieties) and the condition for exact linear aggregation ($\alpha_v = 0$ for noncontinuing varieties). As a result, it would not make sense to use parameter estimates from columns 2 and 3 to calculate price indexes using Eq. 8. The unrestricted estimates of ζ_{vt} also suggest off-equilibrium firm behavior in a small number of cases. Use of the estimates from column 5 is debatable. On the one hand, the conditions implied by aggregation and partial equilibrium are satisfied through the use of constraints. On the other, the Homescan data are clearly fighting the model with respect to β_s^E , which may suggest exact linear aggregation is inappropriate. Therefore, it is not clear from this data that a subsequent estimate of 8 is to be strictly preferred to a formula that implicitly assumes homotheticity.

4.2.2 Potential sources of bias

While the differences between this replication’s parameter estimates and HM’s may plausibly be attributed to data differences, it is worth highlighting a few theoretical and practical challenges. To begin with, the consistency argument from Feenstra (1994) is based on the number of time periods growing to infinity while the number of varieties is held fixed.²⁹ This may be less appropriate for scanner data, where the number of varieties is typically orders of magnitude larger than the number of time periods. While there is no concrete rule for how many time periods are necessary for the asymptotic approximation to be reasonable, Soderbery (2010, 2015) finds that the Feenstra method may be significantly biased when the number of time periods is relatively small.³⁰ Moreover, GMM with many or

²⁸In the standard CES case with monopolistic competition, ε_{vt} reduces to σ^s , which is both the elasticity of substitution and the price elasticity of demand.

²⁹See also Feenstra (1991).

³⁰What “relatively small” means depends on what one is willing to assume about the underlying data; in some simulations, Soderbery finds that $T = 50$ is still insufficient for the bias to be relatively small.

weak moment conditions can be severely biased and asymptotically non-normal (Stock, Wright, and Yogo, 2002). Soderbery (2015) suggests limited information maximum likelihood (LIML) has better finite sample properties for estimating the simpler CES model using the Feenstra method, but solutions are less common for nonlinear GMM. Moreover, one may be concerned that the double-differenced supply and demand shocks—which correspond to shifts in the intercepts of the supply and demand curves—are correlated. This would invalidate the moment conditions in Eq. 18. I calculated Hansen’s J Statistic following estimation with optimal GMM (column 3), and found the null hypothesis of valid overidentifying restrictions was rejected in 103 out of 113 product groups. This suggests failure of Eq. 18 for at least some varieties in each product group, which could be due to simultaneity.³¹ Use of instrumental variables is another possible identification strategy that could overcome the simultaneity problem, though this is not without controversy either.³²

A related concern is potential functional form misspecification. For example, a CES-type aggregator that nests items according to their characteristics assumes constant substitution parameters for different varieties within the same grouping. This is a convenient stylization that makes estimation of the demand system tractable, but misses potential heterogeneity. Before any price index formula like Eq. 8 could be considered by the BLS for use in official statistics, it would be prudent to explore the consequences of assuming this particular functional form. Finally, there is also a practical challenge to using the Feenstra method because it relies on observing a long, varied, time series, yet typical datasets tend to cover relatively short time spans due to definitional or classification differences in the case of trade data (Soderbery, 2015), or UPC churn in the case of scanner data (Melser and Syed, 2016). This is potentially a thorny issue when implementing the HM model in longer, highly disaggregated datasets, where we need to classify varieties as “continuers” and “non-continuers”, but might expect the proportion of “continuers” to fall (perhaps to zero) as the number of time periods grows very large.

5 Conclusion and next steps

This paper investigates HM’s extension to the Feenstra method for estimating consumer demand using heteroskedasticity across product varieties, which is used to estimate the parameters of a preference structure that allows for non-homotheticity. While HM use supplier trade data, I replicate their method using consumer scanner data. While my numerical results are somewhat different than what HM find, this is to be expected due to differences in the data and time periods studied. The immediate next step for this paper is to calculate HM’s proposed cost of living formula using the parameter estimates that conform to theory. Comparing unrestricted and restricted estimates, however, implies that the underlying data may not fit the model well, as binding constraints are required to achieve estimates that can be used in aggregate price indexes, so HM’s proposed price index should be treated with caution when estimated with the Homescan data. Note that the β_s^C and β_s^E govern just one source of nonhomotheticity. Cross-sector nonhomotheticity arises in the model even if all of the subsistence quantities are found to be zero. In fact, HM find this drives the majority of their cross-income differences in import price inflation. This suggests relaxing homotheticity through the use of group-specific sectoral expenditure shares in price indexes. This is an area of active research at BLS.

References

- Broda, Christian and David E Weinstein (2006). “Globalization and the Gains from Variety”. In: *The Quarterly Journal of Economics* 121.2, pp. 541–585.
- (2010). “Product Creation and Destruction: Evidence and Price Implications”. In: *American Economic Review* 100, pp. 691–723.
- Brown, Murray and Dale Heien (1972). “The S-branch utility tree: a generalization of the linear expenditure system”. In: *Econometrica: Journal of the Econometric Society*, pp. 737–747.

³¹As a robustness exercise, HM repeat their estimation on data from multi-product firms only, where the double differences have been taken with respect to a variety within the same firm, removing firm-time fixed effects as well as sector-time fixed effects. I replicated this exercise by taking differences relative to UPCs within the same brand module. I then estimate the model with optimal GMM. I also find the parameter estimates to be largely unchanged, but Hansen’s J test still rejects in 108 out of 113 cases.

³²See Hausman (1996) including Bresnahan’s comment.

- Bureau of Labor Statistics (2013). *Average annual expenditures and characteristics of all consumer units, Consumer Expenditure Survey, 2006-2012*. Tech. rep. URL: <https://www.bls.gov/cex/2012/standard/multiyr.pdf>.
- Bureau of Labor Statistics (2017). *Table 1110: Deciles of income before taxes: Annual expenditure means, shares, standard errors, and coefficients of variation*. Tech. rep. URL: <https://www.bls.gov/cex/2016/combined/decile.pdf>.
- Cage, Robert A, Thesia Garner, and Javier Ruiz-Castillo (2002). “Constructing household specific consumer price indexes: An analysis of different techniques and methods”. In:
- Deaton, Angus and John Muellbauer (1980). *Economics and consumer behavior*. Cambridge university press.
- Feenstra, Robert C (1991). “New goods and index numbers: US import prices”. Working paper.
- (1994). “New product varieties and the measurement of international prices”. In: *American Economic Review*, pp. 157–177.
- Groshen, Erica L et al. (2017). “How Government Statistics Adjust for Potential Biases from Quality Change and New Goods in an Age of Digital Technologies: A View from the Trenches”. In: *Journal of Economic Perspectives* 31.2, pp. 187–210.
- Handbury, Jessie (2013). “Are poor cities cheap for everyone? non-homotheticity and the cost of living across us cities”. Working paper.
- Hansen, Lars Peter (1982). “Large sample properties of generalized method of moments estimators”. In: *Econometrica: Journal of the Econometric Society*, pp. 1029–1054.
- Hausman, Jerry A (1996). “Valuation of new goods under perfect and imperfect competition”. In: *The Economics of New Goods*. University of Chicago Press, pp. 207–248.
- Hottman, Colin J and Ryan Monarch (2018). “Estimating Unequal Gains across U.S. Consumers with Supplier Trade Data”. Unpublished. Version dated Jan. 17, 2018.
- Hottman, Colin J, Stephen J Redding, and David E Weinstein (2016). “Quantifying the sources of firm heterogeneity”. In: *The Quarterly Journal of Economics* 131.3, pp. 1291–1364.
- Ivancic, Lorraine, W Erwin Diewert, and Kevin J Fox (2011). “Scanner data, time aggregation and the construction of price indexes”. In: *Journal of Econometrics* 161.1, pp. 24–35.
- Melser, Daniel and Iqbal A Syed (2016). “Life Cycle Price Trends and Product Replacement: Implications for the Measurement of Inflation”. In: *Review of Income and Wealth* 62.3, pp. 509–533.
- National Research Council (2002). *At What Price?: Conceptualizing and Measuring Cost-of-Living and Price Indexes*. Ed. by Charles Schultze and Christopher Mackie. National Academies Press.
- Oulton, Nicholas (2008). “Chain indices of the cost-of-living and the path-dependence problem: An empirical solution”. In: *Journal of Econometrics* 144.1, pp. 306–324.
- Redding, Stephen J and David E Weinstein (2016). “A unified approach to estimating demand and welfare”. Unpublished.
- Soderbery, Anson (2010). “Investigating the asymptotic properties of import elasticity estimates”. In: *Economics Letters* 109.2, pp. 57–62.

— (2015). “Estimating import supply and demand elasticities: Analysis and implications”. In: *Journal of International Economics* 96.1, pp. 1–17.

Stock, James H, Jonathan H Wright, and Motohiro Yogo (2002). “A survey of weak instruments and weak identification in generalized method of moments”. In: *Journal of Business & Economic Statistics* 20.4, pp. 518–529.

White, Halbert (1980). “A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity”. In: *Econometrica: Journal of the Econometric Society*, pp. 817–838.