

NEW CHARACTERISTICS AND HEDONIC PRICE INDEX NUMBERS

Ian Crawford and J. Peter Neary*

Abstract—Changes in product characteristics on the extensive margin (the addition of new features and the removal of old ones) are an important and hitherto neglected dimension of quality change. Standard techniques for adjusting price indices for new goods cannot handle such changes satisfactorily, and this leads to an economically and statistically significant bias in the measurement of prices and real output. We combine insights from the theories of exact index numbers and demand for characteristics to develop a new method for incorporating changes on the extensive characteristic margin. Applied to U.K. data on new car sales, our method leads to revisions in estimated inflation rates for this commodity group that are both plausible and quantitatively important.

I. Introduction

FAILING to adjust price indices for the arrival of new goods has long been known as an important potential source of mismeasurement in prices and real output.¹ A recent survey by Feldstein (2017) argues that these failures persist and that the resulting overestimate of the rate of inflation implies systematic underestimation of changes in living standards and aggregate output and may be a significant source of the “productivity puzzle,” the slowdown in measured productivity growth in advanced economies over recent decades.

In this paper we describe, and provide a simple solution to, a problem that is directly analogous to the new-goods problem: the new-characteristics problem. We first argue that product innovation often occurs on the extensive margin (products acquire new features) as well as the intensive margin (existing features are upgraded). We then show that while hedonic methods provide a good solution to the new-goods problem when the new goods comprise recombinations of or upgrades to existing characteristics (Pakes, 2003; Ackerman et al., 2007), they do not when a new good acquires an entirely novel characteristic or feature that has never previously been available in that class of products. We then show how to adjust, subject to a functional form assumption, the

hedonic index to account for changes in product specifications on the extensive margin. The method requires only a minor extension to the hedonic methods already in use by national statistical agencies and so it could easily be adopted in practice.²

The new-goods problem arises largely through the use of the matched-model method in index-number construction. National statistical institutes, like the Bureau of Labor Statistics (BLS, 2020) in the United States and the Office for National Statistics (ONS, 2017) in the United Kingdom, gather prices by visiting the same shops repeatedly and forming price-relatives (price-ratios) for the same “matched” good over pairs of visits. These price-relatives are then averaged to form an index. As Pakes (2003) points out, the fundamental problem with the matched-model method is that “it is simply not defined” when a new good appears in the comparison period that did not exist in the base period.

The consequence of this is that “the matched model index is constructed by averaging the price changes of the goods that do not disappear; so it selects disproportionately from the right tail of the distribution of price changes.”³ This results in an upward bias to the resulting price index.

Pakes (2003) proposes a solution based on hedonic methods. The essential idea is to shift the focus away from product space to characteristics space. The hedonic price function (which relates the prices of different products to their characteristics) can be estimated in each period using the products available in that period. It can then be used to construct fixed-weight hedonic price indices, which give bounds to the true Konüs cost-of-living index in a way that is directly analogous to the canonical results for cost-of-goods indices: a base-weighted index giving a Laspeyres-type upper bound and a comparison-period-weighted index giving a Paasche-type lower bound to their corresponding Konüs indices.⁴ The crux of this (in the case of, for example, the hedonic Paasche) is that the later period’s hedonic price regression is used to

Received for publication August 20, 2019. Revision accepted for publication April 13, 2021. Editor: Benjamin R. Handel.

*Crawford: University of Oxford; Neary: University of Oxford, CEPR, and CESifo.

We are very grateful to our editor, Ben Handel, to three anonymous referees, to Erwin Diewert, Andrew Dilnot, Rolf Färe, Steve Redding, and Karl Schlag, and to seminar participants at a Nuffield College/ONS Workshop, the Irish Economic Association Annual Conference, the ESCoE/ONS/BoE Conference on Economic Measurement, the Econometric Society European Meetings in Cologne, the Royal Statistical Society/Royal Economic Society/ONS Workshop at the Warwick Business School in London, and the University of Huddersfield for their very helpful comments.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01079.

¹For example, the report of the Boskin Commission (Boskin et al., 1996) estimated that the U.S. consumer price index overestimated the true rate of inflation by 1.1% per year; of this, failing to account for new goods and quality change in existing goods together accounted for over half: 0.60 percentage points. The other sources of bias were substitution bias (largely eliminated, according to Gordon, 1999, by a subsequent shift from a fixed weight to a variable weight index-number formula), and outlet-substitution bias, arising from insufficiently accounting for shifts to lower price outlets.

²Introduced in Griliches (1961), hedonic quality adjustment of prices is now a firmly established method used by national statistical offices and institutes the world over. In the United Kingdom, hedonic methods were introduced in the CPI in 2003 (for personal computers), and subsequently their use has expanded to include digital cameras, laptops, tablet PCs and handsets for pay-as-you-go cell phones. The United States started using hedonics to adjust rental values as far back as the first half of the last century and now uses it for clothing, footwear, refrigerators, washing machines, clothes dryers, ranges and cooktops, microwave ovens, TVs, and DVD players. Sweden uses the approach for twenty clothing and twelve footwear items and Canada for PCs, laptops, printers, monitors and Internet services. Other countries that use these methods include Australia, New Zealand, Denmark, Finland, Germany, the Netherlands, and Switzerland. See Wells and Restieaux (2013) for a review of practice in the United Kingdom and internationally.

³Pakes (2003), p. 1578.

⁴These provide two-sided bounds when preferences are homothetic and the cost-of-living index is therefore independent of the utility level.

backcast the prices of the goods not available in the base period (see Pakes, 2003, equation 4, p. 1582). The method allows product characteristics to change on the intensive margin through upgrades to, or recombinations of, existing features.⁵

However, changes to product characteristics on the extensive margin, through the introduction of new features or the removal of existing features, pose difficulties for hedonic-based methods. This is because, just as product-based indices are based on “matched models,” the hedonic approach is based on “matched characteristics” and hence suffers a problem that we term “new-characteristics bias.” This is precisely akin to the new-goods problem; the only difference is that the arena in which this occurs is shifted from product space to characteristics space.

A number of authors have commented on this in general terms. Hausman (2003) points out that proper quality adjustment in price indices should allow for changes in the range of product attributes, and Varian (2018) discusses how the new features included in smartphones are a source of quality change not captured by current methods.⁶ To date, the conceptual framework for understanding this problem has not been laid out, and there has been no systematic attempt at quantifying it.

This paper aims to fill this gap. We provide a conceptual framework for understanding new-characteristics bias and propose and implement an explicit solution for it. In more detail, we show that changes to the attributes of existing products on the extensive margin (changes in the range of features incorporated into a product or service) are potentially as important as changes on the intensive margin (upgrades of existing features). We show that the method of hedonic quality-adjustment currently in most widespread use does not properly account for changes on the extensive margin of characteristics and that this is a source of systematic bias in price measurements. We introduce a new method to correct for this that based on the linear characteristics model developed by Gorman in 1956 (Gorman, 1980; Lancaster, 1966), and the mix-adjusted extension of the Sato-Vartia price index (Sato, 1976, and Vartia, 1976) developed by Feenstra (1994). Finally, to illustrate this method, we show the effect of allowing for changes on the extensive margin of characteristics on the quality-adjusted price index for new car sales in the United Kingdom. Our approach is theoretically well founded and yet straightforward to apply in practice: national statistical agencies could easily adopt it, since it can be applied without necessarily requiring any additional estimation steps over and above those normally required for hedonic methods. A price of this simplicity is that we make a functional form

assumption: we assume that preferences are of the generalized (asymmetric) constant elasticity of substitution (CES) form following Feenstra (1994).⁷ An offsetting advantage is that it delivers a single value for the price index as opposed to upper and lower bounds as in other hedonic methods, such as Pakes (2003).

For a recent concrete example of innovation on the extensive margin, consider camera phones.⁸ Now, most mobile telephones feature a digital camera, but it was not always so. There is some debate about whether the first camera phone was the Samsung SCH-V200 (released in South Korea in June 2000) or the J-SH04 by Sharp (released in Japan in November of the same year).⁹ The Samsung was able to take photos at a resolution of 0.35 megapixels (MP), but its storage was limited to just twenty photographs.¹⁰ The Sharp solution was not as good (0.11 MP), but it allowed users to transfer and transmit their photos electronically. The first U.S. camera phone was the Sanyo SCP-5300, released in November 2002. It cost \$400 and had a 0.3 MP capability. It also had a basic flash, white balance control, self-timer, digital zoom, and various filter effects like sepia, black and white, and negative colors. By the end of 2003, camera phones were taking off in the United States, and over 80 million had already been sold worldwide. Over half of the phones sold worldwide in the first nine months of 2004 had cameras in them, and two-thirds of all the phones shipped in the third quarter were camera phones. Leading the way was Finnish manufacturer Nokia (the Nokia N90 had a 2 MP camera, Carl Zeiss optics, autofocus, an LED flash, and a rotating screen mimicking a camcorder).

The next phase of innovation started to move toward the intensive margin in a race for ever greater resolution. The 1.3 MP arrived in the United States with the Audiovox PM8920. It had a dedicated camera button and a number of preset settings for different lighting conditions, a multishot feature, and the ability for users to record their own shutter sound. It was available for \$299 RRP, although as part of various bundles, it could be had for as little as \$190. The Sony Ericsson K800i was released in 2006 with a 3.2 MP camera with autofocus, image stabilization, and a xenon flash. Nokia naturally retaliated with models like the 3.2 MP N73, and then in 2007 with the first 5 MP camera phone in the Nokia N95. This model also represented an innovation on the extensive margin: it could record video at 30 frames per second. Unfortunately (for Nokia) the smartphone era was about to arrive. The original iPhone was launched a few months after the N95, in June 2007. It had a mere 2 MP camera, no flash or autofocus, and could not take video. Camera phones

⁵Although, even in this case, the ability of characteristics-based demand systems to handle new goods can only be regarded as informative to the extent that the new goods “are not too new,” in the words of Akerberg et al. (2007), p. 11; that is, a new product outside the set of all tuples of characteristics that were convex combinations of the characteristics of existing products requires extrapolation outside of the observed range of data.

⁶See also Pakes (2005).

⁷More recently also adopted by Broda and Weinstein (2006) and Redding and Weinstein (2019).

⁸See <https://www.digitaltrends.com/mobile/camera-phone-history/>, from which some of what follows is taken.

⁹The point at issue regarding priority is that the camera and the phone inside the Samsung were in fact separate devices that were simply accommodated in the same housing, whereas the Sharp was fully integrated.

¹⁰They had to be downloaded, using a hard link, to a PC if the user wanted to take more.

continued to improve, but the race stalled somewhat as smartphones took off. The iPhone, with its relatively basic camera, proved that there were more important features than the camera.

Not all innovations turned out to be popular with consumers. Some notable misses include attempts by HTC and LG to interest consumers in 3D. They released phones with dual stereoscopic 5 MP cameras, but there was no real demand, and they were dropped. Manufacturers were more successful with shifting their focus to software features beginning with motion control (since resolutions were so high that users without tripods were unable to hold the camera steady enough to take advantage) and then in-built filters and image-processing effects.

This piece of recent economic history serves to make three main points. First, product innovation occurs on both margins: the race for ever higher resolution exemplifies the intensive margin, while the huge expansion of features and the capabilities of handsets relates to the extensive margin. Second, in the case of handsets, the really significant innovations came via the changes on the extensive margin. Third, some characteristics fail and are withdrawn so changes on the extensive margin are not simply the accretion of ever more features; other features, like 3D, die away. We return to these points in our empirical application below.

It is important to mention that there is a little-used version of the hedonic model that can accommodate changes on the extensive margin. This is the time-dummy method.¹¹ This method works by pooling data across products and periods and regressing (log) prices on a set of product attributes and a sequence of time dummies. The attributes control for quality-related price changes, and the coefficients on the time dummies pick up the pure period-to-period price change. Since the regression is run over data pooled across time periods, any product characteristic that features in at least some good in some period can be included even if it is not present in all periods. This makes it straightforward to accommodate product characteristics that come and go. However, the time-dummy method is not used by statistical offices because it suffers from a simple but major drawback: it requires that the historical consumer price index series be completely revised every time the series is updated.¹² This effectively kills off the time-dummy method as a practical option for statistical agencies, though it provides a useful contrast to our approach. (We discuss it further in appendix B.)

¹¹The time dummy method is based on the repackaging model of Fisher and Shell (1968).

¹²In the United States, both of the two main CPI measures are considered final on publication. The chained index for All Urban Consumers is revised for three quarters (as seasonal weights are revised) but is final after ten to twelve months (see BLS, 2020). In the United Kingdom, the CPI is considered revisable (its predecessor, the RPI, was not), but the Office for National Statistics (ONS) would only do so in “exceptional circumstances” (which we take to mean the discovery of manifest and economically significant mistakes), and in practice it has never been revised (see ONS, 2017). See Ball and Allen (2003) for a discussion of this drawback with the time dummy method.

The more standard hedonic method works by running a sequence of separate, within-period regressions of prices on characteristics. The estimated regression coefficients represent the implicit or shadow price for each of the included characteristics in each period, and these are used to price bundles of characteristics. Applying the shadow prices from period $t + 1$ to the specification of the product in period t generates a period- t product price that values the quality of the product (as measured by its period- t characteristics) at the new period- $(t + 1)$ shadow prices. By contrast with the time-dummy method, this approach simply adds newly estimated shadow price coefficients to the end of the series and does not require any revisions; it is therefore the method that statistical offices use. However, this method of quality adjustment does not account for changes in economic welfare associated with product innovation on the extensive margin of characteristics. To do so would require an estimate of the shadow prices of the characteristic in both the base and comparison periods, but if a characteristic or feature is simply not present in the market in one of those periods, then its shadow price cannot be recovered for that period; there is no variation with respect to that feature to exploit in the price-on-characteristics regression. If such innovations are important, then hedonic methods, as currently applied in practice, will overstate the true rate of price changes and consequently underestimate the rate of growth of real GDP and consumption.

The plan of this paper is as follows. In section II, we describe the linear characteristics model, the economic theory underlying the indirect method of hedonic adjustment, and in section III, we develop a method of accounting for changes in characteristics based on the theory of household behavior under rationing. In section IV, we show how this method can be operationalized using a hedonic version of the price index developed by Sato (1976) and Vartia (1976), and subsequently extended to allow for changes in the range of goods consumed by Feenstra (1994). Section V illustrates how our approach can be applied in practice. Section VA introduces the data we use in our empirical application, section VB, explores the implications of defining a “product” in different ways, and section VC, presents the results. Section VI concludes. The online appendixes present extensions and robustness checks.

II. Characteristics Models and Hedonic Price Indices

The characteristics model was developed in 1956 by Gorman (1980) and subsequently popularized by Lancaster (1966). There are K varieties of products with quantities denoted by x^k and prices by p^k . These market goods are differentiated by J product characteristics with quantities denoted by z^j . It is typically assumed that there are fewer characteristics than market goods, $J < K$, so that the characteristics model normally entails a degree of dimension reduction relative to the preferences-for-goods model.

In the linear characteristics model, the total amount of a given characteristic present in a bundle of varieties $\{x_t^1, \dots, x_t^K\}$ observed in period t is

$$z_t^j = \sum_k a_t^{kj} x_t^k, \quad (1)$$

where a_t^{kj} represents the amount of characteristic j present in one unit of product k according to the product specification in period t . Prices, budgets, and the product specification are time varying (indexed by t), but they are assumed not to be choice variables for the individual consumer.

The model hypothesizes that the price-taking consumer with an exogenous budget has preferences over characteristics rather than products per se, but that she can only purchase characteristics by purchasing the products that embody them. Thus, she selects her preferred mix of products subject to her budget constraint and the requirement that the selected varieties map to the characteristics according to the current specifications given by a_t^{kj} . Preferences over characteristics are represented by a continuously differentiable utility function $v(\cdot)$. Thus the optimizing model is

$$\begin{aligned} \max_{x^1, \dots, x^K} v(z^1, \dots, z^J) \text{ subject to } \sum_k p_t^k x^k &= y_t \\ \text{and } z^j &= \sum_k a_t^{kj} x^k, \end{aligned} \quad (2)$$

where y_t is the consumer's budget.

Given these assumptions about the linear characteristics structure and the consumer's preferences, maximizing behavior implies the first-order conditions

$$\sum_j a_t^{kj} v_t^j = \lambda_t p_t^k \quad \forall k, t, \quad (3)$$

where λ_t is the marginal utility of income, v_t^j denotes the marginal utility of the j th characteristic in period t , and where we assume for simplicity that the consumer is either consuming each good or just indifferent between consuming and not consuming it.¹³ Following Gorman (1980), the shadow prices of characteristics are defined as the ratio of the marginal utility of the characteristics to the marginal utility of income:

$$\pi_t^j \equiv \frac{1}{\lambda_t} v_t^j. \quad (4)$$

We will assume that the marginal utilities of characteristics are weakly positive and therefore, given a positive marginal utility of income, the shadow prices are also weakly positive:

$$\pi_t^j \geq 0 \quad \forall j, t. \quad (5)$$

The first-order conditions then become

$$p_t^k = \sum_j a_t^{kj} \pi_t^j \quad \forall k, t. \quad (6)$$

¹³See Blow et al. (2008) for a discussion of corner solutions.

This is the hedonic pricing equation, which says that the market price of each product in a given period is a linear function of its characteristics, weighted by the consumer's willingness to pay for each characteristic.¹⁴ This implies that the prices are inside the column space of the technology matrix,¹⁵ and it therefore follows that for the linear characteristics model, the budget constraint adds up to total expenditure whether it is expressed in terms of the prices and quantities of market goods or in terms of the shadow prices and quantities of the underlying characteristics.¹⁶

$$\sum_k p_t^k x_t^k = \sum_j \pi_t^j z_t^j = y_t. \quad (7)$$

Consequently there are two ways to think about the consumer's choice problem: one as a demand-for-products problem in product space,

$$\{x^1, \dots, x^K\} = \arg \max_{x^1, \dots, x^K} \left\{ u_t(x^1, \dots, x^K) \mid \sum_k p_t^k x^k = y_t \right\}, \quad (8)$$

and the other as a demand-for-characteristics problem in characteristics space subject to a budget constraint expressed in terms of shadow prices and the same total budget:

$$\{z^1, \dots, z^J\} = \arg \max_{z^1, \dots, z^J} \left\{ v(z^1, \dots, z^J) \mid \sum_j \pi_t^j z^j = y_t \right\}. \quad (9)$$

The two are linked by the linear mapping between goods and characteristics: $u_t(x^1, \dots, x^K) = v(\sum_k a_t^{k1} x^k, \dots, \sum_k a_t^{kJ} x^k)$. These two representations of the consumer's problem are interchangeable, but there is an important difference between them. The consumer's preferences over characteristics represented by $v(\cdot)$ are fixed so do not depend on time. By contrast, the induced preferences over products represented by $u_t(\cdot)$ depend on the time-varying characteristics, coefficients a_t^{kj} . Changes in the quality of goods arising from changes in the characteristics embodied within them show up as changes in preferences for products. This is why the induced preferences $u_t(\cdot)$ have a time subscript. This contrasts with Redding and Weinstein (2019) wherein time-varying preferences for products arise instead from error terms in an empirical demand system; these are then interpreted as taste shocks, which are equivalent to multiplicative

¹⁴As Blow et al. (2008) show, the linear-in-parameters form of the hedonic pricing equation is a general property of the characteristics model that is true of both linear- and nonlinear characteristics models alike. It is not driven by any parametric assumptions on $v(\cdot)$. We thank an anonymous referee for highlighting this.

¹⁵That is $\text{rank}(\mathbf{p}_t : \mathbf{A}_t) = \text{rank}(\mathbf{A}_t)$ where \mathbf{p}_t is the K -vector of prices of goods in period t and \mathbf{A}_t is a $K \times J$ matrix with element a_t^{kj} .

¹⁶See Blow et al. (2008).

price shocks under their CES specification.¹⁷ Redding and Weinstein (2019) propose a particular method for using these shocks to correct a preferences-for-goods cost-of-living index for taste changes. Our approach, by switching focus to time-invariant preferences-for-characteristics and working with the demand-for-characteristics version of the model, is closer to Pakes (2003, 2005).

Dual to the direct utility function, $v(\cdot)$ is the hedonic expenditure function

$$c(\pi_t, v) = \min_{z^1, \dots, z^J} \sum_j \pi_t^j z^j \text{ subject to } v(z^1, \dots, z^J) = v, \quad (10)$$

where π_t denotes the J -vector of shadow prices in period t . This allows us to define our key benchmark for the true cost of living, the constant-utility Konüs hedonic price index:

$$P_K(\pi_0, \pi_1, v) = \frac{c(\pi_1, v)}{c(\pi_0, v)}. \quad (11)$$

Note that this formulation allows for substitution between different combinations of characteristics. Standard index number theory (see Diewert, 1981, for an authoritative account) then provides a number of ways of either approximating or, subject to functional form assumptions regarding $v(\cdot)$, exactly computing the hedonic index. For example, denoting by ω_t^j the hedonic budget share for the j th characteristic $\pi_t^j z_t^j / y_t$, the hedonic Laspeyres,

$$\Pi_L = \sum_j \left(\frac{\pi_1^j}{\pi_0^j} \right) \omega_0^j, \quad (12)$$

and Paasche indices,

$$\Pi_P = \left(\sum_j \left(\frac{\pi_1^j}{\pi_0^j} \right)^{-1} \omega_1^j \right)^{-1}, \quad (13)$$

may be regarded as, respectively, either upper and lower approximations to the hedonic Konüs or as exact hedonic indices if preferences for characteristics are assumed to be either linear or Leontief.¹⁸ The hedonic Sato-Vartia index,¹⁹

$$\Pi_{SV} = \prod_j \left(\frac{\pi_1^j}{\pi_0^j} \right)^{\omega^j} \quad (14)$$

(where the weights ω^j are the normalized logarithmic means of the period 0 and period 1 hedonic budget shares), is ex-

act if preferences are CES, and it approaches the hedonic Paasche and Laspeyres depending on the degree of substitution between characteristics in the limit. The hedonic Fisher index,

$$\Pi_F = \sqrt{\Pi_L \Pi_P}, \quad (15)$$

which is the geometric mean of the hedonic Paasche and Laspeyres, is exact for quadratic preferences over characteristics. Moreover, unlike the Sato-Vartia, it provides a second-order approximation to any twice continuously differentiable homothetic expenditure function and is therefore described as “superlative.” (See Diewert, 1976.)

Finally, to connect theory to practice, the indirect method used by statistical agencies appends an econometric error to the hedonic pricing equation,

$$p_t^k = \sum_j a_t^{kj} \hat{\pi}_t^j + e_t^k, \quad (16)$$

and identifies the shadow prices by regressing prices on product characteristics using within-period cross-variety covariation in product specifications and product prices. The assumption that preferences are defined over fewer characteristics than there are products, $J < K$, is material here as it is required in order to identify the shadow prices. The resulting shadow prices and the appropriate characteristics are then substituted into the chosen index number formula. A sequence of separate hedonic regressions is run on each period of data. As data for new periods become available, changes in product characteristics on the intensive margin can be handled by changing the value of a_t^{kj} to reflect the current product specification, and changes on the extensive margin can be allowed for by changing the set of characteristics in the regression and simply adding or deleting characteristics as regressors. For these reasons, its flexibility and the fact that it does not need to be revised as new observations become available, this method is the one favored by all statistical offices that use hedonic methods.

The main point of this paper is to note that all of these hedonic price indices are functions of the shadow price-relatives $\frac{\pi_1^j}{\pi_0^j}$ and hence require knowledge of the shadow price of each characteristic in both the base and comparison periods, that is, matched pairs of characteristics. This is a direct counterpart of the price-relatives/price-ratios $\frac{p_1^k}{p_0^k}$ required by traditional product-based, matched-model price indices. Just as standard indices are based on matched-models, so the hedonic index is based on matched-characteristics. When characteristics cannot be matched, hedonic methods run into the same problems that matched-model price indices suffer. Product characteristics that existed in period 0 but disappear from the market in period 1 cannot be priced in the later period, and characteristics that are introduced in the later period but did not exist in the base period cannot be priced in the base period. As a result either the denominator (in the case of new

¹⁷The general approach is essentially an application of Barten scaling (Barten, 1964). We discuss this further in section IV.

¹⁸An index number is said to be “exact” if its numerical value always corresponds to that of a Konüs cost-of-living index but its calculation depends only on the observed price and quantity data and does not require knowledge of preference parameters.

¹⁹Sato (1976); Vartia (1976).

characteristics) or the numerator (in the case of disappearing characteristics) of the shadow price-relative is missing.

III. Adjusting for New Characteristics: General

How can we deal with this problem of taking account of changes in the set of characteristics? The standard response is to ignore it by omitting any changes on the extensive margin from the pricing equation and the index. While this stance cannot sensibly be maintained forever without the product specification becoming out-of-date, it can be adapted to form a chained-hedonic index calculated on the basis of the characteristics common to each pair of adjacent periods. Chaining in this manner is routinely used in preferences-for-goods contexts to allow for the entry and exit of products, and it can equally be applied to hedonic price indices to allow for the entry and exit of characteristics.²⁰

However, two potential problems remain—one econometric, the other economic. The econometric problem is omitted variable bias; if the new characteristic is correlated with other existing characteristics (e.g., if a new feature is introduced into varieties with existing high-end specifications), this correlation will bias the estimates of the shadow prices of the other characteristics. The economic problem is that changes on the extensive margin affect economic welfare directly. For example, if a new characteristic becomes available, then it becomes cheaper, other things being equal, for the consumer to reach a given level of utility. This is because something akin to a rationing constraint has been removed. This is analogous to the new-goods problem in preference-for-products index numbers (see Hausman, 1997). Ignoring changes in the extensive margin of characteristics will therefore bias the hedonic price index, even in a chained index. To deal with this bias, we show in this section how changes in the range of characteristics can be incorporated using results from the theory of rationing, following Neary and Roberts (1980).

To see this, consider the constrained hedonic expenditure function in which characteristic 1 is fixed, in period t , at the level \bar{z}_t^1 :

$$c(\pi_t, v, \bar{z}_t^1) = \min_{z^2, \dots, z^J} \pi_t^1 \bar{z}_t^1 + \sum_{j=2}^J \pi_t^j z^j \quad \text{subject to } v(\bar{z}_t^1, \dots, z_t^J) = v. \quad (17)$$

Note that whereas the unconstrained hedonic expenditure function is defined effectively for all v , the constrained hedonic expenditure function is defined only at utility levels that are attainable under the additional constraint on the char-

acteristic. Given this, the constrained expenditure function satisfies Shephard's lemma

$$\begin{aligned} \frac{\partial c(\pi_t, v, \bar{z}_t^1)}{\partial \pi^1} &= \bar{z}_t^1, \\ \frac{\partial c(\pi_t, v, \bar{z}_t^1)}{\partial \pi^j} &= z^j(\pi_t, v, \bar{z}_t^1), \quad \forall j = 2, \dots, J. \end{aligned}$$

The relationship between the constrained and unconstrained hedonic expenditure functions is given by Neary and Roberts (1980):

$$\begin{aligned} c(\pi_t, v, \bar{z}_t^1) &= \pi_t^1 \bar{z}_t^1 + \sum_{j=2}^J \pi_t^j z^j(\pi_t, v, \bar{z}_t^1) \\ &= \pi_t^1 z^1(\tilde{\pi}_t, v) + \sum_{j=2}^J \pi_t^j z^j(\tilde{\pi}_t, v) \\ &= c(\tilde{\pi}_t, v) + (\pi_t^1 - \tilde{\pi}_t^1) \bar{z}_t^1, \end{aligned}$$

where $\tilde{\pi}_t^1$ is the virtual price of the constrained characteristic and $\tilde{\pi}_t = [\tilde{\pi}_t^1, \pi_t^2, \dots, \pi_t^J]$ is the shadow support price vector. The virtual price is defined implicitly as a function of the constraint level, the shadow prices of the unconstrained characteristics, and the utility level. The shadow support prices are those that would precisely induce the consumer to choose the constrained characteristics vector:

$$\bar{z}_t^1 = z^1(\tilde{\pi}_t, v). \quad (18)$$

The welfare cost of the constraint on the characteristic is therefore

$$\begin{aligned} \frac{\partial c(\pi_t, v, \bar{z}_t^1)}{\partial \bar{z}_t^1} &= \frac{\partial c(\tilde{\pi}_t, v)}{\partial \pi_t^1} \frac{\partial \pi_t^1(\pi_t, v, \bar{z}_t^1)}{\partial \bar{z}_t^1} \\ &\quad - \frac{\partial \pi_t^1(\pi_t, v, \bar{z}_t^1)}{\partial \bar{z}_t^1} \bar{z}_t^1 + (\pi_t^1 - \tilde{\pi}_t^1) = \pi_t^1 - \tilde{\pi}_t^1, \end{aligned} \quad (19)$$

where the last equality follows from Shephard's lemma. Thus, the change in welfare associated with a change in the constraint is the difference between the shadow price of the characteristic and its virtual counterpart.

To look at the implications of this in the particular case of changes in the extensive margin of product characteristics, consider two adjacent periods $t \in \{0, 1\}$ and suppose that the first characteristic is not available in the base period ($\bar{z}_0^1 = 0$) but is added to the product specification in the following period. Then in the base period, the constrained expenditure function is

$$c(\pi_0, v, \bar{z}_0^1) = c(\tilde{\pi}_0, v) + (\pi_0^1 - \tilde{\pi}_0^1) \bar{z}_0^1, \quad (20)$$

which becomes

$$c(\pi_0, v, 0) = c(\tilde{\pi}_0, v). \quad (21)$$

²⁰ Keynes argued strongly in favor of the chaining method: "We are not in a position to weigh the satisfactions for similar persons of Pharaoh's slaves against Fifth Avenue's motor cars, or dear fuel and cheap ice to Laplanders against cheap fuel and dear ice to Hottentots. . . . We cannot hope to find a ratio of equivalent substitution for gladiators against cinemas, or for the conveniences of being able to buy motor cars against the conveniences of being able to buy slaves" (Keynes, 1930).

Consequently, the Konüs hedonic cost-of-living index can be expressed conveniently as the ratio of unconstrained hedonic expenditure functions where the base prices are given by the virtual prices:

$$\Pi_K(\pi_0, \pi_1, v) = \frac{c(\pi_1, v)}{c(\pi_0, v)}. \quad (22)$$

The fact that everything can be expressed in terms of the unconstrained expenditure function means that all of the approximate, exact, and superlative formulations for the index described above are available. But all of them require that the shadow price-relative π_1^1/π_0^1 is known.

In principle, recovering the virtual shadow price of the characteristic in the base period is a matter of solving $z^1(\pi, v) = 0$ for π^1 . This is the approach taken in the preference-for-goods context by Hausman (1997) who fits an empirical demand system to observations in which all goods are available (and therefore not subject to the constraint) and then evaluates them at the point where demand for the good of interest is driven down to zero. This approach could be adapted to the characteristics model, but it would be econometrically demanding. Our alternative approach, which is in the tradition of exact and superlative index numbers, is to make structural assumptions about preferences that make the objects of interest (hedonic cost-of-living indices in this case) largely known functions of observables.

In the new-goods context, the most attractive specification of this kind starts with the exact price index for CES preferences due to Sato (1976) and Vartia (1976) and extends it to allow for changes at the extensive margin, following Feenstra (1994).²¹

However, as we have seen, the preference-for-goods approach yields biased results when the set of available characteristics changes. Therefore, we must turn to a different approach.

IV. New Characteristics: A Structural Approach

To develop a practical framework for incorporating new characteristics into the true cost-of-living index, assume that preferences for characteristics exhibit a general (asymmetric) constant elasticity of substitution (CES) form. Thus,

$$v(\mathbf{z}) = \left(\sum_j (\beta^j z^j)^\theta \right)^{\frac{1}{\theta}}, \quad (23)$$

where $\theta \in [-\infty, 1]$. The preference parameters β^j reflect tastes for different characteristics.²² Their presence implies

²¹This approach has been applied to U.S. import demand by Broda and Weinstein (2006), and has been extended by Redding and Weinstein (2019), who attribute all unexplained variation to idiosyncratic taste shocks.

²²Baldwin and Harrigan (2011) call this multiplicative specification “box-size quality”: the utility derived from two boxes of unit-perceived quality, $\beta^j = 1$ equals that from a single box with quality equal to two. The same

that preferences need not be symmetric across characteristics.²³ The corresponding expenditure function is

$$c(\pi, u) = uc(\pi) \text{ where: } c(\pi) \equiv \left(\sum_j (\pi^j / \beta^j)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}, \quad (24)$$

where σ is the elasticity of substitution: $\sigma = (1 - \theta)^{-1}$ and $\sigma \in [0, \infty]$. Because preferences are homothetic, the unit expenditure function $c(\pi)$ is independent of the level of utility. It equals the true price level in a given period, and the ratio of these price levels for two periods gives the Konüs true price index between those periods. However, $c(\pi)$ depends on the unobservable taste parameters, β^j . To eliminate these, we extend the approach of Sato and Vartia, as adapted to allow for changes in the number of goods by Feenstra, from goods space to characteristics space. This allows us to express the true index in terms of observable budget shares.

By the logarithmic version of Shephard’s lemma, the log hedonic budget share for the j th characteristic is

$$\ln \omega^j = (1 - \sigma) (\ln \pi^j - \ln c(\pi) - \ln \beta^j). \quad (25)$$

Now consider two periods $t \in \{0, 1\}$ and suppose that the set of characteristics changes from \mathcal{J}_0 in the base period to \mathcal{J}_1 in the comparison period. The intersection $\mathcal{J} \equiv \mathcal{J}_0 \cap \mathcal{J}_1 \neq \emptyset$ contains the set of characteristics that are available in both periods. Taking the difference in log budget shares for the j th characteristic between periods gives

$$\ln \omega_1^j(\mathcal{J}_1) - \ln \omega_0^j(\mathcal{J}_0) = (1 - \sigma) ((\ln \pi_1^j - \ln \pi_0^j) - (\ln c(\pi_1) - \ln c(\pi_0))), \quad (26)$$

where the taste parameter β^j cancels. Now take a weighted sum (with weights ω^j to be determined) over the subset of characteristics available in both periods \mathcal{J} and rearrange to express the difference in log unit expenditure:

$$\ln c(\pi_1) - \ln c(\pi_0) = \sum_{j \in \mathcal{J}} \omega^j (\ln \pi_1^j - \ln \pi_0^j) - \frac{1}{\sigma - 1} \sum_{j \in \mathcal{J}} \omega^j (\ln \omega_1^j(\mathcal{J}_1) - \ln \omega_0^j(\mathcal{J}_0)). \quad (27)$$

specification is embodied in the equivalence scale model of household composition of Barten (1964).

²³The CES nonetheless assume, that characteristics have a constant elasticity of substitution. This could be relaxed by adopting a nested-CES approach as suggested in the preferences-for-goods context by Broda and Weinstein (2006). In the case of a characteristics model, this would amount to grouping characteristics by their general type (e.g., under such headings as “safety,” “comfort,” or “performance”) and applying a CES utility function to each separable group of characteristics and using an overall CES to also express preferences over groups. We are grateful to an anonymous referee for this observation.

Define the characteristic budget shares with respect to expenditure on the goods in this common set as

$$\omega_t^j(\mathcal{J}) = s_t^j(\mathcal{J})\eta_t, \quad (28)$$

where

$$\eta_t = \frac{\sum_{j \in \mathcal{J}} \pi_t^j z_t^j}{\sum_{j \in \mathcal{J}_t} \pi_t^j z_t^j} \quad (29)$$

is the share of shadow expenditure in period t spent on common characteristics. This gives

$$\ln c(\pi_1) - \ln c(\pi_0) = \sum_{j \in \mathcal{J}} \frac{\mu_j}{\mu} (\ln \pi_1^j - \ln \pi_0^j) - \frac{1}{\sigma - 1} (\ln \eta_1 - \ln \eta_0), \quad (30)$$

$$\mu_j = \frac{\omega_1^j(\mathcal{J}_1) - \omega_0^j(\mathcal{J}_0)}{\ln \omega_1^j(\mathcal{J}_1) - \ln \omega_0^j(\mathcal{J}_0)}, \quad (31)$$

so that the weights are

$$\omega^j = \frac{\mu^j}{\sum_{j' \in \mathcal{J}} \mu^{j'}}, \quad (32)$$

Finally, expressing the difference in log unit costs in levels gives the mix-adjusted Sato-Vartia index that we term the Characteristics-Sato-Vartia-Feenstra Index, or Characteristics SVF Index for short:

$$\Pi_{SVF} = \frac{c(\pi_1)}{c(\pi_0)} = \underbrace{\left(\frac{\eta_1}{\eta_0}\right)^{\frac{1}{\sigma-1}}}_{\text{Bias}} \underbrace{\prod_{j \in \mathcal{J}} \left(\frac{\pi_1^j}{\pi_0^j}\right)^{\omega^j}}_{\Pi_{SV}}. \quad (33)$$

Our new index is the product of two terms. The second, Π_{SV} , is a conventional Sato-Vartia index, extended to hedonic comparisons as in equation (14), so it is defined not over the goods but over the characteristics observed in both periods. The first term is an adjustment factor that has a simple interpretation as the proportional bias, Π_{SVF}/Π_{SV} , arising from ignoring new characteristics in the construction of the index and calculating it over the overlapping set \mathcal{J} . It takes into account the changing set of characteristics on the extensive margin in two ways. On the one hand, it depends on the share of new relative to old characteristics in the bundle and how that changes between periods. If, for example, many new characteristics are introduced, then η_1 will tend to be small and so the hedonic price index will be lower. Removing existing characteristics will have a countervailing effect on the denominator. On the other hand, the bias correction depends on the substitutability between characteristics, as governed by the σ parameter. The bias is less important the higher is the substitution parameter σ , that is, ignoring new characteristics matters less if they are

closer substitutes for existing ones. Equally, the bias is more important whenever the new characteristics are less easily substitutable with existing characteristics.

V. An Application: New Cars and New Characteristics

To show how our approach can be implemented, we apply it to a market-level U.K. data set on new cars and their characteristics. Section VA describes our data, section VB discusses the definition of products, and section VC provides the main results. Further results and robustness checks are given in the appendixes.

A. Data Sources and Description

The data were collected by the U.K. Office for National Statistics from the magazine *What Car*, and give detailed information on all major new car models that became available over the period 1988 to mid-1995. The data on market shares and sales are from the Competition Commission.²⁴

Basic descriptive statistics are given in table 1. Prices and specifications are measured quarterly. Prices are recorded in current, nominal pounds. The specifications listed from “Air conditioning” to “Driver’s airbag” are dummy variables. The acceleration variable records the time taken in seconds from 0 to 97 kilometers per hour (kph). Fuel consumption is measured in liters per 100 km at a speed of 90 kph. Engine size is recorded in cc’s. Brake horsepower and torque are measured at a fixed engine speed, and torque is measured in newton meters (NM). Finally, length is in meters.

Figure 1 shows how the distribution of prices evolved over the period. Panel a shows the density of prices in each quarter (later periods are indicated by lighter line shading). The price distribution moves to the right as prices rise over time. There is also a bimodal characteristic in the earlier years that is somewhat lost in later years where the market is less obviously split into normal/family versus luxury cars. Panel b shows a contour plot of the evolution of the price distribution over time. This illustrates the general upshift in new car prices over the period. The solid line overlaid on the contours shows the quarterly average price. During the period, average prices of new cars rose by about 50% from £7,350 in the first quarter of 1988 to £10,853 by the end of the second quarter in 1995. The average price of new cars over the period as a whole was somewhat over £9,000.


The specifications of new cars generally changed twice per year in the United Kingdom during the period studied: at the beginning of the fourth quarter, which marks the time when

²⁴http://webarchive.nationalarchives.gov.uk/20111202195250/http://competition-commission.org.uk/rep_pub/reports/2000/439cars.htm. As noted by a referee, statistical agencies may not always have market share data available. Under these circumstances, they typically produce equally weighted indices based just on the price quotes in their sample (the Carli, Dutot, and Jevons indices are the most popular). Consonant with this practice, a statistical agency without access to weights would set $w_t^k = \frac{1}{|\mathcal{J}_t|}$ and $\omega_t^j = \frac{1}{|\mathcal{J}_t|}$ and $h_t = \frac{|\mathcal{K}|}{|\mathcal{K}_t|}$ and $\eta_t = \frac{|\mathcal{J}|}{|\mathcal{J}_t|}$.


TABLE 1.—DESCRIPTIVE STATISTICS

Statistic	N	Mean	S.D.	Min	Max
Year	3,160	1991	2.12	1988	1995
Price (£)	3,160	9,168	2,629	4,215	17,409
Air-conditioning	3,160	0.01	0.09	0	1
Central locking	3,160	0.48	0.50	0	1
Manual sunroof	3,160	0.29	0.46	0	1
Electric sunroof	3,160	0.09	0.29	0	1
Electric front windows	3,160	0.28	0.45	0	1
All-electric windows	3,160	0.04	0.19	0	1
Powered (heated) seat	3,160	0.14	0.35	0	1
Headlamp cleaners	3,160	0.01	0.12	0	1
Electric mirrors	3,160	0.14	0.35	0	1
Trip computer	3,160	0.01	0.07	0	1
Split-fold seats	3,160	0.65	0.48	0	1
Seat adjustment	3,160	0.29	0.45	0	1
Radio	3,160	0.11	0.31	0	1
Radio cassette	3,160	0.77	0.42	0	1
CD player	3,160	0.003	0.06	0	1
Steering adjustment	3,160	0.20	0.40	0	1
Driver's airbag	3,160	0.11	0.32	0	1
Acceleration (secs, 0 to 97 kph)	3,160	13.42	2.72	8.70	21.40
Fuel consumption (L/100 km at 90 kph)	3,160	5.42	0.61	3.62	7.06
Engine size (cc)	3,160	1,391	256	988	1,998
Brake horsepower	3,160	73	21	41	118
Torque (NM)	3,160	81	19	50	134
Length (m)	3,160	3.98	0.35	3.05	4.93
French	3,160	0.17	0.37	0	1
German	3,160	0.14	0.35	0	1
Italian	3,160	0.08	0.28	0	1
Japanese	3,160	0.08	0.28	0	1
Spanish	3,160	0.03	0.16	0	1
Swedish	3,160	0.03	0.16	0	1
U.K.	3,160	0.14	0.35	0	1
U.S.	3,160	0.33	0.47	0	1

FIGURE 1.—THE DENSITY OF NEW CAR PRICES OVER TIME



(a) Density of Prices in each Quarter
(Lighter shading denotes later periods)



(b) Evolution of Price Distribution Over Time
(Solid line shows the quarterly average price)

new annual registration (number-plate) letters were issued in the United Kingdom, and the beginning of the second quarter, which normally coincides with changes in vehicle excise duty and taxes on motor fuels in the government's annual finance bill.

Looking at the intensive margin first, we see that quality improved in many dimensions: acceleration times of 0 to 60 mph improved over the period by 6.5%, while fuel economy improved by 2.8%. Both torque and brake horsepower improved by 5.2% on average (torque and BHP are proportional at a fixed engine speed). The availability of features also improved: central locking was present in a quarter of cars at the start of the period; this rose to nearly two-thirds by the end. Manual sunroofs went from being available on 17.6% of models to 25%. Electric front windows were uncommon at the beginning with an 8% market share, but by the end, 40% of new models had them. Electrically operated door mirrors went from a similar level to being available on a quarter of new cars, and split-fold rear seats increased from being specified in around half of new cars to nearly three-quarters.

Of more interest from our perspective, on the extensive margin, there were a number of changes to the set of characteristics over the period. In terms of new characteristics, electronic seat and steering adjustability became available in the fourth quarter of 1988; electric sunroofs were introduced in the first quarter of 1990; air-conditioning was introduced in the fourth quarter of 1993, as were powered (heated) seats and drivers' airbags. CD players were introduced in the second quarter of 1994. These appeared alongside radio-cassette players, which were available throughout, but by the time CD players were introduced, simple radio-only models had gone; they disappeared from the new car market by the fourth quarter of 1992. In the case of in-car entertainment, there was a clear sequence from radios (gone by the end of 1992), via radio-cassettes, toward CD players (from 1994). Other features were short-lived and simply came and went: headlamp cleaners were available in some cars (Volvos) at the start of our period of observation but disappeared at the start of 1990, only to reappear briefly from mid-1993 to mid-1994 (Alfas and Volvos again) before disappearing again. Similarly early trip computers became available (also on Alfas) but disappeared only to be offered briefly by Rover before again disappearing. Trip computers are now completely standard, but they did not take off until after the end of our period of observation.

B. *The Definition of Products*

In the case of differentiated product markets, the researcher has a choice regarding the definition of a product, and that choice matters.²⁵ There are many options, but the following three seem to encompass the range of possibilities. First, we could take a narrow view and define a product as any good

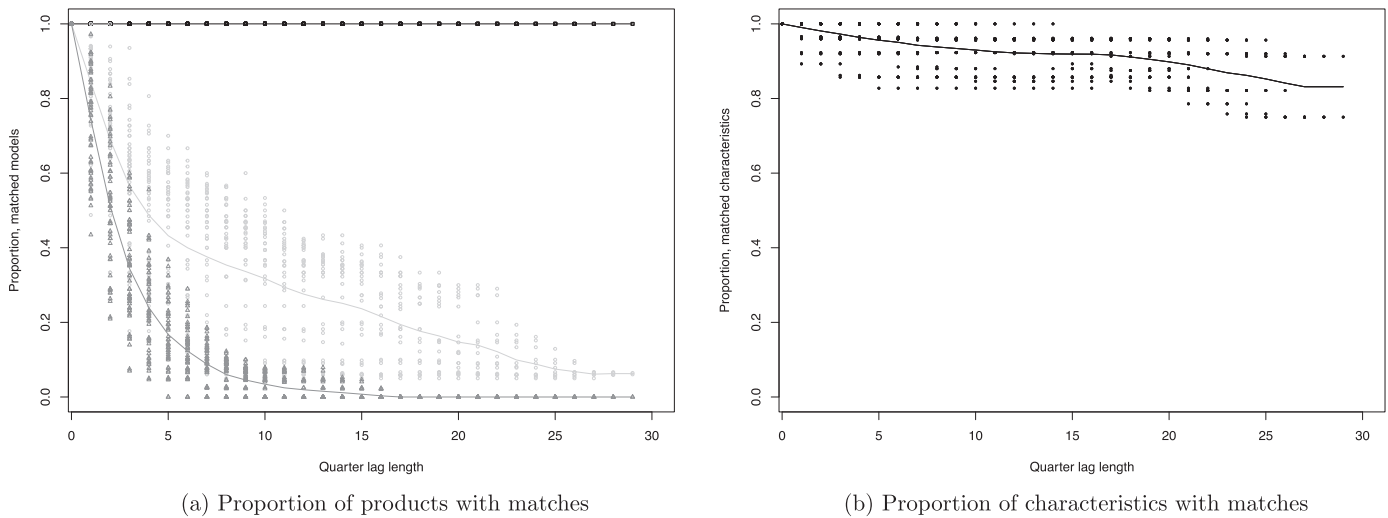
with a fixed set of characteristics on both the extensive and intensive margins. On this definition, neither the range nor the specification of a product's features would ever change, and so the quality of a product would be fixed from the time of its launch. For example, if Toyota improved the fuel economy of the Corolla, this would entail the death of an old product and the birth of a new one: the older product would disappear and would be replaced. A drawback of this view is that the introduction of new products and the withdrawal of old ones would likely mean that the rate of churn of products could be rapid. As a result, we would expect that product-space indices would suffer greatly from a lack of matched models, and so the classic new-goods problem would be exacerbated. Second, we could instead take a broad view in which products are defined broadly enough that there is no entry or exit of products at all. So, for example, changes in the specifications of the car marketed as a Toyota Corolla would all be classed as variants with different characteristics but Toyota Corollas nonetheless. In this case, there would be no new goods but lots of changes in the characteristics of products on both margins. As a result, matched-model product-space indices would be easy to construct, though these would suffer from quality bias, and hedonic matched-characteristic indices would suffer from the new-characteristics problem. Third, we could adopt an intermediate approach, defining products on the basis of the presence or absence of a fixed list of embodied characteristics. Thus, existing products would be allowed to have upgrades, but as soon as a new feature was added or removed, it would be classed as a new product. So, for example, a Toyota Corolla with air-conditioning would be classed as a different product from one without but one with improved fuel efficiency would not. The entry or exit of new characteristics and new products would thus go together. In our narrowest definition (where a product is defined by its exact specifications), there are 353 products in our data; in our broadest (where a product is defined by its model name), there are 36; and in our intermediate definition, there are 171.

To illustrate some of the implications of these choices, we have constructed the data for each of these different approaches.²⁶ Figure 2 illustrates the prevalence of matched models and characteristics in our data for different time separations (lag-lengths) between base and comparison periods for the three approaches to the definition of products. The left-hand panel (a) shows the proportion of products that are available in a base period that can be matched to products in a comparison period against the number of quarters between the base and comparison period. The three different scatters and lines correspond to the different product definitions

²⁵We are grateful to the editor and referees who made us think more carefully about this point.

²⁶We do this by first creating product dummies according to each definition and then averaging both prices and characteristics by product \times time period. For example, in the case of the Toyota Corolla, for the broadest definition of the product, the price and the characteristics in the t th period are the means over all variants available in that period; for the intermediate definition, the price and the characteristics are the means over all Corollas that share the same set of characteristics on the extensive margin in that period, while for the narrowest definition, no aggregation is necessary.

FIGURE 2.—PROPORTION OF MATCHED PRODUCTS AND CHARACTERISTICS, BY LAG-LENGTH AND PRODUCT DEFINITION



(broad, intermediate, narrow). For each lag length between periods, the scatterplots show the proportion of matched models for all possible base-period and comparison-period pairs, while the lines show the average proportion. Both the data points and the lines equal 1 in the top left-hand corner where the lag length is 0. (When the base period and comparison period are the same, the match between the products available in each is, of course, perfect). As we move to the right, the lag length between base and comparison period grows, and there are progressively fewer matched products, except for the broad definition of products where the match remains perfect as the list of models remained constant throughout the period, whatever the time difference between base and comparison period. However, for the intermediate and the narrow definition of goods, the proportion of available matches drops rapidly. The proportion of products that can be matched drops below 0.5 within a year on both definitions and drops to 0 at a five-year separation in the case of the narrow definition; on this definition, no products survive the entire period. Expressing the same information as market shares by value shows a similar pattern.²⁷

The right-hand panel b shows the same information for characteristics. Since the way we define characteristics is independent of the way we define products, the proportion of matched characteristics does not vary with the granularity of products. Panel b shows that characteristics are longer-lived than products (on all but the broadest definition). Nonetheless, it is not possible to construct hedonic price indices composed of fully matched characteristics because of changes in the extensive margin of characteristics. It is this feature that gives rise to the issue of new-characteristics bias.

This discussion has two implications for our approach. On the one hand, it shows that the problems of new-goods bias

and new-characteristics bias, discussed in section I and formalized in sections II to IV depend on how goods are defined: on the broad definition of goods, there can be no new-goods bias by assumption. On the other hand, it suggests important practical considerations in implementing either the narrow or the intermediate definitions of products. Because of the frequent churning of products they imply, constructing consistent index numbers over time faces a common-goods problem: there may be no or very few comparable goods in periods we wish to compare. (In terms of the notation we introduced in section IV, the intersection $\mathcal{K} \equiv \mathcal{K}_0 \cap \mathcal{K}_1$ of the set of goods available in both periods may be empty.) Of course, in principle, the broad definition of products faces an analogous common-characteristics problem. However, the empirical fact that the range of characteristics changes much less frequently than that of goods in the narrow definition suggests that the broad definition is operational, and the methods we have developed in previous sections allow us to deal with this problem. On both a priori and practical grounds, therefore, the narrow definition of goods seems the most plausible. However, to illustrate the alternatives, we present results for all three definitions.

C. Results

In this section, we present the results of calculating our new index for the data on U.K. cars. Since our solution is a modified hedonic Sato-Vartia index, we focus on the Sato-Vartia index and its extensions throughout. Other index number formulas (Laspeyres, Paasche, Fisher) have also been calculated: the matched-model and matched-characteristic Fisher indices are virtually indistinguishable from the Sato-Vartia, while the Paasche and Laspeyres indices exhibit their classic substitution-biased behavior. The full range of indices is shown in appendix A for the broad definition of products.

²⁷Further details are available from us.

TABLE 2.—THE SATO-VARTIA INDEX AND ITS EXTENSIONS

	Product Based	Hedonic
Constant Mix	$P_{SV} = \prod_{k \in \mathcal{K}} \left(\frac{p_1^k}{p_0^k} \right)^{w^k}$	$\Pi_{SV} = \prod_{j \in \mathcal{J}} \left(\frac{\pi_1^j}{\pi_0^j} \right)^{\omega^j}$
Mix-Adjusted	$P_{SVF} = \left(\frac{h_1}{h_0} \right)^{\frac{1}{s-1}} \prod_{k \in \mathcal{K}} \left(\frac{p_1^k}{p_0^k} \right)^{w^k}$	$\Pi_{SVF} = \left(\frac{\eta_1}{\eta_0} \right)^{\frac{1}{s-1}} \prod_{j \in \mathcal{J}} \left(\frac{\pi_1^j}{\pi_0^j} \right)^{\omega^j}$

They are also available from us for the intermediate and narrow definitions. Appendix A also illustrates the confidence intervals on the hedonic indices. In the main results below, the confidence intervals are omitted to avoid cluttering the presentation. Finally, fixed-base indices have been calculated for all possible bases. The ones shown here are based on the first quarter of our data; the others are also available from us.

Using the Sato-Vartia as our basic index-number formula, the four classes of index number we present are given in table 2. Here, P_{SV} is the standard Sato-Vartia index for goods, and P_{SVF} is its extension by Feenstra (1994) to allow for changes in the product mix consumed. The notation is standard: s is the elasticity of substitution between goods, the weights w^k are the normalized logarithmic means of the period 0 and period 1 budget shares, and h_t is the share of expenditure in period t spent on goods commonly available in both periods. Similarly, Π_{SV} is the Sato-Vartia index applied to a fixed bundle of characteristics, as given in equation (14), while Π_{SVF} is its extension to changes in characteristics at the extensive margin, introduced in equation (33).

To construct our hedonic index numbers, we first estimate the linear hedonic pricing equation (16) by regressing goods prices on characteristics to obtain estimates of the shadow prices. The estimation is done with weighted least squares, using market shares as weights, subject to the restrictions from the underlying theory that shadow prices are weakly positive ($\hat{\pi}_t^j \geq 0$), and that the hedonic demands satisfy adding-up:

$$\sum_k p_t^k x_t^k = \sum_j \hat{\pi}_t^j z_t^j. \quad (34)$$

This is carried out for each of our three definitions of products in every period. The hedonic regressions are estimated using the characteristics listed in table 1.²⁸ Where a characteristic is missing, either because it is yet to be introduced or because it has been withdrawn, it is omitted from the regression model. Otherwise, the specification is the same in all periods.

Given the estimated shadow prices and the observed characteristics $\{\hat{\pi}_t^j, z_t^j\}$, we recover the elasticity parameter re-

quired to adjust the Sato-Vartia index by fitting the asymmetric CES demand system,

$$\hat{z}_t^j = \beta_j^\sigma \left(\frac{\hat{\pi}_t^j}{P_t} \right)^{-\sigma} \frac{y_t}{P_t}, \quad (35)$$

where $P_t = \left(\sum_j \beta_j^\sigma \hat{\pi}_t^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$, to the characteristics by nonlinear least squares. Our procedure therefore does not account for any endogeneity, and this should be borne in mind when assessing the results.²⁹ Once again, we carry out these procedures separately for each of our three definitions of products (narrow, intermediate, broad). Confidence intervals for the price indices constructed from these estimates, and for all the indices that use hedonic estimates for quality adjustment, are calculated using the jackknife procedure. Constructing a mix-adjusted, matched-model, product-based index (i.e., a nonhedonic index) according to the Sato-Vartia-Feenstra method also requires an estimate of the substitution elasticity. We calculate this in a similar way using product prices and quantities instead of shadow prices and characteristics. The results are reported in appendix C.

Fixed-Base Indices. Figure 3 shows the main results for fixed base indices (1988q1 = 1). The index numbers fall into five groups. There are two groups of product-based indices: matched-model (indicated by \diamond) and mix-adjusted (indicated by $+$). There are two groups of hedonic indices: matched-characteristics (indicated by $*$) and our characters-mix-adjusted index, which allows for extensive-margin changes (\bullet). Finally, for comparison, we include a group of indices based on the time-dummy method (indicated by \times). Recall that this method is not used by statistical agencies because it needs to be constantly retrospectively reestimated and revised as each new year is added. The effects of these revisions and the economic model underlying the index are detailed in appendix B. Figure 3 only shows the very final index of the sequence of revised indices. While the time-dummy approach is not used in practice, we include it in the figures as a useful sanity check on the other results.³⁰

Within each of these five groups, three indices are shown. These correspond to each of our three alternative ways of defining products: the broad definition with 36 products (solid line), the intermediate wherein there are 171 products (dotted line), and the narrow definition wherein there are 353 products (dot-dash). Each group of indices is collectively shaded so that the general trends for each group can be more easily picked out.


Looking first at the product-based, matched-model (\diamond) group, we note the following. First, on the narrow definition of products (353 products with no within-product quality change as measured by changing characteristics), the index fails to cover the period. It ends by early 1992, as at that

²⁸ We do not report the results of these 3×30 regressions here. Instead we focus on the substantive results: the price index numbers that result. The full set of regression results, the code used to generate them, and the rest of the empirical work in the paper are available from us.

²⁹ We are grateful to a referee for pointing out this caveat.

³⁰ We are grateful to an anonymous referee for suggesting this.

FIGURE 3.—FIXED-BASE PRICE INDICES



1988q4: adjustable seats and adjustable steering introduced; 1990q1: electric sunroofs introduced; 1990q2: headlamp cleaners withdrawn; 1991q1: trip computers introduced; 1992q1: trip computers withdrawn; 1992q3: trip computers reintroduced for one quarter; 1992q4 radio-only and trip computers withdrawn; 1993q2: headlamp cleaners reintroduced; 1993q4: air-conditioning, power seats, and driver's airbags introduced; 1994q2 CD players available; 1994q4: headlamp cleaners finally withdrawn.

point, no products that existed in 1988q1 still remain in the sample; they have all exited so can no longer be matched. On the intermediate and broad definition of products, the indices at least cover the period. However, while the index for the intermediate (171 product) definition spans the period, figure 2 showed that within a few quarters of the start of the series, half of the products already cannot be matched and are excluded. This index should therefore be considered unreliable for the reasons set out in Pakes (2003). Furthermore, both of these indices suffer from quality bias as the characteristics of the products are changing over the period, and neither of these indices (being nonhedonic) allows for this.

Looking next at the (+) group of product-based indices, which are mixed-adjusted using the Feenstra (1994) method, we see a number of features. Focusing first on the broad definition of goods, the index is identical to the non-mix-adjusted counterpart. This is simply because on the broad definition, there was no entry or exit of products over the period, hence no need to mix-adjust. The indices are the same as a result; hence, the remarks about quality bias made above apply to both. Turning to the intermediate and the narrow definitions, we can see that the indices that result from applying the Feenstra (1994) method to products give unusably unstable results. The estimates of the elasticity parameter are not the source of the instability; their values are generally not too far from the estimates for the hedonic models and these are stable, as we shall see. Rather, the instability arises from the fact that the index is driven almost entirely by the mix-adjustment factor,


because there are so few matched models when products are defined in this way. Note too that mix adjustment in product-based indices can work only when there exists an index to adjust: when the index based on the narrow definition runs out in 1988q1, so does the mix-adjusted version.

Summarizing the results for the product-base indices, we conclude that these indices suffer from a range of problems depending on the way goods are defined: failure to cover the period, instability or unreliability, and quality bias. It should be emphasized that these results do not apply to the Redding and Weinstein (2019) approach wherein error terms in an empirical demand system are interpreted as time-varying preferences or preference shocks and then adjusted for. Doing so may improve the product-based indices in terms of both stability and neglected quality-bias.³¹

We now turn to the hedonic indices. The general tenor of the results is as follows. First, the matched-characteristic, hedonic indices (*) show a lower rate of price growth than the (broad) matched-model, product-based index (◇). This reflects the standard finding that adjustment for quality improvements on the intensive margin pushes down the measured rate of price growth or, equivalently, that failure to adjust for quality improvements causes an upward bias. Second, making the adjustment for entry and exit of characteristics (●) tends to push down the index more; or, equivalently, matched-characteristic hedonic indices do not go far enough

³¹We are grateful to an anonymous referee for making this observation.

FIGURE 4.—FIXED-BASE ANNUAL INFLATION RATES



1988q4: adjustable seats and adjustable steering introduced; 1990q1: electric sunroofs introduced; 1990q2: headlamp cleaners withdrawn; 1991q1: trip computers introduced; 1992q1: trip computers withdrawn; 1992q3: trip computers reintroduced for one quarter; 1992q4 radio-only and trip computers withdrawn; 1993q2: headlamp cleaners reintroduced; 1993q4: air-conditioning, power seats and driver's airbags introduced; 1994q2 CD players available; 1994q4: headlamp cleaners finally withdrawn.

because they do not allow for changes on the extensive margin. The impact of the changes on the extensive margin (notably, important comfort and safety features: air-conditioning and airbags) which all appeared together in 1993 show up clearly. Third, these general results for hedonic indices are not particularly dependent on the definition of products, and similar ranking and patterns emerge regardless.

Finally we consider the results from the time-dummy method (green group). As discussed in section I, this index must be revised every time another period is added and the indices in figure 3 represent only the final revision.³² The results from the time-dummy model track the other hedonic methods until 1993 when, as we have seen, the extensive margin changes begin to accrue more rapidly. The model underlying the index differs from the Gorman-type (1956) characteristics model in a number of important ways,³³ but perhaps the most relevant here is the fact that the model assumes that products are perfect substitutes. Recall that the higher the degree of substitutability, the lower is the effect of product innovations on the extensive margin. The perfect substitutability assumed between products, and hence the implied substitutability be-

tween the characteristics embodied in these products, therefore means that the index will not fully capture the effects of extensive margin changes. This is indeed what we observe in the resulting figure.


Figure 4 shows the annual inflation rates for the fixed base indices. We omit the results for the mixed-adjust product-based indices because of the way they were observed to behave in figure 3. Some additional instability in the product-based, matched-model indices shown in figure 3 is much more apparent when the data are presented as rates-of-change in figure 4. In terms of inflation rates, the two hedonic approaches track each other reasonably well except during this period in which the burst of the extensive-margin changes occurred.

As we saw in section IV, the adjustment factor $(\eta_1/\eta_0)^{1/(\sigma-1)}$ to the index can be interpreted as the proportional bias arising from ignoring extensive-margin changes in characteristics. The bias therefore depends on both the empirical rate of entry and exit of characteristics and the estimate of the elasticity of substitution. Given the empirical rate of entry and exit, the more substitutable the entering and exiting characteristics are with respect to the other characteristics, the lower the bias caused by ignoring them. Figure 5 shows the sensitivity of the bias term to alternative values of σ . It is measured in the final period of the data shown in figure 3. There are, in fact, three curves illustrated, each corresponding to one of the three definitions of products. However, the curves for the narrow ($K = 353$) and the intermediate ($K = 171$)

³²See appendix B for the dynamic effects of revisions in our data.

³³The underlying model assumes that products are perfect substitutes; there is a functional form difference in the hedonic regression itself; the shadow prices of characteristics are assumed time-invariant, and as a result, unlike the classical hedonic indices reviewed in section II, the index is independent of any reference characteristics. See appendix B for a full description of the model.

FIGURE 5.—NEW-CHARACTERISTICS BIAS AS A FUNCTION OF THE ELASTICITY OF SUBSTITUTION



definitions are visually indistinguishable. The three points illustrated on each curve are the bias values corresponding to the estimated value of σ , and the curves show what the bias would be if the elasticity of substitution were different.³⁴ As the figure shows, the results are potentially sensitive to the elasticity of substitution if it were much lower than the values we estimate and use. For the values of σ that we estimate, the implied values of the bias for different product definitions are very close to each other.

Chained indices. The benefit of fixed-base, as opposed to chained, index numbers is that it is clear what is being compared to what; either the set of products in the case of product-based, matched-model index numbers or the set of characteristics in the case of matched-characteristic, hedonic indices, is held constant across the entire period. However, one of the findings from the fixed-base results was that the churn in products makes even fixed-base, matched-model index numbers that are based on anything but broadly defined product categories unusable. It is to be expected that such indices suffer in the fixed-base approach as it gets harder and harder, and in the end impossible, to find matches as the time between the base and the comparison period gets longer. Chained indices, by contrast, have the significant advantage of allowing the mix of products or characteristics to change over time. With chained index numbers, matches only need


to be found and maintained between adjacent periods. Such matches are relatively abundant, as we saw in figure 2, and so fewer products have to be dropped. This is a major benefit associated with chaining and partly explains their wide use. While chaining is usually argued for on the basis that it helps to address the new-goods problem, it clearly also applies in characteristics space: just as chained product-based indices help with new goods, chained hedonic indices should help with the new-characteristics problem. A second benefit of chaining is that it allows preferences to change between links in the chain. It is therefore a very simple way of adjusting for taste change with respect to either products or characteristics.

Figure 6 illustrates the chained indices. Once again, as with figure 3, we have five groups of indices: matched-model and mix-adjusted, product-based indices (indicated by \diamond and $+$ respectively); matched-characteristic and our hedonic indices (indicated by $*$ and \bullet); finally the time-dummy index (\times). Note that the chained time-dummy index is, necessarily by construction, identical to its fixed-base counterpart shown in figure 3.

We note that the chained matched-model indices (\diamond group) now behave relatively well; they all cover the period studied and are relatively stable. However, since they make no adjustment for product improvements, we regard both indices based on the broad and intermediate definitions of goods as suffering from a degree of upward quality bias. The index based on the narrow definition of goods does not suffer to

³⁴See appendix C for further details.

FIGURE 6.—CHAINED PRICE INDICES



1988q4: adjustable seats and adjustable steering introduced; 1990q1: electric sunroofs introduced; 1990q2: headlamp cleaners withdrawn; 1991q1: trip computers introduced; 1992q1: trip computers withdrawn; 1992q3: trip computers reintroduced for one quarter; 1992q4 radio-only and trip computers withdrawn; 1993q2: headlamp cleaners reintroduced; 1993q4: air-conditioning, power seats, and driver's airbags introduced; 1994q2 CD players available; 1994q4: headlamp cleaners finally withdrawn.

the same degree, but since it necessarily omits products that change specification between adjacent periods, its representativeness suffers in periods in which many products undergo such changes. The mix-adjusted, product-based indices (+) are, slightly, improved through chaining in the sense that, compared to figure 3, we can see that all of the indices do at least now cover the whole period and they are somewhat less unstable. However, they are still either unusable (narrow and intermediate definitions of products) or unnecessary (broad definition where no mix adjustment is needed). The chained hedonic indices track each other well until the period where the bulk of extensive margin changes occur. As this point the matched-characteristic hedonic indices (*) fail to capture the effect of these new characteristics, as shown by the mix-adjusted hedonic index (•).


Figure 7 shows the corresponding annual rates of inflation. Again we have omitted the inflation rates for the mix-adjusted, product-based indices (+ group) for the same reasons as previously. Both figures 6 and 7 show that the chained hedonic mix-adjusted index does not depart significantly from the matched-characteristics indices during the early part of the study period. The quarterly changes in specifications on the extensive margin were evidently too minor to matter greatly. However, later on it is clear that although chaining helps, it does not always provide an empirical fix for the new-characteristics problem. The burst of innovations that occurred together in 1993 (air-conditioning, power (heated) seats, and drivers' airbags) are not ignorable even over the

short term, and hence our mix-adjusted hedonic index adjusts to accommodate that. Comparing figures 3 and 6, we can see that by the end of our period of study, the overall effects of adjusting for the characteristics mix are similar regardless of whether we choose a fixed-base or a chained index approach.

VI. Conclusion

Hedonic methods are now an established part of national statistical offices' tool kit. Since the time-dummy method based on the repackaging model cannot be used by statistical offices because of the revisions it entails, the only method in widespread use is the indirect method or variations on it. Hedonic indices based on the indirect method suffer from a problem akin to new-goods bias when there are changes in product specifications on the extensive margin (new features are introduced, old features withdrawn). We have shown that the approach developed by Feenstra (1994) for mix-adjusting the Sato-Vartia index can be adapted to take account of changes in the mix of characteristics. The resulting index, which we call the Characteristics-Sato-Vartia-Feenstra Index, provides an effective and theoretically consistent method of allowing for such changes. We have also shown in our application to U.K. new car prices that it makes a quantitatively important difference to estimated inflation rates. Very roughly, if we take the Boskin Commission's estimate of the bias in measured inflation from new goods and quality change of 0.6 percentage points per year, our results suggest that allowing

FIGURE 7.—CHAINED ANNUAL INFLATION RATES



1988q4: adjustable seats and adjustable steering introduced; 1990q1: electric sunroofs introduced; 1990q2: headlamp cleaners withdrawn; 1991q1: trip computers introduced; 1992q1: trip computers withdrawn; 1992q3: trip computers reintroduced for one quarter; 1992q4 radio-only and trip computers withdrawn; 1993q2: headlamp cleaners reintroduced; 1993q4: air-conditioning, power seats, and driver's airbags introduced; 1994q2 CD players available; 1994q4: headlamp cleaners finally withdrawn.

for changes in the extensive margin of characteristics could add half as much again to that, at least for those commodity groups in which proliferation of product features is important.

Two potential drawbacks of our approach need to be borne in mind. The first is that it is based on the assumption that preferences for characteristics are asymmetric CES. Constant elasticity of substitution preferences can be derived as the aggregation of the choices of individual consumers with extreme-value-distributed idiosyncratic preferences, as shown in Anderson et al. (1992). This provides some justification for using the CES with market-level data as we do here. Nonetheless, the asymmetric CES is “not-quite-superlative”;³⁵ that is, it is not a flexible functional form, and therefore we cannot necessarily be confident that the approximation to true preferences will be good (although see Hill, 2006). This means that the use of the Sato-Vartia may involve a misspecification error if true preferences are not close to the generalized CES. In this respect, we draw some comfort from the empirical results that show that the not-quite-superlative Sato-Vartia index is, in our data, empirically indistinguishable from the superlative Fisher index. The Fisher index *is* superlative and so does not (to second order) suffer from misspecification. Since the Sato-Vartia empirically approximates the Fisher so closely, we conclude that the misspecification in the Sato-Vartia is not empirically significant in our application. This is similar to a finding in Redding and

Weinstein (2019)³⁶ where it is also used to justify the use of the CES. Of course, this argument is an empirical one and may not necessarily extend to other contexts and applications. In general, it would therefore seem a sensible, cautious approach to also always compute superlative indices and check whether they diverge from the CES-based index. If they do, then the misspecification is cause for concern. Should this be the case, then a further possible avenue would be to combine our characteristics-based approach with the translog preferences considered by Feenstra and Weinstein (2017).³⁷

The second drawback is that just as the Sato-Vartia index is not-quite-superlative, the Sato-Vartia-Feenstra index is not-quite-exact. Like the Sato-Vartia-Feenstra index, ours requires the value of the elasticity of substitution parameter σ . We estimated this by fitting CES demand-for-characteristics functions. This represents the one additional piece of estimation that a statistical office would need to do over and above its normal hedonic estimation work. An alternative would be to select suitable parameter values for σ from the literature and conduct a sensitivity analysis.

Despite these shortcomings, our method buys the ability to address the new-goods/new-characteristics problem without any of the difficult econometric work involved in solving estimated characteristics-demand equations for reservation prices. This is, we would argue, a practical plus, and our method is one that statistical offices could easily adopt.

³⁵W. Erwin Diewert, personal communication.

³⁶See p. 507.

³⁷We are grateful to an anonymous referee for suggesting this procedure.

REFERENCES

- Akerberg, D., C. L. Benkard, S. Berry, and A. Pakes, "Econometric Tools for Analyzing Market Outcomes" (pp. 4171–4276), in J. J. Heckman and E. E. Leamer, eds., *Handbook of Econometrics*, vol. 6, pt. A (Amsterdam: Elsevier, 2007). 10.1016/S1573-4412(07)06063-1
- Anderson, S. P., A. De Palma, and J.-F. Thiss, *Discrete Choice Theory of Product Differentiation* (Cambridge, MA: MIT Press, 1992).
- Baldwin, R., and J. Harrigan, "Zeros, Quality, and Space: Trade Theory and Trade Evidence," *American Economic Journal: Microeconomics* 3:2 (2011), 60–88. 10.1257/mic.3.2.60
- Ball, A., and A. Allen, "The Introduction of Hedonic Regression Techniques for the Quality Adjustment of Computing Equipment in the Producer Prices Index (PPI) and Harmonised Index of Consumer Prices (HICP)," *Economic Trends* 592 (2003), 30–36.
- Barten, A. P., "Family Composition, Prices, and Expenditure Patterns," in P. Hart, G. Mills, and J. K. Whitaker, eds., *Econometric Analysis for National Economic Planning* (London: Butterworth, 1964).
- Blow, L., M. Browning, and I. Crawford, "Revealed Preference Analysis of Characteristics Models," *Review of Economic Studies* 75:2 (2008), 371–389. 10.1111/j.1467-937X.2008.00485.x
- BLS, *Handbook of Methods: Consumer Price Index* (Washington DC: Bureau of Labor Statistics, 2020).
- Boskin, M. J., E. Dulberger, R. Gordon, Z. Griliches, and D. Jorgenson, "Toward a More Accurate Measure of the Cost of Living," Final Report to the Senate Finance Committee (December 4, 1996), <http://www.ssa.gov/history/reports/boskinrpt.html>
- Broda, C., and D. E. Weinstein, "Globalization and the Gains from Variety," *Quarterly Journal of Economics* 121:2 (2006), 541–585. 10.1162/qjec.2006.121.2.541
- Diewert, W. E., "Exact and Superlative Index Numbers," *Journal of Econometrics* 4:2 (1976), 115–145. 10.1016/0304-4076(76)90009-9
- , "The Economic Theory of Index Numbers: A Survey" (pp. 163–208), in A. Deaton, ed., *Essays in the Theory and Measurement of Consumer Behaviour in Honour of Sir Richard Stone* (Cambridge: Cambridge University Press, 1981).
- Feenstra, R. C., "New Product Varieties and the Measurement of International Prices," *American Economic Review* 84:1 (1994), 157–177.
- Feenstra, R. C., and D. E. Weinstein, "Globalization, Markups, and U.S. Welfare," *Journal of Political Economy* 125:4 (2017), 1040–1074. 10.1086/692695
- Feldstein, M., "Underestimating the Real Growth of GDP, Personal Income, and Productivity," *Journal of Economic Perspectives* 31:2 (2017), 145–164. 10.1257/jep.31.2.145
- Fisher, F. M., and K. Shell, "Taste and Quality Change in the Pure Theory of the True Cost of Living Index" (pp. 97–139), in J. N. Wolfe, ed., *Value, Capital and Growth* (Edinburgh: Edinburgh University Press, 1968).
- Gordon, R. J., "The Boskin Commission Report and Its Aftermath," *Monetary and Economic Studies* 17:3 (1999), 41–68.
- Gorman, W. M., "A Possible Procedure for Analysing Quality Differentials in the Egg Market," *Review of Economic Studies* 47:5 (1980), 843–856. 10.2307/2296916
- Griliches, Z., "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change" (pp. 173–196), in Z. Griliches, ed., "Hedonic Price Indexes for Automobiles: An Econometric Analysis of Quality Change" (pp. 173–196), in NBER, ed., *The Price Statistics of the Federal Government* (New York: NBER, 1961).
- Hausman, J., "Sources of Bias and Solutions to Bias in the Consumer Price Index," *Journal of Economic Perspectives* 17:1 (2003), 23–44. 10.1257/089533003321164930
- , "Valuation of New Goods under Perfect and Imperfect Competition" (pp. 209–247), in T. F. Bresnahan and R. J. Gordon, eds., *The Economics of New Goods* (Chicago: Chicago University Press, 1997).
- Hill, R. J., "Superlative Index Numbers: Not All of Them Are Super," *Journal of Econometrics* 130:1 (2006), 25–43. 10.1016/j.jeconom.2004.08.018
- Keynes, J. M., *A Treatise on Money*, vol. 1: *The Pure Theory of Money* (London: Macmillan, 1930).
- Lancaster, K. J., "A New Approach to Consumer Theory," *Journal of Political Economy* 74:2 (1966), 132–157. 10.1086/259131
- Neary, J. P., and K. W. Roberts, "The Theory of Household Behaviour under Rationing," *European Economic Review* 13:1 (1980), 25–42. 10.1016/0014-2921(80)90045-8
- Office for National Statistics, *Revisions Policy for Consumer Price Inflation Statistics* (Newport, UK: Office of National Statistics, 2017).
- Pakes, A., "A Reconsideration of Hedonic Price Indexes with an Application to PC's," *American Economic Review* 93:5 (2003), 1578–1596. 10.1257/000282803322655455
- , "Hedonics and the Consumer Price Index," *Annales d'Économie et de Statistique* 79–80 (2005), 729–749.
- Redding, S. J., and D. E. Weinstein, "Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences," *Quarterly Journal of Economics* 135:1 (2019), 503–560. 10.1093/qje/qjz031
- Sato, K., "The Ideal Log-Change Index Number," this REVIEW 58:2 (1976), 223–228.
- Varian, H., "Measurement Challenges in High Tech: Silicon and Statistics," presented at ESCoE/ONS/BoE Conference on Economic Measurement, Bank of England, London (2018).
- Vartia, Y. O., "Ideal Log-Change Index Numbers," *Scandinavian Journal of Statistics* 3:3 (1976), 121–126.
- Wells, J., and A. Restieaux, "Review of Hedonic Quality Adjustment in U.K. Consumer Price Statistics and Internationally," Office for National Statistics, UK., mimeograph (2013).