

BLS WORKING PAPERS



U.S. Department of Labor
U.S. Bureau of Labor Statistics
Office of Prices and Living Conditions

A New Vehicles Transaction Price Index: Offsetting the Effects of Price Discrimination and Product Cycle Bias with a Year-Over-Year Index

Brendan Williams, U.S. Bureau of Labor Statistics
Erick Sager, U.S. Bureau of Labor Statistics

Working Paper 514
May 2019

A New Vehicles Transaction Price Index: Offsetting the Effects of Price Discrimination and Product Cycle Bias with a Year-Over-Year Index^{*}

Brendan Williams[†] and Erick Sager[‡]

May 2019

Abstract

Using a new transaction level dataset on new vehicle purchases from J.D. Power, this paper documents violations of assumptions underlying cost-of-living theory and constructs a new price index that address these violations. First, we show that there are significant differences in consumer characteristics at different points in a typical product cycle, which suggests profit maximizing behavior through intertemporal price discrimination. When the composition of consumers changes each month, observed prices do not represent the purchases of a typical consumer at each point in time. Consumer segmentation violates assumptions necessary for measuring cost-of-living and can introduce bias into matched model price indexes. Second, we construct a price index that uses product age to control for consumer heterogeneity at different points of a product's life cycle. Specifically, we compare the price of a vehicle at a particular point in its life cycle to its previous-year's model at the same point in its life cycle, which we implement by assuming stability in the pattern of model year introduction and comparing each vehicle to its model year equivalent as of 12 months prior. Finally, this year-over-year price index does not convey information about short-run fluctuations in new vehicle prices, so we reincorporate high-frequency price change into the index by using a novel filtering technique. This new index averages 0.24 percentage points lower growth than the current CPI for new vehicles on an annual basis.

JEL Codes: C43, E31, L62

Keywords: price measurement, quality change, product life cycle, intertemporal price discrimination.

^{*}The authors thank Alberto Cavallo, Erwin Diewert, John Greenlees, Bart Hobijn, Emi Nakamura, Leonard Nakamura, Mick Silver, Marshall Reinsdorf, and participants of the Conference on Review of Income and Wealth (CRIW) at the 2018 NBER Summer Institute, the 2018 Meeting of the Group of Experts on Consumer Price Indexes in Geneva, the November 2017 meeting of the Bureau of Labor Statistics Technical Advisory Committee (BLSTAC), and the 2016 Annual Conference of the Society of Government Economists, for insightful comments and discussions. The authors also thank staff at the Bureau of Labor Statistics for their valuable contributions to earlier iterations of this research, especially Ralph Bradley, Austin Frazer, Emily Hunter, Tod Reese, Leonard Sabetti, Bill Thompson, and Randy Verbrugge, and thank Rob Cage, Thesia Garner, David Friedman, and Crystal Konny for innumerable discussions and comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Board of Governors, the Federal Reserve System, the Bureau of Labor Statistics, or the US Department of Labor.

[†] Corresponding author: Williams.brendan@bls.gov.

Bureau of Labor Statistics, Division of Consumer Prices and Price Indexes.

[‡]Erick.R.Sager@frb.gov. Federal Reserve Board of Governors, Division of Research and Statistics.

This research was performed while Erick Sager was a Research Economist in the Division of Price and Index Number Research of the Bureau of Labor Statistics.

1. Introduction

This paper presents research behind the first index based entirely on transaction data proposed for incorporation into the US Consumer Price Index. Transaction data provide real-time weighting information necessary to calculate a superlative index, which approximates a cost-of-living index, the conceptual objective of the CPI.¹ Product life cycles, the patterns of behavior from the introduction to the exit of a version of a product, complicate the measurement of price change. Previous research has shown a salient feature of the new vehicle product life cycle is the tendency for prices to decline during the model year. Aizcorbe, Bridgman, and Nalewaik (2010) suggest that intertemporal price discrimination (price skimming) might explain this pattern. A variety of other explanations for the product cycle's pricing behavior have been proposed going back at least to Pashigian, Bowen, and Gould (1995), who attributed the pricing pattern to fashion effects. Aizcorbe et al. (2009) conclude with the suggestion that year-over-year (YoY) price indexes could be used to hold the characteristics of buyers constant. We analyze transaction-level data from J.D. Power and confirm one of the regularities seen in prior research: the tendency for substantial declines in prices over the course of a model year. We reject chain drift as the source of this movement. We discuss implications of intertemporal price discrimination in terms of cost-of-living theory and the difficulties it creates for constructing a "representative consumer." The current CPI uses item replacement to offset within-model-year price change to mitigate these effects, but we show this method does not work well with dynamically weighted data. The impact of product life cycles and item replacement methodology have not been thoroughly explored in the price index literature. We consider a YoY index to achieve a measure of long-run price change that avoids complications from the product cycle regardless of whether the underlying source of these patterns is intertemporal price discrimination or other effects. We show that ignoring item replacements and showing product life cycle effects implies an annual bias of 2.81% in the CPI for new vehicles and 4.84% in the J.D. Power-based index. These rates are multiples larger in magnitude than estimates of cumulative bias in the CPI, suggesting that these elements of index methodology deserve serious attention.² Our final index shows a nearly 6% increase since December 2007 compared to an almost 8% increase in the CPI for new vehicles, a difference in the annual rate of about 0.24%.

This low-frequency measurement from YoY price indexes comes at the expense of information on the short-run behavior of the economy. Thus, we use time-series filtering techniques to extract the high-frequency component from a monthly price index. This monthly index is unsuitable for measuring long-run price change because the application of an item replacement procedure introduces a bias. We construct a hybrid index with a trend based on the YoY index and a high-frequency component extracted from the monthly index reflecting short-run changes in the industry while also maintaining an accurate long-run measurement of price change. The high-frequency component of the monthly index based on transaction data is correlated with the same detrended component from the CPI for new vehicles. The close relationship between extracted components based on different data sources suggests that the detrended patterns contain information about the same data generating process. Our method of hybrid index construction has applications in other areas where price index construction may require trade-offs between long-run accuracy and short-term detail.

The remainder of the paper is organized as follows. Section 2 provides a literature review that contextualizes our index within the larger body of work on transaction data, price indexes, and quality

¹ See, for example, "Chapter 17. The Consumer Price Index," p 2-3.

² For example, the Boskin Commission in 1996 estimated a total bias of 1.10%, and Brent Moulton (2018) estimates 0.85%.

adjustments. Section 3 describes the J.D. Power data and documents observations on the behavior of vehicle prices and quantities. Section 4 discusses the standard formulation of price indexes and how the J.D. Power data are seemingly inconsistent with standard theory. Section 5 discusses how a year-over-year price index can resolve these deviations from standard theory. Section 6 develops the methodology of our proposed a year-over-year price index, while Section 7 details the behavior of the resulting price index. Section 8 develops an index that represents price change across different frequencies by combining the high-frequency cyclical component of a monthly index with a year-over-year trend. Section 9 discusses specific biases that arise in our monthly index. Section 10 concludes.

2. Related Literature

This paper is related to an established literature on price dynamics for durable goods, with a subset of papers specifically focusing on the market for new vehicles – see Corrado, Dunn, and Otoo (2006), and Aizcorbe, Bridgman, and Nalewaik (2010), and Copeland, Dunn, and Hall (2011). Several empirical observations arise from previous research, most notably that vehicle prices exhibit strong and persistent within-model-year price declines. A second empirical finding is that vehicle sales are hump-shaped over the duration of the model year, with peak sales occurring around the middle of the year. This paper also arrives at the same conclusion about price declines and hump-shaped sales.

This paper’s confirmation of existing empirical observations helps validate the informativeness of previously available vehicle price data. Several previous papers have used data from J.D. Power, as we do in this paper. However, these papers use aggregated transaction data from J.D. Power’s PIN Explorer lookup tool. While our data comes from the same source, our dataset contains the transaction-level data that underlies the aggregated data in PIN Explorer, as well as details of every transaction including finance terms and trade-in information.

The literature also provides various explanations for the observed model year price declines. Aizcorbe et al. (2010) match self-reported consumer demographics and income from NOPWorld to the sales data from J.D. Power and find that high-income consumers purchase vehicles early in the model year while low-income consumers tend to delay their purchases well into the model year. Aizcorbe et al. (2010) interpret this finding as evidence of *intertemporal price discrimination*, positing that low-income consumers may be more sensitive to prices and dealerships can maximize profits by incrementally lowering prices to attract lower income consumers over time.

Corrado, Dunn, and Otoo (2006) focus on the role of rebates and interest rate subvention (concessionary finance) and construct an approximation to the actual transaction price that consumers face after negotiated adjustments to the list price. Corrado et al. (2006) attribute model year price declines to a “seasonal” pattern driven by *obsolescence*, or a *loss of newness*, and note that these drops do not reflect persistent changes in the “actual price of new vehicles.” They suggest that these early purchases are by “fashion-oriented” or “price-inelastic” consumers might explain within-model-year price declines, but they find no evidence that there is a declining pattern in consumer age, their proxy for income.

Finally, Copeland, Dunn, and Hall (2011) focus on *inventory patterns’ role* in vehicle pricing dynamics by using data from Ward Communications on American manufacturers’ vehicle inventories. From the lens of a model of vehicle price setting, Copeland et al. (2011) find that dealers and manufacturers intentionally keep prices high at the beginning of the year to accumulate inventory. Dealers maintain a high inventory to sales ratio in order to offer customers variety and maintain high prices at the beginning of a model year rather than lose stock. Only in the later part of the model year do demand shifts—such as intertemporal price discrimination and fashion effects—drive downward price movement.

For supermarket products, Melser and Syed (2016) also study pricing along a product's life cycle. Unlike this paper, which attempts to create a real-time price index, Melser and Syed (2016) define a product life cycle as not only product age but also time until product exit. After estimating the effect of the product life cycle on price, they find that the net aging effect is negative (older products that are closer to exiting tend to have lower prices). The authors then construct a monthly price index that adjusts prices for quality change at item replacements and does so in a way that does not attribute life cycle price variation to quality. This paper takes a different approach, by creating a year-over-year price index that effectively prices all vehicles in units of their age 1 price (or put differently, we control for product age by comparing two versions of an item at the same point in each product's life cycle).

This paper finds additional evidence supporting consumer heterogeneity consistent with intertemporal price discrimination. In particular, we use data on consumer financing to show that consumers rely more heavily on external financing for purchases made at the end of the model year.

These proposed explanations for model year price declines – such as price discrimination, inventory management, and product obsolescence – are difficult to separately identify from the data and likely operate simultaneously on price dynamics. Furthermore, these explanations should not be reflected as cost-of-living declines over the long term since they each represent a violation of the theory behind a cost-of-living index. In constructing a price index, this paper is more agnostic about what drives model year price declines and instead takes the stance that we can control for these various explanations through their correlation with product age. Accordingly, we use twelve-month (year-over-year) price changes to construct our proposed price index, as year-over-year price changes provide a measure of long-run price change that does not require directly estimating effects of price discrimination, inventory control, fashion effects, or obsolescence of a model year over time.

With respect to price index construction, none of this prior research has produced a long-term, monthly price index series from transaction data as this paper does. The closest instance is Corrado et al. (2006) who construct long-run measures of price change such as the year-over-year price change from one December to the next between two model years, but they do not create a monthly price index.

Finally, this paper is also related to a literature on quality adjustments applied to price change after products are substituted within the CPI. Bils (2009) shows that durable products in the CPI tend to exhibit trend price declines before exiting the market, at which time the BLS identifies a replacement item for inclusion in the CPI. Bils (2009) uses price changes across the two versions of the item (pre and post replacement) to estimate that two-thirds of the price change at product replacement is due to increased product quality, which implies that the CPI misses quality growth and overstates inflation. This paper uses Bils (2009) results by comparing prices across model years in order to control for reasons for price change that are related to product age. Using our dataset to improve quality adjustment construction at the date of model year changeover is the subject of ongoing research.

3. Data

The BLS purchased transaction-level data from the J.D. Power Group.³ J.D. Power receives vehicle transaction records from participating dealerships across the United States. These records provide the raw information for J.D. Power's aggregated "Power Information Network" (PIN) analytics. The data

³ Disclaimer from J.D. Power: "The information supplied by Power Information Network, a business division of J.D. Power and Associates ("PIN") is based on data believed to be reliable but is neither all-inclusive nor guaranteed by PIN. Without limiting the generality of the foregoing, specific data points may vary considerably from other information sources. Any opinions expressed herein reflect the judgment of the authors at this date and are subject to change."

used in this paper covers the time period between January 2007 and March 2015 and represent about a third of the new vehicle purchases in the United States. Data coverage improves over time with increasing dealership participation. Each observation in the dataset includes a transaction price as well as a set of 40 variables showing the value of manufacturer rebates, vehicle characteristics, information on finance terms, and, sometimes, the cost of the vehicle to the dealer.

Most of the characteristic information in the data are indicated and derived from the “squishVIN.” The squishVIN is a shortened version of the vehicle identification number (VIN) commonly used for vehicle registration. VIN numbers are based on an international standard, ISO 3779, that allows certain features of a vehicle, such as its make, model, and origin, to be determined from standardized usage of certain digits. We only receive a squishVIN, the shortened portion of the VIN that identifies features and omits the serial number identification component of the VIN and the ninth “check digit.” The fourth through eighth digits identify vehicle attributes, but the specific meaning of these digits varies by manufacturer. For example, one manufacturer might use the seventh digit to denote drive train type while another might use the fourth digit.

Generally, the squishVIN identifies vehicles down to the trim level and usually provides some additional information on major equipment packages. The specificity of the squishVIN varies by manufacturer. Some manufacturers use squishVINs that seem to denote a specific vehicle model and equipment configuration and others use a single squishVIN to describe a set of vehicles with different packages and transmissions.

When it collects data for the Consumer Price Index, the BLS utilizes statistical methods to draw observations with the goal of creating a representative sample for our survey. In contrast, J.D. Power’s objective is not to select a representative sample of the market, but to capture as many sales as possible. Despite this difference between the BLS survey and J.D. Power’s convenience sample, internally conducted comparisons of J.D. Power and CPI data show similar market shares by brand and a similar geographic distribution of vehicles sales. Publicly available information indicates there has been a shift to trucks away from cars in the past several years.⁴ As of December 2017, the CPI’s sample was evenly split between cars and trucks with a slight majority of trucks. J.D. Power data were much closer to other publicly available sources that indicate trucks outsold cars two-to-one (in terms of the number of vehicles, not expenditures).⁵ CPI’s breakdown may be due to a lag in sample rotation that reflects this market shift. The J.D. Power data may not be designed as a representative sample, but the data appear to provide a similar level of representativeness when compared to the CPI while giving a more timely picture of the current state of the market.

In the CPI survey, prices are observed at the dealer level and price relatives reflect price change within a dealership. In J.D. Power, dealerships are not identified and our analysis reflects consumer activity across dealerships within a geographic area. The methodology currently used by the BLS relies, in part, on outlet level weights. The J.D. Power data do not tell us the specific point-of-sale for a transaction, so the BLS’s exact methodology cannot be applied both in terms of using the point-of-sale as in the weighting scheme and with regard to defining a product in terms of the location where it is purchased.


⁴ “What’s Moving: US Auto Sales.” *Wall Street Journal*. 5 December 2017.
http://www.wsj.com/mdc/public/page/2_3022-autosales.html.

⁵ “Sales and Share of Total Market By Manufacturer.” *Wall Street Journal*. 5 December 2017.
http://www.wsj.com/mdc/public/page/2_3022-autosales.html.

3.1 Empirical Observations


Our analysis of the J.D. Power data reproduced the fundamental patterns seen in previous research on vehicle prices. Over the course of a model year, we observe that a typical vehicle's price declines steadily while its expenditures rise at first and then decline after the vehicle has been on the market about seven months. The within-model-year price behavior leads to problems in measuring long-run price change.

3.1.1 Expenditures: We observe a humped pattern in the sales of vehicles over the course of a model year, seen in Figure 1. Upon being introduced to the market, a typical vehicle's sales increase until reaching a peak about seven months after introduction. Sales steadily decrease thereafter with few vehicle sales occurring 24 months or more after a vehicle's first sale in a geographic area. Copeland, Dunn, and Hall (2011) find a similar pattern in their analysis of older J.D. Power data. They attribute the pattern to an initial production ramp up and dealership inventory accumulation process followed by a managed end-of-life strategy for late model vehicles that remain in inventory as new model year vehicles arrive. This pattern contrasts with the static weighting of the CPI, and, as further discussed in later sections, leads to complications when trying to offset within-model-year price decline with cross model year price change. Details of creating such a month-to-month index are covered in Section 8.1.



3.1.2 Prices: A price index can be constructed using J.D. Power’s data by simply plugging prices and quantities (or expenditure shares) into a price index formula; however, this would ignore the price dynamics at the entry and exit points of a model’s life cycle. We have found that the life cycle dynamics of a vehicle’s price show steady downward movement within a model year. Prices have a strong tendency to decline from the point when a model enters the market to the point where it exits, consistent with previous analysis discussed in the background literature. When a price index constructed on this data only shows matched-model price change (and never a change between model years), the index will reflect a decline.

Our analysis of the data finds that within-model-year price change tends to be strongly negative. Figure 2 plots average price change for model years between 2009-2015, using Tornqvist indexes to measure price change within a given model year. Indexes are normalized to 100 at the beginning of the model year. Figure 2 shows the first 36 months of matched-model price change for all models in a given model year. A few models may have sales after 36 months, but, at this point, three years into a model year, extremely few transactions occur. Our results are very similar to Figure 5 in Corrado et al. (2006) which shows price changes for model years 1990-2004.




* First 36 months of model year sales (Tornqvist)

Over the course of the lifespan of a model year, vehicle prices generally decrease more than 10% over a three-year span. When a simple matched-model index is constructed where matches are only made

within the model year (i.e. there are no “item replacements” where the relative would reflect price change between different model year iterations of the same vehicle), the index drops precipitously. Figure 3 shows that a matched-model Tornqvist index drops almost 30% over the entire course of the dataset.

The results of the matched-model index seem suspect given a clear increase in nominal new vehicle prices over the same time frame. The Dutot price index in Figure 3 below represents the percentage change in average prices in the J.D. Power dataset. The Dutot price index shows prices increasing more than 15% over the same time frame in which the matched-model index drops 30%. The Dutot price index reflects the change in average price between the vehicles sold in time t and $t-1$ according to the following index formula:

$$P_{Dutot,t} = \frac{\frac{1}{N_t} \sum_j^{N_t} Price_{j,t}}{\frac{1}{N_{t-1}} \sum_i^{N_{t-1}} Price_{i,t-1}}$$




The divergence between the Dutot index and strict matched-model index could reflect changing quality in vehicles. Consumers might be paying more in nominal terms, but getting more, in terms of shifting purchases to higher end cars or getting improved cars across the board and paying more for the additional features, but other factors could be at play. Intertemporal price discrimination could lead to continuously declining price relatives. Matched-model indexes in the presence of intertemporal price discrimination are susceptible to downward bias since indexes show the drop in price for a given model over time as prices are cut to attract more price sensitive consumers and never show upward price

movement when a new model is introduced and sold to high willingness-to-pay consumers. More generally, this matched-model index only reflects price change as a specific version gets older. Any aging effects will be comingled with actual price change, and there is no chance for a comparison of an end-of-life vehicle to a new one to offset within-model-year declines.


Given a pure, matched-model approach, the index will not have the chance to show any price change from one iteration of a model to the next. Price increases generally only appear in comparisons across model years. In order to show price increases from one model year to the next, the matched-model approach must be adapted to show price change from an old model year to its successor. In the current CPI, an item replacement or “changeover” process is used in a monthly index to incorporate this price change across model years. As discussed in Section 5.1, this process also offsets product cycle effects. For our index based on J.D. Power data, we instead use a year-over-year measure discussed in Section 5.2 to address these issues.

3.1.3 Consumer Heterogeneity: We find evidence of consumer heterogeneity that would be consistent with intertemporal price discrimination influencing the price dynamics of new vehicles. This evidence supports the more direct evidence of intertemporal price discrimination found in Aizcorbe et al. (2010). Unlike Aizcorbe et al. (2010), we have no measure of consumer income, but the J.D. Power transaction data include certain details on how an individual purchase was paid for and financed. We analyze the changes over the model year in several financing variables: the portion of cash versus financed transactions, the portion of the vehicle cost that is financed, and the APR of the loan. These are shown in Figure 4, Figure 5 and Figure 6 below, graphed against the “age of squishVIN,” which is the number of months since a squishVIN first appeared. We observe that the ratio of vehicle transactions that are paid for in cash rather than financed decreases through the model year, and, of those transactions that are financed, the portion of the transaction financed (generally, the amount of a financed deal not covered by a down payment) tends to increase. APR rates have a more complicated relationship with vehicle age. Percentiles for APR diverge over a vehicle lifespan as the most qualified buyers receive 0% financing incentives for older vehicles while, at the same time, lower-credit consumers appear to enter the market later in the model year, which drives up the higher level APR percentiles.

The J.D. Power data show that buyers of newly debuted vehicles are more likely to pay in cash. These buyers may actually be obtaining financing from a non-dealer source, but show up in the data records as “cash” purchasers. Whether or not a customer brings third-party financing or actually pays cash straight from their bank account, a cash sale is indicative of a customer with greater access to financial resources who would be expected to pay more under an intertemporal price discrimination regime. Analyzing data from January 2009 to March 2015, Figure 4 shows that cash transactions accounted for 43% of expenditures during the debut month of a vehicle. The portion of expenditures represented by cash transactions falls as a vehicle gets older and reaches a nadir of 23% at 17 months. At this point in a vehicle’s lifespan, the vast majority of expenditures have occurred and further sales are sporadic. The red line in Figure 4 depicts the cumulative portion of cash transactions, which is defined as the portion of cash transaction up to the month in question, and shows that it stays around 29%.




In cases where customers take dealer financing, J.D. Power offers data indicating how much of a transaction was covered with financing (the amount of a financed transaction not covered by the down payment). Early purchasers tend to put more money down, so less of the remainder of their expenditure is financed. As seen in Figure 5, in the first sales month of a vehicle, about 89% of the vehicle price is financed. As a vehicle ages, customers tend to borrow more or put less money down relative to the purchase price. A peak is reached 18 months into a vehicle’s life cycle when 97% of the transaction is financed.



Product age has a more complicated relationship with APR than it does with the other variables analyzed here. The central tendency for APR remains mostly flat at 4% over the time period examined (January 2009 to March 2015); however, APR percentiles reveal a more complex story. As seen in Figure 6, they show a divergence in APR rates as customers become more heterogeneous through the model year. Consumers with good credit ratings get lower interest rates. A few months after introduction, those who qualify for concessionary interest rates often get zero percent financing, an effective price cut. The 5th and 10th percentiles for APR are zero for much of the model year. At the upper end of the APR distribution (above the 75th percentile), interest rates increase steadily as the model gets older.

To summarize, measures of consumer finances demonstrate that consumers who purchase vehicles at the beginning of the model year differ from consumers who purchase at the end of the model year. The empirical pattern shows that consumers with weaker financial positions—who need auto loans and who are willing to accept a high APR—delay their purchases until later in the model year when prices are lowest on average. This observation is consistent with the notion that intertemporal price discrimination on the part of sellers generates substantial heterogeneity with high-income consumers buying at the beginning of the model year and less affluent consumers buying near the end.



4. Product Cycle Theory

A variety of factors may explain why we see persistent price declines like those in the strict matched-model index discussed in the previous section. From a certain point of view, this index could reflect a real decline in consumers’ cost-of-living. The within-model-year price change show actual declines in vehicle prices, and the tendency for new model year vehicle prices to exceed those of their predecessors could indicate a price premium on the features and quality of newer vehicles. Given the large difference between the behavior of vehicle average prices and price relatives, we do not think the strict matched-model index reflects the true change in cost-of-living. These declines could be an artifact of price index construction known as “chain drift.” Other explanations for these apparent price declines include intertemporal price discrimination and other factors related to the product cycle.

4.1 Chain drift

Chained price indexes often experience “drift” in one direction. As usually defined, chain drift arises from frequent price oscillations or “bouncing” (see Ivancic (2011)). The vehicle prices we see here actually tend to be relatively stable compared to the extreme price changes often observed in retail scanner data. The drift appears to arise from consistent downward price change within-model-year rather than oscillating prices. Greenlees and McClelland (2010) observed similar behavior in apparel prices and found that “the relentless downward march of prices completely overwhelm the chain drift issue” in their data.


We find further support in dismissing conventional chain drift in the results of the RYGEKS index (Rolling year GEKS, index named for Gini, Eltetö, Köves, and Szulc). Ivancic, Diewert, and Fox (2011) proposed the RYGEKS index as a means of reducing chain drift. The pure GEKS index eliminates chain drift associated with non-transitivity. The RYGEKS is a more practical adaptation of the GEKS. The GEKS is recalculated completely as new data are introduced over time. The RYGEKS starts with a base GEKS calculated over one year and then updated with a “chain link.” The RYGEKS is not fully transitive, but it reduces chain drift from non-transitivity.

We calculate multilateral indexes for J.D. Power data aggregated to the national level.⁶ The RYGEKS formula starts with a thirteen-month GEKS index:

$$GEKS_{0,12} = \prod_{l=0}^{12} \left(\frac{P_{0,l}}{P_{12,l}} \right)^{1/13}$$

where $P_{t,l}$ is the bilateral index between time period t and link month l . For the thirteenth month, the GEKS is updated with a mean splice “chain link”:

$$RYGEKS_{0,13} = GEKS_{0,12} * \prod_{t=1}^{13} \left(\frac{P_{12,t}}{P_{13,t}} \right)^{1/13}$$



⁶ To calculate multilateral indexes, we use the IndexNumR R package from Graham White (2018), R package version 0.1.0. The package can be downloaded from <https://CRAN.R-project.org/package=IndexNumR>.

In Figure 7, we find little difference between a chained Tornqvist and RYGEKS matched-model indexes. All of the matched-model indexes drop precipitously over time. Extending the length of the rolling window from thirteen to eighteen months makes little difference on the index, while extending the window to twenty-four months (at which point there are few matched-models) leads to a slightly faster rate of decline. Differences due to splicing method (using window or movement instead of mean splicing) are also negligible. Were chain drift the cause of the drop, we would expect multilateral methods to attenuate the issue. Problems with these indexes seem to stem from the fact that the monthly price relatives for matched-model goods have a strong tendency to be less than one and not from the oscillation issue addressed by the RYGEKS. Chain drift, as conventionally defined, does not appear to explain the downward price tendencies, nor does this appear to be more generally an index formula issue.⁷

4.2 Intertemporal price discrimination

Stokey (1979) identifies several characteristics of markets that could lend them intertemporal price discrimination. Among other cases, she discusses markets where customers lack pricing information or where the seller has limited capacity, both of which could characterize the new vehicle market. Given the extensive negotiation process generally involved in purchasing a car in the U.S., the new vehicle market lacks transparent pricing. Moreover, secondary sellers have little ability to arbitrage away price differentials since vehicles famously lose much of their value once they become “used”—in part because of information asymmetries in the used vehicle market (Akerloff (1970)). Limited dealer inventory capacity plays an important role in the model developed by Copeland, et al. (2011). As discussed in Section 2 and Section 3, Aizcorbe, et al. (2010) found evidence of intertemporal price discrimination that is supported by our own findings of consumer heterogeneity. Intertemporal price discrimination presents a problem for price index calculation from a cost-of-living perspective, and, more specifically, when ignored, leads to an apparent violation of the assumption of homothetic preferences and the ability to characterize market behavior with a single, representative consumer (see Shaprio and Aizcorbe (2010)).

Cost-of-living theory relies on an assumption of stable preferences (i.e. the same utility function persists over time) to justify making statements regarding the cost of obtaining a certain utility level based on observed expenditures and prices. While this assumption is often violated by constantly shifting preferences in the market, intertemporal price discrimination leads to systematic shifts that bias indexes downward. Intertemporal price discrimination can bias price indexes.

⁷ Researchers in this field often focus on addressing problems by getting the “right” index formula, but the problem of offsetting within-model-year price change seems to necessitate a different approach to relative construction.

When price $\frac{P_{i,t}}{P_{i,t-1}}$ relatives overwhelmingly show price declines (having values less than 1), these relatives cannot be weighted to show a price increase (an index relative greater than 1). Consider a matched-model index where a specific model is denoted by the subscript i , price relatives are constructed as $\frac{P_{i,t}}{P_{i,t-1}}$ and weights for the relatives are designed as $w_{i,t,t-1}$ (this formula can encompass several common index formulas depending on how $w_{i,t,t-1}$ is constructed):

$$P_{Index,t} = \prod_{i=1}^n \left(\frac{P_{i,t}}{P_{i,t-1}} \right)^{w_{i,t,t-1}}$$

In order for predominately declining relatives to produce an increasing index, $P_{Index,t}$ negative weights would have to be used, an impossibility in conventional price indexes.

Konüs laid out the original cost-of-living approach to index theory. He defined a cost-of-living index ($P_{Konüs}$) as the ratio of the minimum cost of attaining a certain level of utility, u , between two time periods given time-specific price vectors, p^t :⁸

$$P_{Konüs}(u, p^0, p^1) = \frac{C(u, p^1)}{C(u, p^0)}$$

where $C(u, p)$ is the function describing the minimum expenditures required to attain utility level u given prices p .

An index of this form is a cost-of-living index regardless of the choice of a reference utility level of u as long as u is constant between the two time periods; otherwise, the resulting index reflects both a change in prices and a change in preferences. Accordingly, a threat to the construction of a cost-of-living index emerges when sellers engage in intertemporal price discrimination. These sellers exploit differences in consumer utility between time periods, which creates a difference in consumer utility, u , between the time periods. For example, consumers may have preferences both between competing vehicle choices (of which they only plan on buying one) and the amount of time they are willing to wait to make the purchase. In this case, the new vehicle seller may list high prices early in the product cycle to take advantage of impatient consumers' desire for a vehicle sooner and may list lower prices later in the product cycle to encourage the most patient consumers to purchase the vehicle before the new model year is released.

Utility may change over time as a result of intertemporal price discrimination based on heterogeneous preferences, in which case changes in u result from differing utility function, or due to differences in consumer budgets, in which case the level of u differs. Either way, the resulting relative is not necessarily a Konüs index since there is no reference level of utility:

$$\frac{C(u_1^1, p^1)}{C(u_0^0, p^0)}$$

Discrimination based on income would be less of a concern if the consumer preferences are homothetic, meaning that a consumer's preferences over how much of each product to purchase scale linearly with the consumer's income. This would allow us to use either time period as a reference utility.⁹ The homothetic preferences assumption is also important for extending cost-of-living theory to a real-world use through aggregation by allowing us to create a fictitious representative consumer. Given that cost-of-living theory assumes that all consumers purchase products in the same proportions, a single consumer (our representative agent) who possesses all of an economy's income also purchases products in these same proportions. In this sense, the single consumer represents the preferences of all agents.

The segmentation of purchases by consumer type into different time periods leads to a violation of homothetic preferences, even in cases where the underlying utility function may be homothetic. Given two customer types that vary by income and, correspondingly, willingness to pay for a good x , customer a differs from b by having a higher reserve price, therefore $p_a > p_b$. Given perfect intertemporal price discrimination, customer a will consume x in the first period (t_1), and customer b will consume x at t_2 . If we try to assume that each period is a snapshot representing the same utility function, the utility function cannot be homothetic. Homothetic utility implies that consumer demands will be proportional to income.

⁸ See the discussion in "The Economic Approach to Index Number Theory: The Single-Household Case." Chapter 17. *Consumer Price Index Manual: Theory and Practice*. 2004.

⁹ For a discussion, see the ILO Handbook (2004), pp 316-317.

When b demands x in t_2 and a has no demand for x , despite a having higher income, the utility function implied by only looking at t_2 is not homothetic. No single time period is representative of aggregate consumer behavior because of time-based market segmentation. When markets can profit from differences in consumer's characteristics, consumer heterogeneity better characterizes the reality of consumers' market behavior than a representative consumer can.

The implications of intertemporal price discrimination for measures of price change have not been well explored in the price index literature. Car dealers using intertemporal price discrimination exploit differences in consumer preferences by selling to high willingness-to-pay consumers at the beginning of the model year and low willingness-to-pay consumers at the end of a model year's life cycle. Resulting price changes reflect both any change in cost-of-living and changes in customer preferences. Intertemporal price discrimination always leads to price declines, a possible explanation for the behavior of within-model-year price.

4.3 Product cycle bias

We consider intertemporal price discrimination along with other possible explanations for the new vehicle price cycle behavior generally as "age bias." In the presence of age bias, the nominal price of an item will reflect both change in price due to the aging of a vehicle and any change in the cost-of-living. Let P_{ijat} be the price of a product i , with an age a , at time t . The price can be decomposed into a time fixed effect θ_t , a product fixed effect ξ_i , an age-characteristic fixed effect κ_{aj} , and a residual \tilde{p}_{ijat} :

$$p_{iat} = \theta_t + \xi_i + \kappa_{ai} + \tilde{p}_{ijat}$$

The change in price from the period prior to t can then be expressed in terms of changes to these underlying components with the product fixed effect cancelling out since it is unchanged by definition:

$$\Delta p_{iat} = \Delta \theta_t + \Delta \kappa_{ai} + \Delta \tilde{p}_{ijat}$$

For our price index, we are interested in capturing the price change due to θ_t , which means controlling or adjusting for changes in age. If the change in age leads to a change in price, a constructed price relative will be biased. This bias occurs at the formation of the first price relative, so bias is not caused by chaining. Although we do not explicitly implement panel regression techniques in constructing a price index, the principle behind controlling for age effect remains.

5. Addressing Product Cycle Behavior

First, we discuss how the current CPI uses "item replacement" to create cross model year price comparisons that tend to offset the within-model-year price declines that characterize the product cycle for new vehicles. However, this method relies on a fixed-weight index formula that is not compatible with the dynamic weighting used in a superlative index, which would be recommended under a cost-of-living approach. We then explore using a year-over-year index to create price comparisons that are robust in the presence of intertemporal price discrimination and other age-related factors. The year-over-year methodology is applied to the J.D. Power data in later sections.

5.1 Item replacement

In order to show price change from one iteration of a model to the next in a matched-model price index, a price relative has to be constructed that shows the change from the old to the new model. In the CPI's sampled survey, item replacement is a built-in part of the process to maintain a sample of a certain size as model years come and go from the market, but it also serves to capture price change related to the introduction and exit of goods. Omitting this cross-version price change can introduce a bias into the index. Generally, this bias is attributed to staggered price setting where sellers wait for a new product


version to introduce price change.¹⁰ However, product replacement can also be used to offset aging and other product life cycle effects. For decades, the BLS has seen a general pattern where new vehicle introductions are associated with upward price change relative to flat or decreasing same-model-year price change. Moulton and Moses (1996) found that item replacements had a large, upward impact on the CPI—accounting for 1.09 percent of the 2.16 percent price increase in 1995 for the items they studied. Recognizing the implications of the replacement relative being more inflationary than same version price change, class-mean imputation was introduced for new vehicles in October 1989 so that discontinued vehicles without comparable replacements would be imputed by price change from other replaced vehicles (see Reinsdorf, Liegey, and Stewart (1996)).

In the current CPI methodology for new vehicles, a “substitution” is made every year when the old model year is replaced once sales of the new exceeds those of the old. For every substitution, a “changeover” relative is constructed that shows the relative price change between the new model year version of a vehicle and its predecessor last month. The relative is adjusted for any observed quality difference between the old and new model year versions of a vehicle. If a vehicle is discontinued or undergoes a major redesign, the changeover relative is imputed based on the mean of other vehicles undergoing substitutions that month in the same or similar geographic area. As discussed in detail below, this changeover relative has the tendency to be a large positive change that offsets within-model-year price declines.

An experimental index based on CPI new vehicles data in Figure 8 shows the impact of removing the changeover process from CPI calculations. Like the matched-model index in Section 3, the resulting index (in red) reflects only within-model-year price change. From January 2007 to December 2014,¹¹ the index without changeovers decreased 15% while the official increased 6.5%. The decrease is not as substantial as the 30% drop in the matched-model Tornqvist index based on J.D. Power data, which may be in part due to the difference in index formula. The J.D. Power-based, matched-model Tornqvist index uses real-time weighting information, while the CPI indexes below follow official CPI procedures and use a statically weighted geomean formula at the lower-level. Both use a Laspeyres formula at the upper level for geographic aggregation using BLS geographic weighting information.

¹⁰ For example, see Nakamura and Steinsson (2012).

¹¹ The research index calculator used to create this index only runs until December 2014 due to changes to our estimation system implemented in January 2015.



The use of real-time weights complicates the current use of the changeover process to offset downward bias in the CPI. Consider the price, $p_{i,a,t}^m$, of a particular vehicle i at date t , where that vehicle belongs to model year m and has been on the market for a months (in other words, has age a). If we consider the component of a price index associated with a specific item with a share weight of s_i , we can see the result of the changeover process in a fixed-weight index. The price change between any two consecutive months may be biased since the comparison reflects a difference in age, but, when a comparison with the next model year iteration $m+1$ is forced in month n , the intervening months cancel and the overall impact of item i on the index reflects a price comparison between two goods of age a . As long as there is no quality change between the m and $m+1$ versions of i , the resulting relative reflects a measure of price change from month t to $t+n+1$ that is not biased by age, but remains biased if the index formula does not accurately capture changes in expenditure consistent with a cost-of-living index.

$$P_{i,t}^{Fixed} = \left(\frac{p_{i,a+1,t+1}^m}{p_{i,a,t}^m} \right)^{s_i} \times \left(\frac{p_{i,a+2,t+2}^m}{p_{i,a+1,t+1}^m} \right)^{s_i} \dots \left(\frac{p_{i,a+n,t+n}^m}{p_{i,a+n-1,t+n-1}^m} \right)^{s_i} \times \left(\frac{p_{i,a,t+n+1}^{m+1}}{p_{i,a+n,t+n}^m} \right)^{s_i} = \left(\frac{p_{i,a,t+n+1}^{m+1}}{p_{i,a,t}^m} \right)^{s_i}$$

The above example approximately characterizes what we observe in the statically weighted, geomean lower-level indexes used for the current CPI (ignoring complicating effects from sample rotation and other practical survey concerns). Once we introduce dynamic weights, denoted $w_{i,t}$ for vehicle i 's weight at date t , such as those used for Tornqvist, Fisher, and other superlative formulas, interpretation of the chained relatives is not as clear:

$$P_{i,t}^{Dynamic} = \left(\frac{p_{i,a+1,t+1}^m}{p_{i,a,t}^m} \right)^{w_{i,t}} \times \left(\frac{p_{i,a+2,t+2}^m}{p_{i,a+1,t+1}^m} \right)^{w_{i,t+1}} \dots \left(\frac{p_{i,a+n,t+n}^m}{p_{i,a+n-1,t+n-1}^m} \right)^{w_{i,t+n-1}} \times \left(\frac{p_{i,a,t+n+1}^{m+1}}{p_{i,a+n,t+n}^m} \right)^{w_{i,t+n}}$$

Prices no longer cancel-out across months and the resulting price change between ages is mixed with the changes in weights of good i relative to other goods. The “changeover” process for reflecting price change from one model year to the next does not translate well to the dynamic data we have in J.D. Power. BLS research on the changeover process, such as Reinsdorf, Liegey, and Stewart (1996), did not consider the ramifications of using this procedure in an index with dynamic weights.

Note that the bias is not the result of the typical “chain drift” bias. When age contributes bias (whether the source of this bias is ultimately intertemporal price discrimination, fashion effects, or another product cycle effect), the unchained relative $\frac{p_{i,a+1,t+1}^m}{p_{i,a,t}^m}$ already reflects bias. In the case of static weights with model year item replacement, chaining can be seen as a corrective to the age bias. In the case of dynamic weights, the impact on bias of chaining is ambiguous, and bias may be exacerbated rather than attenuated.

5.2 Year-over-year

Year-over-year (YoY) indexes have been used since the nineteenth century to ensure price comparisons reflected true price change instead of seasonal fluctuation.¹² Application of this method to new car data, even J.D. Power data, is not new. Aizcorbe, et al. (2009) and Corrado, et al. (2006), assess long-run price change with year-over-year measures. Prices (and weight information if the index formula uses both current and prior period weights) are compared in one time period (be it a week, month, quarter, etc.) to the corresponding period one year ago. The International Labor Organization CPI Manual chapter on seasonality notes that separability conditions on consumer preferences (Diewert 1996 and 1999) would justify year-over-year indexes from an index number theory perspective.¹³ In the article where Konüs (1939) introduced cost-of-living theory, he argued that year-over-year treatment of price change may be necessary to measure changes in cost-of-living since “...there is no possibility of comparing standards of living in the summer and winter months of any year since conditions of life differ as between summer and winter. It is possible to compare only the standards of living in the summer months of one year with those of the summer months of another year.”¹⁴

The new vehicle market does not have a seasonal pattern per se, but an individual vehicle model usually undergoes an annual cycle of price and expenditure patterns related to the annual introduction of new model years. Reasons for the price decline are ambiguous, but the YoY index helps maintain like-to-like comparisons even in the presence of several of the product cycle drivers discussed in the related literature. Given a relatively stable pattern of new vehicle releases every twelve months, the YoY index compares similar points in a vehicle’s life cycle. For instance, this method helps mitigate fashion effects since the price change reflects the difference in price of vehicles with similar “newness.” The price of a brand new, updated vehicle is compared to the point in the previous model year when the vehicle was at the height of its newness for that model year.

Similarly, the YoY index also helps mitigate the potential effects of intertemporal price discrimination. Month-to-month indexes may show the price change between high-willingness to pay consumers at the beginning of a product cycle and more budget-minded consumers who tend to buy

¹² See the ILO Handbook (2004), p 396.

¹³ ILO (2004), p 396.

¹⁴ Konüs (1939), 13.

models being cleared out of dealer inventory. The YoY relatives reflect price comparisons between transactions made by similar consumer types. Under the assumption that consumer preferences are consistent in the same calendar month from one year to the next, expenditures by like consumer types can be compared in terms of their cost-of-living. We assume $u^t = u^{t+12}$ so that:

$$\frac{C(u^{t+12}, p^{12})}{C(u^t, p^0)} = \frac{C(u, p^{12})}{C(u, p^0)}$$

Unlike a direct month-to-month comparison, the utility function does not change, so changes in prices can be interpreted as changes in the cost-of-living.

6. Transaction Data Index Methodology

In order to find a measurement of price change that is not contaminated by product cycle effects, we construct a year-over-year index. This index captures the long-run price trend but does not show short-run price changes including seasonal patterns and short-term shocks. We aggregate all individual transactions for a squishVIN within an area into unit prices and expenditures. Each squishVIN-based “model” is matched to its prior model year equivalent. Quality differences introduced between model years are adjusted based on the cost-based estimates of quality currently used in the CPI for new vehicles. These price relatives are then aggregated into area level indexes based on a Tornqvist index formula. A national level index is then created using the same Laspeyres-type formula and geographic weighting system currently used in the CPI.

6.1 Model definition

We define a vehicle model using a “squishVIN” that is given with every transaction in the J.D. Power data. While the squishVIN might not fully identify a homogenous item, the squishVIN is our best means of defining a “model” in terms of our matched-model requirements. Instead of trying to extract detailed characteristic information from the squishVIN, we simply allow the squishVIN to define a specific good that we track over time. J.D. Power also provides additional specification information, some of which is based on interpreting the squishVIN, but we choose not to use this in the process of defining a model since the squishVIN provides the most accurate means available to us of identifying a consistent set of features to define a constant-quality vehicle.

While the squishVIN provides a clear definition of a model within a model year, we also need to match vehicles across model years. We use a system of backward matching where we try to find a match for a vehicle given its model year in the set of vehicles in the previous model year. Finding the corresponding model in the previous model year is not always straightforward. We first search for a current period squishVIN by simply changing the eighth digit, the one that corresponds to model year, to the character denoting the previous model year. (For example, XXXXXX8X represents a 2008 model, so we would first attempt to match this to the corresponding 2007 model denoted XXXXXX7X). When this process does not succeed, the analyst manually looks at the descriptive information in the J.D. Power data to find a squishVIN in the previous model year that has comparable features, even if there is no match based on squishVIN.¹⁵

¹⁵ The 2010 model year involved a revision to the VIN standard, so all of the 2010 model year vehicles had to be manually matched to corresponding 2009 vehicles. In addition to this mass recoding based on a change in the VIN standard, individual manufacturers periodically change how VIN coding works on their vehicles, which breaks continuity in matching vehicles across model years.

6.2 Unit price

Before we apply a price index formula, we run a first stage aggregation of transactions into a single unit price. Let $p_{v,j,t}^m$ be the j -th transaction price for a new vehicle with characteristics set v (defined by squishVIN), at date t , and model year m . For any new vehicle of model year m with characteristics v at date t , there are $N_{v,t}^m$ total purchases, which are indexed by $j=1, \dots, N_{v,t}^m$. Vehicles have a countable number, V_t , of discernible vehicle configurations at any given date, that are indexed by $v=1, \dots, V_t$. Characteristics for each v include the area in which the vehicle was purchased, the vehicle's make, model, and other detailed characteristics that are recorded by a *squishVIN*. The number of varieties may change over time as new vehicles enter the market or older models exit.

To construct a price index, we compute unit prices for each vehicle $v=1, \dots, V_t$ at each date. Denote such a unit price as $P_{v,t}^m$. Unit prices are constructed as the geometric mean across individual transactions:

$$P_{v,t}^m = \prod_{j=1}^{N_{v,t}^m} (p_{v,j,t}^m)^{(1/N_{v,t}^m)}$$

for each $v=1, \dots, V_t$ and $t=1, \dots, T$.

6.3 Quality adjustment

Our YoY index uses the same cost-based adjustment that the BLS has used since the early 1960s. The BLS currently receives information on the cost of vehicle quality improvements from manufacturers. Certain vehicle manufacturers provide us with estimates of the cost for feature improvements between model year iterations of their products. We receive cost estimates for changes such as the cost of making electronic vehicle stability control standard or changing the styling of a headlight assembly. We use this information, and a markup factor, to calculate quality adjustments. Mechanically, to quality adjust prices, we first compute the unit cost of production for each new model year version m of vehicle v as a geometric mean over transactions:

$$C_{v,t}^m = \prod_{j=1}^{N_{v,t}^m} (c_{v,j,t}^m)^{(1/N_{v,t}^m)}$$

where $c_{v,j,t}^m$ is the production cost of the j -th transaction of model year m of vehicle type v at date t .

Next we compute the retail markup, denoted $\mu_{v,t}^m$ as:

$$\mu_{v,t}^m = \begin{cases} \frac{P_{v,t}^m}{C_{v,t}^m} & \text{If cost data available} \\ 1 & \text{If no cost data available} \end{cases}$$

We obtain manufacturer cost estimates for the estimated costs of improvements, ϕ_v^m , to a vehicle v in m compared to corresponding version of the vehicle in $m-1$, and multiple the cost by the markup. Lastly, we adjust the price by subtracting the cost-based quality adjustment:

$$\hat{P}_{v,t}^m = P_{v,t}^m - \mu_{v,t}^m \phi_v^m,$$

When a completely new vehicle enters, a vehicle is redesigned, or when BLS determines that we do not have information capable of adjusting for large quality improvements between two model years, the price change is treated as “non-comparable.” For the YoY index, this means the vehicle is omitted from the calculation.

6.4 Second stage aggregation

The CPI-U uses a two-stage index aggregation process. The first (elementary) stage is defined as an index for a specific item stratum in a specific geographic area. In the CPI, a geometric mean formula is used for this first stage for most categories—including new vehicles. The second stage aggregates across different items and different geographic areas. Second stage calculations use a Laspeyres-type formula.

The J.D. Power and CPI indexes discussed here use various elementary level index formulas but all use Laspeyres-type aggregation in the second stage to aggregate across areas to calculate a national level index. In the CPI, weights for these first stage indexes come from the Consumer Expenditure Survey and are used to produce biennial “aggregation weights” found by dividing expenditure estimates for an item type in an area by the respective first stage component price index. The aggregation weights at the second stage are updated biennially based on the Consumer Expenditure Survey’s estimates of the amount of money spent on new vehicles. We calculate new estimates of the aggregation weight by dividing the “costweight” found in the production by our new index estimates. Area level weights are thus based on the Consumer Expenditure Survey and not derived from J.D. Power expenditure data. More detail on the two-stage aggregation used in the CPI can be found in Chapter 17 of the BLS Handbook (2018).

6.5 Index formula

This paper focuses on a year-over-year index with a Tornqvist formula. The current new vehicle CPI uses a geometric mean formula, which is exact for a Cobb-Douglas utility function. Lower level indexes, like our area level new vehicles indexes, are commonly considered “elementary indexes,” which typically do not have weights from observed expenditures. The J.D. Power transaction dataset allows us to weight our indexes with observed expenditures instead of, as is currently the case with most CPI components, relying on the assumption of constant expenditure shares to justify a geometric mean index.¹⁶

6.6 Year-over-year index

For each area, the year-over-year (YoY) index shows the annual trend inflation from one calendar year to the next, and one model year to the next. We match each model of vehicle to a similar model in the previous model year. The YoY formula below constructs a price relative ($R_{v,t}^m$) using the price of given model in the current period ($P_{v,t}^m$) and matching it to the price of the corresponding model in the previous model year, 12 months ago ($P_{v,t-12}^{m-1}$) for the set of all model years available. Once a squishVIN match has been made, the twelfth root of the ratio of the matched prices is used as the monthly price relative:¹⁷

$$R_{v,t}^m = \left(\frac{P_{v,t}^m}{P_{v,t-12}^{m-1}} \right)^{\frac{1}{12}}$$

¹⁶ See ILO (2004), p. 212, for a discussion.

¹⁷ A similar method is currently used for the housing component of the CPI. The sixth root of a six-month relative is taken to get a monthly measure of price change. See “Chapter 17. The Consumer Price Index,” p. 21.

These price relatives are incorporated into a Tornqvist formula. We calculate a final index relative based on the geometric mean of these relatives weighted by the average of their current and prior period (as of one year ago) expenditure shares. The weight, $W_{v,t}^m$, for vehicle with characteristics v for model m at time t is constructed as follows,


$$W_{v,t}^m = \frac{1}{2} \left(\frac{I_{m,t} E_{v,t}^m}{\sum_{v=1}^{V_t} I_{m,t} E_{v,t}^m} + \frac{I_{m,t} E_{v,t-12}^{m-1}}{\sum_{v=1}^{V_{t-12}} I_{m,t} E_{v,t-12}^m} \right),$$

where $I_{m,t}$ equals 1 if the model m was available in the data in both the current month, t , and a year ago, $t-12$, and equals zero otherwise. Therefore, expenditures for vehicles are not counted when their price relatives are not used in index calculation. This excludes expenditures for vehicles without transactions in both t and $t-12$ and for vehicles that do not have matches across model years. Thus expenditure shares sum to one. The final index uses the relatives and weights for each of the matched vehicles in time t into an index for new vehicles in a given area:

$$P_t^{\text{YOY}} = \prod_{v=1}^{V_t} (R_{v,t}^m)^{W_{v,t}^m}$$

7. Results

From December 2007 to March 2015, the YoY index shows a 5.9% increase in new vehicle prices, as seen in Figure 9. This compares with an 8.0% increase in the CPI index for new vehicles over the same period. The YoY index could be expected to run lower than the CPI since it uses real-time expenditure information for weights, while the CPI has fixed weighting and over-weights price change at model year changeover. As expected, the YoY index presents a very smooth indicator of price change since higher frequency price fluctuations are omitted in YoY calculations. A method for recovering the high-frequency component of price change is discussed in Section 8.



Bootstrapped confidence intervals and standard errors for the index levels suggest the index levels based on the J.D. Power are estimated with high precision. Looking at the year ending in December 2008, the first year of the YoY index, the bootstrapped standard error for the annual percentage change is 0.012%. According to estimates from BLS’s Price Statistical Methods Division, the standard error for the annual percent change in the CPI for new vehicles for the same period was 0.390%.¹⁸ The variance in the estimator of the CPI price index is calculated using a stratified random groups method that captures the contribution to variance from sampling different geographic localities. Unlike the current CPI, the J.D. Power data are not sampled and the YoY index is calculated across all available areas. The differences in contributions to and estimates of variance differ between the published CPI for new vehicles and our transaction index. The comparison of standard errors are not directly comparable; however, given the relative size and coverage of the J.D. Power dataset, we expect a more precise estimate of price change from this transaction data compared to the current production index in the CPI.


Current lower level indexes in the CPI use a geometric mean formula, which approximates a cost-of-living index. The BLS uses the geometric mean formula for lower level indexes because it has not had access to timely expenditure information.¹⁹ The geometric mean is intended to approximate a cost-of-living index for a constant expenditure utility function. J.D. Power data allows us to use real-time expenditure information to calculate superlative indexes that approximate cost-of-living across a more

¹⁸ Methods for current production estimates of standard error are covered in Chapter 17 of the BLS Handbook pages 37-41.


¹⁹ The Boskin Commission (1996) “...lacking quantity or expenditure information at the lower level, a good approximation to the underlying cost-of-living index is obtained from a geometrically weighted average formula.”

general set of utility functions than the Cobb-Douglas function approximated by the geometric mean index.

While we propose a Tornqvist formula for the YoY index, alternative index formulas can be used to construct cost-of-living bounds for the YoY index as demonstrated in Figure 10. The chained Laspeyres index, which does not reflect substitution effects, presents an upper bound according to the cost-of-living theory set out by Könus. The chained Paasche index reflects perfect substitution and forms the lower bound for the true cost-of-living index. The geometric average of the Paasche and Laspeyres is a Fisher index. The Fisher index (represented with blue squares in the chart below) almost perfectly overlays the YoY (Tornqvist) index. Over the entire seven plus years examined, the Fisher increases 0.05% more than the Tornqvist.



In addition, the impact of quality adjustment was examined by looking at a YoY index where no quality adjustments were made. All quality adjusted relatives were treated as “comparable.” The resulting index might be thought of as an upper bound estimate of the impact of quality adjustments since the analyst would have treated many of the larger quality adjustments as non-comparable and omitted them had quality adjustment values not been available. As shown in Figure 11, over the seven-plus years examined, the non-quality adjusted index rose 1.59% more than the YoY index with quality adjustment. Overall, the cost-based quality adjustments used for our YoY index had a moderate impact on the index level—comparable in magnitude to the difference between the official CPI index for new vehicles and the quality adjusted YoY index.



8. Incorporating High-Frequency Price Change

While the Year-over-Year index helps resolve bias to trend price change, it essentially omits short-term changes in prices. In this section, we extract the short-run fluctuations from an index constructed from *monthly* price changes and combine it with the YoY trend to create an index that represents both short- and long-run price change. The strategy is similar to the CPI’s seasonal adjustment procedure – but in reverse. Instead of removing seasonal volatility through smoothing, we introduce short-run fluctuations into a smooth index.

In section 8.1, we develop an index that combines the detrended portion of a monthly price index with a corrected long-run trend. The monthly index reflects many aspects of the methodology currently used in the CPI. Notably, we extend the CPI’s item replacement (or “changeover”) process to transaction data and encounter a resulting bias in index levels. The band pass filter proposed in Christiano and Fitzgerald (2003) is used to extract the high-frequency component from this monthly index. The two components are combined into a hybrid index that provides information on both short- and long-term price change based on transaction prices and real-time expenditures from the J.D. Power dataset. This process can be viewed as using the YoY measurement as a trend correction for the bias that arises in the monthly index.

Section 8.2 discusses the resulting monthly and hybrid YoY indexes. We investigate bias in the monthly index. Smaller areas tend to show higher price trends than larger areas. Through a resampling

exercise, we confirm that this is the result of bias and not true economic phenomena. Changeovers receive a large amount of weight in small area indexes and small amounts of weight in large areas. The long-run component of the monthly index is biased by area size and does not produce a reasonable measure of price change, but the detrended component of the monthly index shows a relationship with results generated by running the same filter on the new vehicles CPI. Moreover, the detrended components of monthly indexes with different index formulas and other methodological changes also correlate with the new vehicles CPI. The relationship of high-frequency fluctuations across different data sources and different methodologies suggests that the monthly fluctuations capture actual economic activity and represent more than random noise.

8.1 Month-to-month index methodology

Our monthly index roughly follows current CPI methodology and mimics the current production process with respect to item changeover between model years, applying quality adjustments, using a geometric mean index formula, and the upper-level aggregation process. Unlike the CPI, time-varying expenditure weights and transaction prices are used.

The exact CPI methodology cannot be directly implemented on the J.D. Power vehicle transaction data. The current CPI implicitly weights the sample probabilistically selecting dealers through the Telephone Point-of-Purchase Survey (TPOPS) and, within dealerships, probabilistically selecting an individual vehicle model based on dealership sales information. The J.D. Power data only indicate the geographic area of a sale, not a specific outlet, so the outlet selection procedure cannot be replicated. We forgo any sampling process.²⁰ Additionally, the use of transaction rather than dealer estimated prices means that we often do not have observations for vehicles in consecutive months.

We mimic the “changeover” process used to replace old model year vehicles with new ones in the CPI to show price change across model years and to compensate for bias due to within-model-year price declines. The current CPI avoids chaining downward, within-model-year prices by periodically comparing prices between different model years and calculating a changeover relative.²¹

In the CPI, the construction of a changeover relative is part of the item replacement process, and is necessary, in part, because maintaining a fixed sample requires discontinued goods to be replaced with new goods. Here, we construct indexes that allow the sample to vary each month. We imitate the CPI item replacement process by matching each new version of a vehicle to its model year predecessor. When a new model year version of a vehicle is observed, we attempt to match it back to a prior period model, as discussed in Section 3. Once the expenditure in the area on the new model exceeds the expenditure on the old, we compare the price of the old to the new and create a changeover relative. The changeover relative is constructed for a vehicle v in monthly t and model year m :

$$R_{v,t}^m = \frac{p_{v,t}^m}{p_{v,t-1}^{m-1}}.$$

If we have information on the difference in manufacturing cost between the m and $m-1$ version of a vehicle, we apply the quality adjustment process described in Section 6.3 to adjust $p_{v,t}^m$. If the vehicles are thought to be comparable and we do not have cost information, no quality adjustment is made and the

²⁰ Appendix B explores indexes created using J.D. Power data with a fixed sample methodology that adheres to the CPI’s current methodology as closely as possible.

²¹ Reinsdorf, Liegey, and Stewart (1996) discuss the changeover process and the use of changeover relatives in class-mean imputation in the CPI. This process is used for many items in the CPI and described in Chapter 17 of the CPI Handbook (see CPI Handbook (2018), p. 19-20).

prices are directly compared to create a changeover relative. When a significant change to a vehicle's design has been made between model years, the changeover relative is computed based on a weighted average of the observed changeover relatives in the same geographic area. Observations from the prior model year that have been replaced are no longer used in the monthly index after $t-1$. Similarly, observations from succeeding model years are not used until the month in which their sales exceed those of their predecessors.

Price relatives for a given vehicle v , not undergoing a changeover, are constructed using the geometric average price for that vehicle in a given area, denoted by $p_{v,t}$, and the corresponding value in the prior month $p_{v,t-1}$, such that the price relative is given by,

$$R_{v,t}^m = \frac{p_{v,t}^m}{p_{v,t-1}^m}.$$

Weights for the monthly index are constructed from current period expenditures. The denominator used in calculating expenditure shares only reflects expenditures from vehicles that have a price relative used in the final index calculation. This means expenditures for a vehicle v at time t will not be used in calculating the denominator unless vehicle v has corresponding vehicle sales in time $t-1$ (or has a relative imputed during the changeover process).

$$W_{v,t} = \frac{E_{v,t}}{\sum_{v=1}^{V_t} E_{v,t}}$$

Price relatives and monthly expenditure weights for each vehicle i are aggregated into a current period geometric mean index for a geographic region.

$$P_t^{Monthly} = \prod_{v=1}^{N_t} (R_{v,t}^m)^{W_{v,t}}$$

Note that this formula uses expenditure share weights observed in time t while the current CPI uses a geometric mean formula with fixed expenditure shares over time. The use of the geometric mean formula is based on the assumption that consumers can substitute across products upon observing price changes.²²

The impact of the “changeover” relatives on the overall index is dependent on how changeovers are weighted. If sales went to zero for an old item before its new counterpart were introduced, the changeover would receive zero weight and have no effect at all. The case we see in the J.D. Power data is less extreme, but expenditures tend to be relatively low during the changeover period for new vehicles where the expenditures on the old item in time $t-1$ and the new item in t are used to calculate weights. This diminishes the impact of the changeover on the overall index compared to a fixed-weight index. We see an additional issue where differences in area size lead to bias because of a relationship between area size and changeover weighting.

8.2 Construction of the hybrid index

We construct an index that combines the YoY index with elements of a month-to-month index. A time series can be decomposed into trend and high-frequency components. The YoY index captures the

²² See Dalton, Greenlees, and Stewart (1998) for a discussion of the geometric mean in the CPI.

trend component of a new vehicle price change but omits the high-frequency fluctuations that represent short-run price variation, an issue we remedy in our index with the addition of a high-frequency term extracted from a monthly index.

Several well-known statistical filters can be used to separate time series into components that reflect the long-run trend and short-run fluctuations. The trend component of the month-to-month index, T_t^M , is biased by area size (see the discussion in Section 4) and complicated by within-model-year price change. However, the short-run component, C_t^M , reflects information about the short-run behavior of the new vehicle market that is useful in describing price behavior in the short-run even if, by construction, the short-run component does not impact long-run levels. The filter is run on a logged month-to-month index, $\log P_t^{Monthly}$, to obtain estimates for the two components:

$$\log P_t^{Monthly} = \log T_t^M + \log C_t^M$$

The year-over-year index, P_t^Y , does not capture short-run changes but provides a means of measuring long-run price trend. We replace the trend estimate from the month-to-month index with the year-over-year index to create our hybrid index. In log terms:

$$\log I_t = \log P_t^Y + \log C_t^M$$

The hybrid index can be thought of as a trend-adjusted month-to-month index. With a month-to-month index in time t denoted as P_t^M , the long-run trend in that index denoted as T_t , and the year-over-year index as P_t^Y , we can define a trend adjustment:

$$\tau_t = \log P_t^Y - \log T_t$$

Therefore the hybrid index is:

$$I_t = e^{\tau_t} P_t^{Monthly}$$

Notice that when the trend in monthly price change (T_t) is similar in value to the trend in annual price change (P_t^Y), the adjustment does not impact the index. The trend term would only have an impact on the hybrid index if the trend of the month-to-month index deviated from the YoY trend.

8.3 Choice of time series filter

We use the band pass filter developed in Christiano-Fitzgerald (2003) to extract a high-frequency component from the monthly index.²³ Since we intend to use the filter in real time production, endpoint estimation (e.g., the current month) is particularly important. Nilsson and Gyomai (2011) find that the Christiano-Fitzgerald band pass filter has less revision (estimates of the trend are less sensitive to being at or near endpoints) when run in subsequent months, although they find that the Hodrick-Prescott has an advantage in identifying turning points in the trend component. We choose the Christiano-Fitzgerald band pass filter to extract our high-frequency component since we are interested in simulating “real-time” use of this methodology.

Our implementation of the band pass filter is an asymmetric band pass filter on a non-stationary time series (a series with a unit root). Using this band pass filter, we extract a specific range of

²³ The band pass filter is widely used and has better performance in real time compared to other common alternatives such as the Hodrick-Prescott filter.


frequencies. We extract the high frequencies, specifically those between 2 and 24 months. These frequencies represent the higher frequency range we miss when only looking at the 12-month changes of the YoY filter. Of most interest is the calculation of the endpoint, the high-frequency component for the current period that would actually be incorporated into the estimate of index change published on a monthly basis.

8.4 Cyclical component

While the month-to-month indexes fail to provide unbiased measures of long-run price change, the indexes do capture important information about short-run fluctuations in the market. These short-run fluctuations are not captured in the YoY index, which smooths over high-frequency price changes. Using a band pass filter, we extract short-run cyclical components from the month-to-month indexes. The short-run indexes are highly correlated with each other, suggesting the choice of index formula has little impact on measurements of high-frequency fluctuation.

When we apply the same filter to extract the high-frequency component of the CPI for new vehicles during the same time period, as in Figure 12, we see that the two cyclical series are closely related. On a monthly basis, they have a correlation coefficient of 0.43. This implies the short-run cyclical fluctuation in both indexes is more than noise. Since indexes based on two different data sources result in correlated series, their fluctuations most likely represent meaningful information about the real world. Moreover, the relationship between components persists across changes in methodology used on the J.D. Power data as shown in Figure 12 and Table 1.

Figure 12: Extracted Cycle Components by Index Formula



The extracted cyclical component also does not appear to depend on the changeover process, which can have a large impact on index trends. In Figure 13, we observe that a simple, matched-model index with a Tornqvist formula used at the elementary level (the same index presented in Section 3) produces a cyclical pattern that is still related to the CPI for new vehicles and the J.D. Power indexes with changeovers. This indicates that the cyclical pattern observed in the CPI and in the above J.D. Power indexes is not an artifact of implementation of the changeover process but a pattern evident in simple index construction.

Figure 13: Extracted Cycle Components by Index Formula

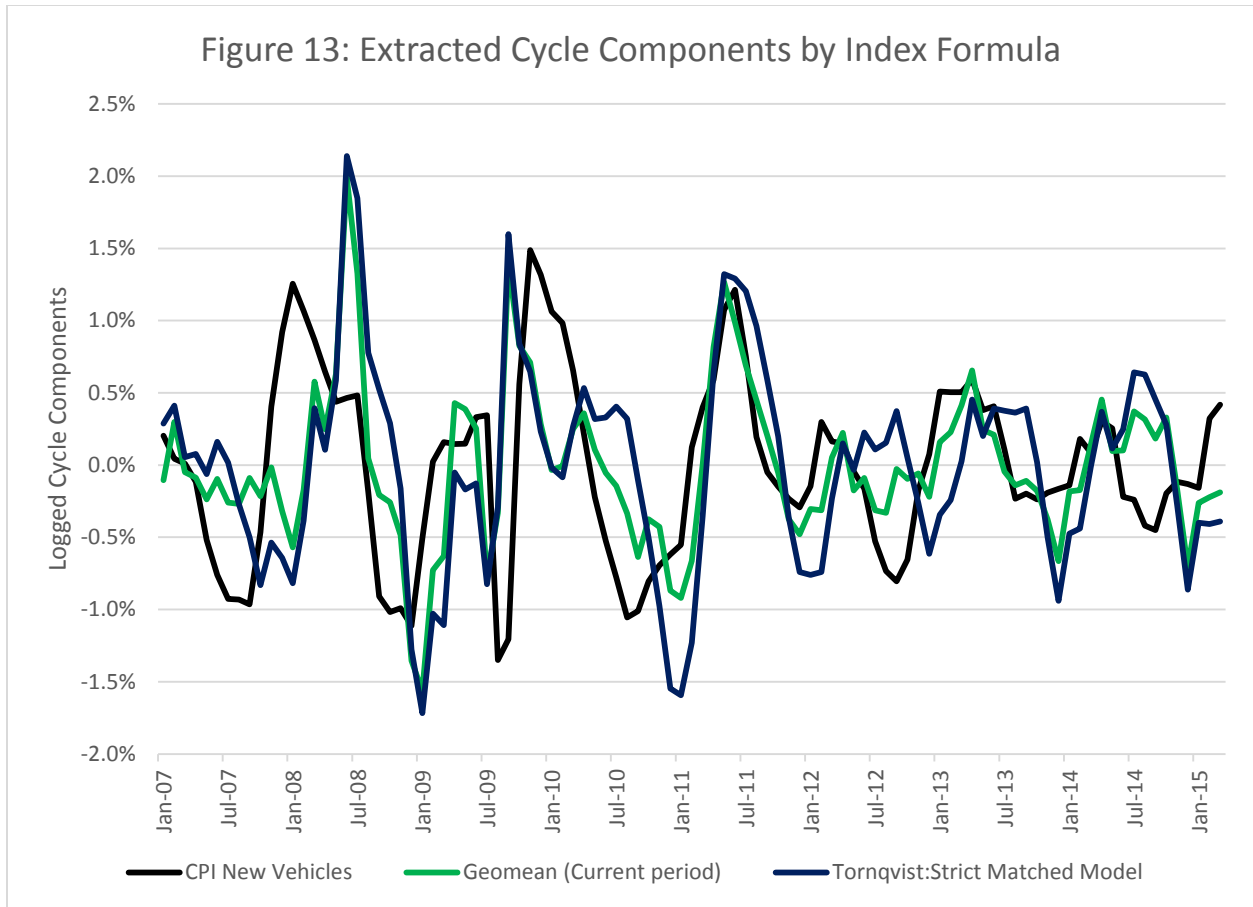


Table 1: Cyclical Behavior and Cross-Correlations*


	STD DEV	CORR +2	CORR+ 1	CORR	CORR- 1	CORR- 2	CORR- 3	CORR- 4	CORR- 5
<i>Correlation with Logged CPI for New Vehicles in time t</i>									
GEOMEAN	0.005	0.185	0.416	0.435	0.418	0.470	0.455	0.351	0.174
LASP	0.006	0.469	0.585	0.481	0.339	0.311	0.251	0.138	0.004
TORN	0.006	0.461	0.593	0.510	0.380	0.338	0.272	0.150	-0.004
TORN:SMM	0.007	-0.141	0.078	0.197	0.327	0.516	0.612	0.591	0.465
CPI NV	0.006								

*The extracted cycles in Table 1 are created by running the band pass filter over the entire series.

8.5 Hybrid YoY Index

In Figure 14, the YoY index combined with the high-frequency component of the monthly index is compared to the CPI. Area level YoY Tornqvist indexes are combined with cyclical components extracted from area level monthly indexes. These area level hybrid YoY indexes are aggregated to a national level index using the aggregation procedure currently used in the CPI. Initially, the cyclical

components are found using a band pass filter run on the first three years of data. Subsequently, the filter is run on a month-by-month basis, as it would in monthly, real-time production.²⁴



Based on simple correlation tests on the changes in the series, the hybrid YoY index based on J.D. Power data appears to lead the CPI, as seen in Table 2. For the entire period from January 2008 to March 2015, the index percentage changes for the two indexes (CPI_t and hybrid YoY) have a correlation coefficient of 0.136. During the same period, the correlation between the percent changes in the CPI and lagged J.D. Power index (CPI_t and hybrid YoY_{t-1}) was 0.494. These results suggest the J.D. Power data in time $t-1$ have predictive power for the CPI in time t . However, the comparison is made difficult by the fact that the early periods of this series use a filter that incorporates information on future data. Correlations for the period after the initialization of the cycle (December 2010 to March 2015) suggest the hybrid YoY index has a stronger contemporaneous correlation with the CPI. The correlation coefficient for the same period index changes is 0.531 and the correlation between the time t CPI and $t-1$ (lagged) hybrid YoY data is 0.558. This later period also reflects a more stable period for the new vehicle market. The earlier period covers the recession and subsequent recovery.²⁵

²⁴ The index in Figure 14 differs from those in the previous section, 8.4, where filtering was run only on the national level index and across the entire time-span of the final index instead of on a month-by-month basis, which is subject to revision in the filter estimates.

²⁵ A Granger Causality test on the differenced logs of the hybrid YoY and new vehicle CPI indexes (Augmented Dickey-Fuller tests suggest that these transformed series are stationary) shows that each series has some predictive power for the other. For a test based on a lag length of 2 (selected based on the Schwartz Information Criteria for a

Table 2: Predictability of CPI using J.D. Power Index

1/2008-3/2015	CPI	CPI (L1)
HYBRID YOY	0.136071	-0.072310
HYBRID YOY (L1)	0.493905	

12/2010-3/2015	CPI	CPI (L1)
HYBRID YOY	0.530582	0.337157
HYBRID YOY (L1)	0.558370	

9. Robustness: Why a Year-over-Year Index?

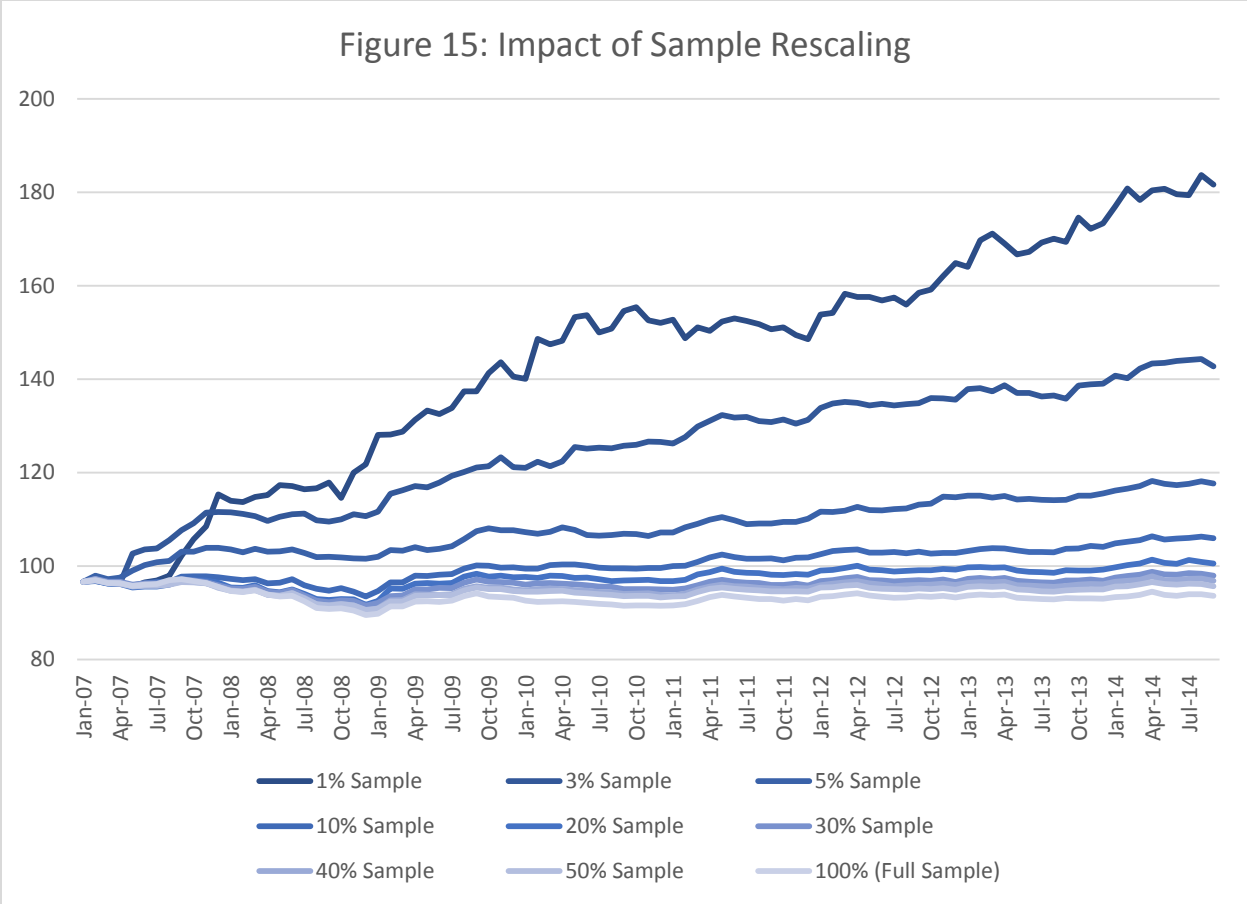
The monthly indexes produced with the changeover methodology, described in Section 8.1, exhibit biased trends yet represent information about the short-run fluctuations in the new vehicle market. In this section we characterize properties of this bias through a counterfactual exercise. We show that resizing an area through resampling traces the source of the bias to differences in the amount of weight put on changeover relatives compared to the weight placed on same-model-year price relatives. Even though index trends are biased, high-frequency price change is both similar to the CPI for new vehicles and robust to a large number of changes in underlying methodological choices. This suggests that combining the extracted high-frequency information from a monthly index with the YoY trend to produce an index that conveys accurate information across different time frequencies.

9.1 Area size simulation

We find that month-to-month indexes for small areas (in population and expenditure terms) tend to exhibit much higher index levels than those in large areas. Small areas have relatively high weight put on price relatives undergoing a “changeover” compared to price relatives reflecting within-model-year price change. Since changeovers tend strongly upward while within-model-year price relatives tend downward, increasing the relative weight on changeovers induces an upward effect on new vehicle price indexes.

To investigate the relationship between changeover weights and area size and confirm that the index level and size correlation did not represent a true economic phenomena, we created simulated areas of different sizes by randomly drawing a subset of transactions from the same population, the largest area in the CPI geography, X300 (medium-sized cities in the South). Since these simulated areas are randomly drawn from the same sample, they should produce similar indexes with some variance based on sampling error. Rather than merely increasing sampling error, we found that small samples showed larger price increases while large areas showed flat or declining price trends as shown in Figure 15.

VAR of the two series), the null hypothesis that hybrid YoY Granger causes the CPI is rejected at the 1% level while the null hypothesis that the CPI index Granger causes the hybrid YoY index is rejected at the 5% level but not the 1%.



Further analysis shows that the weight placed on changeovers increased as we decreased simulated area sizes. On the other hand, the magnitudes of the relatives in the resampling were not correlated with the size of a resampling.²⁶ We analyzed these simulations in terms of this decomposition and found that the difference in the resampling levels occurs through the changeover weight channel, as seen in Table 3. Smaller resamplings result in higher indexes because they put more weight on changeover relatives (generally positive) relative to within-model-year price change (generally negative). This is due to the fact that a changeover is constructed during every model year transition (approximately every 12 months), but, in small areas, the same vehicle might not be observed consistently every month, so within-model-year price relatives can only be constructed for a few months out of the year.

²⁶ Appendix A shows how an index can be decomposed into changeover and continuing relatives and changeover weight share.

Table 3: Area-Size Simulations and Changeover Characteristics

Share of X300	Index Change 1/2007 to 9/2014	Average Monthly Value		
		Changeover share	Changeover relative	Continuing Relative
100%	-3.1%	4.5%	1.05079	0.99613
50%	-1.0%	6.4%	1.05296	0.99656
40%	0.2%	6.5%	1.05218	0.99664
30%	1.4%	6.8%	1.05283	0.99661
20%	4.1%	7.4%	1.05183	0.99666
10%	9.6%	8.7%	1.05640	0.99658
5%	21.8%	10.8%	1.04996	0.99672
3%	47.7%	13.3%	1.05923	0.99635
1%	88.0%	25.2%	1.05662	0.99467

The area size bias we found confounds the deeper underlying issues of weighting item replacements and dealing with price change during the product cycle. Area size differences would not be associated with index levels if every vehicle were observed every month. However, bias from the product cycle may not be fully offset when weights change over time (see Section 5 for a related discussion). We would expect about one-twelfth (8.3%) of the weight to be on changeovers given the annual model year change in new vehicles. In large areas, changeovers receive much less weight. In X300, changeovers receive only 4.5% of the weight on average, each month. Weights for changeover relatives are based on the expenditures from the end (and, in Tornqvist indexes, beginning) of a model year’s sales. Expenditures follow a hump-shaped expenditure curve over the product cycle of a vehicle (see Section 3.1.1), which means the within-model-year relatives tend to receive substantially more weight than changeover relatives in large areas. In large areas, specific vehicles tend to be observed more often. There are fewer cases of “missing” within year observations and the weight on changeover relatives is lower than in small areas.

Fundamentally, when changeover relatives are used alongside same-model-year price relatives, the results are very sensitive to how the changeover relatives are weighted. In the indexes calculated here, the area size (through the mechanism of fewer “missing” observations within year) largely determines the relative weighting of changeover relatives. In a case with no missing observations, area bias may not result, but the weights at the time of vehicle changeover would be too low to offset the intra-model year price declines. For an item that has a product cycle pattern that leads to relatively high weights during the changeover period, this process would overcompensate and thus we conclude that caution should be taken when item replacements (changeovers) are used in a dynamically weighted index.

If we abandoned dynamically weighted index formulas, item replacement could be used without introducing area size bias. This would mean using a fixed-weight formula instead of one of the superlative formulas recommended by cost-of-living theory as well as introducing mechanisms for sampling, fixing weights, and rotating samples. In Appendix B, we create a fixed-weight index that mimics the CPI’s current methodology as closely as possible. While area size bias is no longer a problem in the fixed-weight indexes, the indexes remain very sensitive to the amount of weight placed on changeover relatives. We show that the length of time between sample rotations leads to differences in index levels driven by the impact of sample rotation on the frequency of model year changeover in the

sample. Therefore, we consider the year-over-year measure to be more accurate and robust than the indexes explored here that more closely approximate current BLS methods.

10. Conclusion

This work represents one of the first efforts to incorporate real-time transaction data with expenditure weights into the US Consumer Price Index. One of the persistent features of the new vehicle market is a product cycle characterized by the tendency for a model year vehicle to be introduced at a high price that then declines throughout the model year. The background literature and our own data analysis suggest that a variety of factors might drive the behavior of vehicle prices over the product life cycle. We pay particular attention to the possibility that intertemporal price discrimination may play a role. The methods currently used in the CPI offset this decline as part of the item replacement process, but this method is not compatible with real-time expenditure weights available in transaction data. When weights vary over time, a changeover relative will not necessarily offset within-model-year price changes and may lead to an upward bias if the changeover relative is weighted too heavily or a downward bias if it is underweighted. Thus, we proposed a year-over-year index as an alternative means of measuring price change that is not sensitive to product cycle effects.

Consistent with Greenlees and McClelland (2010), we find that transaction price indexes may be susceptible to downward drift that differs from the commonly discussed "chain drift" resulting from high frequency price change. We eliminated weight "bouncing" chain drift as an explanation for the matched-model index's price declines leading us to explore other means of measuring price change. Matched-model price indexes, especially those based on scanner data, often exhibit strong, downward tendencies that appear to be associated with chain drift. Our work suggests apparent "drift" may result from factors other than weight fluctuations. A distinct product cycle for a product may lead to bias in the index. Intertemporal price discrimination may also explain the tendencies for price decline.

In new vehicles, this drift may be the result of intertemporal price discrimination where matched-model price change may reflect differences in consumer characteristics: Showing the price change from the high willingness-to-pay consumers who pay for new model year cars to the budget-minded consumers who buy the same car for less at the end of the model year. An index that only reflects the price change from price inelastic to elastic consumers, and, moreover, chaining this price change over several production cycles, will be biased and not reflective of a cost-of-living index. Similarly, the downward price movement over a model year may reflect a change in the perceived value of having a new vehicle versus a late model year one (without necessarily having a change in consumer characteristics).

We find that the item replacement process used to offset downward drift in the CPI is ill-suited for transaction data. The weighting of this relative is usually not considered since official statistics have generally been limited to using elementary indexes formulas with equal weights. In dynamically weighted indexes, this method may fail to properly offset product life cycle effects. In our case, indexes based on this method have a bias associated with area size.

Instead of directly adjusting for within-model-year price declines, we focus on a measure of long-term price change that is not susceptible to within-model-year price change: A year-over-year index. The year-over-year (YoY) index represents a general means of ensuring correct price measurement in the presence of price change related to a strong annual product cycle. Regardless of the specific mechanism behind the product cycle behavior of vehicle prices, the YoY method provides a means for measuring accurate price change. More precise corrections should be the subject of future work where the availability of consumer data or more advanced econometric techniques might allow for identifying and adjusting for the effects of intertemporal price discrimination, fashion effects, and other factors that drive

price change over the product cycle. The YoY provides an accurate measure of long-run price change but sacrifices information on short-run market conditions.

We explore using time series filtering techniques to combine a cyclical component extracted from a trend biased by product cycle effects with a trend based on the YoY index. Viewing price indexes as combinations of different frequencies of price measurement could have broader applications in cases where there are trade-offs in price index time horizons, such as indexes for seasonal goods. Our hybrid, year-over-year plus cycle index could be used to incorporate transaction data into the CPI and benefit from the enhanced accuracy of a larger dataset, transaction prices, and weights based on real-time expenditure information.

References

- Aizcorbe, A., B. Bridgman, and J. Nalewaik. 2009. "Heterogeneous Car Buyers: A Stylized Fact." *Finance and Economics Discussion Series* (Board of Governors of the Federal Reserve System). <http://www.federalreserve.gov/pubs/feds/2009/200912/200912pap.pdf>.
- Aizcorbe, A., B. Bridgman, and J. Nalewaik. 2010. "Heterogeneous Car Buyers: A Stylized Fact." *Economic Letters* 109 (1): 50-53.
- Akerlof, George. n.d. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics* (The MIT Press) 3 (84): 488-500. doi:10.2307/1879431.
- Bils, Mark. 2009. "Do Higher Prices for New Goods Reflect Quality Growth or Inflation." *The Quarterly Journal of Economics* 124 (2): 637-675. <http://www.jstor.org/stable/40506240>.
2018. "Chapter 17. The Consumer Price Index." In *BLS Handbook of Methods*. <https://www.bls.gov/opub/hom/pdf/homch17.pdf>.
- Christiano, Lawrence J., and Terry J. Fitzgerald. 2003. "The Band Pass Filter." *International Economic Review* 44 (2): 435-465. doi:10.1111/1468-2354.t01-1-00076.
- Copeland, A., W. Dunn, and G. Hall. 2011. "Inventories and the automobile market." *The RAND Journal of Economics* 42 (1): 121-149. <http://www.jstor.org/stable/23046792>.
- Corrado, Carol, Wendy Dunn, and Maria Otoo. 2006. "Incentives and Prices for Motor Vehicles: What Has Been Happening in Recent Years?" (Federal Reserve Board Divisions of Research & Statistics and Monetary Affairs).
- Dalton, Kenneth V., John Greenlees, and Kenneth J. Stewart. 1998. "Incorporating a geometric mean formula into the CPI." *Monthly Labor Review* 121 (10): 3-7. <https://www.bls.gov/opub/mlr/1998/10/art1full.pdf>.
- Diewert, W.E. 1999. "Index Number Approaches to Seasonal Adjustment." *Macroeconomic Dynamics* 3: 48-69.
- Diewert, W.E. 1996. "Seasonal Commodities, High Inflation and Index Number Theory." (Department of Economics, University of British Columbia).
- Greenlees, J., and R. McClelland. 2010. "Superlative and Regression-Based Consumer Price Indexes for Apparel Using U.S. Scanner Data." St. Gallen, Switzerland: Conference of the International Association for Research in Income and Wealth.
- International Labour Organization (ILO). 2004. *Consumer Price Index Manual: Theory and Practice*. Edited by Peter Hill. Geneva: International Labour Organization. http://www.ilo.org/wcmsp5/groups/public/---dgreports/---stat/documents/presentation/wcms_331153.pdf.
- Ivancic, Lorraine, W. Erwin Diewert, and Kevin J. Fox. 2011. "Scanner data, time aggregation and the construction of price indexes." *Journal of Econometrics* 161 (1): 24-35. doi:10.1016/j.econom.2010.09.003.

- Konüs, A. 1939. "The Problem of the True Index of the Cost of Living." *Econometrica* 7 (1): 10-29. doi:10.2307/1906997.
- Melser, Daniel, and Iqbal A. Syed. 2016. "Life Cycle Price Trends and Product Replacement: Implications for the Measurement of Inflation." *Review of Income and Wealth*, September: 509-533. doi:10.1111/roiw.12166.
- Moulton, Brent. 2018. "The Measurement of Output, Prices, and Productivity: What's Changed Since the Boskin Commission." (Hutchins Center on Fiscal and Monetary Policy Report). <https://www.brookings.edu/wp-content/uploads/2018/07/Moulton-report-v2.pdf>.
- Moulton, Brent, and Karin Moses. 1997. "Addressing the Quality Change Issue in the Consumer Price Index." *Brookings Papers on Economic Activity* 1: 305-349.
- Nakamura, Emi, and Jón Steinsson. 2012. "Lost in Transit: Product Replacement Bias." *American Economic Review* 102 (7): 3277-3316. doi:10.1257/aer.102.7.3277.
- Nilsson, R., and G Gyomai. 2011. "Cycle Extraction: A Comparison of the Phase-Average Trend Method, the Hodrick-Prescott and Christiano-Fitzgerald Filters." *OECD Statistics Working Papers*. doi:10.1787/5kg9srt7f8g0-en.
- Pashigian, Peter B., Bowen, Brian, and Eric Gould. 1995. "Fashion, Styling, and the Within-Season Decline in Automobile Prices." *Journal of Law and Economics* 38 (2): 281-309. <http://www.jstor.org/stable/725687>.
- Reinsdorf, Marshall B., Paul Liegey, and Kenneth Stewart. 1996. "New Ways of Handling Quality Change in the U.S. Consumer Price Index." *BLS Working Paper 276*. <https://www.bls.gov/osmr/pdf/ec960040.pdf>.
- Shapiro, Adam Hale, and Ana Aizcorbe. 2010. "Implications of Consumer Heterogeneity on Price Measures for Technology Goods." BEA Working Papers 0062.
- Stokey, Nancy L. 1979. "Intertemporal Price Discrimination." *The Quarterly Journal of Economics* 93 (3): 355-371. <http://www.jstor.org/stable/188316>.

Appendix A

This appendix presents a useful way to decompose an index into the relatives and weight shares of the subclasses.

For a sample of n items and an expenditure share W and a price relative Rel we can represent a class of price indexes with this formula:

$$P_{Index} = \prod_{i=1}^n Rel_i^{W_i}$$

When W is the average of t and $t-1$ expenditures shares on i , P_{Index} is a Tornqvist index. When W is the t expenditure share, P_{Index} is a geometric mean index. By taking the natural logarithm of both sides we can convert the formula to arithmetic summation:

$$\ln(P_{Index}) = \sum_{i=1}^n W_i \cdot \ln(Rel_i)$$

We can then split the formula into an addition between components. Here we will split the sample into two groups, model year replacement (changeovers) and continued vehicles, but the general approach will work for any set of classes (cars vs trucks, top sellers vs. the remainder, etc.). The original n items can be broken down into $n = r + c$, the total sample (n) as the combination of the sets of item replacements (r) and continuing vehicles (c). Weights are represented in terms of overall total expenditure, so $\sum_{j=1}^{n_r} W_j + \sum_{k=1}^{n_c} W_k = 1$.

$$\ln(P_{Index}) = \sum_{j=1}^{n_r} W_j \cdot \ln(Rel_j) + \sum_{k=1}^{n_c} W_k \cdot \ln(Rel_k)$$

We can show that the index is a weighted average of the relative calculated on item replacements, REL_r , and the relative based on continuing prices, REL_c . We can replace the individual relatives with the relative for their respective sub class. The relative for the subclass relative for model year replacement can be calculated as follows (all weights, w , are in terms of overall sample):

$$\ln(REL_c) = \frac{\sum_{i=1}^n W_i}{\sum_{k=1}^{n_c} W_k} \sum_{k=1}^{n_c} W_k \cdot \ln(Rel_k)$$

This can be rearranged to show that the sum of the weighted relatives of subclass c is equivalent to the product of the weight on subclass c , $\sum_{k=1}^{n_c} w_k$ ($\sum_{i=1}^n w_i = 1$, so all weight shares can be expressed without explicitly showing the summation for total weights) and the overall relative for subclass c . The same relationship holds for the other subclass, r :

$$\ln(REL_c) \cdot \sum_{j=1}^{n_c} W_j = \sum_{j=1}^{n_c} W_j \cdot \ln(Rel_j)$$

We can substitute for the weighted relatives in the price index equation to obtain:

$$\ln(P_{Index}) = \ln(REL_r) \cdot \sum_{j=1}^{n_r} W_j + \ln(REL_c) \cdot \sum_{k=1}^{n_c} W_k$$

The above equation shows that the overall index relative can be expressed in terms of the weights and relatives of the subclasses. The weights of the subclasses sum to one. Since we only have two subclasses here, the weights of one subclass determine the other. We can substitute based on the relationship, $\sum_{k=1}^{n_c} W_k = 1 - \sum_{j=1}^{n_r} W_j$

$$\ln(P_{Index}) = \ln(REL_r) \cdot \sum_{j=1}^{n_r} W_j + \ln(REL_c) \cdot \left(1 - \sum_{j=1}^{n_r} W_j \right)$$

Thus, all price indexes of this form (including Tornqvist and geometric mean) can be viewed as the outcome of three inputs: the model year replacement or changeover relative (REL_r), the continuing price relative (REL_c), and, the changeover weight ($\sum_{j=1}^{n_r} w_j$).

Appendix B

The indexes discussed in this paper do not maintain fixed samples over time and allow the sample to fluctuate between new and discontinued varieties of vehicles. Interest has been expressed in seeing indexes based on a fixed sample drawn from the J.D. Power data with procedures adhering as closely as possible to those used in the CPI currently. In this appendix, we develop indexes maintaining many of the key features of the current CPI methodology for new vehicles: four-year sample rotation, expenditure disaggregation for item selection, fixed-based period weight geometric mean formula, annual substitutions (item replacements) between model years with class-mean imputation for non-comparable substitutions, and cell relative imputation for unavailable items. The J.D. Power data do not identify individual dealerships so the process of outlet selection through the telephone point-of-purchase-survey data cannot be applied.

As discussed in Section 5.1, the use of a fixed-weight index should allow item replacement to be used to offset product cycle effects. However, a sample cannot be held fixed indefinitely.²⁷ The specifics behind implementing sample rotation can lead to differences in index levels and the relative weighting of changeover relatives. Shorter rotation lengths lead to lower indexes with less weight placed on changeover price relatives. In the shortest rotation length we tried, one year, many of the vehicles rotated out of the sample without having a changeover. In effect, short rotation lengths lead to a censoring of changeover frequency. The longer the time between rotations, the less likely a changeover is “missed” because of sample rotation.

The current CPI is constructed using a four-year, fixed sample rotated on a staggered basis. Entire PSUs (primary sampling units, our most fundamental geographical group) are rotated at a time. Rotations occur on a rolling basis with one-eighth of the sample rotating at once. Rotations occur in February or August. No price relatives are created between items rotating in and out. Item replacements between model years occur annually, but each observation generally reflects the same make, model, and trim over the life of the sample. Fundamental differences also remain related to the CPI’s price collection and J.D. Power’s transaction prices. CPI procedures collect a vehicle’s offer price for the sampled configuration regardless of whether that specific configuration has been sold or is even available as long as the same make, model, and trim vehicle has been sold within a month. In order for a vehicle to be priced in the J.D. Power indexes, the specific configuration denoted by the squishVIN must be observed in a transaction.

Data collectors select an item for the sample based on expenditure information provided by dealers. Implementation of the disaggregation procedure is driven by the ability and willingness of the dealership to provide expenditure information for different categories of vehicles. Disaggregation selects a specific model for the sample by sequentially selecting among various categories of vehicles probabilistically selecting a category based on the information provided by the dealer. Typically, a disaggregation for a vehicle would start with selecting a car or truck, then move on to a specific make, model, and then a final set of options. First stage disaggregation usually relies on either a ranking or percentage expenditure share of the selection categories. Procedures for disaggregation do not specify a “recall period” for this expenditure share information. Generally, the shares reflect the dealer’s interpretation of “current.” For our fixed sample experiments with the J.D. Power data, we set weights

²⁷ Vehicles exit the marketplace and, even in cases where old models can be replaced with matched successors, the new vehicles may share little with predecessors besides a name and branding—drawing into question how “fixed” a sample can be.


based on the expenditure share of vehicles in the year prior to the rotation. This reflects a one year recall period.

As previously discussed, the J.D. Power data do not indicate the dealership of purchase so we cannot implement the first stage of creating a sample observation, selecting a dealership to survey. We essentially skip this stage of sampling and pool expenditures across all dealerships within a geographic area. We take all vehicles observed in the initiation period and weight them according to their expenditure share within the area during that period. The CPI does not observe expenditures (or expenditure shares) for vehicles in the sample, but vehicles are sampled based on the expenditure information used in disaggregation. We use explicit weights instead of selection probabilities, but, fundamentally, we are still getting the result as targeted in CPI sampling: vehicles weighted proportionally to their sales at the time of initiation. A vehicle with X% of the expenditures in an area would be expected to be represented by X% of the sampled observations in that area in the CPI but will be represented by a single unit with X% of the weight in the indexes constructed using J.D. Power.


Like the CPI, we construct changeover item replacements between model year versions of a vehicle when the sales of a new version of a vehicle exceed those of the old version. The same cost-based, quality adjustment procedures and values are used here as in the production CPI and the year-over-year indexes. When an item replacement is non-comparable, the relative is imputed from other item replacements in the same area (or nearby area if none are available from the same area during that month).

Initially, we constructed fixed sample indexes that required an exact match for an item in the sample to appear in order for an item replacement to occur. When a vehicle has no exact match, it is imputed until the sample rotates. Holding the sample “fixed” by only allowing vehicles to be replaced by direct successors created a high rate of attrition. By the end of the four-year rotation cycle, almost 80% of the sample by weight was imputed. We do not face this sort of attrition in the CPI currently because dealers can alert data collectors to discontinued vehicles and help choose successors.

Sample Attrition: Share of Imputes Fixed-Weight Index with Four-Year Rotation



Working with the J.D. Power data, we have to turn to data-driven methods of item replacement since we need a process that works without dealer input. We mimic these item replacement procedures by allowing replacement by another vehicle of the same make and model if a specific variation in the prior model year is no longer sold. When a non-exact match is made, a non-comparable changeover is implemented. We experimented with different rotational lengths (one-, two-, three-, or four-year rotations) and both “exact” and “broad” item replacement methodologies.



The indexes based on the exact match substitution methodology show a decline while those calculated allowing for broader substitutions show a slight increase. Both methodologies show an impact of sample rotation frequency on index levels. Looking at these indexes in terms of the decomposition discussed in Appendix A illustrates the mechanisms at work. However, these components do not directly reflect final index levels due to differences in time and geographical aggregation. Aggregating monthly weights and area relatives over the entire time span examined (December 2007 through March 2015) shows that indexes based on the “broad” match have much higher shares of weight on changeovers than those based on “exact” matches.

Rotation Length (Years)	Model Year Match	Changeover Share	Changeover Relative	Continuing Relative
1	Exact	6.22%	1.0458	0.9964
2	Exact	6.37%	1.0497	0.9960
3	Exact	6.44%	1.0503	0.9961
4	Exact	6.76%	1.0518	0.9954
1	Broad	7.49%	1.0458	0.9964
2	Broad	8.10%	1.0500	0.9961
3	Broad	8.37%	1.0507	0.9962
4	Broad	8.45%	1.0527	0.9956

For the exact match indexes, we found about 6.2 to 6.8% of the weight was placed on changeovers. We would expect about 8.3% (one-twelfth) of the weight to be on changeover relatives for a fixed-weight index with annual changeovers. Under the exact match methodology, vehicles without successors do not have a changeover in the final year. A vehicle without a successor will show a within-model-year price drop during its final year but not an offsetting changeover. This leads to lower weight on changeovers and lower index levels.

The “broad” match indexes have fewer cases where vehicles fail to have changeovers because they do not have successors (“missed” changeovers will still occur if a make and model is no longer sold). This adaptation led to higher shares of weight on changeovers ranging from 7.49% to 8.45% depending on the length of time between rotations. In both the exact and broad match indexes, the changeover relative becomes larger with longer rotation cycles. In addition to the mechanical impacts related to the frequency of sample rotation, differences in index levels may be due to compositional differences of the fixed samples. Older samples may contain vehicles that reflect older technologies or that consumers have shifted expenditure away from. Compositional effects may be related to the larger changeover price relatives observed in the longer-length rotations

Like the indexes with monthly weights and item replacement discussed earlier in this paper, the results of the fixed-weighted indexes demonstrate potential complications that arise when using item replacement as a method to offset product cycle price change. Indexes are very sensitive to the relative weighting of changeover price relatives compared to those based on “continuing” (within-model-year) price change. Where differences in the frequency at which vehicles are observed in different sized areas led to an area size bias in the indexes with monthly weights, differences in rotation length created variation in index levels for fixed sample indexes. Sample rotations introduce a mechanical problem into the index where the timing of rotation may lead to items not showing the correct balance between within-model-year and changeover price change. Instituting longer periods between rotations appears to reduce this censoring of changeover price changes but leads to samples that are less representative of current market conditions.

The overarching issue we have confronted in this research project is how to measure price change in a product with a strong product cycle. We have proposed year-over-year measurement as a means of addressing this issue. The current production CPI relies on an item replacement process to offset the tendency for prices to drop within a model year. Applying the same item replacement method to transaction data requires making a variety of sacrifices and assumptions. When applied to dynamically weighted indexes, the effects of item replacement can be ambiguous, and, in our specific case, the

combination of item replacement and dynamic weights creates a bias associated with area size. In order to viably use item replacement, we turned to the fixed-weighted indexes discussed in this appendix. Introducing fixed weighting requires abandoning the superlative index formulas best suited for a cost-of-living index when transactions data sets with dynamic weights are available. The geometric mean formula can be considered an approximation for a cost-of-living index under the restrictive assumption of a Cobb-Douglas utility function, which implies fixed expenditure shares. For new vehicles, this assumption does not hold since expenditures (and expenditure shares) vary over the product cycle of the vehicle. Moreover, implementing a fixed-weight index requires constructing a fixed sample, which introduces additional issues regarding sample attrition, delayed entry of new goods, and, as the sample ages, a general loss of representativeness. A fixed sample cannot be held forever, and introducing sample rotation introduces additional complexities including censoring the frequency of item replacement and limiting its effectiveness.