

Product Creation and Destruction: Evidence and Price Implications

By CHRISTIAN BRODA AND DAVID E. WEINSTEIN*

This paper describes the extent of product creation and destruction in a large sector of the US economy. We find four times more entry and exit in product markets than is found in labor markets because most product turnover happens within firms. Net product creation is strongly procyclical and primarily driven by creation rather than destruction. We find that a cost-of-living index that takes product turnover into account is 0.8 percentage points per year lower than a “fixed goods” price index like the CPI. The procyclicality of the bias implies that business cycles are more volatile than indicated by official statistics. (JEL E31, E32, L11, O31)

Over the last 20 years, economists have dramatically improved our theoretical understanding of how product innovation influences major aspects of macroeconomic performance. Not only has research explored the potential role that product creation and destruction has for explaining business cycle fluctuations (e.g., Andrei Shleifer 1986; Ricardo J. Caballero and Mohamad L. Hammour 1994; Fabio Ghironi and Marc J. Melitz 2005 among others) but economists have also examined the key role played by new and better products for long run growth (e.g., Paul M. Romer 1987; Gene M. Grossman and Elhanan Helpman 1991; Philippe Aghion and Peter Howitt 1992; Jakob Klette and Samuel Kortum 2004 among others). Despite the vast theoretical implications of product creation and destruction, the empirical analysis on the aggregate behavior of product turnover lags far behind its theoretical counterpart. This gap has emerged largely because of data availability. Even the recent emergence of scanner databases specific to particular stores is not useful for understanding product creation and destruction because store-specific data are not appropriate to analyze the extent of creation and destruction of products for the consumer. In this paper, we document the nature, extent and cyclicity of product entry and exit in the United States with special attention to the implications that it has for the measurement of prices. In particular, we quantify the biases that arise in “fixed goods” price indexes like the CPI because they largely ignore the changes in overall product quality available to consumers that occur as new products replace outdated ones.

We introduce a unique dataset that contains the universe of products with bar codes purchased by tens of thousands of households in sectors that cover around 40 percent of all expenditures on goods in the CPI. We first explore the vast amount of information about product creation and destruction that is lost if researchers have access only to firm level data. By matching bar codes

*Broda: University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637 (e-mail: cbroda@chicagobooth.edu); Weinstein: Department of Economics, Columbia University, 420 W. 118th St. MC3308, New York, NY 10025 (e-mail: david.weinstein@columbia.edu). The authors wish to thank the NSF (grant SES-0452460) and the Federal Reserve Bank of New York for providing critical financial support for this project. We would like to thank the Director of Research at the FRBNY, Joseph Tracy, for his early support of this project. We would also like to thank ACNielsen's Director of Research Frank Piotrowski, Ivan Rocabado, and Maura Elhbretch for their careful explanation of the data, and Olivier Blanchard, Jeff Campbell, and Steve Davis for very useful comments. Nonetheless, the views expressed here do not reflect those of the Federal Reserve System, and all errors remain our own. Alexis Antoniadis provided us with outstanding research assistance. In addition, we would like to thank the Global Financial Markets Initiative at the University of Chicago GSB for research support.

with firms, we document the multiproduct nature of the firm and show that the vast majority of product creation and destruction happens within the boundaries of the firm. In particular, we find four times more entry and exit in product markets than that found in establishment and labor market data (e.g., Timothy Dunne, Mark J. Roberts, and Larry Samuelson 1988, 1989; Steven J. Davis and John C. Haltiwanger 1992). In a typical year, 40 percent of household expenditures are on goods that were created in the last 4 years, and 20 percent of expenditures are in goods that disappear in the next 4 years.

We also document the cyclical patterns of product creation and destruction. We find that net creation is strongly procyclical, with more products being introduced in expansions and in product categories that are booming. Destruction of goods is countercyclical, although its magnitude is quantitatively less important. This is suggestive of models where firms have an incentive to defer implementation of the product until aggregate demand is relatively high (as in Jacob Schmookler 1962 and Shleifer 1986).¹ While early studies of the labor market suggest that job destruction responds more to cyclical movements than job creation (c.f. Olivier J. Blanchard and Peter Diamond 1990; Jeffrey R. Campbell 1998; Davis and Haltiwanger 1996), we find the opposite to be true in product markets (in line with more recent findings in labor markets by Robert E. Hall 2005 and Robert Shimer 2005). That is, product creation moves more than product destruction with business cycle fluctuations.

The economic significance of product turnover can be measured by estimating the quality bias in conventional “fixed goods” price indexes like the CPI. The fact that the quality available for consumers rises as a result of the creative destruction process is a central feature of Schumpeter’s work.² We show that since most product creation and destruction is unobserved by the BLS, there remains a substantial bias arising from new and higher quality goods in the CPI. This upward bias in measured inflation averages between 0.6 and 0.9 percentage points per year depending on the aggregation methodology in our sample of goods. This implies that inflation was around seven percentage points lower than suggested by the CPI over the period studied (1994–2003). The bias measures the compensating variation needed to keep consumers indifferent between the set of goods available in 1994 and those in 2003. It implies that consumers are willing to pay around seven percent of their income to access the set of goods available in 2003 relative to those available in 1994. We find that innovations in sectors like electronics, drugs, prepared foods, and new and improved household appliances (e.g., microwaves and dishwashers) were particularly important for consumer welfare.

We also find the bias to be procyclical, which suggests that business cycles are more pronounced than is typically reported in official statistics. For example, during the 2001 recession, real consumption of the products in our sample fluctuated by 0.4 percentage points more than suggested by national statistics. This cyclicity arises in part because 30 cents of each additional dollar of consumption is spent on new, higher quality goods that are ignored by conventional price indexes. It

¹ As noted by Shleifer (1986), Joseph A. Schumpeter (1939) thought the innovation process to be essentially autonomous and independent of market demand. By contrast, Schmookler (1962, 1966) believed that high demand periods were conducive to large profits from innovation. Thus, innovation would be concentrated in booms. Judd (1985) also suggests cyclical patterns of innovation that are driven by a different mechanism. In periods where variety per capita is low, profits to innovators are high and innovation thrives. However, as the patents from new products expire, prices of recently introduced products drop and profits for innovators fall. Only after an extended period of growth does innovation restart the process. In this paper we present evidence of Schmookler’s view of innovation.

² Throughout *Capitalism, Socialism and Democracy* (1942), Schumpeter refers to new goods as increasing the quality available to consumers. Page 84 provides a clear reference: “in capitalist reality (the type of competition) which counts (is) the competition from the new consumer goods, the new technology, the new source of supply, the new type of organization—competition which commands a (...) decisive quality advantage.” This aspect of the creative destruction process is also an important part of endogenous growth models like Grossman and Helpman (1991) and Aghion and Howitt (1992).

is important to highlight that while we study the cyclical patterns of hundreds of product categories, our data is limited to a single aggregate business cycle around the 2001 recession.

Our work is related to several different literatures. It complements the seminal work by Dunne, Roberts, and Samuelson (1988) and Davis and Haltiwanger (1992) by directly exploring a major element of the Schumpeterian creative destruction process: *product* entry and exit. This is the key mechanism through which the creative process has an impact on the welfare of consumers. We also document the basic cyclical properties of product creation and destruction in ways similar to the literature on job turnover (e.g., Davis, Haltiwanger, and Scott Schuh 1996; Campbell 1998; Caballero and Hammour 1994). This is closely related to the literature on innovation cycles and in particular to the work of Shleifer (1986). In his model, although inventions arrive evenly over time, they are implemented in waves. The waves arise because firms have an incentive to defer implementation until aggregate demand is relatively high. While our evidence is consistent with the work of Shleifer (1986), we leave a systematic examination of models of the innovation cycles to future work.

Second, our work is related to the papers that study the implications of quality change. We can examine the quality growth inherent in the process through which new products replace outdated ones. This is a key aspect of Schumpeter's process of innovation. Moreover, most of the studies examine the welfare impact of the introduction of specific products and fall short of computing aggregate price biases.³ A notable exception is the work of Mark Bils and Peter J. Klenow (2001) that used the US Consumer Expenditure Survey to quantify the quality bias in 66 consumer durable goods. The Advisory Commission to the CPI (1996) used a few studies that estimate these biases for specific sectors to extrapolate to the entire CPI and argue that the quality bias was around 0.6 percent per year. David E. Lebow and Jeremy B. Rudd (2003) survey the recent improvements on price measurement and conclude that the estimates of the quality bias found by the Advisory Commission were based on "at least a moderate degree of hard evidence" for only ten percent of the CPI. The rest comes from either inadequate evidence or is entirely subjective. Our paper uses detailed price and quantity data at the bar code level to estimate changes in quality and all the parameters necessary to compute an exact aggregate price index for almost half of the goods in the CPI.

I. Data Description

A. Overview

An important contribution of this paper is to bring a new dataset to bear on price measurement. Since the ACNielsen Homescan database has not been used extensively in other studies, it is worth spending some time describing its features. ACNielsen provides handheld scanners to approximately 55,000 households which then scan in the purchases of every good with a bar code.⁴ These households represent a demographically balanced sample of households in 23 cities in the United States.⁵ Bar codes are concentrated in grocery, drugstore and mass-merchandise sectors. Overall, the database covers around 40 percent of all expenditure on goods in the CPI.

³ An important paper that quantifies the extent of quality growth is Bils (2004). He examines the impact of product substitution in durable goods inflation by matching sales data on cars and 18 other consumer durables with the BLS's price data. The paper suggests that quality growth for durables has averaged about five percent per year in recent periods.

⁴ The 1994 data are based on a sample of 40,000 households; the data for 1999–2001 encompass 55,000 households, and the data for 2002 and 2003 represent 61,000 households.

⁵ The cities are Atlanta, Baltimore, Boston, Buffalo, Charlotte, Chicago, Columbus, Dallas, Denver, Detroit, Houston, Los Angeles, Miami, Minneapolis, New York, Philadelphia, Phoenix, Sacramento, San Antonio, San Francisco, Seattle, St. Louis, and Tampa.

The dataset is ideal for understanding how prices evolve for a large share of consumption expenditures for a number of reasons that we explain in this section. First, instead of relying on a small sample of goods we observe virtually the entire universe of goods purchased by households in the sectors we examine. We observe vastly more information about goods in these sectors than is observed by the Bureau of Labor Statistics (BLS) and other statistical agencies. Our database covers approximately 700,000 different goods purchased at some point by our household sample, while the BLS sample for the entire CPI covers only approximately 85,000 goods.⁶ This actually overstates the relative size of the BLS sample because our data covers only a subset of all consumption expenditure categories. In particular, the goods that underlie the CPI are spread over 305 expenditure categories called “Entry Level Items” or ELIs. Our sample covers roughly 104 of the 305 ELIs, which suggests that for the same set of expenditure categories that we examine, the BLS is working with a sample that is less than 5 percent as large.⁷

A second distinctive feature of our database is that the data collection point is the household and not the store. This is an advantage relative to more easily available store-specific scanner databases, where the researcher cannot distinguish whether a product that is new to the store is truly new to the consumer or whether the price or products at that store are representative. Our data circumvents these limitations by using data directly collected by a representative set of households. In particular, since the product information comes from actual purchases and is not restricted to particular outlets, our basket of goods is more in line with those of the representative household in the United States than that of other studies or even the BLS survey of prices. As long as one household out of the 55,000 that are surveyed weekly consumes a product, it becomes part of the “universe” of products for which we have information.

A third crucial characteristic of this database is that along with prices of each of the products, quantities of the same products are also collected at the same frequency. One must collect quantity data in order to correctly account for quality changes in the measurement of prices. However, when BLS field agents survey outlets, they observe the prices only of the products they sample. Thus, we have a unique opportunity to measure the extent of the quality bias in the CPI for a large number of categories included in the consumption basket. Moreover, in contrast to studies that use more aggregated industry or sector information, the data allow us to explore the multi-product nature of firms more precisely.

The data we use is collected in the following manner. With the scanners provided by ACNielsen, households scan the items they purchased at the conclusion of every shopping trip. If the purchase was made at large retail stores, the price is automatically downloaded from the store’s database. Otherwise the household enters the price paid and records any deals used that might affect the price (only a small fraction of the transactions are recorded this way). The matched price and quantity data means that we do not have to estimate the weights for prices we use; we know exactly how many units were purchased.

The data we obtained includes the average price paid and total quantities of all products with Universal Product Codes (UPCs, what we will also refer to as “bar codes”) purchased by the representative household at the quarterly frequency for six years: 1994 and 1999–2003. The data were weighted by ACNielsen to correct for sampling error. For example, if the response rate for a particular demographic category is low relative to the census, ACNielsen reweights the averages so that the price paid and the quantity purchased is representative of the United States as a whole.⁸

⁶ As a benchmark, the USDA estimates the total number of food UPCs to be 320,000, and the total number of UPCs in an average store in the United States to be 30,000 (see Anthony Gallo 1995).

⁷ A sample list of 400 general categories included in the database appears in the appendix of Christian Broda and David E. Weinstein (2007).

⁸ We also ran all of our tables and regressions with the raw (unadjusted) data, and the main results of the paper are unaffected. One concern about the database is that results might be driven by the growth in the number of bar codes *per se*.

Although it is difficult to enforce how a company uses a bar code, most industry experts strongly caution firms not to use the same bar code on more than one product.⁹ Doing so could cause confusion among retailers who would have trouble knowing what they were selling and for consumers whose receipts would not match their actual purchases. Similarly, firms typically do not use multiple UPCs for the same product because that makes it very difficult for retailers to reorder out of stock items. As a result, manufacturers tend to use other bar code systems for internal use and reserve the UPC for tracking products that are identical to the consumer. Therefore, it is reasonable to assume that all goods with different UPCs differ in some way that might cause consumers to pay a different price for them and that it is rare for a meaningful quality change to occur that does not result in a change of UPC. For example, changing the slogan on a Heinz ketchup bottle does not require a new bar code, but changing the size of the bottle does. In other words, it is safe to assume that if the bar code changes, it is likely that some noticeable characteristic of the product has changed.

As we proceed with our data analysis, it will be necessary to keep track of three levels of aggregation. It is easiest to understand these levels of aggregation by means of an example drawn from our data. At the lowest level we have a product which we identify using the UPC. For example, a box of “100-count Centrum Multi-Vitamins From A-to-Zinc in tablets” has a UPC of 030005-423936. Each UPC in turn belongs to a “brand module,” i.e., the brand “Centrum” within the module “Multi-Vitamins.” In the first quarter of 2001, there were 16 different goods (UPCs) marketed under this brand module. A manufacturer, in this case Wyeth, may have several brands within a module—e.g., “Centrum,” “Centrum Silver,” and “Centrum Performance”—and each of these would constitute a different brand module. At the highest level, we have a “product group,” which in this case would be “Nutritional Supplements,” which contains not just multivitamins but also other modules like “Kids’ Vitamins.”¹⁰

Table 1 presents descriptive statistics of the number of UPCs at the different levels of aggregation just defined. The table shows that there are roughly 650,000 different UPCs sold in each year in around 51,000 different brand modules which can be aggregated into 1,094 modules or 122 product categories. In other words, the average product group contains about nine modules, the average module contains 46 brands, and the average brand module contains about 34 products.

B. Stylized Facts

In this section we describe the extent, nature, and cyclicity of product creation and destruction in a large sector of the US economy. We present four stylized facts that document the main characteristics of product creation and destruction.

Stylized Fact 1: The Importance of Multiproduct Firms.—The first fact concerns the vast loss of information about creation and destruction that occurs when the unit of analysis is the firm rather than the product. This is important if one wants to compare our results with

In Appendix A of Broda and Weinstein (2007), we show two facts that suggest that this is not the case. First, the share of goods purchased with a bar code as a share of total shopping expenditures has remained constant at around 0.85 throughout this period. Second, we observe the average real expenditure per UPC to be constant over time. This suggests that over our sample period the growth in total sales has been approximately proportional to the number of UPCs.

⁹ In order to obtain a bar code, a company must register with the Universal Code Council, which costs around \$750 dollars plus an annual maintenance fee of \$150 per UPC. This means that the financial costs of registering new products are not likely to present an important obstacle for products entering and remaining in our database. Prices quoted from http://www.cummingsdesign.com/bar_codes101_UCC_App.htm. We also checked that in our data there are no two identical UPCs with different product descriptions, or two different UPCs with the same product description.

¹⁰ In Appendix C of Broda and Weinstein (2007) we provide another example of the levels of aggregation in our database that helps explain how we identify the elasticities we estimate in Section IV.

TABLE 1—DESCRIPTIVE STATISTICS

Year	Number of UPCs				
	Total	By brand module		By firm	
		Median	Average	Median	Average
1994	532,789	8	46.1	8	36.0
1999	651,343	5	33.6	8	39.6
2000	666,445	5	32.5	8	40.4
2001	658,055	4	31.4	8	39.9
2002	690,036	4	31.7	8	41.7
2003	697,312	4	31.5	8	42.0
Average	649,330	5	34	8	40
Per-year growth					
2003/1994	1.03	0.93	0.96	1.00	1.02
2003/1999	1.02	0.95	0.98	1.00	1.01
Year	Number of brand modules			Number of modules	Number of product groups
	Total	By firm		Total	Total
		median	Average		
1994	31,884	1	2.9	1,097	118
1999	51,881	2	3.9	1,096	122
2000	53,579	2	4.0	1,095	122
2001	54,213	2	4.1	1,093	122
2002	57,146	2	4.2	1,092	122
2003	58,135	2	4.2	1,093	123
Average	51,140	2	4	1,094	122
Per-year growth					
2003/1994	1.07	1.08	1.04	1.00	1.04
2003/1999	1.03	1.00	1.02	1.00	1.01

those of firm level studies of entry and exit. The first six digits of the UPC is the manufacturer identifier number, which can be matched to the parent company. For example, the manufacturer of Centrum (“030005”), Wyeth, also produces in other product groups, e.g., “Snack Bars” and “Medications, Remedies, Health Aids,” and using other six-digit brand identifiers (e.g., Advil and Robitussin).

By matching the manufacturer identifier numbers with firm codes, we can quantify the extent to which the world at the level of the UPC differs from viewing the world at the level of the firm. In our data, the typical firm sells eight different UPCs in two different brand modules. The distribution of UPCs per firm is highly skewed, however, with a large number of firms having a small number of products. As a result, the average firm sells 40 UPCs under 4 different brands in 3 modules, which, in turn, are contained in 2 product groups.

Table 2 highlights the multiproduct nature of firms in these markets. It describes firm characteristics by sales size. Only the smallest of firms sell in a single brand and product group. Over 60 percent of the sales in the fourth quarter of 2003 come from firms that sell over 700 UPCs in over 35 different brands and 19 different product groups. The bottom line is that the bulk of output in these sectors is produced by firms marketing hundreds of different products under dozens of brands in a variety of markets. While this is consistent with previous work that has documented the extent of *industry* diversification in US firms or plants (e.g., Mary Streitweiser 1991; Boyan Jovanovic and Richard J. Gilbert 1993, and Andrew B. Bernard et al. 2006), aggregation up to the firm level results in a substantial loss of information about the

TABLE 2— AVERAGE FIRM CHARACTERISTIC BY FIRM SIZE (2003:IV)

Total firm sales (in \$)	Average number of				
	UPCs	Brands	Modules	Product groups	Market share (1)
1–100,000	3	1	1	1	0.00
100,000–1,000,000	10	2	3	2	0.02
1,000,000–5,000,000	33	5	6	3	0.05
5,000,000–10,000,000	69	7	10	4	0.04
10,000,000–50,000,000	139	11	20	7	0.17
50,000,000–100,000,000	386	22	50	14	0.10
100,000,000–500,000,000	713	37	71	18	0.33
500,000,000–1,000,000,000	1,191	68	110	28	0.12
> 1,000,000,000	3,431	182	246	54	0.16

Note: (1) The market share is based on total value of UPCs for firms in 2003:IV within a bin compared to total value in 2003:IV.

process of creative destruction.¹¹ As we will see below, most of the product entry and exit happen within existing firms.

Stylized Fact 2: The Extent of Product Entry and Exit.—Before analyzing the extent of product entry and exit in our data, we define the statistics we use to describe the data and discuss some potential problems that can arise due to the nature of our database. We report the following measures of product creation and destruction:

$$(1) \quad \text{Entry Rate}(t,s) = \frac{\# \text{New UPCs}(t,s)}{\# \text{All UPCs}(t)} ; \quad \text{Exit Rate}(t,s) = \frac{\# \text{Disappearing UPCs}(t,s)}{\# \text{All UPCs}(s)}$$

$$(2) \quad \text{Creation}(t,s) = \frac{\text{Value of New UPCs}(t,s)}{\text{Total Value}(t)} ;$$

$$\text{Destruction}(t,s) = \frac{\text{Value of Disappearing UPCs}(t,s)}{\text{Total Value}(s)}$$

The entry rate is defined as the number of new goods in period t relative to period s as a share of the total number of products purchased in period t . Note that a new product is one that was consumed in period t by at least one household but was not part of the consumption basket of any household in period s . The exit rate is defined in a similar way. Creation and destruction are the weighted analogues of the entry and exit rates. Instead of simply counting the number of new and disappearing goods we use their value in consumption. Thus, creation is the share in total expenditure on those goods that were consumed in period t but were not available in period s .

Goods with a strong seasonal or fashion cycle tend to exhibit large price reductions as the seasons change. To prevent these factors affecting our measures of product turnover we define the periods t and s as years. We report the extent of product creation and destruction for the case in which t and s are one, four and nine years apart. It is exceedingly rare for a product to have positive sales in one year, zero sales in the following year, and then to have positive sales again

¹¹ Though marginal to this paper, it is interesting to note that a one percent increase in the overall market share of a firm is associated with 0.6 percent more UPCs per firm. This elasticity results from a cross-sectional regression of the log of the number of UPCs per firm on the log of the market share of the firm in a module.

TABLE 3—PRODUCT ENTRY AND EXIT IN THE UNITED STATES

Period	9-year 1994–2003	4-year 1999–2003	1-year median
Entry rate	0.78	0.50	0.25
Creation	0.64	0.37	0.09
Entrant relative size	0.49	0.56	0.30
Exit rate	0.72	0.46	0.24
Destruction	0.37	0.18	0.03
Exiter relative size	0.23	0.23	0.09

Notes: Entry rate = Number of new UPCs (t) / total number of UPCs (t)

Exit rate = Number of disappearing UPCs ($t - 1$) / total number of UPCs ($t - 1$)

Creation = Value of new UPCs (t) / Total value (t)

Destruction = Value of disappearing UPCs ($t - 1$) / Total value ($t - 1$)

Entrant relative size = Average sales of entrant (t) / Average firm sales (t)

in the following year, which suggests that product creation and destruction is not being driven by goods drifting in and out of the sample. Although not shown in the table, less than two percent of the products are products that reappear after being destroyed. In value terms, these products compose less than 0.2 percent of the sample. If we drop products that reenter after a period of being out of our sample from our calculation, the creation and destruction measures are effectively unchanged.

Finally, notice that these measures of product turnover include *any* change in products, including those driven by changes in the size of products, their flavor, or other characteristics that can be secondary for the consumer. We describe the importance of some of these characteristics in explaining the extent of product turnover below. In the following sections, we assess the economic significance of this product turnover and provide specific examples about how our methodology values these types of bar code changes.

Table 3 summarizes the extent of product creation and destruction using weighted and unweighted measures at different frequencies. The first column presents data on entry and exit rates between 1994 and 2003, the second column reports the same numbers between 1999 and 2003, and the third column presents the median annual rates for each year between 1999 and 2003. Column 2 reveals that almost 50 percent of the products that existed in 2003 were not around in 1999. These new products composed 37 percent of expenditures in 2003. The value of disappearing UPCs, that is those that existed in 1994 but did not exist in 2003, was much smaller: 18 percent of expenditure in 1994. The fact that creation is larger than destruction suggests that new products are systematically displacing market share from existing products. As will become apparent in the following sections, this displacement is indicative of biases in conventional price indexes that ignore the effects of changing quality.¹²

It is useful to compare our results with those of studies of plant and establishment turnover (c.f. Dunne, Roberts, and Samuelson 1988 and Davis and Haltiwanger 1992). Table 2 in Dunne et al. is directly comparable to our Table 3. Using the census of manufacturing firms collected every five years, they find that on average ten percent of all manufacturing output in a particular census year came from plants that were not present in the previous census, while plants that

¹² In Table 3 of Broda and Weinstein (2007) we also reported the average size of the entering and exiting UPCs relative to the average size of common UPCs. New UPCs tend to be larger than disappearing ones but both are smaller than common UPCs by a large margin.

were present in one census but disappeared in the following census accounted for 14 percent of output.¹³ While the total amount of output linked to entry and exit of plants in all manufacturing is 24 percent, Dunne et al. (1988) find that in the “Food Processing” sector the extent of plant turnover is roughly 15 percent, or 30 percent smaller than in the average manufacturing sector. This is less than a quarter of the market share of new and disappearing products over a four-year period in Table 3 (i.e., 56 percent). That is, relative to this study, we find that there is four times more product creation and destruction than plant creation and destruction. Davis and Haltiwanger present similar numbers for the importance of entry and exit of establishments, but weighted by employment rather than output. They find that over a one-year period, roughly three percent of current employment came from the entry of new establishments, while 2.5 percent of past employment was in establishments that disappeared in the following year. Altogether this amounts to 5.5 percent of employment, which is 2.5 times smaller than the market share due to new and disappearing products at the same frequency.

It is typical for firm level studies to drop the smallest firms in their samples as this is the group of firms with the largest measurement error. Similarly, in our data some of the bar code products are purchased by a small number of households. In order to show that the main results on the extent of product turnover are not driven by these products, we replicated Table 3 excluding those UPCs that were purchased by fewer than 20 households in a given year. The levels of entry and exit are marginally smaller than those in Table 3, suggesting that this correction has only a minor effect on the level of product turnover. We present this table in the appendix to Broda and Weinstein (2007).

For roughly 20 percent of the products that were purchased in the fourth quarter of 2003, we have detailed information about the characteristics of the UPC, including the package size and the flavor of the product. This allows us to proxy the extent of product creation that is driven primarily from changes in volume and flavors of existing products. For example, a new UPC might differ from an existing UPC only in the number of vitamin tablets in the bottle. This would be characterized as a volume change. Thus, we can calculate how much of overall creation is due to innovations in volume and how much is due to innovations in flavor. First, we verified that the subsample of UPCs with volume and flavor information has a similar creation rate as the whole sample. The rate of creation for this subsample is 35 percent over a 4-year period, which is very close to the 37 percent rate of the full sample in Table 3. The rate of creation due to new sizes is 1.9 percent or roughly 5 percent of overall creation. The rate of creation due to new flavors was only 2.3 percent, which also is only a small fraction of total creation. Thus, we conclude that the vast majority of the creation of goods in our sample is not due to volume or flavor changes.

Stylized Fact 3: Product Turnover is Concentrated Within Firms and Sectors.—In this section, we focus on the characteristics of the creation and destruction of products. We describe how much of product creation and destruction occurs within firms, within brands and across different types of products.

Interestingly, we observe most of the product entry and exit occurring within the boundaries of the firm. To more closely compare our results with those at the plant or establishment level, we examine the extent of product turnover within a “manufacturer identification number.” A firm that owns several plants can have several manufacturer identification numbers. Thus, the extent of manufacturer turnover will always be larger than that of firm turnover. Table 4 reveals that the extent of manufacturer entry and exit is much less than that of products. At one-year frequencies,

¹³ Because we are interested in comparing plant turnover to product turnover, we report only the magnitudes presented by Dunne et al. (1988) that correspond to “new firms, new plants” and “diversifying firms, new plants.” We exclude from this comparison the category “diversifying firm, product mix” as this represents existing plants that simply change the industry in which they sell.

TABLE 4—CREATION AND DESTRUCTION OF BRANDS AND MANUFACTURER ID NUMBERS

Period	Brand turnover			Manufacturer ID turnover		
	9-year 1994–2003	4-year 1999–2003	1-year median	9-year 1994–2003	4-year 1999–2003	1-year median
Entry rate	0.73	0.35	0.16	0.51	0.26	0.11
Creation	0.30	0.12	0.01	0.18	0.07	0.01
Exit rate	0.50	0.27	0.15	0.29	0.15	0.09
Destruction	0.18	0.03	0.00	0.06	0.02	0.00

Notes: Brand entry rate = Number of new brands (t) / total number of brands (t)

Brand exit rate = Number of disappearing brands ($t - 1$) / total number of brands ($t - 1$)

Manufacturer ID entry rate = Number of IDs (t) / total number of IDs (t)

Manufacturer ID exit rate = Number of disappearing IDs ($t - 1$) / total number of IDs ($t - 1$)

only one percent of all consumption expenditures come from establishments with a new manufacturer identification number. Comparing the extent of firm creation to that of product creation (Table 3) suggests that 92 percent of product creation happens within existing manufacturers, and 97 percent of product destruction happens within existing manufacturers. At four-year frequencies the comparable numbers are 82 and 87 percent, respectively. This implies that over a four-year period, 18 percent of the value of overall consumption is coming from products of completely new manufacturers, and 13 percent of product exit is happening because manufacturers disappear. In our data, product entry and exit is six to 30 times as important as manufacturer entry and exit depending on the time frame used. Table 4 also shows the summary statistics for both entry and exit of brands. Not surprisingly, most product entry and exit happens within existing brands.

One obvious question is whether the product categories that exhibit a lot of turnover correspond to our priors about which products are knowledge intensive. To explore this question we focus on the one hundred modules with the largest sales values, since there are many modules that have trivial market shares and very few UPCs (e.g., “retort pouch bags”). In Table 5, we report the ranks of the top ten and lowest ten modules in terms of turnover, where we define turnover as the sum of creation and destruction in the module. The fact that it is easier to be innovative when developing prerecorded video recordings, cameras, and computer software than when developing new forms of granular sugar, frankfurters and butter, is suggestive that these measures are capturing meaningful innovations for the consumer.

The aggregate patterns described until this point mask important relationships which exist at the level of the product. In Table 6, we examine the behavior of destruction by UPC characteristic. Here we divide up the goods that existed in 2002 into age bins. One-year-old goods are goods that existed four quarters earlier but not eight quarters earlier; two-year-old goods are goods that existed eight quarters earlier, but not twelve quarters earlier; and so on. The first panel shows that among younger UPCs, the share of exiting UPCs is larger than among older UPCs. The lower panels show that destruction is also systematically related to size. For bins of small UPCs (in terms of share in overall expenditure), destruction rates are higher than for bins of large UPCs. Thus, destruction rates are higher for smaller and younger UPCs. The lower panel shows that destruction does not monotonically fall with the size of the brand. Small and large brands have higher rates of destruction than middle sized brands.

Stylized Fact 4: Creation is Procyclical and Destruction is Weakly Countercyclical.—In this section we address how product creation and destruction covaries with aggregate measures of

TABLE 5—1-YEAR CREATION AND DESTRUCTION RATES BY MODULE RANKED BY TURNOVER

Module description	Ranking (1)	Creation (2)	Destruction (3)	Net creation (2) – (3)	Turnover (2) + (3)
Video products—prerecorded	1	0.56	0.09	0.47	0.65
Cameras	2	0.44	0.18	0.26	0.62
Computer software	3	0.32	0.19	0.13	0.51
Telephone and accessory	4	0.25	0.11	0.14	0.36
Vacuum and carpet cleaner appliances	5	0.25	0.14	0.11	0.39
Candle and candle in holder	6	0.30	0.15	0.14	0.45
Disposable diapers	7	0.17	0.06	0.12	0.23
Storage and space management	8	0.19	0.10	0.09	0.29
Kitchen utensil and gadget	9	0.18	0.10	0.08	0.28
Nutritional supplements	10	0.14	0.06	0.08	0.20
Sugar—granulated	100	0.01	0.01	0.01	0.02
Frankfurters—refrigerated	99	0.01	0.01	0.00	0.02
Butter and spreads	98	0.02	0.00	0.01	0.02
Cheese—cottage	97	0.01	0.01	0.00	0.02
Sausage—breakfast	96	0.02	0.01	0.01	0.03
Seafood—tuna—shelf stable	95	0.02	0.01	0.02	0.03
Eggs—fresh	94	0.03	0.01	0.02	0.04
Dinners—frozen	93	0.04	0.00	0.03	0.04
Marinara sauce	92	0.03	0.01	0.02	0.04
Pizza—frozen	91	0.04	0.01	0.03	0.04

Notes: The top 100 modules in terms of sales were ranked in terms of turnover (i.e., the sum of the creation and destruction rates). Creation (module) = Value of new UPCs (t) of module m /Total value (t) of module m . Destruction (module) = Value of disappearing UPCs ($t - 1$) of module m /Total value ($t - 1$) of module m . Net creation equals creation minus destruction.

consumption. For expositional ease (and slightly abusing notation), we will redefine creation in terms of the total value in period s rather than the total value in t (as in equation (2)). This transformation is useful as we can divide total sales growth into the sales of UPCs that survive and those that are new or disappearing. It is useful for us to define a UPC as “common” in periods s and t if the UPC was purchased in both periods. The variable COM_t is the total expenditures in period t on all UPCs that existed in periods t and s . The total expenditure in period s on the set of UPCs that existed in period s but did not exist in period t , i.e., the disappearing UPCs, are denoted by D_s ; and the total expenditure in period t on the set of UPCs that existed in period t but did not exist in period s , i.e., the new UPCs, is denoted by N_t . We make use of the relationships that $V_s \equiv COM_s + D_s$ and $V_t \equiv COM_t + N_t$ to obtain:

$$(3) \quad \frac{V_t - V_s}{V_s} \equiv \underbrace{\frac{COM_t - COM_s}{V_s}}_{\text{Growth due to common UPCs}} - \underbrace{\frac{D_s}{V_s}}_{\text{Destruction}} + \underbrace{\frac{N_t}{V_s}}_{\text{Creation}}.$$

We call D_s/V_s the rate of product “Destruction” and N_t/V_s the rate of product “Creation.” V_t , COM_t , and N_t are adjusted for inflation so that real dollar values are compared over time.¹⁴

¹⁴ We use the CPI Food index to deflate these series.

TABLE 6—1-YEAR DESTRUCTION RATES BY UPC CHARACTERISTIC

Age in years	Age	
	Destruction	Market share (in $t - 1$)
1	0.141	0.115
2	0.119	0.094
3–7	0.059	0.342
8+	0.026	0.449

Share ($t - 1$)	UPC share (in its module)	
	Destruction	Market share (in $t - 1$)
0–0.00001	0.209	0.014
0.00001–0.0001	0.096	0.110
0.0001–0.001	0.024	0.336
0.001–0.01	0.008	0.469
0.01–0.1	0.006	0.071
0.1+	0.004	0.000

Number of UPCs per brand	Brand size	
	Destruction	Market share (in $t - 1$)
1–9	0.119	0.150
10–99	0.062	0.550
100–999	0.068	0.271
1,000–9,999	0.332	0.030

Notes: Destruction (bin) = Value of disappearing UPCs ($t - 1$) of bin / Total value ($t - 1$) of bin. Market share is the sales share of the UPCs in a particular bin.

In order to quantify the cyclicalities of product creation and destruction, we present the patterns of net creation (creation less destruction), as well as creation and destruction separately. While we have only 20 consecutive quarters of data, the period studied includes the 2001 recession. Figure 1A plots net creation and the growth in overall sales of the ACNielsen sample. The pattern that emerges is procyclical. ACNielsen sales are weakest during the recession of 2001, and this is also the period during which net creation of goods reaches its trough. In the later years (2002 and 2003) sales and net creation pick up. Figure 1B shows the procyclicality of creation. That is, product creation is largest in periods where sales growth is strongest. Figure 1C shows a clear countercyclical pattern of destruction.

We should be cautious when interpreting the results of figures 1A–1C because there are only 24 quarters over which the data are collected.¹⁵ However, the cyclicalities of creation and destruction can be examined at the product group level. For each product group, we examine how creation and destruction covaries with the overall consumption in that product group. Exploiting the advantages of the accounting identity in equation (3), we can separately run regressions for net creation, creation, and destruction on the overall growth (the left hand side of (3)) of each product group in a particular period. Following Caballero and Hammour (1994), we also present results for periods where consumption growth is above and below average separately.¹⁶ We restrict the

¹⁵ We observe more creation and destruction in these plots than in Table 3. This is because goods created or destroyed in the last three quarters also count towards creation and destruction at the four-quarter frequency, while they do not at an annual frequency.

¹⁶ We do not include leads and lags of consumption growth as most of these coefficients are small and insignificant.




FIGURE 1A. SALES GROWTH AND NET CREATION (*Four-quarter growth rates*)FIGURE 1B. SALES GROWTH AND CREATION (*Four-quarter growth rates*)FIGURE 1C. SALES GROWTH AND DESTRUCTION (*Four-quarter growth rates*)

TABLE 7—CYCLICALITY AT THE PRODUCT GROUP LEVEL

	Full sample			Product groups with above average consumption growth			Product groups with below average consumption growth		
	Net creation	Creation	Destruction	Net creation	Creation	Destruction	Net creation	Creation	Destruction
Consumption growth	0.351	0.299	−0.053	0.384	0.375	−0.009	0.362	0.241	−0.121
(Product group level)	[0.018]**	[0.020]**	[0.013]**	[0.040]**	[0.048]**	[0.029]	[0.044]**	[0.043]**	[0.031]**
Observations	1,815	1,815	1,815	875	875	875	940	940	940
Number of rpg	122	122	122	119	119	119	119	119	119
R ²	0.18	0.12	0.01	0.11	0.07	0	0.08	0.04	0.02

Notes: Creation = Real value of new UPCs (t) / Total value ($t - 1$). Value of new UPCs (t) is adjusted by inflation using the CPI Food and Beverages index between t and $t - 1$. Destruction = Value of disappearing UPCs ($t - 1$) / Total value ($t - 1$). Net creation is the difference between Creation and Destruction. Consumption Growth is the growth rate of aggregate consumption in the product group. All regressions include product group fixed effects. Standard errors in brackets.

** Significant at the 1 percent level.

* Significant at the 5 percent level.

impact of consumption on creation and destruction to be the same across all groups, but we allow for a different constant across groups.

Table 7 shows the results from regressing net creation, creation, and destruction on total consumption growth. Obviously, because there is an accounting identity linking these variables, we are not looking for a structural relationship but simply to describe the comovement of each of these variables with overall growth across sectors. The first column shows that net creation rises significantly in periods where consumption growth is high. A one percentage point increase in sales growth is associated with a rise in net creation of 0.35 percentage points. This is suggestive evidence for models where firms have an incentive to defer implementation of the product until aggregate demand is relatively high as in Schmookler (1962) and Shleifer (1986). These models differ from the traditional Schumpeterian creative process which is independent of market demand.

It is also interesting to note that while net creation is strongly procyclical, it is primarily driven by the procyclicality of creation rather than the countercyclicality of destruction. The coefficient under the column “creation” can be interpreted as how much creation moves with sectoral consumption growth. An additional one percentage point growth in consumption of a particular product group is associated with 0.3 percentage point increase in the share of new goods. Creation is strongly procyclical and covaries more with demand in expansions than in contractions. The opposite is true for destruction rates. Destruction is countercyclical, but destruction responds more strongly in recessions than in booms. As suggested by Figure 1, however, most of the procyclicality of net creation comes from the procyclicality of creation.¹⁷ The rapid response of

¹⁷ Caballero and Hammour (1994) find the coefficient on job destruction rates to be almost double that of job creation, which is opposite to what we find in product markets. The larger response of destruction in recessions than in booms is similar to the pattern observed for job destruction. They interpret the process of creative destruction as one where new production units are being created, and outdated ones are being destroyed. However, they acknowledge that the creation process could also be interpreted as one of product innovation. We can use their model to interpret the implications of the cyclical patterns of product flows. Product creation has a strong “insulating” effect over product destruction because falls in demand are largely accommodated by a fall in the creation rate. This effect reduces the impact of a recession on existing products and helps dampen the response of product exit to cycles. This predicted

product creation to the cycle is suggestive of either a small cost of producing new products or that “product innovations” are stored until market conditions are suitable for their implementation.

II. Price Measurement

Our exploration of the data yielded several important stylized facts that will help us understand problems in price indexes and the measurement of welfare. First, the extent of product turnover within firms means that the bulk of the creative destruction process is driven by within-firm product creation and destruction. Second, new products systematically displace market share from existing products. Finally, net creation is strongly procyclical. These facts suggest that our ability to compare standard of living of consumers over time or to measure the severity of business cycles is likely hampered by the fact that conventional price indexes do not systematically adjust to product creation and destruction. We now turn to quantifying these forces.

Conventional methods of calculating price indexes are limited in the way they capture the impact of new and disappearing products. Consumer price indexes are based on price quotations for a fixed set of products over a long period of time (four to five years in the United States). Even when products are rotated into or out of the sample, the BLS does not correct for the quality upgrades that happen between rotations. The BLS methodology for performing these “scheduled rotations” involves field agents collecting price quotations for both the old and new sample in an overlapping month. For example, if cameras are scheduled to be resampled in January 2008, then the BLS field agents will collect prices for the old and new sample of cameras that month. For January, they will use the price inflation of the old sample. For February, they use the price inflation of the new sample. Any quality difference between samples is completely ignored. If three-megapixel cameras have replaced most two-megapixel cameras in the market, the BLS will not take this quality change into account because their new sample would not contain many two-megapixel cameras (whose prices were presumably falling). In particular, as will become clearer in the next subsection, without a measure of the quantities of the different cameras purchased in both periods, it is hard to make any correction for quality.

The BLS also allows products to be replaced within the scheduled sample period if the product disappears *from the particular store being surveyed*. These are called “forced substitutions.” Goods that may drop out of the sample because of forced substitutions may still be available to the consumer in many other stores, but they will be replaced in the CPI survey because the product is not available in the particular store visited by the field agent. There are two reasons why forced substitutions will miss the type of product creation and destruction we observe in the overall population of goods. First, a good that drops out of the sample is likely to be replaced with a good of comparable vintage rather than a new good. This is because the BLS agent has instructions to find a good whose characteristics match as closely as possible that of the exiting good. For example, if Nikon decides to replace a particular three-megapixel camera with a four-megapixel one, the BLS agent will attempt to find another three-megapixel camera with which to replace the discontinued camera. This implies that the new four-megapixel camera is likely not to be used in the updated CPI. Second, when substituting the new three-megapixel camera into the sample, the BLS typically uses a direct comparison method to compare the prices of the two different goods. This method treats any difference in prices between the two different

response of product destruction is consistent with the patterns observed in the data. Thus, the cyclical characteristics of product turnover suggest that the large insulating effect of product creation implies that recessions may play only a small role in terms of “cleansing” outdated or relatively unprofitable products. Not surprisingly, this is the opposite conclusion from that found when examining the implications of job flows.

three-megapixel cameras as an actual price change, even if it may reflect quality differences between products (see BLS 2005 and Bils 2004).¹⁸

Finally, explicit quality corrections by the BLS (e.g., using hedonics) are not standard in any of the categories of goods included in our data. The only exception in our sample is “Computer Software,” for which the BLS started using hedonic adjustments since January 1998. For these reasons, we believe that the quality biases that we measure in this paper are not accounted for in conventional price indexes like the US CPI.¹⁹

Before presenting the utility based framework we use to derive the bias in the CPI, we use the methodology developed by Davis and Haltiwanger (1992) as a way to illustrate the implications that product turnover has for the measurement of prices. This methodology allows us to examine price changes of all goods including those that are new or disappeared. We begin by expressing movements in prices in terms of mean growth rates. If we denote the price of UPC u in period t by p_{ut} , then the mean growth rate is defined as

$$(4) \quad g_{ut} \equiv \frac{p_{ut} - p_{us}}{\frac{1}{2}(p_{ut} + p_{us})}.$$

The mean growth rate is a monotonic transformation of the conventional growth rate.²⁰ For goods that were sold in both periods, i.e., “common” goods, we observe prices for both periods and therefore have no problem computing price changes in (4). This is not the case for new or disappearing goods. Without loss of generality, assume that the price of goods when they are not sold in the market (i.e., its “reservation” price) is infinite.²¹ In this case, mean growth rates are confined to the range of -2 to 2 . This allows us to nicely capture John Hicks’ (1969) simple intuition regarding new and disappearing products. Since Hicks we have interpreted the introduction of new products as a price *decline* from its reservation price to its market price, and a disappearing product is a product whose price *increases* from the market price to its reservation price.

In Figure 2, we plot the histogram of price changes for all goods over a four-year period, 1999 to 2003. We weight each bin in the histogram by its weight in consumption relative to the total value of consumption in 1999 *plus* 2003. For example, the height of the bar at “ -2 ” is given by the value of new goods relative to the total consumption of goods in 1999 and 2003.²² Notice that the height of the bar at “ -2 ” is roughly half the creation rate we found in Table 3 for the same time period. The same relation is true for the products that disappeared over this period and the destruction rate. The main feature of the figure that we want to emphasize is that to the extent that the value of goods created exceeds the value of goods that disappeared, a greater share of goods experience unobserved price declines than unobserved price increases. This fact is important because it establishes that the mean of the full distribution lies *to the left* of the mean of the

¹⁸ In particular, if a field agent is forced to substitute a 16-ounce bottle of Coke for a 32-ounce bottle, the BLS is not likely to treat that as a price increase if the 16-ounce bottle costs more than half as much, even though it is common for manufacturers to offer volume discounts. In this sense, the BLS methodology only crudely corrects for size changes.

¹⁹ The contribution of computer software to the overall quality bias we find is tiny (less than one percent) so we ignore this adjustment when presenting the main results.

²⁰ The conventional growth rate can be written as $2g/(2 - g)$.

²¹ If we had assumed a different value for the reservation price, it would affect only the location of the bars denoting birth and death but not their height. For example, if the reservation price were three times higher than the observed price, the bars would have been located at -1 and 1 . The graph would imply a qualitatively similar picture of the data.

²² Formally, the height of the bar at “ -2 ” in Figure 2 is given by $N_t/(V_s + V_t)$.




FIGURE 2. WEIGHTED PRICE GROWTH DISTRIBUTION
(4-year differences)

distribution of price changes of common goods. In other words, a price index that does not take birth and death into account is likely to be biased *upwards*.²³

Of course, the figure makes a number of simplifying assumptions that amount to treating all product introductions and disappearances symmetrically. As we already discussed, new bar codes can be introduced if secondary characteristics of the product are modified (like the volume of a soda can) or if the new product has truly distinct characteristics (digital versus film cameras). In the next section, we move away from the simplifying assumptions behind this figure and provide an exact calculation of this bias in the CPI—or equivalently, the welfare impact of product turnover allowing for a richer demand structure. In particular, we examine how different types of product creation and destruction affect a price index that is exact for a nested CES utility function in the presence of quality upgrading.

A. Creative Destruction and Quality Upgrading: An Exact Price Index

We now turn to formally deriving the bias present in conventional price indexes that ignore product turnover. This is equivalent to quantifying the impact that quality change, through product innovation, has on the welfare of consumers.

An important feature of our methodology is that we allow for a different impact of within-brand module product creation and destruction (i.e., new UPCs of an existing brand module) from that of across brand creation and destruction (i.e., new brands). In principle, we could allow different impacts in all brand modules, but this would be impossible in practice because many brand modules have just a few UPCs in them, and many product groups are composed of just a few brand modules. As a result, we decided to assume the same structure of substitutability within and across brand modules for each product group but allow it to vary across

²³ As will be clear in the next section, the bias does not necessarily need to be upward. As in Bart Hobijn (2002), we find that the quality bias in the current CPI can be downward if the prices (per unit of quality) of new goods are higher than those of the goods that disappear.

product groups. In particular, we model the impact of varieties on utility using a three-tiered CES aggregator. The first level describes how UPCs within a brand module in a particular product group enter the subutility function describing the representative consumer utility. The second describes how brands within a product group subutility function, and the last aggregates product groups.

This imposes a number of restrictions on the data. First, we constrain the within-brand module elasticity of substitution to be the same within any product group. This is, perhaps, easiest to understand in the context of an example. Consider the product group of “Crackers,” which contains brand modules like “Nabisco–Premium-Flaked Soda Crackers” and “Pepperidge Farm Goldfish–Cheese Crackers.” Our first restriction forces the elasticity of substitution *across* different brand modules within the same product group to have the same elasticity of substitution, but we allow the across-brand module elasticity to vary across product groups. Thus, two brand modules in a different product group (e.g., “Halls–Cough Drops” and “Herbon Glacial–Cough Drops”) will have a different elasticity of substitution than that of “Nabisco–Premium-Flaked Soda Crackers” and “Pepperidge Farm Goldfish–Cheese Crackers.” Our second restriction is that *within* brand module elasticities are also constrained to be the same as *within* brand module elasticities for other brand modules in the same product group. For example, all UPCs within the brand module “Nabisco–Premium-Flaked Soda Crackers” are equally substitutable within each other, and the elasticity is the same as the one for UPCs contained in the brand module “Pepperidge Farm Goldfish–Cheese Crackers,” but different than the elasticity of substitution for UPCs that compose Halls Cough Drops.

We now write down these restrictions formally. For expositional purposes, we begin by specifying the upper level utility function as:

$$(5) \quad U = \left(\sum_{g \in G} (C_{gt})^\rho \right)^{\frac{1}{\rho}},$$

where product groups are indexed by g , $\sigma = \rho/\rho - 1$ is the elasticity of substitution across product groups and G is the set of all product groups.²⁴ The set G is fixed over time ($G_t = G \forall t$), and so ρ plays no role in the analysis that follows.

We model the two lower tiers as follows:

$$(6) \quad C_{gt} = \left(\sum_{b \in \Psi_g} (c_{bgt})^{\rho_g^a} \right)^{\frac{1}{\rho_g^a}},$$

where c_{bgt} is the total quantity consumed of brand module b in product group g at time t , $\sigma_g^a = \rho_g^a/(\rho_g^a - 1)$ is the elasticity of substitution *across* brand modules within product group g , and Ψ_g is the set of all possible brand modules within a product group g . The set of existing brand modules in period t is a subset of this set, i.e., $B_{gt} \subset \Psi_g$, and can vary in each period. For future reference, it is useful to define the set of brands within group g that exist throughout all the time period (i.e., the “common brands”) as B_g where $B_g = B_{gt} \cap B_{gs}$.

Sales of a brand module, say, “Nabisco Premium-Flaked Soda Crackers” are aggregates of the different UPCs that make up the brand module:

²⁴ We do not need to assign quality weights to the aggregate consumption goods in our analysis because these weights would be redundant given that we later assign a quality weight to each individual variety (see equation (7)). Increasing the representative consumer’s preference for an aggregate good can always be represented as increasing the consumer’s preference for every variety of that aggregate.

$$(7) \quad c_{bg} = \left(\sum_{u \in \Omega_{bg}} (d_{ubgt} c_{ubgt})^{\rho_g^w} \right)^{\frac{1}{\rho_g^w}},$$

where c_{ubgt} is the consumption of UPC u of brand b of product group g in period t , $\sigma_g^w = \rho_g^w / (\rho_g^w - 1)$ is the elasticity of substitution *within* brands of brand b and product group g , and Ω_{bg} is the set of all possible UPCs that can exist in a particular brand module in product group g . The parameters d_{ubgt} play a crucial role in the analysis as they capture the different quality of UPCs that can exist in the market of a particular brand module in product group g . For example, “Nabisco Premium Unsalted Crackers” and “Nabisco Premium Multi-Grain Crackers” are two different UPCs of the brand module “Nabisco Premium-Flaked Soda Crackers,” and each of these UPCs have their unique quality parameter, d_{ubgt} . We define $U_{bgt} \subseteq \Omega_{bg}$ as the set of all new UPCs that have positive sales. It is also useful to define the set of UPCs in a brand module that are common over time as U_{bg} where $U_{bg} = U_{bgt} \cap U_{bgs}$.²⁵

As we show below, this three-tier specification allows for the introduction of a UPC to have a different impact on the price index for two reasons. First, as noted above, each UPC has its own quality level. Second, if the UPC belongs to a new brand, the elasticity of substitution used to value its introduction, σ_g^a , is different than the introduction of a new UPC within an existing brand, σ_g^w . In the data we expect that $\sigma_g^a < \sigma_g^w$ because within brand-group UPCs should be more substitutable than new brands.

The intuition for how we will measure the impact of quality changes can most easily be garnered from the unit cost functions. The minimum unit cost function of the subutility function in (7) is given by the following expression:

$$(8) \quad P_{bgt} = \left(\sum_{u \in U_{bgt}} \left(\frac{p_{ubgt}}{d_{ubgt}} \right)^{\sigma_g^w} \right)^{\frac{1}{\sigma_g^w}},$$

where p_{ubgt} is the price of UPC u of brand b in product group g in period t . For simplicity, define $\tilde{p}_{ubgt} = p_{ubgt}/d_{ubgt}$ as the quality-adjusted price. Here, it is important to remember that equation contains only those UPCs with positive sales in time t and that the exact price index of a brand module depends on the quality-adjusted prices of the UPCs contained within it. These properties are common to a number of different models, including the translog case.

Analogously, the minimum unit cost function of (6) can be denoted by

$$(9) \quad P_{gt} = \left(\sum_{b \in B_{gt}} (P_{bgt})^{\sigma_g^a} \right)^{\frac{1}{\sigma_g^a}}.$$

And the overall price index is given by

$$(10) \quad P_t = \left[\sum_{g \in G} P_{gt}^\sigma \right]^{\frac{1}{\sigma}}.$$

Equations (8)–(10) constitute the main building blocks for the calculation of exact aggregate price indices that follows.

We can now understand how a change in quality will be measured by our index. The difficulty in correcting for changes in quality stems from the fact that quality-adjusted prices are unobserved. However, by observing quantities purchased together with prices, we can use the

²⁵ More generally, we could define U_{bg} as any nonempty subset of common goods, but we always choose the largest subset.

information in the demand system to uncover the quality parameters. In particular, the CES demand system provides a simple way of recovering quality adjusted prices from observed consumer purchases. For UPCs of equal price, those with a higher quality will result in a lower quality-adjusted price and a higher market share. That is, the share of consumption of UPC u will depend directly on the quality-adjusted price:

$$(11) \quad s_{ubgt} = \left(\frac{p_{ubgt}/d_{ubgt}}{P_{bgt}} \right)^{1-\sigma_g^w}.$$

Suppose that UPC u' with positive sales at time $t - 1$ is replaced at time t by a higher quality UPC, u , i.e., $d_{ubgt} > d_{ubgt-1}$. If the higher quality UPC has a *lower* quality-adjusted price, then this implies that it will also have a *larger* market share. Alternatively, we can write the quality-adjusted price as

$$(11a) \quad \ln \frac{p_{ubgt}}{d_{ubgt}} = \frac{\ln s_{ubgt}}{1 - \sigma_g^w} + \ln P_{bgt}.$$

When we write the equation this way, it becomes immediately apparent that if one UPC is replaced by another with a larger market share, the new UPC must have a lower quality adjusted price. In other words, if we know the elasticity of substitution, we can infer the difference in quality-adjusted prices of two UPCs from the difference in their market shares. The seminal insight of Robert C. Feenstra (1994) was that one can use the market share of entering and disappearing goods to eliminate the quality parameters from the price index in (8) and write it *only* in terms of prices and market shares even when goods are constantly being replaced. We will now generalize and formalize this simple intuition.

It will be useful to keep track of two different sets of goods that have been introduced earlier in this section. First, the share of common goods in that particular brand module (i.e., $u \in U_{bgt}$) is defined as $s_{bgt}^{Com} = \sum_{u \in U_{bg}} s_{ubgt}$. Second, B_g is the set of brands in product group g that were consumed in both periods t and s . That is, $b \in B_g$ is a brand that was sold both in periods t and s , i.e., a brand that is common over this period. Similarly, we can define $s_{gt}^{Com} = \sum_{b \in B_g} s_{bgt}$ where $s_{bgt} = \sum_{u \in U_{bgt}} s_{ubgt}$.

The nested CES structure embedded in (5)–(7) implies that the exact price index that allows for product creation with different quality levels can be estimated using price and share data from individual UPC purchases. We are now ready to present the main proposition in the paper that defines the difference between an exact price index with quality change, EPI, that takes all price changes into account, including those of new and disappearing products, and a “conventional” exact price index, CEPI, that captures price changes only of the set of goods that existed on both periods (i.e., the common set of goods). Given the current methods used by the BLS of rotating products and adjusting for quality, we can use the CEPI as a benchmark for the US CPI.

PROPOSITION 1: For $g \in G$, if $d_{ubgt} = d_{ugbs}$ for $u \in U_{bg} = U_{bgt} \cap U_{bgs}$, $U_{bg} \neq \emptyset$ and there exists $b \in B_g$, $B_g \neq \emptyset$, then the exact price index for product group g with new and disappearing brands and UPCs is given by

$$\text{EPI}(\mathbf{p}_t, \mathbf{p}_{t-1}, \mathbf{x}_t, \mathbf{x}_{t-1}, B, U) = \prod_{g \in G} \left[\text{CEPI}_g \left(\frac{s_{gt}^{Com}}{s_{gs}^{Com}} \right)^{\frac{1}{\sigma_g^w - 1}} \prod_{b \in B_g} \left(\frac{s_{bgt}^{Com}}{s_{bgs}^{Com}} \right)^{\frac{w_{bg}}{\sigma_g^w - 1}} \right]^{w_g},$$

where weights w_{bgt} and w_{gt} are log-ideal CES Kazuo Sato (1976) and Yrjö Vartia (1976) weights defined as follows:

$$w_{gb} = \frac{\frac{s_{bgt} - s_{bgs}}{\ln s_{bgt} - \ln s_{bgs}}}{\sum_{b \in B_g} \left(\frac{s_{bgt} - s_{bgs}}{\ln s_{bgt} - \ln s_{bgs}} \right)} \quad \text{and} \quad w_g = \frac{\frac{s_{gt} - s_{gs}}{\ln s_{bs} - \ln s_{gs}}}{\sum_{g \in G_g} \left(\frac{s_{gt} - s_{gs}}{\ln s_{gt} - \ln s_{gs}} \right)}.$$

This result states that the exact price index with quality change is equal to the “conventional” exact price index, $\text{CEPI}_g(B_b, U_{bg})$, i.e., the exact price index of the UPCs that existed in both periods, multiplied by two ratios of the share of common goods over time. The first share ratio, $s_{gt}^{com}/s_{gs}^{com}$, captures the impact of creation and destruction of brands on the shares of common brands within a product group, and the second share ratio, $s_{bgt}^{com}/s_{bgs}^{com}$, captures the impact of creation and destruction of UPCs on the shares of common UPCs within a brand.

When the share of new brands in period t is larger than the share of disappearing brands in period s , this ratio is smaller than 1. In this case, the EPI will be smaller than the CEPI because the new brands have a lower price per unit quality than the disappearing brands; see equation (11). The conventional “fixed-good” index will miss this increase in quality because it occurs through the entry and exit of brands. In a similar way, the second share ratio, $s_{bgt}^{com}/s_{bgs}^{com}$, captures the role played by new and disappearing UPCs *within* brand modules that were common in both time periods. s_{bgt}^{com} equals the expenditure share in period t of UPCs that are available in both periods in a brand module (i.e., $u \in U_{bg} = (U_{bgt} \cap U_{bgs})$). Thus, $s_{bgt}^{com}/s_{bgs}^{com}$ is just the ratio of the share of common goods within a brand module in period t relative to share of common goods in period s . In this case, if one UPC is introduced and one is discontinued within a particular brand, then the EPI will fall relative to CEPI only if the incoming UPC has a lower quality-adjusted price than the outgoing UPC. The lower quality-adjusted price will imply that the new goods have a higher market share than the disappearing good used to have. This implies that the share ratio will be smaller than unity because the share of the common goods will have fallen. A share ratio smaller than one means that the EPI will be lower than the CEPI, and hence the conventional price index is biased upwards.

It is also important to recognize the importance of the exponent in this formula. The term w_{bg} simply captures the fact that one cares more about quality upgrading in brand modules that have larger market shares than ones with smaller shares. However, σ_g^w and σ_g^d play more subtle roles. As the elasticity of substitution rises, a given movement in the share of common goods over time will have a smaller effect on the inflation bias. The intuition is simple. If goods are highly substitutable, then the introduction of a new high quality good will have a big impact on the prices and quantities of existing goods. This means that the conventional price index will not be very biased since most of the welfare gain from the introduction of the new good can be elicited from examining what happens to common goods. In the limit as the elasticity of substitution approaches infinity, the inflation bias goes to zero because all quality changes are captured in price and quantity changes of existing goods.

In principle, we could use Proposition 1 to examine quality/new goods biases over any time horizon; however, two factors make some time horizons more sensible than others in practice. First, it makes sense to define periods t and s in years to prevent seasonal factors from driving product turnover. Thus, UPCs will be considered destroyed only if they were not purchased at any time during a yearlong period. Second, we need to decide how many years should separate the two periods. While this choice is inherently arbitrary, we decided to present biases calculated over a four-year period (1999–2003) and nine-year period (1994–2003) for two reasons. Because

goods tend to remain fixed within the CPI over four- to five-year periods, this is an appropriate time frame to use when comparing a fixed good index with our new-good adjusted index. Moreover, using a long time difference implies that the methodology is robust to product life cycle considerations driven, for instance, by marketing or fashion trends. Below we explain this fully with the use of two examples.

Proposition 1 requires that the taste or quality parameters for common goods must remain constant in start and end years of the sample. However, d_{ubgt} can vary over short horizons due to anything that might affect demand (e.g., marketing or fashion considerations). Indeed, in the next section, we will explicitly use these fluctuations in the quality parameters to achieve identification of elasticities of demand. The reason we assume immutable preferences *over long time horizons* when deriving our price indexes is that if the utility function is changing over time for either exogenous reasons (e.g., fashion) or endogenous reasons (e.g., marketing) then one cannot make sensible statements about how price changes affect welfare, nor can one derive exact price indexes because identical price vectors will yield different utility levels at different times. Thus, we are forced to assume that there is some time horizon over which preferences are fixed, as are the qualities of the UPCs in the set U_{bg} .

One can better understand the implications of our choice of time horizon by considering two examples of how the proposition captures the impact of different types of creation and destruction. First, let's consider the case of a new type of sunscreen that replaces an earlier type. If the new sunscreen is just a repackaging of last year's sunscreen without a noticeably different quality or price, then, *ceteris paribus*, the new sunscreen will have a market share equal to that of the old sunscreen. If this is true, then the share of common goods, i.e., $s_{bgt}^{com}/s_{bgs}^{com}$, will be unchanged and our measured quality bias from the replacement of the old model would be zero.

If, instead, the new sunscreen is priced identically but is of a higher quality than the old model, then, *ceteris paribus*, its market share will rise. This result comes directly from (11) because the new sunscreen will have a lower price per unit quality (p_{ubgt}/d_{ubgt}) than the old sunscreen. If this is the case, the higher share of the new good relative to the old good implies that $s_{bgt}^{com}/s_{bgs}^{com} < 1$. Our methodology will then imply that the EPI will be smaller than the CEPI, i.e., the conventional price index will have a positive quality bias. The extent of this bias will depend on the share of this brand in consumption, and how substitutable the two sunscreens are. Notice that if the new product has a higher price and a lower share, the price per unit quality (p_{ubgt}/d_{ubgt}) of the new model might be higher than that of the old model even if quality of the new good is higher. If this were the case consumer welfare would *fall* because the price per unit quality would have risen. Thus, the price index that takes quality into account may well be higher than the one that does not incorporate quality changes. This example highlights the importance of collecting quantity data to assess from market observables the price per quality of products.

We now consider a different example that captures the role played by the product life cycle. Imagine that new goods have a high market share when they are initially introduced but that the market share falls to two-thirds its initial level in the second year, one-third its initial level in the third year, and to zero in the fourth and subsequent years. For simplicity, assume that there is a constant flow of new goods whose market shares follow this pattern.

How does our methodology capture the welfare gains from these new products? Since we are examining the bias over a four-year period (1999–2003), our methodology will treat any good that was not available in 1999 but was consumed in 2003 as new. Proposition 1 indicates that the relevant share of new products used in the calculation of the bias between 2003 and 1999 is the share of these goods in 2003. Notice, however, that in 2003 the set of new goods includes goods that are three years old (and have low market shares) as well as those that are one and two years old (and have higher market shares). This implies that the actual share of new goods used is not the high share we observe at the beginning of the life cycle of new products but rather an *average* share over

the product life cycle of all new goods.²⁶ The same is true for disappearing goods. The share of disappearing goods will not be the low share that these products have exactly before they leave the market, but rather their average share during their life cycle. Thus, if the set of new goods has an average share over their life cycle that is larger than that of the disappearing products, we should expect to see a smaller quality-adjusted price index than the conventional price index. As discussed in the case of the sunscreens the opposite might be the case if consumers perceive that the price per quality of the new set of goods is not as low as that of the old set of goods.

III. Estimating the Elasticities

In order to compute the bias we need to obtain estimates of the “within” and “across” brand module demand elasticities which can then be used to estimate the relationship implied by (6) and (7). We rely closely on the methodology derived by Feenstra (1994) as extended by Broda and Weinstein (2006). Rather than providing a detailed description of the methodology, we present the intuition for it here and sketch out a few of the key equations.

The basic problem that we face is that we want to obtain a demand and supply equation using only information on prices and quantities. Obviously, we face the standard endogeneity problem for a given UPC. We do not know the slopes of the demand and supply equations that generated this data. The key is to realize that although we cannot identify supply and demand, the data does tell us something about the joint distribution of supply and demand parameters.

This intuition can be explained graphically. Wassily Leontief (1929) suggested that these data could be described by a set of demand (σ) and supply (ω) elasticities that can be represented by a hyperbola. This hyperbola can be built by finding the supply elasticity that minimizes the sum of squared errors for each demand elasticity. For example, assume σ_0 is the true elasticity of demand. Then if the price and quantity data points are as in Figure 3, it suggests that ω_0 may be a better fit to explain the data with the assumed demand elasticity than ω_1 . However, if we had assumed σ_1 is the true elasticity of demand, then ω_1 would likely be the best estimate. Because there are many σ s that can be the true demand parameter, we have not solved this problem, but the key insight, which we plot in Figure 4, is that the set of elasticities that maximize the likelihood function satisfy the hyperbola suggested by Leontief.

This is where we use Feenstra’s main insight. The panel nature of the data allows us to obtain a different hyperbola for each UPC. Figure 5 shows what happens if we take another UPC within that same brand. We can repeat the previous steps to compute that UPC’s own hyperbola. As long as demand and supply shocks are not drawn from the same distribution as the previous UPC, the new hyperbola will be different than the one in Figure 4. This illustrates the importance of two of our key identifying assumptions: first, that the elasticities of supply and demand are the same for each UPC within a product group, and second, that the relative variances of demand and supply shocks differ across bar codes so that the hyperbolas will also differ in shape. Given these assumptions, we can identify ω and σ by locating the point at which the two hyperbolas intersect. With more than two UPCs, the hyperbolas will not intersect at a single point in general, and so in this case we would pick the $\hat{\omega}$ and $\hat{\sigma}$ that minimizes a weighted distance from the intersections of the hyperbolas.

²⁶ Note that the share of new goods enters into Proposition 1 given that $s_{bgt}^{com} = 1 - s_{bgt}^{new}$ where s_{bgt}^{new} is the share of new goods within a brand module in period t . The share is exactly the average only in the special case where product life cycles are identical across new products, the pace of innovation is constant, and the life cycle of the product is less than four years.




FIGURE 3.

Formally, we do this by first modeling the supply and demand conditions for each good within a brand module cell. We estimate the demand elasticities, using the following system of differenced demand and supply equations:

$$(12) \quad \Delta^{k_{bg}} \ln s_{ubgt} = -(\sigma_g^w - 1) \Delta^{k_{bg}} \ln p_{ubgt} + \varepsilon_{bgt}^{k_{bg}}$$

$$(13) \quad \Delta^{k_{bg}} \ln p_{bgt} = \frac{\omega_g^w}{1 + \omega_g^w} \Delta^{k_{bg}} \ln s_{bgt} + \delta_{bgt}^{k_{bg}}.$$

Equation (12) is a transformation of the demand for a given UPC u in brand module b and product group g , derived from the CES demand in (11), and equation (13) is derived from the supply curve of that UPC. Both are expressed in terms of shares, where s_{ubgt} is the share of UPC u in brand module b within product group g . The equation for each UPC u is differenced with respect to time and a benchmark UPC of the same module, brand and product group. More specifically the difference operator we use for the shares and domestic prices is defined as $\Delta^{k_{bg}} x_{ubgt} = \Delta x_{ubgt} - \Delta x_{k_{bg}bgt}$. In this setup, the k^{th} good always corresponds to the largest selling UPC marketed in a particular brand module. The parameter ε_{ubgt} represents demand shocks to a particular UPC that might cause demand for that UPC to move relative to other UPCs marketed under the same brand module.²⁷ Obvious examples of such shocks are seasonal shifts in demand such as holidays, weather changes or diet changes that cause consumers to favor particular goods within a brand over others. Supply shocks are represented by δ_{bgt} and can be thought to include assembly line shocks that affect some UPCs within a firm's product mix

²⁷ Note that these shocks can be interpreted as high frequency taste shocks, i.e., changes in d_{ubgt} . Our only requirement is that these taste shocks are mean zero.




FIGURE 4.




FIGURE 5.

but not others. Both enter the expressions above in differenced form: $\varepsilon_{ubgt}^{k_{bg}} = \varepsilon_{ubgt} - \varepsilon_{k_{bg}bgt}$ and $\delta_{ubgt}^{k_{bg}} = \delta_{ubgt} - \delta_{k_{bg}bgt}$.

The k -differencing is critical to understanding our identification strategy. Any brand module level shocks—e.g., advertising, firm level supply shocks, or general demand shocks—are purged from the data and cannot affect our estimates. We are left with pure within-brand module variation that is likely to render $\varepsilon_{ubgt}^{k_{bg}}$ and $\delta_{ubgt}^{k_{bg}}$ uncorrelated, i.e., $E_t(\varepsilon_{ubgt}^{k_{bg}} \delta_{ubgt}^{k_{bg}}) = 0$. The second identifying assumption is that σ_g^w and ω_g^w are restricted to be the same over time and for all UPCs of a given brand module–product group (but is allowed to vary over product groups). As will become clear below this provides enough conditions to identify the main parameters of interest.

The derivation of the key moment conditions for identification has been explained in detail in Broda and Weinstein (2006a, b) so here we just provide an intuition for the main identification strategy. As in Feenstra (1994), it can be shown that using the panel nature of the dataset

and the assumption that demand and supply elasticities are constant over UPCs of the same product group we can obtain identification of the “within” demand elasticities. In particular, we can define a set of moment conditions for each brand module and product group by using the independence of the unobserved demand and supply disturbances for each UPC over time, i.e.,

$$(14) \quad G(\beta_g) = E_t(v_{ubg}(\beta_g)) = 0 \quad \forall u, b, \text{ and } g,$$

where $v_{ubg} = \varepsilon_{ubg}\delta_{ubg}$ and $\beta_g = \begin{pmatrix} \sigma_g^w \\ \omega_g^w \end{pmatrix}$. For each product group, g , all the moment conditions that enter the GMM objective function can be stacked and combined to obtain Hansen’s (1982) estimator:

$$(15) \quad \hat{\beta}_g = \arg \min_{\beta_g \in B} G^*(\beta_g)'WG^*(\beta_g) \quad \forall g,$$

where $G^*(\beta_g)$ is the sample analog of $G(\beta_g)$ stacked over all varieties u of a good g , W is a positive definite weighting matrix, to be defined below, and B is the set of economically feasible β_g which is common across importers and goods (i.e., $\sigma_g^w > 1$ and $\omega_g^w > 0 \quad \forall g$). Note that this implies that there are as many moment conditions as the number of UPCs in a particular product group g . We follow Broda and Weinstein (2006) in the way we implement this optimization.²⁸ Standard errors are obtained by bootstrapping.

The problem of measurement error in average purchase prices motivates our weighting scheme. In particular, there is good reason to believe that average prices calculated based on large numbers of purchases are better measured than those based on small numbers of purchases.²⁹ The use of the between estimate coupled with our need to estimate σ_g^w , ω_g^w and a constant means that we need data from at least four different UPCs in each product group and at least two time differences to identify β . Estimates of demand elasticities across brand modules are obtained using a similar procedure to the one just described. Instead of using UPC level data, we use market shares and unit prices at the brand level across modules and assume that the across brand elasticities in all modules within a product group are the same. We aggregate prices to form brand level prices by using the exact brand price index implied by the CES. Thus, we can obtain estimates for the two sets of elasticities per product group that are key to estimating the impact of product turnover. In the next section we describe these two sets of elasticities separately.

²⁸ We first use Feenstra’s approximate (15) to solve for β_g . In around 85 percent of the product groups this produces estimates in the feasible set. If this procedure renders imaginary estimates or estimates of the wrong sign we use a grid search of β s over the space defined by B . In particular, we evaluate the GMM objective function for values of $\sigma_g^w > 1$ and $\omega_g^w > 0$ at intervals that are approximately five percent apart. For computational easiness, we performed the grid search over values of σ_g and γ_g where γ_g is related to ω_{ig} in the following way: $\omega_g = \gamma_g/(\sigma_g(1 - \gamma_g) - 1)$. The objective function was evaluated at values for $\sigma_g \in [1.05, 131.5]$ at intervals that are 5 percent apart, and for $\gamma_g \in [0.01, 1]$ at intervals 0.01 apart. Only combinations of σ_g and γ_g that imply $\sigma_g^w > 1$ and $\omega_g^w > 0$ are used. To ensure we used a sufficiently tight grid, we cross-checked these grid searched parameters with estimates obtained by nonlinear least squares as well as those obtained through Feenstra’s original methodology. Using our grid spacing, the difference between the parameters estimated using Feenstra’s methodology and ours differed by only a few percent for those σ_{ig} and ω_{ig} for which we could apply Feenstra’s “between” approach.

²⁹ In the appendix of Broda and Weinstein (2007), they show that this requires us to add one additional term inversely related to the quantity consumed and weight the data so that the variances are more sensitive to price movements based on large value UPCs than small ones.

TABLE 8—SHARE RATIOS OF COMMON GOODS AND ESTIMATED ELASTICITIES OF SUBSTITUTION

Percentile	Per year share ratios		Elasticities of substitution	
	Within (average)	Across	Within (σ_g^w)	Across (σ_g^a)
1	0.53	0.52	3.3	1.8
5	0.67	0.74	4.0	3.0
10	0.77	0.83	4.7	3.3
25	0.86	0.92	6.5	4.5
Median	0.91	0.97	11.5	7.5
75	0.96	0.98	29.9	22.8
90	1.22	1.00	50.5	48.6
95	1.47	1.03	65.3	50.5
99	2.19	2.20	265.6	63.4
Number of product groups	122	122	122	122

Note: Per year share ratios are calculated as follows (weights and shares are defined in the text): Within (Average) = $(1/g_b) \sum_{b \in g} (S_{bg2003}^{Com} / S_{bg1999}^{Com})^{1/4}$; Across = $(S_{g2003}^{Com} / S_{g1999}^{Com})^{1/4}$ where g_b is the number of brand modules in a product group.

IV. Results

Our empirical and theoretical section has suggested two main implications of creative destruction for aggregate price measurement. The first concerns the magnitude of bias arising from an index that does not take into account the fact that many prices are moving to and from their reservation levels. The second concerns the cyclical nature of the bias given the cyclical patterns of product creation and destruction observed in the data. We will address each in turn.

A. Quality—New Goods Bias

Figure 2 suggested that the mean of the full distribution of price changes will be smaller than the mean of the price changes of a common goods index. The key question is by how much. We compute the bias relative to a CES price index computed over a set of common goods over time (i.e., the CEPI in Proposition 1). We think this is a reasonable benchmark because for common goods, the CEPI price index yields an almost identical rate of inflation as the Tornqvist index and the chained CPI.³⁰ Moreover, in Appendix F in Broda and Weinstein (2007), we show how a quarterly seasonally adjusted CEPI tracks the BLS's Food and Beverage CPI extremely closely.

Proposition 1 indicates that the magnitude of the bias depends on two factors: the ratios of the share of common goods over time and elasticities. The former indicates the relative importance of quality shifts in the data and the second tells us about how much of these shifts are being picked up in the conventional index. Table 8 presents the distribution of the per-year ratio of the share of common goods within and across brand modules.³¹ Formally, the share ratios we report are those that appear in Proposition 1. For expositional ease, we refer to $s_{gt}^{com} / s_{gs}^{com}$ as the across share ratio, and average of $s_{bgt}^{com} / s_{bgs}^{com}$ within a product group as the within share ratio. The median

³⁰ In Appendix G of Broda and Weinstein (2007), we present the biases of the Laspeyres, Paasche, CPI, chained CPI, and CES price indexes over four-year periods relative to the Tornqvist. As one can see from the table the Chained CPI, CES, and Tornqvist indexes yield inflation rates that differ by less than 1/100th of a percentage point.

³¹ Per year ratios are derived from four-year calculations. The reason for doing this is that in the computation of the quality bias we want to exclude the impact that high frequency turnover (e.g., like that coming from fashion products) can have on the quality bias.

(across all product groups) within per year ratio is 0.93. Over a four-year period the median ratio is 0.75.

This number is easiest to understand in terms of a simple example. Imagine a firm produced two goods in a brand module, one of which had a market share that was double the other. If the firm replaces the low selling good with a new product that now sells as well as the high selling good, this would generate a share ratio of 0.75 (i.e., $s_{bgt}^{com}/s_{bgs}^{com} = 0.5/0.67$). Thus, a ratio of 0.75 suggests that the typical new good has 50 percent higher market share than the good it replaces. The ratio across brands, i.e., $s_{gt}^{com}/s_{gs}^{com}$, is 0.87 over four years, or 0.97 per year, which is, as expected, smaller than $s_{bgt}^{com}/s_{bgs}^{com}$, but still less than one. This suggests that while the typical product group experienced new brand modules that had lower quality adjusted prices than previous brand modules, much of the quality improvement appears to have been happening within the product mix of particular firms.

The third and fourth columns of Table 8 show the distribution of estimated elasticities of substitution. The typical within-brand module elasticity is 11.5.³² The direct implication is that a one percent price decline of a UPC within a brand module causes its sales to rise by 11 percent. This would be the case if the various versions of, say, Nabisco Ritz Crackers are close substitutes with one another—a fairly plausible conjecture. This estimate is slightly higher than the typical demand elasticity found between different products of the same brand in marketing studies that range from four to seven; see Jean-Pierre Dube and Puneet Manchanda (2005) and Alan L. Montgomery and Peter E. Rossi (1999). As a sensitivity test of our benchmark estimate of the quality bias we will also use these typical elasticities found in the marketing literature to calibrate the size of the quality bias.

A second way to assess whether this number seems reasonable is to consider other estimates of the elasticity of substitution. Broda and Weinstein (2006) estimate the elasticity of substitution for US imports at various levels of aggregation. They find a typical elasticity of between three and four for the most disaggregated trade data (ten-digit Harmonized System categories). Clearly, products produced within the same brand module should be a lot more substitutable than imports from different countries within the same ten-digit sector, so it is comforting to see that our typical estimated elasticity is larger than their ten-digit elasticity.

A final reasonability check is to see if the within-brand module elasticities are larger than the across-brand module elasticities. This would be true if we believe that products marketed under a particular brand in a module are more substitutable than products across different brands in the same module. For example, it is likely that different types of Nabisco Ritz Crackers are more similar to each other than they are to Graham Crackers. When we estimate the across brand module elasticity of substitution, we see that the median elasticity of substitution across brand modules is 7.5, which is smaller than the median within brand module elasticity. While we can reject at all conventional levels of significance the hypothesis that the median across and within elasticities are the same, this might not be fully convincing since the median within and across elasticities do not correspond to the same module. In order to test if within brand elasticities are smaller than the across ones in the same module, we take the difference between the estimated within- and across-brand module elasticity in each product group. The median difference is 2.4, and this is statistically different from 0 at the 5 percent level (t -statistic equals 2.04). Thus, not only is the typical within elasticity higher than the typical across elasticity, this is also true within product groups.

The results, thus far, indicate that while there is significant product upgrading, the high levels of substitutability mean that a large share of this upgrading is likely to be captured in the movements of existing prices. However, Proposition 1 enables us to compute the actual bias

³² We could easily compute standard errors for the three-quarters of elasticities that were not grid searched. For these sectors the median within-brand module t statistic was 5.4 and the median across elasticity was 4.5.

TABLE 9—AGGREGATE QUALITY BIAS

Period	CES across brands within product groups, CES within brands				
	1994–2003	1999–2003			
Sigmas	Estimated (1)	Estimated (1)	$\sigma_g^w = \sigma_g^a = 7$	$\sigma_g^w = \sigma_g^a = 4$	$\sigma_g^w = 7$ and $\sigma_g^a = 4$
Per year quality bias	−0.9	−0.7	−1.1	−2.3	−1.5
Cumulative bias	−8.1	−2.8	−4.4	−9.2	−6.0

Notes: (1) Estimated sigmas are the 244 elasticities of substitution computed in this paper. Sigmas in the other columns were taken from marketing studies.

since we know all of the ratios of common shares and elasticities. When we do this, in Table 9, we find the conventional exact price index overstates inflation by 2.80 percentage points over the period 1999–2003 relative to the quality-adjusted index, or 0.69 percentage points per year. Over the period 1994–2003 the estimated bias is 8.1 percentage points or 0.91 percentage points per year.³³

The results of this paper are driven by two opposing forces. First, the high rates of product turnover that we find, and second the relatively high elasticities of substitution that we estimate among UPCs. We perform a number of sensitivity tests of these results in Table 9. We use the elasticities of substitution commonly found in the marketing literature as a way of checking the magnitude of our results. If we use the upper tail of the estimates of elasticities in this literature (which is around seven) for all our within and across elasticities to evaluate Proposition 1, we obtain a quality bias of around 1.1 percentage points per year, or almost 40 percent larger than our benchmark estimate. Moreover, if we use the low end of the typical range of elasticities for both within and across elasticities, the quality bias grows to over two percentage points per year. If instead we used the low estimate for the across elasticities and the high estimate for the within elasticities we get an estimate of the quality bias of around 1.5 percentage points per year. Thus, by using the estimates of the elasticities we compute in this paper, we generate an estimate of the quality bias.³⁴

In equation (5) we explicitly define nests based on the product groups that ACNielsen provides. While this has clear advantages, as one might expect goods within a product group to be more substitutable than outside a product category, this might not always be the case. In addition to the classification of product groups, ACNielsen also provides us with a “module” category that encompasses similar goods of different brands but is defined at a lower level of aggregation. As a way of testing how restrictive is this assumption, we performed the calculation of the bias changing the nests in (5) to be defined over modules as opposed to product groups. With this different assumption, we find the per year bias during the periods 1999–2003 to be 0.6 percentage points, only marginally smaller than our benchmark estimate. This different utility structure has the advantage of imposing a weaker demand structure over the data at the lower level of aggregation. However, we do not use this specification as our benchmark specification for two reasons. First, we do not have enough observations to compute the demand elasticities in about 20 percent of

³³ In Appendix H of Broda and Weinstein (2007), we present the 20 product groups that contribute the most to the quality bias. We do not present a confidence interval for the bias as it tends to produce a very narrow range and is extremely computer intensive to compute. For example, in Broda and Weinstein (2006) we find that the bias in the import price index between 1972 and 2001 was 1.2 percentage points per year with a 10–90 percentile confidence interval having a width of 0.3 percentage points or about plus or minus 12.5 percent. The reason the band is so narrow is that we use a vast amount of data, and the aggregate bias is a weighted average of the individual parameter standard errors. Hence, the confidence interval tends to decline with the number of observations and sectors.

³⁴ In general the bias will move with $1/(\sigma - 1)$, so using an elasticity of 13 will produce a bias estimate half as large as the elasticity of 7 we report.

TABLE 10—CYCLICALITY OF THE QUALITY BIAS AT THE PRODUCT GROUP LEVEL
OVER 4-QUARTER PERIODS

Weights Year and quarter dummies	LHS variable : Q4 bias		
	No No	Yes No	Yes Yes
Product group consumption growth	−0.057*** [0.010]	−0.106*** [0.011]	−0.097*** [0.011]
Constant	−0.019*** [0.001]	−0.017*** [0.001]	−0.012*** [0.002]
Observations	1,236	1,236	1,236
R ²	0.02	0.07	0.11

Notes: Q4 bias is the bias in the conventional exact price index over a four quarter period.
Standard errors in brackets.

*** Significant at the 1 percent level.

the brand modules.³⁵ Second, the average precision of the demand elasticities estimated at the module level falls substantially. Nonetheless, it is reassuring that we find that the main quality bias does not change much.

B. Cyclical Nature of the Bias

In Stylized Fact 4, we highlighted the cyclical patterns of creation and destruction and related it to the work of Shleifer (1986). Table 10 shows that the aggregate pattern is confirmed at the product group level. Here we regress the four-quarter bias by product group against the growth rate of sales in the product group over the same four quarters. We present different specifications (with and without sales weights, with and without year and quarter dummies). The coefficient on sales growth can be interpreted as the elasticity between sales growth at the product group level and the product group four-quarter bias. The last column shows that for every one percent growth in sales in a particular group, the bias in the conventional price index for that group increases by 0.1 percentage point.

Figure 6 shows the four-quarter biases computed on a rolling basis against the sales growth of the entire ACNielsen sample of products. The pattern observed supports the conjecture that the bias is cyclical. The bias moves between 0.32 and 0.71 percentage points over four quarters depending on the extent of sales. In particular, the lowest bias was recorded in the trough of the recession of 2001, and the peak was recorded in the fourth quarter of 2002. This procyclicality in the CPI bias suggests that real consumption is more volatile than is implied by national statistics. In particular, real consumption of the products in our sample was 0.4 percentage points more volatile than official real consumption.³⁶ For food consumption, four-quarter real growth rates during the mild 2001 recession fluctuated from 4.3 percent in 2000:I to 1.37 percent in 2001:II. This suggests that real consumption volatility was at least 10 percent more volatile after controlling for quality.

³⁵ To make sure that the difference in the bias was not driven by different set of modules in our benchmark estimate and in this new estimate, we also performed our benchmark estimate without the modules that were dropped in the new estimation. The per year bias without those modules was still 0.7 percentage points.

³⁶ This is true if the personal consumption expenditure price index (PCE) is used to deflate the nominal consumption series.




FIGURE 6. THE CYCLICAL NATURE OF THE QUALITY BIAS

In order to assess how much these biases might matter, we use the food and beverages component of the CPI from 1990 to 2007 and estimate the implied bias due to cyclical and other factors based on the results in Table 10. Figure 7 summarizes our main findings by combining the benchmark level of the quality bias found in the previous subsection and the cyclical nature of the CPI bias in this subsection. The solid line shows the four-quarter seasonally adjusted inflation rates from the BLS Food and Beverages CPI component from the first quarter of 1990 to the third quarter of 2007. The dashed line shows the same series adjusted for the benchmark quality bias we document in this paper. The average quality bias is assumed to be 0.8 percentage points. We use four-quarter real consumption growth rates for food reported by the Bureau of Economic Analysis (BEA) and the information in Table 10 to compute the cyclical nature of the bias. Formally, we use the following equation to estimate the bias over this period: $\text{Quality Adjusted Inflation} = \text{BLS Food and Beverage CPI Inflation} - 0.8 + (\text{Real Food Consumption Growth} - 1.9) \times 0.1$, where 1.9 is the average percent growth over the period between 1990 and 2007. For the years that overlap with our sample period, we observe the bias moving by 0.35 percentage points, almost the same cyclical nature as we found using the actual bias in Figure 6. In this extended sample, we see the largest bias during the 2000 boom, where the bias reaches over one percent, and the smallest bias at the end of the 1991 recession, where the bias was just below 0.5 percent. We highlight the largest and smallest biases in the graph.

V. Conclusion

Since Schumpeter completed his classic work, economists have known that economic welfare depends heavily on the creation and destruction of products. However, statistical agencies charged with measuring the cost of consumption have used methodologies that largely ignore the creative-destruction process. Our standard inflation indexes are computed using price surveys of existing goods with little or no information about the quantitative importance of these goods.




FIGURE 7. BLS AND QUALITY-ADJUSTED FOOD AND BEVERAGE CPI INFLATION
USA 1994–2003

Note: Quality adjusted inflation = BLS inflation – $(0.8 + (\text{Real food cons. growth} - 1.9) \times 0.1)$.

The inability of statistical agencies to systematically adjust for the destruction of obsolete goods and the creation of new and improved ones means that there are likely to be substantial differences between a true cost of living index and official price indexes.

This paper is the first large scale examination of product creation and destruction and its implications for price measurement. Using a new dataset that covers around 40 percent of all expenditures on goods in the CPI, we conclude that there is an upward bias of around 0.8 percent in this fraction of the CPI. The bias is also strongly procyclical, which suggests that business cycles are more pronounced than is typically reported in official statistics.

Considerably more work needs to be done before we can definitively measure the magnitude of the overall quality bias of the CPI. First, we need to have more and better data about sections of the CPI that are not covered in our sample. In particular, we need to have information on both prices and quantities in these sectors. Second, we need to find ways of testing the implications of functional form assumptions on our estimates of quality bias. In particular, in this paper our choice of the CES aggregator was based on the fact that it is a good approximation of a superlative index for common goods, it is useful in comparing our results with existing theory, and that we have methodological limitations on applying more flexible structures. However, the prominence and cyclicity of quality upgrading indicates that the biases and measurement issues we have documented are likely to be robust features of the data.

REFERENCES

- Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott. 2006. "Multi-Product Firms and Product Switching." National Bureau of Economic Research Working Paper 12293.
- Bils, Mark. 2004. "Measuring the Growth from Better and Better Goods." National Bureau of Economic Research Working Paper 10606.
- Bils, Mark, and Peter J. Klenow. 2001. "Quantifying Quality Growth." *American Economic Review*, 91(4): 1006–30.
- Blanchard, Olivier Jean, and Peter Diamond. 1990. "The Cyclical Behavior of the Gross Flows of U.S. Workers." *Brookings Papers on Economic Activity*, 1990(2): 85–143.

- Boskin Commission.** 1996. "Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index." <http://www.ssa.gov/history/reports/boskinrpt.html>.
- Broda, Christian, and David E. Weinstein.** 2004. "Variety Growth and World Welfare." *American Economic Review*, 94(2): 139–44.
- Broda, Christian, and David E. Weinstein.** 2006. "Globalization and the Gains from Variety." *Quarterly Journal of Economics*, 121(2): 541–85.
- Broda, Christian, and David E. Weinstein.** 2007. *Product Creation and Destruction: Evidence and Price Implications*. National Bureau of Economic Research Working Paper 13041.
- Caballero, Ricardo J., and Mohamad L. Hammour.** 1994. "The Cleansing Effect of Recessions." *American Economic Review*, 84(5): 1350–68.
- Caballero, Ricardo J., and Mohamad L. Hammour.** 1996. "On the Timing and Efficiency of Creative Destruction." *Quarterly Journal of Economics*, 111(3): 805–52.
- Campbell, Jeffrey R.** 1998. "Entry, Exit, Embodied Technology, and Business Cycles." *Review of Economic Dynamics*, 1(2): 371–408.
- Davis, Steven J., and John C. Haltiwanger.** 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Quarterly Journal of Economics*, 107(3): 819–63.
- Dube, Jean-Pierre, and Puneet Manchanda.** 2005. "Differences in Dynamic Brand Competition across Markets: An Empirical Analysis." *Marketing Science*, 24(1): 81–95.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson.** 1988. "Patterns of Firm Entry and Exit in U.S. Manufacturing Industries." *RAND Journal of Economics*, 19(4): 495–515.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson.** 1989. "The Growth and Failure of U.S. Manufacturing Plants." *Quarterly Journal of Economics*, 104(4): 671–98.
- Feenstra, Robert C.** 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review*, 84(1): 157–77.
- Gallo, Anthony.** 1995. "First Major Drop in Food Product Introduction in the Last 20 Years." *Journal of Food Distribution Research*, 26(1): 9–13.
- Ghironi, Fabio, and Marc J. Melitz.** 2005. "International Trade and Macroeconomic Dynamics with Heterogeneous Firms." *Quarterly Journal of Economics*, 120(3): 865–915.
- Hall, Robert E.** 2004. "The Labor Market Is Key to Understanding the Business Cycle." Unpublished.
- Hicks, John.** 1969. *A Theory of Economic History*. Oxford: Oxford University Press.
- Hobijn, Bart.** 2002. "On Both Sides of the Quality Bias in Price Indexes." Federal Reserve Bank of New York Staff Report 157.
- Jovanovic, Boyan.** 1993. "The Diversification of Production." *Brookings Papers on Economic Activity, Microeconomics*, 1993(1): 197–247.
- Lebow, David E., and Jeremy B. Rudd.** 2003. "Measurement Error in the Consumer Price Index: Where Do We Stand?" *Journal of Economic Literature*, 41(1): 159–201.
- Leontief, Wassily.** 1929. "Ein Versuch Zur Statistischen Analyse von Angebot und Nachfrage." *Weltwirtschaftliches Archiv*, 30(1): 1–53.
- Montgomery, Alan L., and Peter E. Rossi.** 1999. "Estimating Price Elasticities with Theory-Based Priors." *Journal of Marketing Research*, 36(4): 413–23.
- Rigobon, Roberto.** 2003. "Identification through Heteroskedasticity." *Review of Economics and Statistics*, 85(4): 777–92.
- Sato, Kazuo.** 1976. "The Ideal Log-Change Index Number." *Review of Economics and Statistics*, 58(2): 223–28.
- Shimer, Robert.** 2005. "Reassessing the Ins and Outs of Unemployment." Unpublished.
- Shleifer, Andrei.** 1986. "Implementation Cycles." *Journal of Political Economy*, 94(6): 1163–90.
- Streitwieser, Mary L.** 1992. "The Extent and Nature of Establishment-Level Diversification in Sixteen U.S. Manufacturing Industries." *Journal of Law and Economics*, 34(2): 503–34.
- US Department of Labor.** 1997. *BLS Handbook of Methods*. Washington, DC: US Government Printing Office.
- Vartia, Yrjo.** 1976. "Ideal Log-Change Index Numbers." *Scandinavian Journal of Statistics*, 3(3): 121–26.