Product Creation and Destruction: Evidence and Price Implications — Source link

Christian Broda, David E. Weinstein

Institutions: University of Chicago, Columbia University

Published on: 01 Jun 2010 - The American Economic Review (National Bureau of Economic Research)

Topics: Price level, Product (category theory), Product market, Price index and Index (economics)

NBER WORKING PAPER SERIES

PRODUCT CREATION AND DESTRUCTION: EVIDENCE AND PRICE IMPLICATIONS

Christian Broda
David E. Weinstein

Working Paper 13041
http://www.nber.org/papers/w13041

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2007

The authors wish to thank the NSF (grant #0214378) 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 vice-president of Pricing 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 Antoniades 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. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

Product Creation and Destruction: Evidence and Price Implications
Christian Broda and David E. Weinstein
NBER Working Paper No. 13041
April 2007
JEL No. E21,E31,E32

ABSTRACT

This paper describes the extent and cyclicality of product creation and destruction in a large sector of the U.S. economy and quantifies its implications for the measurement of consumer prices. We find four times more entry and exit in product markets than is typically found in labor markets because most product turnover happens within the boundaries of the firm. Net product creation is strongly pro-cyclical, but contrary to the behavior of labor flows, it is primarily driven by creation rather than destruction. High rates of innovation are also accompanied by substantial price volatility of products. These facts suggest that the CPI deviates from a true cost-of-living index in three important dimensions. The quality bias that arises as new goods replace outdated ones causes the CPI to overstate inflation by 0.8 percent per year; the cyclicality of the bias implies that business cycles are more volatile than indicated by official statistics; and finally, sampling error is sufficiently large that over the last 10 years policymakers could not statistically distinguish whether quarterly inflation was accelerating or decelerating 65 percent of the time.

Christian Broda
University of Chicago
Graduate School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
cbroda@chicagogsb.edu

David E. Weinstein
Columbia University, Department of Economics
420 W. 118th Street
MC 3308
New York, NY 10027
and NBER
dew35@columbia.edu

I. Introduction

Over the last twenty 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., Shleifer (1986), Caballero and Hammour (1994) and Ghironi and Melitz (2005) among others) but economists have also examined the key role played by new and better products for long-run growth (e.g., Romer (1987), Grossman and Helpman (1991), Aghion and Howitt (1992) and Klette and 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.[1] In this paper, we document the nature, extent and cyclicality of product entry and exit in the U.S. with special attention to the implications that it has for the measurement of prices. In particular, we quantify the biases that arise in "fixed-good" 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 barcodes 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 only have access to firm-level data. By matching barcodes with firms, we document the multi-product 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., Dunne, Roberts and Samuelson (1988, 1989) and Davis and Haltiwanger (1992)). In a typical year, 40 percent of household's expenditures are in 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 pro-cyclical, with more products being introduced in expansions and in product categories that are booming. Destruction of goods is counter-cycle, although its magnitude is quantitatively less important. This is suggestive of models where firms have an

[1] This gap has emerged largely because of data availability. The lack of empirical evidence was not resolved even with the appearance of scanner databases specific to particular stores, as store-specific data are not appropriate to analyze the extent of creation and destruction of products for the consumer.

incentive to defer implementation of the product until aggregate demand is relatively high (as in Schmookler (1962) and Shleifer (1986)).[2] While labor market data suggest that job destruction responds more to cyclical movements than job creation (c.f. Blanchard and Diamond (1990), Campbell (1998), Davis and Haltiwanger (1996)), we find the opposite to be true in product markets. Product creation moves more than product destruction with business cycle fluctuations.[3] The high degree of product churning is also associated with substantial price volatility as products move through their life cycle.

These new facts have important implications for the measurement of prices. They suggest that the CPI may deviate from a true cost-of-living index in three important dimensions. First, the creation and destruction of products can hide quality upgrades that are not captured by indexes that are based on the prices of a fixed set of goods over time.[4] Second, if the quality bias has a pro-cyclical component it can augment the volatility of real variables over the cycle. Finally, since the CPI is based on the prices of a sample of goods available in the market it is subject to sampling error. Given the high degree of price volatility of individual goods that we observe in the data, the potential for sampling error is large. The core of the paper is about quantifying these three implications of product churning for the measurement of the true cost-of-living.

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 averages 0.6 – 0.9 percent per year depending on the aggregation methodology. This implies that inflation was around 7 percentage points lower than suggested by the CPI over the period studied (1994 – 2003). The fact that the quality available for consumers rises as a result of

---

[2] As noted by Shleifer, 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 finds 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. In future work, we will systematically compare the cyclical predictions of different models of product innovation.
[3] In terms of Caballero and Hammour (1994), we find a strong "insulating" effect of creation that dampens the response of destruction to the business cycles.
[4] We discuss the role played by "forced" substitutions later in the paper. The BLS performs regular product substitutions every 4 to 5 years, so except for forced substitutions the CPI is based on a fixed set of goods within each product group over this period.

the creative destruction process is a salient feature of Schumpeter's work.[5] The bias is also strongly pro-cyclical, which suggests that business cycles are more pronounced than is typically reported in official statistics. For example, during the 2001 recession real consumption in the products in our sample fluctuated by 0.4 percentage points more than suggested by national statistics. This cyclicality arises because 30 cents of each additional dollar of consumption is spent on new higher-quality goods which are ignored by conventional ("fixed-goods") price indexes. Lastly, we calculate the standard deviation of inflation rates at different frequencies based on sampling error. We find the sampling error in 1-quarter seasonally adjusted inflation rates to be 0.5 percent (annualized rates). This suggests that over the last 10 years policymakers could not statistically distinguish whether quarterly inflation was accelerating or decelerating 65 percent of the time.

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 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 directly 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

---

[5] 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).

short of computing price biases.[6] 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. Lebow and Rudd (2004) survey the recent improvements on price measurement and conclude that the estimates of the quality bias for less than 10 percent of the CPI were based on “at least a moderate degree of hard evidence”. The rest comes from either inadequate evidence or is entirely subjective. A contribution of our paper is to use 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.

II. Data Description

II. 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 who then scan in the purchases of every good with a barcode.[7] These households represent a demographically balanced sample of households in 23 cities in the US.[8] Barcodes are concentrated in grocery, drugstore and mass-merchandise sectors. Overall, the database covers around 40 percent of all expenditure on goods in the CPI.

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 BLS and other statistical agencies. Our database covers

---

[6] 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 5 percent per year in recent periods. In earlier work, Bils and Klenow (2001) used the US Consumer Expenditure Survey to quantify the quality growth in 66 consumer durable goods.

[7] The 1994 data is based on a sample of 40,000 households; the data for 1999-2001 encompasses 55,000 households, and the data for 2002 and 2003 represents 61,000 households.

[8] 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.

approximately 700,000 different goods purchased at some point by our household sample, while the BLS sample for the entire CPI only covers approximately 85,000 goods.[9] 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 only working with a sample that is less than 5 percent as large.[10]

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 are not restricted to particular outlets, our basket of goods is more in line with those of the representative household in the U.S. than that of other studies or even the BLS survey of prices.[11] Moreover, by treating each product that enters or exits the basket of goods consumed by our sample of 55,000 households as a new or disappearing product in U.S. consumption we can capture even the smallest of product innovations in a large number of sectors.[12] 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.[13]

---

[9] 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 US to be 30,000 (see Gallo, 1997).

[10] A sample list of 400 general categories included in the database appears in Appendix A.

[11] Fenwick et al. (2003) provides an interesting measure of how the lack of quantity data at the product level can make the products sampled by statistical offices not representative of consumer purchases. They compare the sales weights of top selling consumer durable goods (e.g., top 10 models of TVs) with the percentage of price quotes of these goods used to calculate UK's retail price index (RPI). They find large differences between these two measures suggesting that the products used in the calculation of UK's RPI may not be representative of basket of goods consumed.

[12] This information is especially valuable for a number of reasons. From the "experimentation" literature we know that the most active segments of the market are typically small in size. Studies based on product surveys or data based on publicly traded firms typically ignore or cannot fully capture the information in these segments.

[13] As we will show in the next section, it is very infrequent for a product that drops from the sample to reappear in a later quarter, suggesting that the product has truly disappeared from the market place. The few cases in which this happens only affects 1-quarter comparisons, as they are driven by seasonal factors. Even at this high frequency, if we drop products that re-enter after a period of being out of our sample from our calculation, the creation and destruction measures are effectively unchanged.

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 only observe the prices 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, the use of barcode data means that our empirical construct of a good is extremely close to the theoretical concept, in contrast to studies that use more aggregated industry information. This allows 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 a store with Scantrack technology, 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. 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 refer to as "barcodes") purchased by the representative household at the quarterly level for 6 years: 1994 and 1999-2003. The data was 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 re-weights the averages so that price paid and the quantity purchased is representative of the U.S. as a whole.[14, 15]

Before proceeding further, it is worth taking a moment to review precisely what this data measures. Firms have strong incentives to label their products with UPCs because these machine readable identifiers dramatically reduce the costs for the store to stock and sell items. In order to obtain a barcode, a company must register with the Universal Code Council, which costs around

---

[14] We also ran all of our tables and regressions with the raw (unadjusted) data and the main results of the paper are unaffected.

[15] One concern about the database is that results might be driven by the growth in the number of barcodes *per se*. In Appendix B, we show two facts that suggest that this is not the case. First, the share of goods purchased with a barcode as a share of total 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.

$750 dollars plus an annual maintenance fee of $150 per UPC.[16] 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.

Although it is difficult to enforce how a company uses a barcode, most industry experts strongly caution firms not to use the same barcode on more than one product. 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 barcode 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. Changing the slogan on a Heinz ketchup bottle does not require a new barcode, but changing the size of the bottle does. In other words, it is safe to assume that product quality (insofar as it might change the price paid by consumers) will change if and only if the barcode changes.

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 Multi-Vitamins but also other modules like "Kids Vitamins".[17]

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

[16] Prices quoted from http://www.cummingsdesign.com/bar_codes101_UCC_App.htm
[17] In Appendix C we provide a second example of the levels of aggregation in our database that helps explain how we identify the elasticities we estimate in section V.B.

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 category contains about 9 modules, the average module contains 46 brands, and the average brand-module contains about 34 products.

II. B. Stylized Facts

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

Stylized Fact 1: The Importance of Multi-Product 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 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 6-digit brand identifiers (e.g. Advil and Robitussin).

By matching the UPC identifier with firm code, 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 8 different UPCs in 2 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 multi-product 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 4th 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 U.S. firms or plants (e.g., Streitweiser

(1991), Jovanovic (1993) and Bernard, et al (2006)), aggregation up to the firm level results in a substantial loss of information about the process of creative destruction.[18] As we will see below, most of the product entry and exit happen within existing firms.

Stylized Fact 2: The Significance of Product Entry and Exit

Price indexes in all countries are essentially based on prices of individual goods that are common over long periods of time (at least over 4 or 5 years).[19] As a result, only the set of goods that exist in the base period and the current period are included in the index. This is a potential problem for price measurement since goods whose prices move to and from their reservation price are excluded from the analysis but are part of the choice set available to consumers at some point.

We can understand the severity of the problem by employing the methodology developed by Davis and Haltiwanger (1992). We need to keep track of two types of price movements: price changes of common goods and the price movements of goods that are created or destroyed. For common goods, we observe prices for both periods and therefore have no problem computing the price change. For new and disappearing goods, we assume that the reservation price is infinite. This assumption only affects where in the graph we describe below we locate the bars associated with the birth and death of products.

In order to keep track of both of these forces, it is useful to express 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

$$g_{ut} = \frac{p_{ut} - p_{ut-1}}{\frac{1}{2}\left(p_{ut} + p_{ut-1}\right)}. \tag{1}$$

The mean growth rate is a monotonic transformation of the conventional growth rate.[20] Given that new and disappearing goods have infinite reservation prices, mean growth rates are confined to the range of -2 to 2. In the remainder of this section, we will use the term "growth rate" to refer to the rate computed according to equation (1).

---

[18] Though marginal to this paper, it is interesting to note that a 1 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.

[19] In the case of the U.S. CPI regular substitutions are currently scheduled every 4 years. We will comment on the role that "forced" substitutions play in the U.S. CPI below.

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

In Figure 1A, we plot the distribution of one-quarter price changes for all goods in our sample over the period 1999-2003. One of the features of the data that immediately jumps out from this plot is the large number of products that die and are born during a typical quarter. The bar at “-2” indicate that 19 percent of all goods available in both quarters are new (i.e., did not exist the previous quarter) and the bar at “+2” implies that 18 percent of goods disappear (i.e., do not exist the following quarter).[21] The fact that close to 40 percent of goods in a quarter are moving to and from their reservation price and are not being measured in conventional indexes means that strong assumptions are needed in order to provide confidence that the moments of the trimmed distribution of products match those of the full distribution.

An obvious critique of the previous plot is that we weight all goods equally regardless of their importance in consumption and, most likely, utility. A simple correction for this is to weight the data so that the height of the bars represents the share of total consumption rather than the share of total products. We do this in Figure 1B using a simple average of consumption at time t and t-1. Three results are immediately apparent. First, the distribution of price changes of common goods is far more compressed when we weight by value. About 80 percent of all price changes now fall in the range of plus or minus 10 percent. Second, the relative importance of birth and death is diminished substantially. Extreme values are about 5 percent of total consumption.

The third point is more subtle but has very important implications. If we look at the inflation rates of common goods, we see that there is a large amount of variance in individual price changes across goods. If we take the sales weighted average of absolute changes in the prices of common goods, we find that the average percentage change in prices is 8.8 percent in any quarter. This number is very much in line with price changes in the set of goods used in the CPI. Klenow and Kryvtsov (2005), for example, find that the comparable *monthly* number in the CPI is 13.3 percent. This suggests that the high level of price volatility is not a unique feature of the ACNielsen dataset, but rather a more general feature of all samples. Indeed, our dataset seems to exhibit one-third less price volatility than the BLS’s *Commodities and Services Survey*. A more conventional way of expressing this volatility is in terms of confidence intervals and

[21] If we had assumed a different value for the reservation price, it would only affect the location of the bars denoting birth and death but not their height. For example, if the reservation price were 3 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.

standard deviations. Ninety percent of all price changes in our dataset lie between -19 percent and 20 percent in a quarter. However, due to the existence of fat tails in the price distribution, the standard deviation of price changes is 20 percent per quarter. This large variance in quarterly inflation rates will be an important ingredient in our explanation of how sampling error affects the CPI methodology.

Another way of looking at the data is to take long differences (i.e. 1999-2003). Figures 2A and 2B present these results. Over this five year period, products whose prices moved to or from their reservation price account for over 60 percent of the sample. In the figure weighted by sales, the entry and exit of products is an unmistakably feature of the distribution. Clearly any price index based on a common set of goods is likely to be problematic because it ignores a sizable share of the sample whose prices are having large changes. Over 25 percent of the goods (in average sales-weighted terms) either were born or disappeared during the 4-year period. Since the weight used to scale this figure is an average taken over both *t* and *t*-1, new and exiting goods always have a weight equal to zero in one period. Thus, their importance in any one time period is roughly twice as large. In other words, this same data implies that 37 percent of the expenditures in 2003 were on goods that did not exist in 1999, and 18 percent of expenditures in 1999 were on goods that did not survive until 2003. Moreover, 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 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.

Table 3 summarizes the extent of product creation and destruction using weighted and un-weighted measures at different frequencies. The first column presents data on entry and exit between 1994 and 2003, the second column reports the same numbers between 1999 and 2003, and the third column presents the median numbers for each year between 1999 and 2003. In each case, we report the following measures of product creation and destruction: [22]

$$(2)\quad \text{Entry Rate}\left(\text{t}\right)=\frac{\text{\# New UPCs (t)}}{\text{\# All UPCs (t)}} \quad ; \quad \text{Exit Rate}\left(\text{t-1}\right)=\frac{\text{\# Disappearing UPCs (t-1)}}{\text{\# All UPCs (t-1)}}$$

[22] Notice that if the value of overall UPCs does not change over time, the bar for "birth" ("death") in figures 1A and 2A equals half the values for entry rate (exit rate), while the bars for "birth" ("death") in figures 1B and 2B equals half the values for the creation share (destruction share).

$$\text{Creation}\left(\mathrm{t}\right)=\frac{\text{Value of New UPCs (t)}}{\text{Total Value (t)}}\quad ; \quad \text{Destruction}\left(\text{t-1}\right)=\frac{\text{Value of Disappearing UPCs (t-1)}}{\text{Total Value (t-1)}}$$

Table 3 reveals that almost 80 percent of the products that existed in 2003 were not around in 1994. These new products comprised 64 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: 37 percent of expenditure in 1994.[23] This suggests that the new UPCs systematically displaced market share from the UPCs that were common throughout this period. This can also be summarized in terms of the ratio of the shares of common UPCs. The last row in this table shows that this ratio is below unity, i.e., the share of common UPCs in 1994 was larger than in 2003. As we will emphasize in the next section, this is an important indication that new products are of a higher quality than the products that exited.

The table also reports the average size of the entering and exiting UPCs relative to the average size of common UPCs. If we look at the third column of the table, we can see the relative sizes of UPCs that did not exist the year before. Here we see that the market share of the average new UPC is 30 percent as large as the share of common UPCs. For the average exiting UPC, this share is 9 percent. This suggests that new UPCs tend to be larger than exiting ones but both are smaller than common UPCs by a large margin. Over longer time periods the differences diminish somewhat but do not disappear – the share of entering products rises to 56 percent of that of existing products. Interestingly, the share declines at 9-year differences, which is suggestive of a life-cycle pattern for the products in our sample. We will return to this feature of the data when describing the fourth stylized fact.

It is useful to compare our results with that of studies of firm and establishment turnover as 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 5 years, they find that on average 10 percent of all manufacturing output in a particular census year came from plants that were not present in the previous census, while plants that were

---

[23] 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, we can drop the UPCs that were purchased by the smallest number of individual households to understand the impact of measurement error. In Appendix D we replicate table 3 dropping the UPCs that were purchased by less than 20 households. The measures of product creation and destruction are only marginally unaffected. Contrary to firm-level studies, the un-weighted measures of entry and exit are only slightly smaller than the ones including the full sample of UPCs.

present in one census but disappeared in the following census accounted for 14 percent of output.[24] 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 exiting products over a 5-year period in Table 3 (i.e., 56 percent). In other words, there is four times more product creation and destruction than firm 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 1-year period, roughly 3 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.

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. We also report the correlations of entry and exit across modules and across time within modules.

Interestingly, we observe most of the product entry and exit occurring within the boundaries of the firm. Table 4 reports the extent of firm entry and exit. At 1-year frequencies, 1 percent of all consumption expenditures come from new firms. Comparing the extent of firm creation to that of product creation (Table 3) suggests that 92 percent of product creation happens within an existing firm, and 97 percent of product destruction happens within existing firms. At 4-year frequencies the comparable numbers are 82 and 87 percent, respectively. This implies that over a 4-year period, 18 percent of the value of overall consumption is coming from products of completely new firms, and 13 percent of product exit is happening because firms disappear. In our data, product entry and exit is 5 times as important as firm entry and exit. Table

[24] 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.

4 also shows the summary statistics for both entry and exit of brands. Not surprisingly, most product entry and exit happens within existing brands.[25]

One obvious question is whether the product categories that exhibit a lot of market 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 varieties (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. To the extent that one believes that it is easier to be innovative when developing pre-recorded video recordings, cameras, and computer software than when developing new forms of sugar, frankfurters and butter, one can have some confidence that these measures are capturing true innovations rather than simple re-labeling of existing products.

An aspect of product creation and destruction that we explore is how entry and exit co-varies across modules and time. The left panel of Table 6 examines the correlation between entry and exit rates *across* product modules. For expositional ease, we only use 1-year entry and exit variables. The correlation between entry and exit variables *across* modules is positive and high. The average contemporaneous correlation between entry and exit rates is 0.70. In terms of creation and destruction, the contemporaneous correlation is slightly higher at 0.83. This suggests that modules tend to be characterized by either high entry and exit or low entry and exit.[26]

The right panel of Table 6 examines the correlation between entry and exit rates *within* product modules over time (i.e., after subtracting the average entry and exit rate from each module over the four-year period). Here we exploit the panel nature of the dataset to examine whether periods in which entry is high are also periods of relatively high exit. We construct correlations using entry and exit variables that have been purged by module averages over time.

[25] In Appendix E, we further document the characteristics of product turnover. One dimension that is interesting to describe is how much of product creation is the result of new sizes or new flavors of existing products. For roughly 20 percent of the products that were purchased in Q4 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 sizes and flavors of existing products. We find that creation from new sizes is 1.9 percent or roughly 5 percent of overall creation. New flavors add 0.4 percentage points to creation or less than 1 percent.

[26] The simple correlations we find for product entry and exit are similar to those found in terms of firm entry and exit by Dunne et al (1988).

We find that in 2000 and 2003 modules that had creation above their module average were also modules that had destruction above their module averages. This suggests that these are periods where both creation and destruction co-vary positively. By contrast, 2001 and 2002 were years where modules with creation higher than their module average had low destruction relative to their own average destruction. In these years, creation and destruction co-varied negatively. Interestingly, the correlation between creation in period t and the destruction in period $t+1$ is positive, which indicates that higher than average creation in one period is followed by higher than average destruction in the next period. However, overall, there is not a strong intertemporal pattern in creation and destruction rates.

Up until this point, we have been describing the magnitude and overall trends in prices and product creation and destruction. However, these aggregate patterns mask important relationships which exist at the level of the product. We begin by examining the behavior of destruction by UPC characteristic in Table 7. 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.

We now focus our attention in the behavior of UPCs over their life cycle. Table 8 presents information on two sets of products over time, those UPCs that existed in 1994 and those that existed in 1999. We report means and standard deviations across all modules of several statistics. The first row of the first panel shows the average market shares across modules of the set of UPCs available in 1999. By 2000, the set of UPCs available in 1999 comprised 92 percent of the expenditure in the average module. This share falls steadily as one moves closer to the end of our sample. By 2003, the set of UPCs in 1999 were 70 percent of all expenditure in UPCs. The second panel shows the cumulative exit rates of the same sets of goods. For instance, the exit rate of 1999 goods in 2000 was 20 percent. As time passes the cumulative exit rates rise but at a decreasing rate. As can be inferred from following the 1994 set of UPCs over time, the increments in the exit rates remain relatively constant after a few years. Finally, the bottom of Table 8 shows the average size of surviving products relative to all products in the module. UPCs that survive 4 years are typically 23 percent larger than the average existing UPC. Over a 9 year period, those UPCs that survive have market shares that are 50 percent larger than the average UPC in the module.

Stylized Fact 4: Creation is Strongly Pro-Cyclical and Destruction is Weakly Counter-cyclical

In this section we address how product creation and destruction co-varies with aggregate measures of consumption. For expositional ease (and slightly abusing notation), we will re-define creation in terms of the total value in period $t-1$ rather than the total value in t (as in (2)). This transformation is useful as we can divide total sales growth into the sales of UPCs that are common over time and those that are new or disappearing:

(3) $$\frac{V_t - V_{t-1}}{V_{t-1}} \equiv \underbrace{\frac{C_t - C_{t-1}}{V_{t-1}}}_{\substack{\text{Growth in} \\ \text{Common UPCs}}} - \underbrace{\frac{D_{t-1}}{V_{t-1}}}_{\substack{\text{Destruction} \\ \text{(t-1)}}} + \underbrace{\frac{N_t}{V_{t-1}}}_{\substack{\text{Creation} \\ \text{(t-1)}}}$$

where $V_{t-1} = C_{t-1} + D_{t-1}$ is the total value of UPC sales in $t-1$, and is the sum of the value of UPCs that are common between period t-1 and t, C_{t-1}, and those that are discontinued between t and $t-1$, D_{t-1}; and $V_t = C_t + N_t$ is the value of all UPCs in $t-1$, and N_t is the value of all those UPCs that are created between t and $t-1$. We call D_t product "Destruction" and C_t product "Creation". V_t, C_t, and N_t are adjusted for inflation so that real dollar values are compared over time.[27]

In order to quantify the cyclicality of product creation and destruction, we present the patterns of net creation (creation – destruction), creation and destruction as defined in (3). While we have only 20 consecutive quarters of data, the period studied includes the 2001 recession. Panel (a) of Figure 3 plots net creation (creation minus destruction) and the growth in overall sales of the ACNielsen sample. The pattern that emerges is strongly pro-cyclical. ACNielsen sales are weakest during the recession of 2001, and this is also the period during which net

[27] We use the CPI Food index to deflate these series.

creation of goods reaches its trough. In the later years (2002 and 2003) sales and net creation pick up. Panel (b) shows that the cyclicality of net creation is driven mainly by the pro-cyclicality of creation. That is, product creation is largest in periods where sales growth is strongest. Panel (c) shows a clear counter-cyclical pattern of destruction. While the correlation of destruction is strongly negative with overall sales, the magnitude of the fluctuations is small, and hence it explains a small fraction of the cyclicality of net creation.

We should be cautious when interpreting the results of Figure 3 because there are only 24 quarters over which the data is collected.[28] However, the cyclicality of creation and destruction can be examined at the product group level. For each product group we examine how creation and destruction co-varies 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)) for 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.[29] We restrict the impact of consumption on creation and destruction to be the same across all groups but we allow for a different constant across groups.

Table 9 shows the results from regressing net creation, creation and destruction on total sales growth. Obviously because there is an accounting identity linking these variables, we are not looking for a structural relationship but simply to describe the co-movement 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 leads net creation to rise by 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 pro-cyclical, it is primarily driven by the pro-cyclicality of creation rather than the counter-cyclicality of destruction. The coefficient under the column "creation" can be interpreted as how much creation moves with

[28] 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.
[29] We do not include leads and lags of consumption growth as most of these coefficients are small and insignificant.

sectoral consumption growth. An additional 1 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 pro-cyclical and co-varies more with demand in expansions than contractions. The opposite is true for destruction rates. Destruction is counter-cyclical, but destruction responds more strongly in recessions than in booms. As suggested by Figure 3, however, most of the pro-cyclicality of net creation comes from the pro-cyclicality of creation.[30]

We can use Caballero and Hammour's (1994) model to interpret the implications of the cyclical patterns of product flows.[31] The rapid response of 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 (like in Shleifer (1986)). In both cases, 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 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.

III. 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 productivity. 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, the large number of products facing consumers and the price volatility observed in the data means that small sample problems can

[30] In particular, the insulating role of product creation seems larger than that of job creation. They find the coefficient on job destruction rates to be almost double that of job creation while Table 9 suggests the opposite is true when looking at product turnover. The larger response of destruction in recessions than in booms is similar to the pattern observed for job destruction.

[31] Caballero and Hammour (1994) use their model to understand the implications of job flows, 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 use their model to interpret the new facts on product turnover discussed in this section.

result in substantial error in price measurement. Third, the fact that new products seem to systematically displace market share from existing products means that common goods price indexes are likely to have substantial biases because they ignore the effects of quality upgrading. Finally, the cyclicality of creation means that our ability to measure the severity of business cycles is likely hampered by the fact that our price indexes do not systematically adjust to product creation and destruction. We now turn to quantifying these forces.

There are two approaches to quantifying the implications that these facts have for price measurement. The first is to examine the implications that these stylized facts have for an index that is commonly used – in this case the Consumer Price Index. The second is to start with a utility based or axiomatic approach and assess what impact this has on an exact index. The reason to take all these different approaches is that the CPI is an *ad hoc* index. Thus, while we can discuss how sampling error can affect the standard error of the index, we cannot discuss "biases" without having a notion of what the "true" inflation rate is. Exact and Superlative indexes are useful in this respect because they answer the question of what the price level is, or approximately is, for a given utility function or class of utility functions. Superlative indexes in particular have often been used as an unbiased "benchmark" in discussions of CPI bias. The downside of simply focusing on exact or superlative indexes is that they are not the formulas actually used by government statistical agencies, and so we feel it is useful to understand how much the facts we discuss matter in theoretically-grounded indexes as well as the *ad hoc* one that is actually used.

III. A. Common Goods Indexes: Sampling Error in the CPI

Although the CPI originated as close to a pure Laspeyres index, the current CPI contains many modifications. In this section, we provide a brief overview of how the CPI is computed, focusing on the areas that our data allow us to address.[32] We describe the methodology used by the BLS to sample the products whose prices enter the CPI. This is a small subset of the goods' prices collected by ACNielsen's Homescan database. Thus this database provides an invaluable resource to calculate the extent of sampling error in the CPI methodology.[33]

---

[32] Readers interested in many of the details that are not included in this brief description, please refer to BLS [2007]

[33] Richardson (2003) suggests that scanner data can have the potential of reducing the CPI's sampling error considerably. Leaver and Larson (2001) provide estimates of the variance of a "scanner-based" CPI for cereals. More generally, the BLS provides a measure of the sampling error based on combining the price quotes of the

The CPI contains two levels of aggregation. At the *upper* level, 305 Entry-Level Item (ELI) price indexes in each of 38 urban *areas* are combined either using a Laspeyres formula in the case of the standard CPI or a Tornqvist formula when using the chained CPI.[34] Weights from the Consumer Expenditure Survey are used for this upper level. Each of these indexes is in turn based on a *lower* level sample of individual products. The BLS uses on average 7 price *quotations* per item-area (85,000 price quotes all) for the non-shelter component of the CPI. Thus the average ELI is based on 280 price quotes obtained from 38 different urban areas. This lower level price quotation is critical for the BLS's approach to price measurement.[35]

Price quotes are selected into the CPI by first conducting a survey of households and asking them where they purchase items in each ELI or expenditure category (e.g. audio equipment). Once a store is selected through this "Point of Purchase Survey" or POPS, a BLS field representative will visit the store and try to determine the market share of each individual product in the store. This is done by asking the store personnel, measuring shelf space, or assigning equal probability to all items. The market shares for each individual product, when available, are used to perform a "weighted" sampling of products within the store. The products selected using this procedure enter the lower level index.

After the first visit of the BLS agent to the store, monthly or bi-monthly telephone calls are performed to the store to assess the price of the products that were originally picked.[36] Since January 1999, the price quotations from various stores in an elementary item-area have been geometrically averaged to form an item-area (*ia*) sub-index:

(1)
$$\pi_t^{ia} \equiv \prod_{j=1}^{N^{ia}} \left(\frac{p_{jt}^{ia}}{p_{jt-1}^{ia}} \right)^{\frac{1}{N^{ia}}},$$

where *j* are individual goods that belong to the item-area *ia*, and N^{ia} is the total number of prices quotes in item-area area *ia*.[37] As mentioned above, on average $N^{ia} = 7$, and the average number

products sampled in different ways. However, none of these ways include increasing the number of goods being sampled.

[34] The 305 ELIs are aggregated into to 211 "strata" or elementary indexes in the CPI. An elementary index can contain several ELIs.

[35] Almost all of the deviations from the standard Laspeyres index in the US case – hedonics, geometric averaging of prices, sample rotations, etc. – occur at this lower level.

[36] If products are discontinued the BLS has to perform a "forced substitution" of the product. Bils (2005) focuses on these episodes in the case of consumer durable goods that are surveyed by the BLS.

[37] Since the introduction of the lower-level geometric average, the BLS publishes the extent of the correction of the lower-substitution bias driven by the implementation of this new formula. On average the annual lower-level

of price quotes over all areas for an item equals 280, i.e. $\sum_a N^{ia} = 280$. Prior to 1999, the BLS used a Laspeyres index to capture item-area averages. This modification allowed the BLS to reduce the lower-level substitution bias. Note that no weights specific to the individual good *j* appear explicitly in equation (1). This masks the fact that the BLS selects the products in each lower level index using a "weighted" sampling as explained above. If field agents actually could select goods with probabilities equal to their market share, then the lower level aggregator would approximate a sales-weighted geometric average of individual products.

The overall CPI inflation is then defined by aggregating item-area indices as follows:

(2) $$\Delta CPI_{t,r} \equiv \sum_{i,a} s_r^{ia} \pi_t^{ia}\text{, where } s_r^{ia} = \frac{E_r^{ia}}{\sum_{i,a} E_r^{ia}}$$

and E_r^{ia} is the expenditure at the item-area level in the base year *r* obtained from the consumer expenditure survey.[38] A key feature of equation (2) is that the item-area price indexes are weighted by the expenditure shares of the reference period *r*. This creates an "upper-level substitution bias" because although consumers can switch away from high priced items, the fixed-weight index does not allow this. Since August 2002, the BLS has added a chained CPI index that is defined using a Tornqvist formula at the upper level to address this problem. The chained CPI (CCPI) uses a geometric average formula to aggregate all ELI price changes using weights equal to the ELIs past and current expenditure shares. The chained inflation formula is given by the following expression:

(3) $$\Delta CCPI_t \equiv \prod_{i,a} \left(\pi_{t,t-1}^{ia}\right)^{\frac{1}{2}\left(s_t^{ia} + s_{t-1}^{ia}\right)},$$

where weights are expenditure shares from periods *t* and *t*-1. Because current expenditure shares are not available at the time the index is published, the BLS provides interim chained CPI measures contemporaneously to the release of the CPI, which is updated once the new expenditure shares are measured.

The methodology just described gives rise to sampling error. This error arises as only a sample of retail prices is used to compute the CPI, instead of using the complete universe of

---

substitution bias has been around 0.3 percentage points. For a detailed survey of different measures see Lebow and Judd (2004).

[38] The expenditure reference period used by the CPI until December 2001 was 1993-1995 data. Throughout 2002-2003 the reference period was 1999-2000.

retail prices. The fact that ACNielsen records every product and price consumed by their sample of households provides an ideal scenario to quantify the extent of the sampling error. Each BLS ELI encompasses approximately the same amount of expenditure as an ACNielsen product group, but the BLS only samples approximately 280 prices per ELI out of the around 5,300 prices per product group we observe. Thus even if the index were unbiased, there is certain amount of measurement error that makes both the observed level of inflation and changes from that level stochastic. This is something that we quantify in the results section.

III.B. Creative Destruction and Quality Upgrading: An Exact Price Index

The price indexes that we have been discussing thus far are legitimate for cases in which the quality of goods does not change or the set of good is fixed over time. In this section, we work through the theory of how creative destruction affects price measurement using a price index that is exact for a CES utility function in the presence of quality upgrading.

To better measure the impact of creation and destruction, 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 elasticities within and across 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 comprised 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 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 sub-utility function describing the representative consumer utility. The second describes how brands within a product group sub-utility 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 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."

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

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

where product groups are indexed by g, ρ is the elasticity of substitution across product groups and G is the set of all product groups. 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:

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

where c_{bgt} is the total quantity consumed of brand-module b in product group g at time t, ρ_g^a 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 $B_g \subseteq B_{gt} \subset \Psi_g$.

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

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

where c_{ubgt} is the consumption of UPC *u* of brand *b* of product group *g* in period *t*, ρ_g^w 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_{ubg} 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-Flaked Soda Crackers" is one particular brand-module, *b*, of the product group "Crackers," and "Nabisco Premium Unsalted Crackers" and "Nabisco Premium Multi-Grain Crackers" are two different UPCs of the brand-module "Nabisco Premium-Flaked Soda Crackers". Each of these UPCs have their unique quality parameter, d_{ubg}. We define $U_{bgt} \subseteq \Omega_{bg}$ as the set of all UPCs that have non-zero sales. It is also useful to define the set of UPCs in a brand-module that are common over time as U_{bg} where $U_{bg} \subseteq U_{bgt} \subset \Omega_{bg}$. As we show below, this 3-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 individual UPC has its own quality level. Second, if the UPC belongs to a new brand, the elasticity of substitution used to value its introduction, ρ_g^a, is different than the introduction of a new UPC within an existing brand, ρ_g^w. In the data we expect that $\rho_g^a < \rho_g^w$ as within brand-group UPCs are 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 sub-utility function in (5) is given by the following expression:

$$P_{bgt} = \left(\sum_{u \in U_{bgt}} \left(\frac{p_{ubgt}}{d_{ubg}} \right)^{\sigma_g^w} \right)^{\frac{1}{\sigma_g^w}}, \tag{7}$$

where p_{vbgt} is the price of UPC *u* of brand *b* in product group *g* in period *t* and $\sigma_j^i = \rho_j^i / \left(\rho_j^i - 1 \right)$. For simplicity, define $\tilde{p}_{ubgt} = p_{ubgt} \big/ d_{ubg}$ as the quality-adjusted price. Here, it is important to remember that equation (7) only contains 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 (5) can be denoted by

$$P_{gt} = \left(\sum_{b \in B_g} \left(P_{bgt} \right)^{\sigma_g^a} \right)^{\frac{1}{\sigma_g^a}} . \tag{8}$$

And the overall price index is given by,

$$P_t = \left[\sum_{g \in G} P_{gt}^{\ \sigma} \right]^{\frac{1}{\sigma}} . \tag{9}$$

Equations (7) - (9) 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. Suppose that UPC u with positive sales in time $t-1$ is replaced in time t with a higher quality UPC, u', i.e. $d_{ubg} < d_{u'bg}$. If the higher quality UPC has a lower quality-adjusted price, an index that ignores the quality upgrade will generally result in higher measured price change relative to the true quality-adjusted inflation. 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 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:

$$s_{ubgt} = \left(\frac{p_{ubgt} / d_{ubgt}}{P_{bgt}} \right)^{1-\sigma_g^w} E_{bgt} \tag{10}$$

The seminal insight of Feenstra [1994] was that one can use this fact to eliminate the quality parameters from the price index in (7) and write it *only* in terms of prices and market shares even when goods are constantly being replaced.

We turn next to the derivation of the exact price index that allows for the creation and destruction of products with different qualities over time. Let $U_{bgt} \subset \Omega_{bg}$ be the subset of all UPCs of brand b of group g consumed in period t. The nested CES structure embedded in (4)-(6) 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 UPCs purchases. The main proposition in this paper is an extension of Feenstra (1994) and Broda and Weinstein (2006):

PROPOSITION 1: For $g \in G$, if there exists a $b \in B_g = B_{gt} \cap B_{gt-1}$, $B_g \neq \varnothing$, then the exact price index for product group *g* with new and disappearing brands *and* UPCs is given by,

$$EPI\left(\mathbf{p}_t, \mathbf{p}_{t-1}, \mathbf{x}_t, \mathbf{x}_{t-1}, B, U\right) = \prod_{g \in G} \left[CEPI_g \left(\frac{S^c_{gt}}{S^c_{gt-1}} \right)^{\frac{1}{\sigma^a_g - 1}} \prod_{b \in B_g} \left(\frac{S^c_{bgt}}{S^c_{bgt-1}} \right)^{\frac{w_{bg}}{\sigma^w_g - 1}} \right]^{w_g},$$

where $S^c_{bgt} = \dfrac{\sum_{u \in U_{bg}} p_{ubgt} x_{ubgt}}{\sum_{u \in U_{bgt}} p_{ubgt} x_{ubgt}}$, and $S^c_{gt} = \dfrac{\sum_{b \in B_g} \sum_{u \in U_{bgt}} p_{ubgt} x_{ubgt}}{\sum_{b \in B_{gt}} \sum_{u \in U_{bgt}} p_{ubgt} x_{ubgt}}$, $U_{bg} = U_{bgt} \cap U_{bgt-1}$, and $CEPI_g$ is the standard conventional exact price index defined over the common goods and weights w_{bgt}, and w_{gt} are defined in Sato (1976) and Vartia (1976).

This result states that the exact price index with quality change (*EPI*) is equal to the "conventional" exact price index, $CEPI_g\left(B_b, U_{bg}\right)$, i.e. the exact price index of the UPCs without quality change over time, multiplied by two functions of the share of common goods over time, i.e., the S^c ratios.

The first S^c ratio captures the impact that new brand-modules have on the EPI relative to the CEPI. It can be easily interpreted as the ratio over time of the expenditure shares of the *common* brands in product group *g*. That is, when the share of new brands in period *t* is larger than the share of disappearing brands in period $t-1$, this S^c ratio is smaller than 1. In this case, as new brands have a lower price per quality than the disappearing goods (see (10)) the EPI will fall relative to CEPI, since the conventional "fixed-good" index will miss the increase in quality that occurred through the entry and exit of brands. In a similar way, the second S^c ratio captures the role played by new and disappearing UPCs *within* brand-modules that were common in both time periods.[39] S^c_{bgt} equals the expenditure share in period *t* of UPCs that are available in both periods in a brand-module (i.e., $u \in U_{bg} = \left(U_{bgt} \cap U_{bgt-1}\right)$). Thus, the S^c ratio is just the ratio of the share of common goods in *t* relative to share of common goods in *t*-1. In this case, if one UPC is

[39] All of the index numbers used in this paper suffer from the classic "index number problem". In particular, results are dependent on the base year or years used. Since we are examining long-run changes, we use two base years 1972 and 1990.

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 S^c ratio will be smaller than unity because the share of the common goods will have fallen. A S^c ratio smaller than one means that the exact price index will be lower than the conventional price index 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^a play more subtle roles. As the elasticity of substitution rises, then a given movement in the share of common goods over time (i.e., S^c ratio) 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.

IV. 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).

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 demand and supply equations:

$$(11) \qquad \Delta^{k_{bg}} \ln s_{ubgt} = -\left(\sigma_g^w - 1\right)\Delta^{k_{bg}} \ln p_{ubgt} + \varepsilon_{bgt}^{k_{bg}}$$

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

Equation (11) represents the optimal demand for a given UPC *u* in brand-module *b* and product group *g*, and equation (12) represents the supply of that UPC. Both are expressed in terms of shares, where s_{ubgt} is the share of variety *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. 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 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 ε_{ubgt} and δ_{bgt} uncorrelated, i.e. $E_t\left(\varepsilon_{ubgt}^{k_{bg}} \delta_{ubgt}^{k_{bg}}\right) = 0$. The second identifying assumption is that σ_g^w is restricted to be the same over time and for all UPCs of a given brand-module-product group (but varies over product groups).

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.

(13) $$G(\beta_g) = E_t\left(\upsilon_{ubg}\left(\beta_g\right)\right) = 0 \quad \forall u, b, \text{ and } g.$$

where $\upsilon_{ubg} = \varepsilon_{ubg}\delta_{ubg}$ 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:

(14) $$\hat{\beta}_g = \arg\min_{\beta_g \in \mathrm{B}} \mathrm{G}^*\left(\beta_g\right)' W \mathrm{G}^*\left(\beta_g\right) \quad \forall g.$$

where $\mathrm{G}^*\left(\boldsymbol{\beta}_\mathbf{g}\right)$ is the sample analog of $G\left(\boldsymbol{\beta}_\mathbf{g}\right)$ 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 $\boldsymbol{\beta}_\mathbf{g}$ 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.[40] 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.[41] The use of the between estimate coupled with our need to estimate σ_g^w, ω_g^w and a

---

[40] We first use Feenstra's approximate (14) to solve for $\boldsymbol{\beta}_\mathbf{g}$. 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 $\boldsymbol{\beta}$'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 5 percent apartFor 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: $\omega_g = \frac{\gamma_g}{\sigma_g\left(1-\gamma_g\right)-1}$. The objective function was evaluated at values for $\sigma_g \in \left[1.05, 131.5\right]$ at intervals that are 5 percent apart, and for $\gamma_g \in \left[0.01, 1\right]$ at intervals 0.01 apart. Only combinations of σ_g and γ_g that imply $\sigma_g > 1$ and $\omega_g > 0$ are used. To ensure we used a sufficiently tight grid, we cross-checked these grid-searched parameters with estimates obtained by non-linear 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 only by a few percent for those σ_{ig} and ω_{ig} for which we could apply Feenstra's "between" approach.

[41] In the appendix of Broda and Weinstein (2006), 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.

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 as 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 set of elasticities per product group that are key to estimating the impact of product churning. In the next section we describe this two set of elasticities separately.

V. Results

Our empirical and theoretical section has suggested two main implications of creative-destruction for aggregate price measurement. The first concerns the inherent stochastic element of the CPI due to the small sample of prices that are used in the index. The second concerns the magnitude of bias arising from an index that not only deviates from an exact or superlative index but also does not take into account the fact that many prices are moving to and from their reservation levels. We will address each in turn.

V. A. Sampling Error in the CPI

Since the CPI is based on a random sample of prices, it can deviate from the true price index because of sampling error. The current literature on index numbers devotes a great deal of their attention to matters such as choice of formula, weighting schemes, etc., but seem to have made only passing reference to the fact that most price indexes are derived from samples and are thereby subject to sampling error.

A simple way to understand the importance of sampling error in the CPI is to bootstrap the standard errors of the CPI by repeatedly drawing random samples of comparable size and examining how sensitive the price index is to these different samples. In order not to confuse issues of sampling with those of new goods we will only focus on the set of goods that survive over the full period in question. Given that a product group is approximately the size of an ELI, then the BLS is sampling 280 prices per product group. If one conservatively assumes that a product group is more like a stratum, the BLS samples approximately 400 products per sub-index. In either case, this is a small share of the total number of prices.

Unfortunately, we do not know precisely how BLS field agents decide how to weight each good when deciding which goods to include. We therefore will make two extreme assumptions. The first is that the BLS field agents know the true sales shares of all goods in an ELI when they establish the base year (1999); and the second is that they randomly select goods without knowing anything about the sales weights. Obviously the truth lies somewhere between these two extreme assumptions, but this implies that the true degree of measurement error will also lie between these bounds.

Table 10 presents the results from bootstrapping the standard errors for a price index that covers our sample of goods that is constructed using the CPI methodology.[42] We perform the bootstrapping at a quarterly frequency and annual frequency and report the annualized inflation rates and standard errors. In all cases, we bootstrap the distribution using 250 iterations. Since any seasonal factor should affect each sample identically, we should interpret the distribution of inflation measures at the quarterly frequency to be standard errors on seasonally adjusted data.

At a quarterly frequency with 250 iterations and using a "weighted" sampling, the typical inflation rate is 0.75. The standard error is 0.49 indicating when we observe quarterly inflation rates of around 1 percent, the 95th percentile confidence interval for the true inflation is between 0 and 2 percent. If we assume that field agents do not observe sales weights, the 95 percent confidence bands expand to just over 1.1 percentage point. This implies that we can only be sure that the economy is away from deflationary terrain in cases where measured inflation is above 1 percent. Sampling 400 products slightly reduces the sampling variation as expected, but does not qualitatively change the results.[43] This sampling error does not vary much by quarter.[44]

Not surprisingly, moving to a four-quarter frequency reduces the sampling error presumably because the impact of short-term price fluctuations arising from sales and other high

[42] The price index we compute tracks the changes in CPI food and beverage index quite closely over the sample period. See Appendix F for details

[43] At first glance this number seems too large, however, it follows immediately from the volatility of price changes that we described earlier. In the data preview section we highlighted the fact that our data, as well as the data for the CPI, exhibit very high levels of price volatility at the individual good level. More to the point, we documented that the standard deviation of a price change for a common good at a quarterly frequency is 20 percent. Given this level of volatility of individual prices, even with 29,000 price quotes, we would expect the standard deviation of the average inflation rate to be 0.16 at a quarterly frequency or 0.62 annualized. This number is quite close to the numbers we obtain from the bootstrapping procedure and gives the intuition behind the numbers we report.

[44] See Appendix F for details.

frequency events is reduced. Indeed, one-quarter inflation rates (annualized rates) are roughly 3 times more volatile that four-quarter inflation rates. The standard error for four-quarter inflation rates is 0.1 indicating that sampling error means that the index is precise to plus or minus 0.2 percent per year at the annual frequency.

Another important number to compute is the sampling error for changes in inflation rates. This lets us examine how much confidence we have in a movement in inflation. At the quarterly frequency, the standard errors range from 0.76 and 0.88 depending on how well one assumes that BLS field agents sample. This means that unless a seasonally adjusted quarterly inflation rate moves by 1.5 to 2 percentage points in either direction (annualized), one cannot be confident that the fluctuation isn't being driven by measurement error. Once again, at the annual frequency the standard errors fall but they remain around 0.2 indicating that inflation movements of less than plus or minus 0.4 percent could easily be driven by sampling error. Overall, this implies that at quarterly frequencies inflation rates have large standard errors, and policymakers and market participants should be cautious in interpreting changes in inflation trends.

Another way of understanding these numbers is to compare them to the actual inflation rates we observe. In order to make our numbers comparable, we should make two adjustments. First, the BLS uses approximately 29,000 price quotes for the ELIs covered in our sample, but 85,000 quotes for the CPI as a whole. Since the entire CPI averages 2.5 times more prices than we use in our simulation, one would expect the standard error of the CPI to be smaller than that of our simulation. If we assume that the standard error of the CPI falls with the square root of the number of price quotes, this suggests we should divide our estimate of the standard error by 1.58. On the other hand, the work of Klenow and Kryvtsov [2005] suggests that our sample of prices exhibits less price volatility than those in the CPI as a whole. If we think of the Klenow and Kryvtsov find the (weighted) absolute price change in the CPI to be 13.3 percent, whereas the comparable price fluctuations in our data are 8.8 percent, then the standard error of the typical price quote in the CPI would be 1.51 times larger that of a price in our sample. This comparison suggests that the large sample size in the CPI implies that average volatility is 1.58 times *lower* in the CPI, but the higher price volatility in the sectors not covered in our sample implies that one should expect the CPI to have 1.51 times *more* volatility than our sample. The similarity of the implied downward and upward adjustments suggests that the variance of the CPI is, to a first approximation, close to the variance of the sample we are using, so we will compare the implied volatility in our sample with that of the CPI directly.

Figure 4 summarizes our results of the impact of sampling error on innovations in the inflation rate. Here we plot the kernel density of seasonally-adjusted quarterly inflation rate changes between 1994 and 2003. Inflation innovations whose absolute value are greater than 1.5 percent denote those changes that we can be 95 percent confident were not generated by measurement error. As one can see from the graph, about 50 percent of all inflation changes fall within the interval [-1.5, 1.5], which means that we cannot reject the hypothesis that half of all observed quarterly inflation innovations were due to measurement error. Obviously, the smaller standard errors for four-quarter inflation rate innovations means that we can be far more certain about annual inflation rates. However, the large standard errors at the quarterly frequency imply that the measurement error is sufficiently large to make it difficult to determine the direction of inflation changes about half of the time.

V. B. Quality – New Goods Bias

A separate issue concerns the magnitude of the bias in the CPI. We already know from Figures 1 and 2 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. In order to examine the quality bias, we need to leave behind the *ad hoc* structure of the CPI and turn to the magnitude of the bias relative to an exact index. 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 CES price index yields an almost identical rate of inflation at the Tornqvist index and the chained CPI.[45] The role of forced substitutions is discussed below. Explicit quality corrections (like hedonics) are not standard in any of the categories of goods included in our data.[46]

Proposition 1 indicates that the magnitude of the bias depends on two factors: S^c ratios and elasticities. The former indicates the relative importance of quality shifts in the data and the

[45] In Appendix G 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.

[46] The only exception is computer software where the CPI started using hedonics since January 1998. The contribution of computer software to the overall quality bias we find is tiny (less than 1 percent) so we ignore this adjustment when presenting the main results.

second tells us about how much of these shifts are being picked up in the conventional index. Table 11 presents the distribution of the per-year ratio of the share of common goods within and across brand-modules.[47] The median within per-year ratio is 0.93. Over a five-year period this implies a ratio of 0.69. This number is easiest to understand in terms of a simple example. Imagine a firm produced two goods with equal market shares within a brand-module. If they replace one of these goods with a new product that has twice the market share of the common good, that would generate a ratio of 0.66. Thus, a ratio of 0.69 is consistent with substantial quality upgrading by firms. The across ratio is 0.98 per year, or 0.87 over 4-years, which is significantly smaller than the within S^c ratio, but still less than one. This suggests that while the typical product group experienced new brand-modules that were of higher quality 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 11 show the distribution of estimated elasticities of substitution. The typical within-brand-module elasticity is 11.5. The implication is that a one percent price decline of a UPC within a brand-module causes its sales to rise by 12 percent. This would be the case if the various versions of, say, Nabisco Ritz Crackers are fairly 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 4 to 7 (see Dube et al. (2005) and Rossi (1999)). We believe this makes our estimated quality biases conservative estimates.

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 3 and 4 for the most disaggregated trade data (10-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 10-digit sector, so it is comforting to see that our estimated elasticity is larger than their 10-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

[47] Per-year ratios are derived from 5-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 churning (e.g., like that coming from fashion products) can have on the quality bias.

under a particular brand in a module are more substitutable than different brands. 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 convincing because we would like to see that within-brand-module elasticities are smaller than the across ones in the same module. In order to assess this, for each product group we take the difference between the estimated within- and across-brand-module elasticity. 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 since we know all of the ratios of common shares and elasticities. When we do this, 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 8.5 percentage points or 0.91 percentage points per year.[48]

We can now contrast our approach with that of the BLS. There are two ways in which new and disappearing products affect the CPI. First, the BLS regularly schedules product substitutions in 25 percent of the categories every year. This means that with the exception of "forced" substitutions (when the product that is being sampled ceases to be sold at the point of purchase), the same good is priced over a 4-year period. Thus, in general, over any four year period, we can think of the CPI as a fixed-goods price index.[49]

As we just noted, however, the BLS is sometimes forced to substitute a product when it disappears from the particular store being sampled. There are two important reasons why forced substitutions will miss the type of product creation and destruction we observe in the overall

[48] In Appendix H, we present the twenty product groups that contribute the most to the quality bias.
[49] We saw in previous sections how a common-goods CES price index is almost identical to a superlative common-goods price index. At the product category level, we know that the CPI is an superlative index.

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 3-megapixel camera with a 4-megapixel one, the BLS agent will attempt to find another 3-megapixel camera with which to replace the discontinued camera. This implies that the new 4-megapixed camera is likely not to be used in the updated CPI. Second, when substituting the new 3-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 3-megapixel cameras as a price change, which implies that any quality change between products is ignored (see BLS (2005) and Bils (2004)).[50] For these two reasons we believe that the vast majority of product creation and destruction and its price implications are ignored by the CPI.

In Stylized Fact 4 we highlighted the cyclical patterns of creation and destruction. We now turn to assessing whether the bias to the common-goods price index has a cyclical component. Figure 5 shows the four-quarter biases computed on a rolling basis against the sales growth of whole 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 4th quarter of 2002. Despite the clear cyclical pattern, the figure plots only 16 observations in a period with only a mild recession. Table 12 shows that the aggregate pattern is confirmed at the product group level. Here we regress the four-quarter bias 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 10 percent growth in sales in a particular group, the bias in the conventional price index for that group increases by 1 percentage point.

We can now put together all of our results. The average quality bias we find is around 0.8 percent per year. Given this upward bias and the estimates for sampling error computed in the

[50] This is true for the vast majority of CPI product categories in our sample. For other products (airline fares, cable television, personal computers, motor vehicles, telephone services, etc), BLS uses hedonics or other methods to infer quality.

previous section, one cannot be 95 percent certain that the U.S. is not in deflation unless measured quarterly C-CPI inflation rates is higher than 1.8 percent per year.[51] Sampling error at the quarterly frequency is sufficiently high that quarterly rates of chained inflation of 1.8 percent at an annualized rate could still be consistent with deflation (adding the bias and the sampling error together). Roughly 30 percent of the actual quarterly CPI inflation rates reported by the BLS over the period 1994-2003 were below 1.8 percent. This suggests that sampling error coupled with the estimated bias means that approximately one third of all quarterly inflation rates cannot be statistically distinguished from zero [52]

VI. 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. The fact that these surveys only comprise a small fraction of all goods in the market gives rise to substantial sampling error in aggregate indexes. Moreover, 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 biases in price measurement.

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. Moreover, we find that sampling error for the set of goods we examine means that one would need to observe a conventionally measured quarterly inflation rate of close to 2 percent to be 95 percent confident that true inflation is not negative. The bias is

[51] When the comparison is made with the regular CPI, then we have to add the upper level substitution bias to the confidence interval. We estimated the difference between a C-CPI or Tornqvist index and CPI (lower level geometric, and upper level Laspayres) to be 0.2 percent in our data. This implies that the CPI has to be higher than 2 percent for policymakers to be certain that the U.S. is not in deflation.

[52] For space considerations we present the kernel density of seasonally-adjusted quarterly inflation rates between 1994 and 2003 in Appendix I.

also strongly pro-cyclical, which suggests that business cycles are more pronounced than is typically reported in official statistics.

The paper also fills an important gap in the existing literature on creative destruction. An extensive theoretical literature focuses on how the process through which new products replace old ones influences major aspects of macroeconomic performance. Product creation and destruction is at the core of many models of long-run growth and business cycle fluctuations. Despite its theoretical prominence, the empirical analysis on the aggregate implications of product turnover still lags far behind its theoretical counterpart. Our work complements the papers by Dunne et al. (1988) and Davis and Haltiwanger (1992) on firm turnover by documenting a key component of the creative destruction process that remained unmeasured, that of product entry and exit.

Considerable 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 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 cyclicality of quality upgrading indicates that the biases and measurement issues we have documented are likely to be robust features of the data.

References (Incomplete)

Anderson, Simon P., A. de Palma, and J. Thisse (1992), *Discrete Choice Theory of Product Differentiation*, Cambridge, MA: MIT Press.

Barlevy, G., "On the Cyclicality of Research and Development" mimeo Federal Reserve Bank of Chicago, 2006.

Bernard, A., S. Redding and P. Schott (2006) " Multi-product Firms and Product Switching," NBER working paper 12293.

Bilbiie, F., F. Ghironi and M. Melitz, "Endogenous Entry, Product Variety, and Business Cycles," mimeo Harvard University, June 2006.

Bilbiie, F., F. Ghironi and M. Melitz, "Monopoly Power and Endogenous Variety in Dynamic Stochastic General Equilibrium models," mimeo Harvard University, June 2006.

Bils, Mark, 2004. "Measuring Growth from Better and Better Goods."*NBER Working Paper* 10606, June.

Bils, Mark, and Peter J. Klenow, 2001. "Quantifying Quality Growth." *American Economic Review*, 91, 1006-1030.

Blanchard, O. and P. Diamond, "The Cyclical Behavior of the Gross Flows of U.S. Workers." *Brookings Papers on Economic Activity*, 1990, (2), pp. 85-143.

Bresnahan T. and R. J. Gordon, eds., *The Economics of New Goods*. Chicago: University of Chicago Press.

Broda, C. and D. Weinstein, "Globalization and the Gains from Variety", *Quarterly Journal of Economics*, Volume 121, Issue 2 - May 2006

Broda, C. and D. Weinstein, "Variety Growth and World Welfare**",** *American Economic Review*, Vol. 94, No. 2, May 2004

Boskin Commission Report (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.

Caballero, R., and M. Hammour. "On the Timing and Efficiency of Creative Destruction", *Quarterly Journal of Economics*,Vol. 111, No. 3. (Aug 1996)

Caballero, R., and M. Hammour. "The Cleansing Effects of Recessions", *Quarterly Journal of Economics*, Vol. 111, No. 3. (Aug 1996)

Campbell, J. “Entry, Exit, Embodied Technology, and Business Cycles” *Review of Economic Dynamics,* 1(2), 371-408. April 1998.

Chatterjee, S. and R. Cooper. “Entry and Exit, Product Variety and the Business Cycle” NBER working paper no. 4562, December 1993.

Davis, S. and J. Haltiwanger, “Gross Job Creation, Gross Job Destruction, and Employment Reallocation” *The Quarterly Journal of Economics*, Vol. 107, No. 3 (Aug., 1992), pp. 819-863

Dunne, T., M. Roberts and L. Samuelson, “Patterns of Firm Entry and Exit in U.S. Manufacturing Industries” *The RAND Journal of Economics*, Vol. 19, No. 4 (Winter, 1988), pp. 495-515

Dunne, T., M. Roberts and L. Samuelson, “The Growth and Failure of U. S. Manufacturing Plants” *The Quarterly Journal of Economics*, Vol. 104, No. 4 (Nov., 1989), pp. 671-698

Feenstra, R. “A Homothetic Utility Function for Monopolistic Competition Models, Without Constant Price Elasticity” *Economic Letters*, 78, 2003, 79-86.

Feenstra, R. “New Product Varieties and the Measurement of International Prices,” American Economic Review, 84(1), March 1994, 157-177.

Feenstra, R. and M. Shapiro, *Scanner Data and Price Indexes*, 2003, Studies of Income and Wealth, Volume 64, The University of Chicago Press.

Fenwick, D., A. Ball, P. Morgan, and M. Silver, “Price Collection and Quality Assurance of Item Sampling in the Retail Price Index” in *Scanner Data and Price Indexes* (2003), The University of Chicago Press.

Gallo, A. “First Major Drop in Food Product Introduction in the last 20 years” Food Marketing, December 1997, US Department of Agriculture.

Geroski, P. and C. Walters, “Innovative Activity over the Business Cycle,” *The Economic Journal,* July 1995, 105, pp. 916-928.

Ghironi, F. and M. Melitz, “International Trade and Macroeconomic Dynamics with Heterogenous Firms”, *Quarterly Journal of Economics* CXX (August 2005): 865-915

Gollop, F. and J. Monahan, 1991, "A generalized Index of Diversification: Trends in U.S. Manufacturing" Review of Economics and Statistics 43 no.2: 318:230

Hall, Robert E. “The Labor Market Is Key To Understanding the Business Cycle.” Unpublished manuscript, Stanford University, September 23, 2004.

Hausman, J. (2003), "Sources of Bias and Solutions to Bias in the CPI," *Journal of Economic Perspectives*.

Jovanovic, B. and R. Gilbert, “The Diversification of Production”, *Brookings Papers on Economic Activity. Microeconomics*, Vol. 1993, No. 1. (1993), pp. 197-247.

Jovanovic, B. and S. Lach, “Product Innovation and the Business Cycle” *International Economic Review*, Vol. 38, No. 1. (Feb., 1997), pp. 3-22.

Klenow, Peter J. and Oleksiy Kryvtsov, 2005. “State-Dependent or Time-Dependent Pricing: Does It Matter for Recent U.S. Inflation?” Stanford University mimeo.

Leaver, S. and W. Larson, 2001, “Estimating variances for a scanner-based consumer price index,” *Proceedings of the American Statistical Association*, Governmnet Statistics Section.

Lebow, David E., and Jeremy B. Rudd, 2003. “Measurement error in the Consumer Price Index: Where do we stand?” *Journal of Economic Literature*, 41, 159-201.

Moulton, Brent R. and Karin E. Moses (1997), "Addressing the Quality Change Issue in the Consumer Price Index," *Brookings Papers on Economic Activity*, 305-366.

Nordhaus, William D., 1998. “Quality change in price indexes”. *Journal of Economic Perspectives*, Winter, 59-68.

Richardson, D. “Scanner Indexes for the Consumer Price Index,” chapter 2 in *Scanner Data and Price Indexes*, edited by Feenstra, R. and M. Shapiro, The University of Chicago Press.

Shapiro, Matthew D. and David W. Wilcox (1996) “Mismeasurement in the Consumer Price Index: An Evaluation.” *NBER Macroeconomics Annual*, 11, 93-142.

Shimer, Robert. “Reassessing the Ins and Outs of Unemployment.” Unpublished manuscript, University of Chicago, 2005.

Streitweiser, M., 1991, "The extent and Nature of Establishment Level Diversification in Sixteen US Industries" Journal of Law and Economics 34 no.2, pt.2: 503-34

Stokey, N. "Learning-by-Doing and the Introduction of New Goods," *Journal of Political Economy*, 96 (August 1988), pp. 701-717.

U.S. Department of Labor (1997), *BLS Handbook of Methods*, U.S. Government Printing Office.

Table 1
Descriptive Statistics

Year	Number of UPCs					Number of Brand-Modules			Number of Modules	Number of Product Groups
		By Brand-Module		By Firm			By Firm			
	Total	Median	Average	Median	Average	Total	Median	Average	Total	Total
1994	532789	8	46.1	8	36.0	31884	1	2.9	1097	118
1999	651343	5	33.6	8	39.6	51881	2	3.9	1096	122
2000	666445	5	32.5	8	40.4	53579	2	4.0	1095	122
2001	658055	4	31.4	8	39.9	54213	2	4.1	1093	122
2002	690036	4	31.7	8	41.7	57146	2	4.2	1092	122
2003	697312	4	31.5	8	42.0	58135	2	4.2	1093	123
Average	649330	5	34	8	40	51140	2	4	1094	122
Per-year Growth										
2003/1994	1.03	0.93	0.96	1.00	1.02	1.07	1.08	1.04	1.00	1.04
2003/1999	1.02	0.95	0.98	1.00	1.01	1.03	1.00	1.02	1.00	1.01

TABLE 2
Average Firm Characteristic by Firm Size (2003Q4)

Value Bins (in $)	UPCs	Brands	Modules	Product Groups	Share
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	1191	68	110	28	0.12
> 1,000,000,000	3431	182	246	54	0.16

Note: The share is based on total value of UPCs for firms within a bin compared to total value in 2003Q4
Ignore entry and exit of firms, only consider firms that exists in both periods t and t-1

TABLE 3
Product Entry and Exit in the U.S.

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
Ratio Share Common (t/t-1)	0.57	0.77	0.93

Notes: Entry Rate = number of new products (*t*) / total number of products (*t*)
Exit Rate = number of disappearing products (t-1) / total number of products (*t*-1)
Creation = Value(new products, *t*) / Total value (*t*)
Destruction = Value(disapp. products, *t*-1) / Total value (*t*-1)

TABLE 4
Share of Product Entry and Exit within Brands and within Manufacturer ID

Period	Brand Entry and Exit			Firm ID Entry and Exit		
	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

Table 5
Net and Gross Rates by Module, Size-Weighted Averages Ranked by SUM (over 1-year)

MODULE DESCRIPTION	RANKING (1)	CREATION (2)	DESTRUCTION (3)	NET ENTRY (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 APPL	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

Note: Select top 100 MODs by average value and then rank top 10/bottom 10 by SUM

TABLE 6
Correlations between Module Entry and Exit Variables

No Correction for Fixed Module Effects — Correlations between 1-year Entry and Exit Rates

	Exit			
Entry	2003	2002	2001	2000
2003	0.73	0.74	0.69	0.71
2002	0.72	0.70	0.74	0.73
2001	0.65	0.69	0.68	0.75
2000	0.64	0.65	0.67	0.69
	Destruction			
Creation	2003	2002	2001	2000
2003	0.84	0.76	0.77	0.73
2002	0.86	0.80	0.78	0.76
2001	0.89	0.85	0.83	0.81
2000	0.83	0.86	0.87	0.83

Correction for Fixed Module Effects — Correlations between 1-year Entry and Exit Rates

	Exit			
Entry	2003	2002	2001	2000
2003	0.24	0.23	-0.02	-0.07
2002	0.16	0.07	0.20	0.05
2001	-0.38	-0.29	-0.23	-0.10
2000	-0.04	-0.04	0.06	0.11
	Destruction			
Creation	2003	2002	2001	2000
2003	0.55	-0.29	0.04	-0.44
2002	0.45	-0.05	-0.22	-0.24
2001	-0.45	0.33	-0.12	0.37
2000	-0.56	0.12	0.17	0.37

TABLE 7
Destruction Rates (Over 4-Quarters) by UPC Characteristic

	Age	
Age in years	Destruction	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

	UPC Share (in its module)	
Share (t-1)	Destruction	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

	Brand Size	
# of UPCs per Brand	Destruction	Share (in t-1)
1 - 9	0.119	0.150
10 - 99	0.062	0.550
100 - 999	0.068	0.271
1000 - 9999	0.332	0.030

TABLE 8
Market Share and Average Product Size
(Mean and Standard Deviations across 1090 Modules)

	Market Share					
	1994	1999	2000	2001	2002	2003
1999 UPCs		1.00	0.92	0.83	0.76	0.70
			(0.11)	(0.18)	(0.21)	(0.23)
1994 UPCs	1.00	0.58	0.54	0.50	0.47	0.44
		(0.32)	(0.31)	(0.31)	(0.30)	(0.29)
	Cumulative Exit Rates					
	1994	1999	2000	2001	2002	2003
1999 UPCs			0.20	0.30	0.36	0.41
			(0.09)	(0.12)	(0.14)	(0.15)
1994 UPCs		0.47	0.50	0.54	0.57	0.59
		(0.14)	(0.14)	(0.15)	(0.15)	(0.15)
	Average Size of Surviving Products Relative to All Products in the Module					
	1994	1999	2000	2001	2002	2003
1999 UPCs		1.00	1.17	1.18	1.21	1.23
			(0.14)	(0.21)	(0.26)	(0.32)
1994 UPCs	1.00	1.46	1.46	1.46	1.48	1.49
		(0.60)	(0.63)	(0.68)	(0.74)	(0.79)

TABLE 9
Cyclicality at the Product Group Level

	Full Sample			Above Average Consumption Growth			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	1815	1815	1815	875	875	875	940	940	940
Number of rpg	122	122	122	119	119	119	119	119	119
R-squared	0.18	0.12	0.01	0.11	0.07	0	0.08	0.04	0.02

All regressions include product group fixed effects. Standard errors in brackets
* significant at 5% level; ** significant at 1% level

TABLE 10

Inflation under different Sampling Methodologies
2003 - 1994

Sample Size	280	280	400	400
Random Draw	Unweighted	Weighted	Unweighted	Weighted
Iterations	250	250	250	250
Base	99	99	99	99
	1-Quarter Differences (Annualized rate)			
Mean	0.73	0.75	0.76	0.75
St.Dev.	0.56	0.49	0.44	0.37
Coef. of Variation	0.8	0.7	0.6	0.5
	4-quarter Differences			
Mean	1.04	0.95	1.10	0.95
St.Dev.	0.12	0.11	0.10	0.07
Coef. of Variation	0.1	0.1	0.1	0.1

Inflation Changes under different Sampling Methodologies
Annualized Inflation Change (percentage points) 2003 - 1994

Sample Size	280	280	400	400
Random Draw	Unweighted	Weighted	Unweighted	Weighted
Iterations	250	250	250	250
Base	99	99	99	99
	1-Quarter Differences (Annualized rate)			
Mean	-0.16	-0.37	-0.16	-0.37
St.Dev.	0.88	0.76	0.70	0.49
Coef. of Variation	-5.5	-2.1	-4.4	-1.3
	4-quarter Differences			
Mean	-0.52	-0.54	-0.53	-0.54
St.Dev.	0.18	0.17	0.14	0.11
Coef. of Variation	-0.3	-0.3	-0.3	-0.2

TABLE 11

Percentile	Per-Year Ratio of Common Shares		Elasticities of Substitution	
	Within	Across	Within	Across
1	0.82	0.96	3.3	1.8
5	0.84	0.96	4.0	3.0
10	0.87	0.97	4.7	3.3
25	0.90	0.98	6.5	4.5
Median	0.93	0.98	11.5	7.5
75	0.96	0.99	29.9	22.8
90	1.22	1.05	50.5	48.6
95	1.47	1.10	65.3	50.5
99	2.19	1.22	265.6	63.4
Number of Product Groups	122	122	122	122

TABLE 12
Cyclicality of the Quality Bias at the Product Group Level
Over 4-Quarter Periods

	LHS variable : Q4 Bias		
Weights	no	yes	yes
Year & Quarter Dummies	no	no	yes
Product Group Sales Growth	-0.057***	-0.106***	-0.097***
	[0.010]	[0.011]	[0.011]
Constant	-0.019***	-0.017***	-0.012***
	[0.001]	[0.001]	[0.002]
Observations	1236	1236	1236
R-squared	0.02	0.07	0.11

Standard errors in brackets; *** significant at 1%

FIGURE 1 A and B

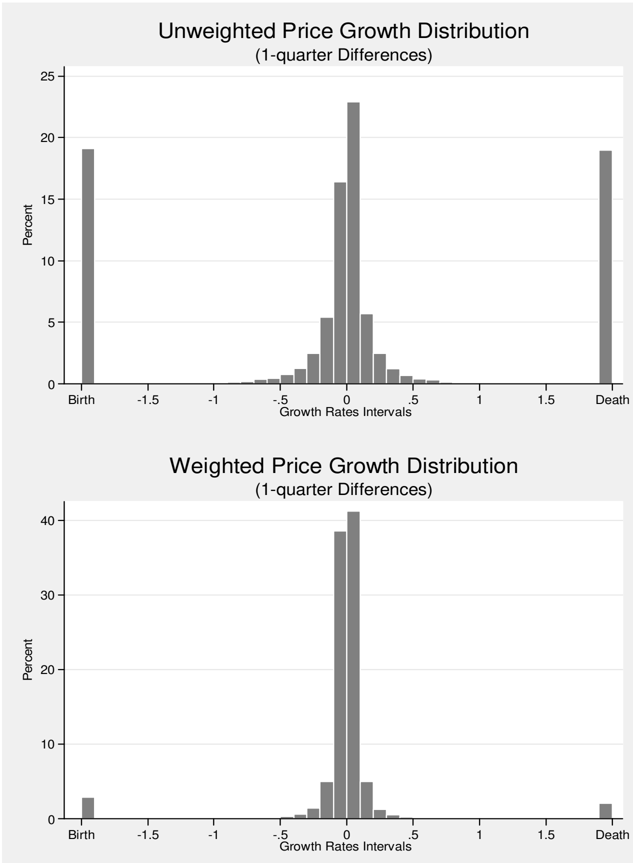

FIGURE 2 A and B

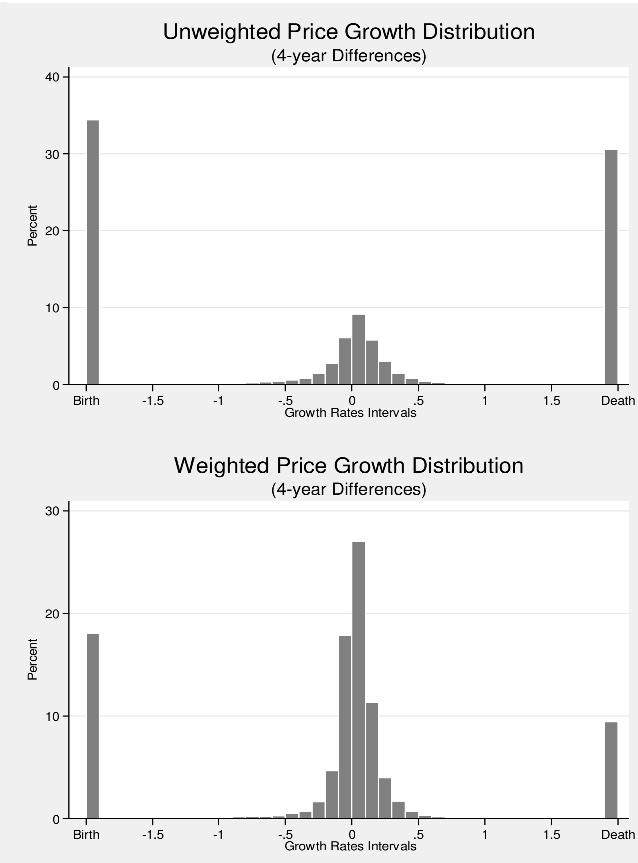

Figure 3A: Sales Growth and Net Creation (Q4/Q4 growth rates)

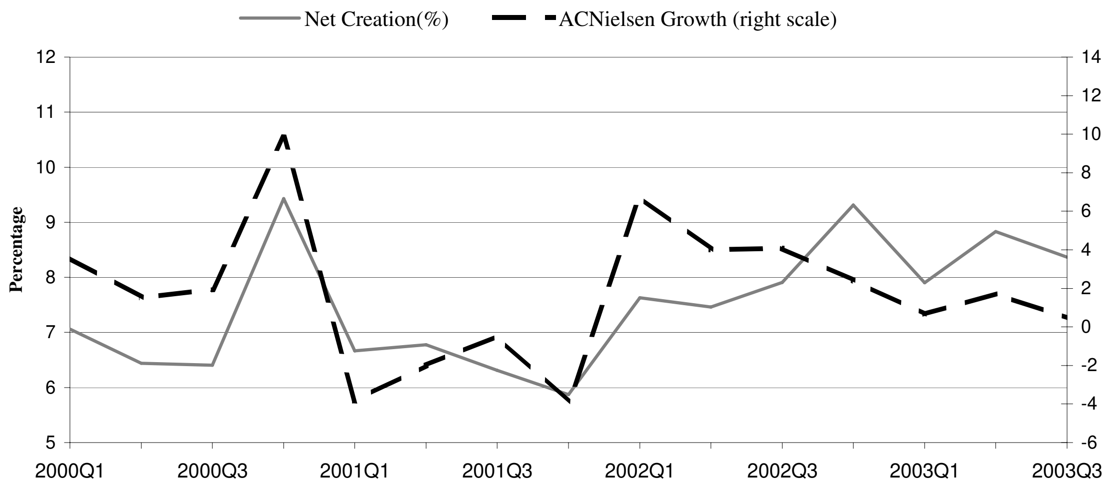

Figure 3B: Sales Growth and Creation (Q4/Q4 growth rates)

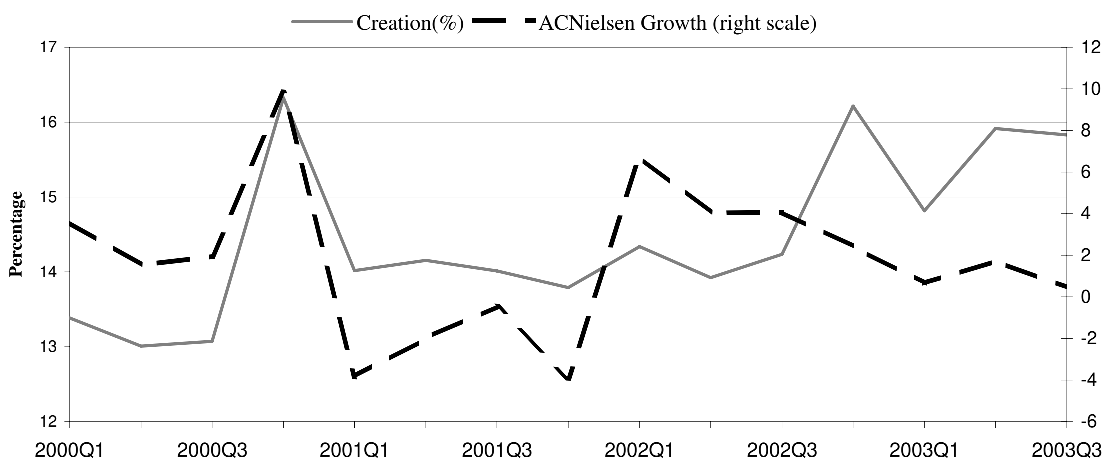

Figure 3C: Sales Growth and Destruction (Q4/Q4 growth rates)

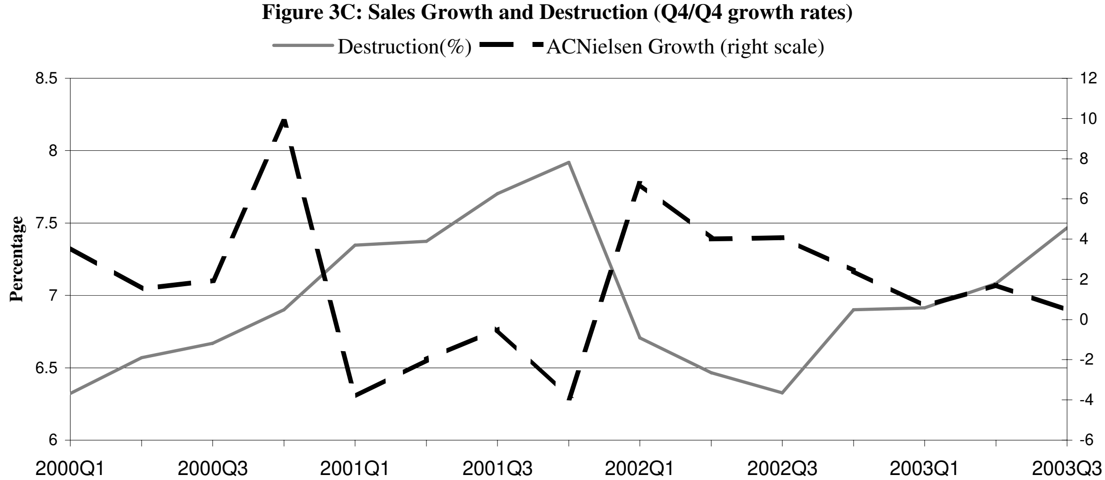

FIGURE 4
Kernel Density of Changes in Seasonally Adjusted 1-Quarter Inflation Rates during (1994 - 2003)

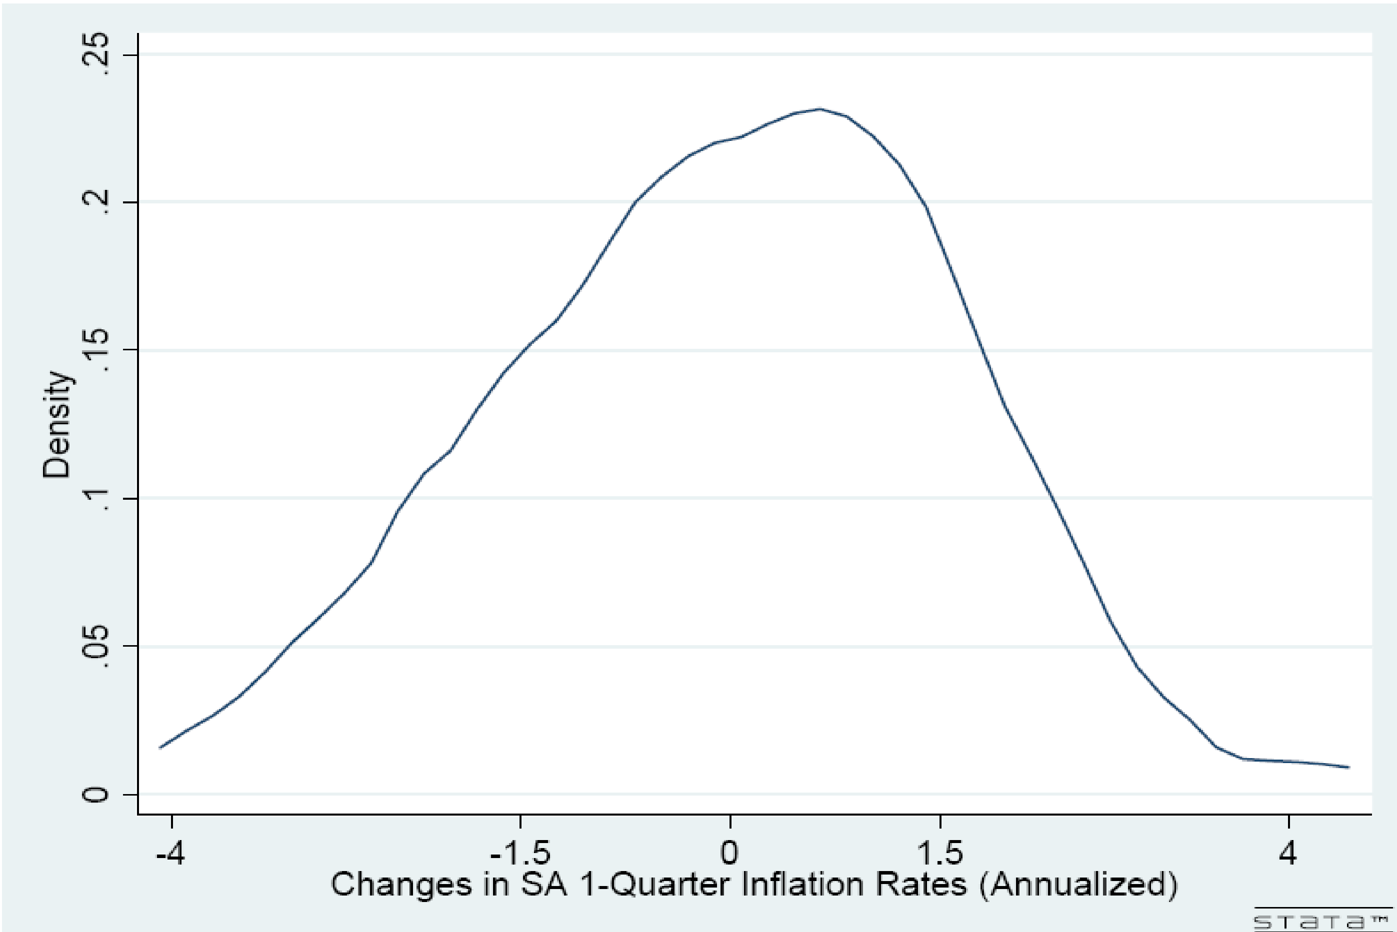

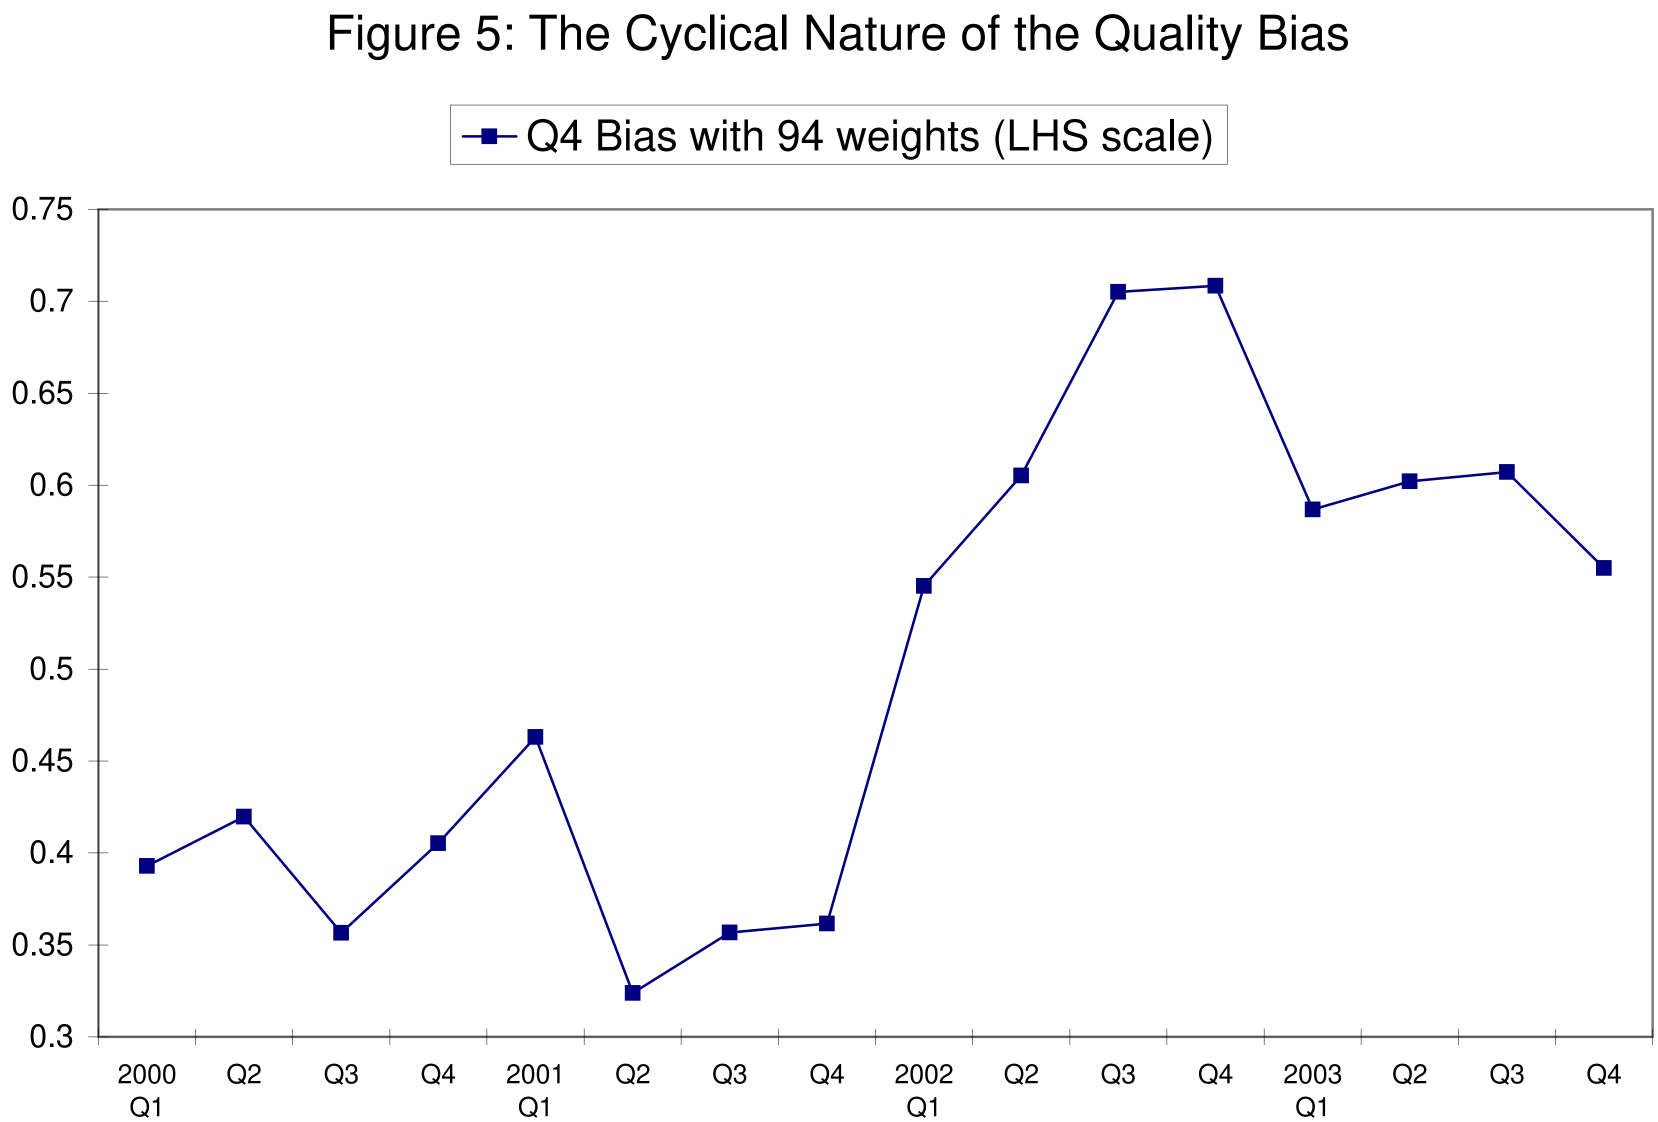
Figure 5: The Cyclical Nature of the Quality Bias
Q4 Bias with 94 weights (LHS scale)
0.75
0.7
0.65
0.6
0.55
0.5
0.45
0.4
0.35
0.3
2000 Q1
Q2
Q3
Q4
2001 Q1
Q2
Q3
Q4
2002 Q1
Q2
Q3
Q4
2003 Q1
Q2
Q3
Q4

Appendix to "Product Creation and Destruction: Evidence and Price Implications"
Broda and Weinstein (2007)

Appendix A: Sample List of 400 Modules (out of over 1000)

Description for 400 modules (out of over 1000)

ABRASIVE CLEANSERS-LIQUID	BUCKETS & PAILS	COOKING SAUCE	DRY DINNERS - REMAINING
ABRASIVE CLEANSERS-POWDERED	BURNER AND RANGE APPLIANCE	COOKING SPRAYS	DRY DINNERS - RICE
ACNE REMEDIES	BUTTER	COOKING WINE & SHERRY	DRY ERASE BULLETIN BOARD AND A
ADHESIVE BANDAGES	BUTTER-FRUIT & HONEY	COOKWARE PRODUCT	DUMPLINGS-CANNED
ADHESIVE BANDAGES - LIQUID - P	CAKE DECORATIONS & ICING	COOLERS - REMAINING	DUSTPANS
ADHESIVE NOTE PADS	CAMERAS	CORDIALS & PROPRIETARY LIQUEUR	DYE AND DYE REMOVER
ADHESIVE TAPE	CAN OPENER APPLIANCE	CORN DOGS-FROZEN/REFRIGERATED	EAR DROPS
ADULT-INCONTINENCE	CANADIAN WHISKEY	CORN MEAL	EGG COLORING KITS/DYE
AFTER SHAVE COSMETICS - MEN'S	CANDLE AND CANDLE IN HOLDER	CORN/POTATO STARCH	EGG MIXES-DRY
AIR CONDITIONER APPLIANCE	CANDLE HOLDER AND ACCESSORY	CORNED BEEF - CANNED	EGGNOG - FRESH & CANNED
AIR PURIFIER AND CLEANER APPLI	CANDY-CHOCOLATE	CORNED BEEF HASH-CANNED	EGGS-FRESH
AIR/SPECIALTY FRESHENERS - REM	CANDY-CHOCOLATE-MINIATURES	CORRECTION FLUID AND ERASERS	ELECTRIC KNIFE AND PEELER APPL
AIR/SPECIALTY FRESHENERS - SOL	CANDY-CHOCOLATE-SPECIAL	COSMETIC AND NAIL GROOMING ACC	ENEMAS-READY-TO-USE
AIR/SPECIALTY FRESHNERS - AERO	CANDY-DIETETIC - CHOCOLATE	COSMETIC KITS	ENGINE TREATMENT AND ADDITIVE
ALCOHOLIC COCKTAILS	CANDY-DIETETIC - NON-CHOCOLATE	COSMETICS - APPLICATOR BRUSHES	ENTREES - ITALIAN - 1 FOOD - F
ALE	CANDY-HARD ROLLED	COSMETICS - CONCEALERS	ENTREES - ITALIAN - 2 FOOD - F
ALUMINUM FOIL	CANDY-KITS	COSMETICS - NONCOTTON APLCTRS/	ENTREES - MEAT - 1 FOOD - FROZ
AMMONIA	CANDY-LOLLIPOPS	COSMETICS-BLUSHERS	ENTREES - MEAT - 2 FOOD - FROZ
ANALGESIC & CHEST RUBS	CANDY-NON-CHOCOLATE	COSMETICS-EYE SHADOWS	ENTREES - MEXICAN - 1 FOOD - F
ANCHOVY PASTE	CANDY-NON-CHOCOLATE-MINIATURES	COSMETICS-EYEBROW & EYE LINER	ENTREES - MEXICAN - 2 FOOD - F
ANTACIDS	CANNED AIR	COSMETICS-FACE POWDER	ENTREES - MULTI PACK - FROZEN
ANTI-GAS PRODUCTS	CAPERS	COSMETICS-FOUNDATION-CREAM AND	ENTREES - ORIENTAL - 1 FOOD -
ANTI-SLEEP PRODUCTS	CAT & DOG LITTER	COSMETICS-FOUNDATION-LIQUID	ENTREES - ORIENTAL - 2 FOOD -
ANTI-SMOKING PRODUCTS	CAT AND DOG LITTER	COSMETICS-LIPSTICKS	ENTREES - POULTRY - 1 FOOD - F
ARTIST AND HOBBY PAINT AND SUP	CAT FOOD - DRY TYPE	COSMETICS-MASCARA	ENTREES - POULTRY - 2 FOOD - F
AUTOMATIC DISHWASHER COMPOUNDS	CAT FOOD - MOIST TYPE	COSMETICS-NAIL POLISH	ENTREES - REMAINING - 1 FOOD -
AUTOMOTIVE COMBINATIONS	CAT FOOD - WET TYPE	COSMETICS-NAIL POLISH REMOVER	ENTREES - REMAINING - 2 FOOD -
BABY ACCESSORY	CATSUP	COSMETICS-REMAINING	ENTREES - SEAFOOD - 1 FOOD - F
BABY BIB AND BURP CLOTH	CEREAL - GRANOLA & NATURAL TYP	COTTON - SWABS/BALLS/ROLLS/APL	ENTREES - SEAFOOD - 2 FOOD - F
BABY BOTTLES & NIPPLES	CEREAL - HOT	COUGH AND COLD THROAT SPRAYS	ENTREES-REFRIGERATED
BABY CARE PRODUCTS-BATH	CEREAL - READY TO EAT	COUGH DROPS	ENTREES/SIDE DISHES - SHELF ST
BABY CARE PRODUCTS-LOTIONS	CHARCOAL	COUGH SYRUPS & TABLETS	EXTRACTS
BABY CARE PRODUCTS-OIL	CHARCOAL/WOOD LIGHTERS	CRACKERS - CHEESE	EYE CARE - REMAINING
BABY CARE PRODUCTS-OINTMENTS	CHEESE - COTTAGE	CRACKERS - FLAKED SODA	EYE DROPS & LOTIONS
BABY CARE PRODUCTS-POWDER	CHEESE - FARMERS	CRACKERS - FLAVORED SNACK	FABRIC FINISHERS
BABY CEREAL & BISCUITS	CHEESE - GRATED	CRACKERS - GRAHAM	FABRIC PROTECTORS
BABY FOOD - JUNIOR	CHEESE - NATURAL - AMERICAN CH	CRACKERS - OYSTER	FABRIC SOFTENERS-AEROSOL
BABY FOOD - STRAINED	CHEESE - NATURAL - AMERICAN CO	CRACKERS - REMAINING	FABRIC SOFTENERS-DRY
BABY JUICE	CHEESE - NATURAL - BRICK	CRACKERS - SANDWICH & SNACK PA	FABRIC SOFTENERS-LIQUID
BABY MILK AND MILK FLAVORING	CHEESE - NATURAL - MOZZARELLA	CRACKERS - SPRAYED BUTTER	FABRIC WASHES - SPECIAL
BABY PACFR/TEETHR & BOTTLE/NIP	CHEESE - NATURAL - REMAINING	CRACKERS - SPRAYED FLAKE	FACE CLEANSERS & CREAMS & LOTI
BACON-BEEF & CANNED	CHEESE - NATURAL - VARIETY PAC	CRACKLINS - REFRIGERATED	FACIAL TISSUE
BACON-REFRIGERATED	CHEESE - PROCESSED - CREAM CHE	CRANBERRIES - SHELF STABLE	FALSE EYELASH AND ACCESSORY
BAGS - FOOD STORAGE	CHEESE - PROCESSED - LOAVES	CRAYONS	FALSE NAIL AND NAIL DECORATION
BAGS - FREEZER	CHEESE - PROCESSED - SNACK	CREAMERS - POWDERED	FAN AND CEILING FAN APPLIANCE
BAGS - LAWN & LEAF	CHEESE - PROCESSED SLICES - RE	CREAMERS-LIQUID	FEATHER DUSTERS
BAGS - OVEN	CHEESE - RICOTTA	CREME RINSES & CONDITIONERS	FEMININE HYGIENE-DEODORANT SPR
BAGS - PAPER - LUNCH	CHEESE - SHREDDED	CROUTONS	FEMININE HYGIENE-DOUCHES
BAGS - SANDWICH	CHEESE - SPECIALTY/IMPORTED	DAIRY - DIP - REFRIGERATED & F	FEMININE HYGIENE-MISCELLANEOUS
BAGS - TALL KITCHEN	CHEESE-NATURAL-MUENSTER	DAIRY - POTATO TOPPING - REFRI	FEMININE HYGIENE-TOWELETTES
BAGS - TRASH/TRASH COMPACTOR	CHEESE-NATURAL-SWISS	DAIRY - SOUR CREAM - REFRIGERA	FILM
BAGS - WASTE	CHEESE-PROCESSED SLICES-AMERIC	DAIRY DIGESTIVE AIDS	FIREPLACE LOGS
BARBECUE SAUCES	CHERRIES-MARASCHINO	DAIRY-BUTTERMILK-REFRIGERATED	FIRST AID - GAUZE - ROLLS
BARBECUED BEEF & PORK-CANNED	CHICKEN - SHELF STABLE	DAIRY-CREAM-REFRIGERATED	FIRST AID - GAUZE PADS
BARLEY-DRY	CHILDREN'S COLOGNE & GIFT SETS	DAIRY-FLAVORED MILK-REFRIGERAT	FIRST AID - HYDROCORTISONES
BATH ADDITIVES - DRY	CHILI SAUCE	DAIRY-MILK-REFRIGERATED	FIRST AID - ICE AND HEAT PACK
BATH ADDITIVES - LIQUID	CHILI-REFRIGERATED	DATES	FIRST AID - THERMOMETERS
BATH OIL - DRY	CHILI-SHELF STABLE	DENTAL ACCESSORIES	FIRST AID - TREATMENTS
BATH OIL - LIQUIDS	CHILIES	DENTAL FLOSS	FISH & SEAFOOD & COCKTAIL SAUC
BATHROOM ACCESSORY	CHOCOLATE CHIPS & MORSELS	DENTURE ADHESIVES	FLASHLIGHT BULBS
BATHROOM SCALE	CIGARETTES	DENTURE CLEANSERS	FLASHLIGHTS
BATTERIES	CIGARS	DEODORANTS - COLOGNE TYPE	FLOOR CARE - WAXES
BATTERY CHARGERS	CLAM JUICE	DEODORANTS - PERSONAL	FLOOR CARE-CLEANERS
BEAN SPROUTS-CANNED	CLEANERS - BATHROOM	DEPILATORIES - MEN'S	FLOUR - ALL PURPOSE - REMAININ
BEANS-DRY	CLEANERS - DISINFECTANTS	DEPILATORIES - WOMEN'S	FLOUR - SINGLE PURPOSE
BEER	CLEANERS - NON-DISINFECTANT	DESSERTS - RTS SINGLE SERVINGS	FLOUR-ALL PURPOSE-WHITE WHEAT
BEER & WINE MAKING KITS	CLEANERS - POWDERS	DETERGENT BOOSTERS	FONDUE SAUCE
BEVERAGE STORAGE CONTAINER	CLEANERS - WINDOW	DETERGENTS - HEAVY DUTY - LIQU	FOOD COLORING
BLANK AUDIO PRODUCT	CLEANERS-HUMIDIFIERS/VAPORIZER	DETERGENTS - LIGHT DUTY	FOOD PROCESSOR AND GRINDER APP
BLANK VIDEO PRODUCT	CLEANERS-METAL	DETERGENTS-PACKAGED	FOOD STORAGE CONTAINERS
BLEACH - DRY	CLEANERS-PASTE & JELLY	DEVILED HAM - CANNED	FOOT COMFORTS PRODUCTS
BLEACH - LIQUID/GEL	CLEANERS-SEPTIC TANK	DIARRHEA REMEDIES	FOOT PREPARATIONS-ATHLETE'S FO
BLENDER APPLIANCE	CLOTH-POLISHING/CLEANING	DIETING AIDS - APPETITE SUPPRE	FOOT PREPARATIONS-REMAINING
BLOOD PRESSURE KIT AND ACCESSO	CLOTHESPINS	DIETING AIDS-COMPLETE NUTRITIO	FRANKFURTERS-REFRIGERATED
BLOOD URINE STOOL TEST PRODUCT	COCKTAIL MIXES-DRY	DINNERS-FROZEN	FRANKS-COCKTAIL-REFRIGERATED
BLUINGS	COCKTAIL MIXES-LIQUID	DIP - CANNED	FREEZER SUPPLIES
BODY MASSAGER APPLIANCE AND AC	COCKTAIL ONIONS	DIP - MIXES	FRUIT DRINKS & JUICES-CRANBERR
BODY WARMERS	COCKTAIL PRODUCTS-BITTERS & HE	DISHWASHER RINSING AIDS	FRUIT DRINKS & MIXES - FROZEN
BORATEEM	COCOA	DISHWASHING ACCESSORY PRODUCT	FRUIT DRINKS - ORANGE - FROZEN
BOUILLON	COCONUT	DISINFECTANTS	FRUIT DRINKS-CANNED
BOURBON-BLENDED	COFFEE - LIQUID	DISK DISKETTE AND DATA CARTRID	FRUIT DRINKS-OTHER CONTAINER
BOURBON-STRAIGHT/BONDED	COFFEE - SOLUBLE	DISPOSABLE CUPS	FRUIT JUICE - APPLE
BRANDY/COGNAC	COFFEE - SOLUBLE FLAVORED	DISPOSABLE DIAPERS	FRUIT JUICE - APPLE - FROZEN
BRATWURST & KNOCKWURST	COFFEE AND TEA MAKER APPLIANCE	DISPOSABLE DISHES	FRUIT JUICE - GRAPE
BREAD-SPECIALTY-CANNED	COFFEE FILTERS - DISPOSABLE	DIURETIC REMEDIES	FRUIT JUICE - GRAPE - FROZEN
BREADING PRODUCTS	COFFEE SUBSTITUTES	DIVIDERS TABS LABELS AND TAGS	FRUIT JUICE - GRAPEFRUIT - FRO
BREADMAKER APPLIANCE	COLD REMEDIES - ADULT	DOG & CAT TREATS	FRUIT JUICE - ORANGE - FROZEN
BREAKFAST BARS	COLD REMEDIES - CHILDREN	DOG FOOD - DRY TYPE	FRUIT JUICE - ORANGE - OTHER C
BREAKFAST DRINKS - POWDERED	COLOGNE & PERFUME-WOMEN'S	DOG FOOD - MOIST TYPE	FRUIT JUICE - REMAINING - FROZ
BREATH FRESHENERS	COLOGNE/LOTION-MEN'S	DOG FOOD - WET TYPE	FRUIT JUICE - UNCONCENTRATED -
BREATH SWEETENERS	COMBINATION LUNCHES	DOUGH PRODUCTS - COOKIES & BRO	FRUIT JUICE-GRAPEFRUIT-CANNED
BREATHING AIDS EXTERNAL	COMPLETE NUTRITIONAL PRODUCTS	DOUGH PRODUCTS-BISCUITS-REFRIG	FRUIT JUICE-ORANGE-CANNED
BRONCHIAL REMEDIES	COMPUTER SOFTWARE	DOUGH PRODUCTS-BREAD-FROZEN	FRUIT JUICE-PRUNE
BROOMS/ MOPS & WAX APPLICATORS	CONFECTIONERY PASTE	DOUGH PRODUCTS-DINNER ROLLS-RE	FRUIT PECTINS
BRUSHES - AUTOMOTIVE	CONTACT LENS SOLUTION	DOUGH PRODUCTS-REMAINING-REFRI	FRUIT PROTECTORS
BRUSHES - MISCELLANEOUS	CONTRACEPTIVES-FEMALE	DOUGH PRODUCTS-SWEET ROLLS-REF	FRUIT PUNCH BASES & SYRUPS
BRUSHES-BATHROOM	CONTRACEPTIVES-MALE	DRAIN PIPE OPENERS	FRUIT SALADS-REFRIGERATED
BRUSHES-KITCHEN & SCRUB	COOKER STEAMER AND DEHYDRATOR	DRIED BEEF-SHELF STABLE	FRUIT SPREADS
BUCKETS & PAILS	COOKIES	DRINKWARE CONTAINER SET	FRUIT-DRIED AND SNACKS
BURNER AND RANGE APPLIANCE	COOKING BAGS W/SEASONING	DRY DINNERS - PASTA	FRUIT-GLAZED

Appendix B
Value per UPC and non-barcode goods

One concern about the database is that results might be driven by the growth in the number of barcodes *per se*. Fortunately, ACNielsen has provided us with data on the overall purchases of products with and without barcodes in the same expenditure categories for each quarter. If it has become easier for firms to obtain a barcode over the period we investigate, we would expect the share of goods with a barcode over total sales to rise over time. In Figure 1, we plot (solid line, right hand side scale) the share of sales of goods with barcodes relative to total shopping expenditures over time. Unfortunately, we cannot do this in 1994 due to an error in the sampling of total purchases. However, the share of goods purchased with a barcode relative to total shopping purchases stood at an almost constant value of 0.85 between the first quarter of 1999 and the fourth quarter of 2003 which indicates that there was no movement in the share of goods with barcodes over this time period.

Another way of checking that the ratio of goods with barcodes relative to those without barcode has remained relatively unchanged over the sample period is to check the average expenditure per UPC over time. Observing a declining expenditure per UPC would be suggestive of a fall in the cost of obtaining a UPC over the sample period. In Figure 1, we plot (dashed line, right hand scale) movements in the average expenditure per UPC deflated by the food and beverages component of the CPI at the quarterly frequency. As one can see, while there are some small fluctuations in the average amount of expenditure per UPC, overall the series is remarkably flat. This indicates that over the four years for which we have data growth in total sales moved approximately proportionally with the number of UPCs.

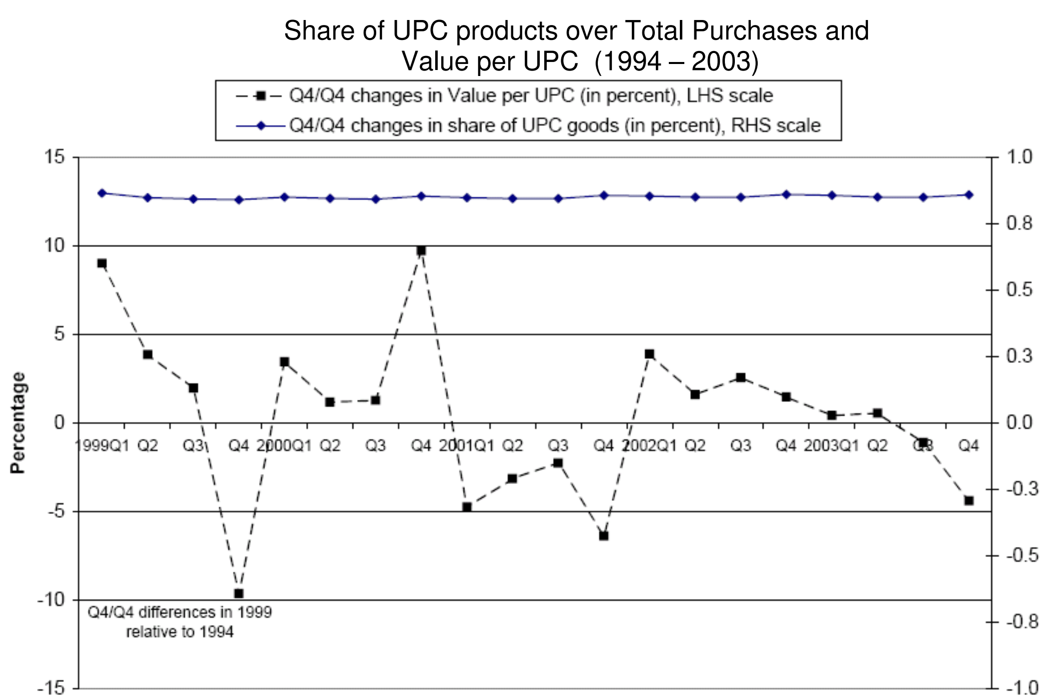

Appendix C
Levels of Aggregation in the ACNielsen Database

Below is an example of UPCs, Modules and Product Groups in our database. UPC numbers are illustrative and not the actual codes.

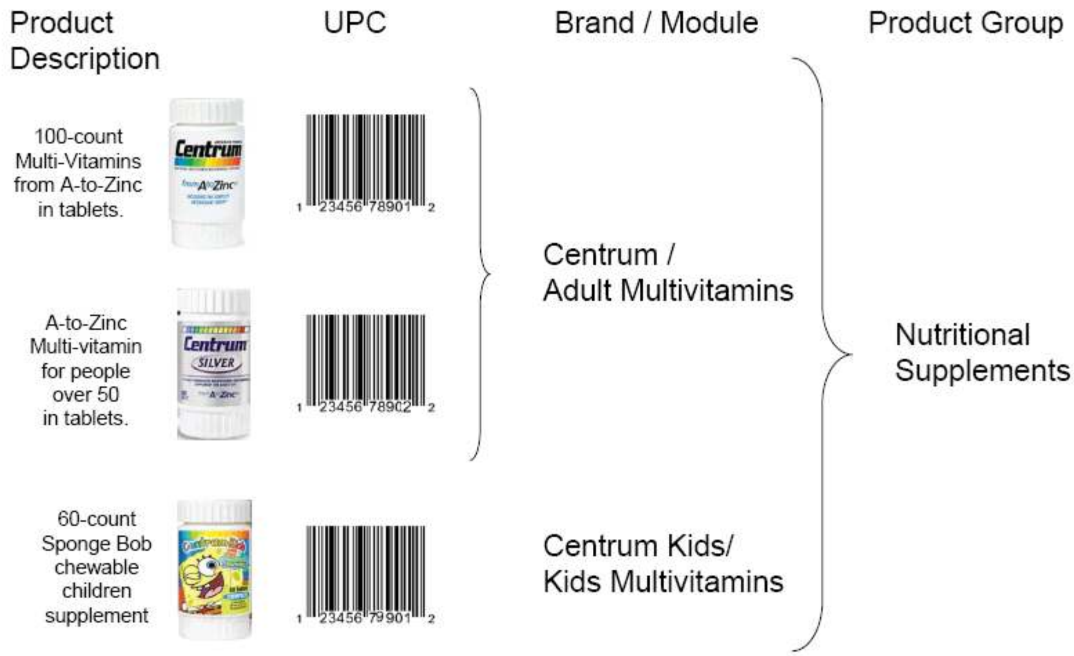

Appendix D
Product Creation and Destruction excluding UPCs with low number of "raw buyers"

Some of the UPCs in our data are purchased by only a small fraction of the overall number of households. In order to show that the main effects in the paper are not driven by these UPCs we replicate Table 3 in the paper but exclude those UPCs with less than 20 households purchasing them. 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 churning.

Appendix Table 3: Product Entry and Exit in the U.S. dropping UPCs with less than 20 "raw buyers"

Period	9-year 1994 - 2003	4-year 1999 - 2003	1-year Median
Entry Rate	0.75	0.49	0.25
Creation	0.64	0.37	0.10
Entrant Relative Size	0.57	0.60	0.35
Exit Rate	0.69	0.43	0.23
Destruction	0.42	0.22	0.05
Exiter Relative Size	0.32	0.32	0.17
Ratio Share Common (t/t-1)	0.63	0.81	0.94

Notes: All UPCs with less than 20 households buying it are dropped from the table.

Appendix E
Product Creation and Destruction due to size and flavor only

For roughly 20 percent of the products that were purchased in Q4 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 sizes and flavors of existing products. For instance, a new UPC can be a "200-count Centrum Multi-Vitamins From A-to-Zinc in tablets" which differs from the UPC described in section 2 only in the amount of vitamins included in the bottle. Thus we can calculate how much of overall creation is due to innovations in size and how much is due to innovations in flavor. We find that for this sample of goods the overall creation is 35.3 percent, which is very similar to that in the overall sample. Creation from new sizes is 1.9 percent or roughly 5 percent of overall creation. Adding new flavors raises creation to 2.3 percent, which still is a small proportion of overall creation.

Share of Entry and Creation due to New Sizes and Flavors

	1999-2003
Creation	35.3%
Creation due to New Size	1.9%
Creation due to New Size&Flavor	2.3%

* This includes a sample of 20% of all goods in Q4 2003 for which size and flavor data is available.

Appendix F
Bootrapped Standard Errors by quarter and Food and Beverage CPI

The figure below shows the 90th confidence interval for inflation using ACNielsen common set of goods and the actual inflation rate from the BLS for the Food and Beverage component of the CPI. As discussed in the paper, the ACNielsen includes a larger set of goods than that included in the CPI's Food and Beverage component.

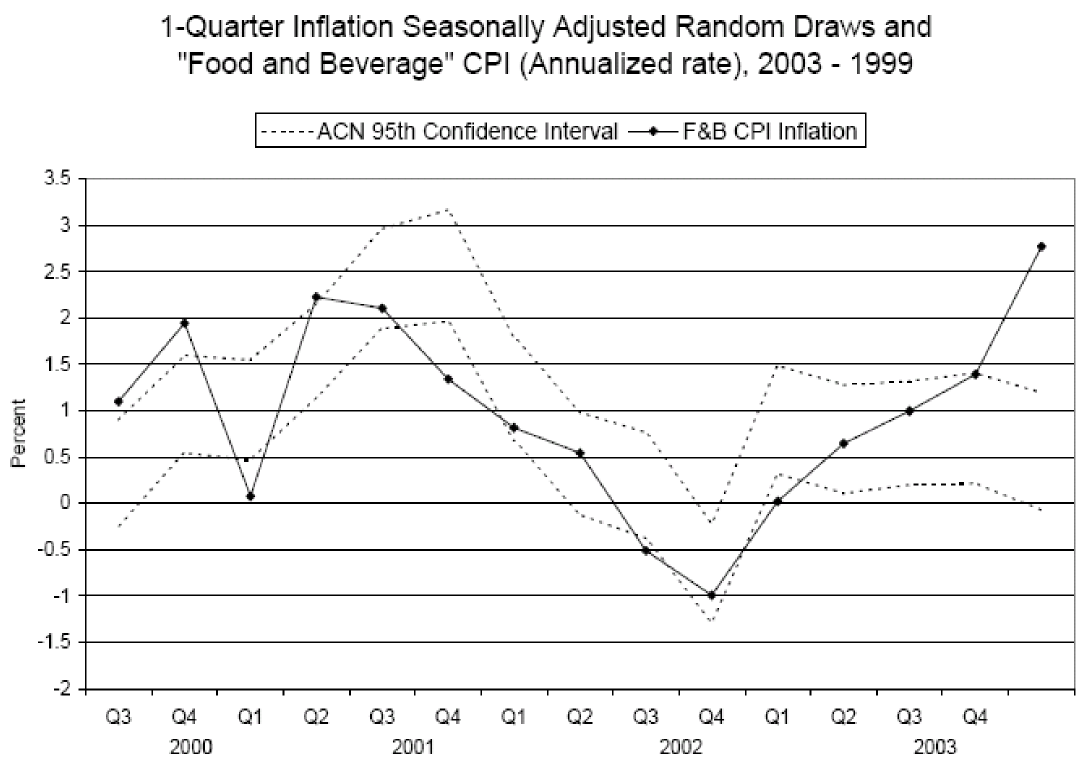

Appendix G
Per year Substitution Bias relative to Tornqvist (in percent)

	Sample and Weights 1994				
Avg Lower	Laspayres	Paasche	Geometric	Geometric	Ideal CES
Avg Upper	Laspayres	Paasche	Laspayres*	Tornqvist**	Ideal CES
03Q4-94Q4	0.36	-0.37	0.14	0.00	0.00
03Q4-99Q4	0.31	-0.14	0.17	0.00	0.00
99Q4-94Q4	0.36	-0.52	0.10	0.00	0.00

	Sample and Weights 1999				
Avg Lower	Laspeyres	Paasche	Geometric	Geometric	Ideal CES
Avg Upper	Laspeyres	Paasche	Laspeyres	Tornqvist	Ideal CES
03Q4-99Q4	0.25	-0.25	0.07	0.00	0.00

* Closest to the CPI methodology.
** Closest to the C-CPI methodolgy

Appendix H
Product Group with the largest contribution to the Quality/New Goods Bias

Product Group Name	Weight in Consumption	Contributin to Quality Bias
ELECTRONICS, RECORDS, TAPES	3.04%	12.40%
PREPARED FOODS-FROZEN-READY TO SERVE	2.80%	7.06%
DRUGS, REMEDIES (NON-PAIN), MEDICAL ACCESSORIES	1.87%	4.75%
SNACKS	2.62%	4.20%
HOUSEWARES, APPLIANCES	1.86%	3.68%
PAIN REMEDIES	0.65%	3.48%
DEODORANT	0.39%	3.15%
LAUNDRY SUPPLIES	0.86%	3.01%
ICE CREAM, NOVELTIES	1.19%	2.47%
COFFEE	0.88%	2.37%
PAPER PRODUCTS	2.55%	2.25%
HAIR CARE	1.13%	2.16%
CEREAL	2.36%	1.92%
FROZEN NOVELTIES	0.55%	1.86%
KITCHEN GADGETS	0.53%	1.84%
HOUSEHOLD SUPPLIES	0.83%	1.72%
CHEESE	2.14%	1.60%
DRESSINGS/SALADS/PREP FOODS-DELI	1.44%	1.60%
SOFT DRINKS - LOW CALORIES	0.61%	1.54%

Appendix I
CPI Inflation (1996-2006)

Extrapolating the quality bias and the sampling error to the entire CPI suggests that over the last 10 years, we cannot sign inflation precisely almost 30 percent of the time. The density below the red density function between the dotted lines (+/- 0.98 percent) is equal to 28 percent.