

Missing Growth from Creative Destruction

Philippe Aghion, Antonin Bergeaud

Timo Boppart, Peter J. Klenow, Huiyu Li*

November 2017

Abstract

Statistical agencies typically impute inflation for disappearing products based on surviving products, which may result in overstated inflation and understated growth. Using U.S. Census data, we apply two ways of assessing the magnitude of “missing growth” for private nonfarm businesses from 1983–2013. The first approach exploits information on the market share of surviving plants. The second approach applies indirect inference to firm-level data. We find: (i) missing growth from imputation is substantial — at least 0.6 percentage points per year; and (ii) most of the missing growth is due to creative destruction (as opposed to new varieties).

*Aghion: Collège de France and London School of Economics; Bergeaud: Paris School of Economics-EHESS and Bank of France; Boppart: IIES, Stockholm University; Klenow: Stanford University (klenow@stanford.edu); Li: Federal Reserve Bank of San Francisco. We thank Raouf Boucekkine, Pablo Fajgelbaum, Colin Hottman and Stephen Redding for excellent discussions and Victoria De Quadros for superb research assistance. Ufuk Akcigit, Robert Feenstra, Chang-Tai Hsieh, Xavier Jaravel, Chad Jones, Per Krusell, Torsten Persson, Ben Pugsley, John Van Reenen and numerous seminar participants provided helpful comments. Opinions and conclusions herein are those of the authors and do not necessarily represent the views of the Federal Reserve System. All results have been reviewed to ensure that no confidential information from the U.S. Census Bureau or U.S. Bureau of Labor Statistics have been disclosed.

1 Introduction

Whereas it is straightforward to compute inflation for an unchanging set of goods and services, it is much harder to separate inflation from quality improvements and variety expansion amidst a changing set of items. In the U.S. Consumer Price Index (CPI), over 3% of items exit the market each month (Bils and Klenow, 2004). In the Producer Price Index (PPI) the figure is over 2% per month (Nakamura and Steinsson, 2008).

The Boskin Commission (Boskin et al., 1996) highlighted the challenges of measuring quality improvements when incumbents upgrade their products. It also maintained that the CPI does not fully capture the benefits of brand new varieties. We argue that there exists a subtler, overlooked bias in the case of creative destruction. When the producer of the outgoing item does not produce the incoming item, the standard procedure at statistical offices is to resort to some form of *imputation*. Imputation inserts the average price growth among a set of surviving products that were not creatively destroyed.¹ We think this misses some growth because inflation is likely to be below-average for items subject to creative destruction.²

Creative destruction is believed to be a key source of economic growth. See Aghion and Howitt (1992), Akcigit and Kerr (2010), and Aghion, Akcigit and Howitt (2014). We therefore attempt to quantify the extent of “missing growth”—the difference between actual and measured productivity growth—due to the use of imputation in cases of creative destruction. Our estimates are for the U.S. nonfarm business sector over the past three decades.

¹U.S. General Accounting Office (1999) details CPI procedures for dealing with product exit. For the PPI, “If no price from a participating company has been received in a particular month, the change in the price of the associated item will, in general, be estimated by averaging the price changes for the other items within the same cell for which price reports have been received.” (U.S. Bureau of Labor Statistics, 2015, p.10) The BLS makes explicit quality adjustments, such as using hedonics, predominantly for goods that undergo periodic model changes by *incumbent* producers (Groshen et al., 2017).

²A similar bias could arise due to creative destruction at times of regular rotation of items in the CPI and PPI samples.

In the first part of the paper we develop a growth model with (exogenous) innovation to provide explicit expressions for missing growth. In this model, innovation may either create new varieties or replace existing varieties with products of higher quality. The quality improvements can be performed by incumbents on their own products, or by competing incumbents and entrants (creative destruction). The model predicts missing growth due to creative destruction if the statistical office resorts to imputation.

In the second part of the paper we use two alternative approaches to estimate the magnitude of missing growth based on our model. For both approaches we use micro data from the U.S. Census on employment at all private nonfarm businesses for the years 1983–2013. For the first approach we look at employment shares of continuing (incumbent), entering, and exiting plants. If new plants produce new varieties and carry out creative destruction, then the inroads they make in incumbents' market share should signal their contribution to growth.

In the second approach, we extend the algorithm from Garcia-Macia, Hsieh and Klenow (2016) to estimate the arrival rates and step sizes of the various kinds of innovations (creative destruction by entrants and incumbents, incumbent own innovation, expanding variety by entrants and incumbents). We then use our accounting framework to calculate missing growth. This second approach allows us to estimate the contribution of each of the different types of innovation to missing growth. It does not assume that only new plants introduce new varieties and carry out creative destruction, but does rely on indirect inference.

Our findings from these two quantifications can be summarized as follows. First, missing growth from imputation is substantial. We estimate that missing growth averages around one-third of true total productivity growth using the two methods. Second, the primary source of missing growth appears to be creative destruction rather than new varieties.

Example: The following numerical example illustrates how imputation can miss growth. Suppose that: (i) 80% of products in the economy experience no innovation in a given period and are subject to a 4% inflation rate; (ii) 10% of products experience quality improvement without creative destruction, with their quality-adjusted prices falling 6% (i.e., an inflation rate of -6%); and (iii) 10% of products experience quality improvement due to creative destruction, with their quality-adjusted prices also falling by 6%. The true inflation rate in this economy is then 2%. Suppose further that nominal output grows at 4%, so that true productivity growth is 2% after subtracting the 2% true inflation rate.

What happens if the statistical office resorts to imputation in cases of creative destruction? Then it will not correctly decompose growth in nominal output into its inflation and real growth components. Imputation means that the statistical office will ignore the goods subject to creative destruction when computing the inflation rate for the whole economy, and only consider the products that were not subject to innovation plus the products for which innovation did not involve creative destruction. Thus the statistical office will take the average inflation rate for the whole economy to be equal to

$$\frac{8}{9} \cdot 4\% + \frac{1}{9} \cdot (-6\%) = 2.9\%.$$

Presuming it correctly evaluates the growth in nominal GDP to be 4%, the statistical office will (incorrectly) infer that the growth rate of real output is

$$4\% - 2.9\% = 1.1\%.$$

This in turn implies "missing growth" in productivity amounting to

$$2\% - 1.1\% = 0.9\%.$$

This is end of the example, which hopefully clarifies the main mechanism by which imputation can miss growth from creative destruction.

Our study relates to several strands of literature. The first is the pioneering work of Abramovitz (1956), Jorgenson and Griliches (1967), Griliches (1996) and Diewert (2000) on the measurement of Total Factor Productivity (TFP).³ Second, our paper builds on the innovation-based endogenous growth literature (Romer, 1990; Grossman and Helpman, 1991; Aghion and Howitt, 1992). Third is the literature on growth, reallocation and firm dynamics (Klette and Kortum, 2004; Akcigit and Kerr, 2010; Acemoglu et al., 2013; Haltiwanger, 2015).⁴

Our paper touches on the recent literature about secular stagnation and growth measurement. Gordon (2012) observes that a rising flow of patented innovations has not been mirrored by an acceleration in measured TFP growth.⁵ He argues that innovation has run into diminishing returns, leading to an irreversible slowing of TFP growth. Syverson (2016) and Byrne, Fernald and Reinsdorf (2016) conclude that understated growth in certain sectors, most notably the ICT sector, cannot account for the productivity slowdown since 2005 because the ICT sector is small relative to the aggregate economy. In contrast to these studies, we look at missing growth for the whole economy, not just from the ICT sector. We find that missing growth has not risen in the past decade.

More closely related to our analysis are Feenstra (1994), Bils and Klenow (2001), Bils (2009), Broda and Weinstein (2010), Erickson and Pakes (2011), Byrne, Oliner and Sichel (2015), and Redding and Weinstein (2016).⁶ We make two contributions relative to these important papers. First, we compute

³See also Hulten (2000) and Lipsey and Carlaw (2000, 2004).

⁴See also Davis, Haltiwanger and Schuh (1998) and Bartelsman and Doms (2000).

⁵Related studies include Jones (1995), Kortum (1997), and Bloom et al. (2017).

⁶Feenstra (1994) corrects for biases in the U.S. import price indices of six manufacturing goods, in particular due to an expanding set of available product varieties. Bils and Klenow (2001) use the U.S. Consumer Expenditure surveys to estimate “quality Engel curves” and assess the unmeasured quality growth of 66 durable goods which account for 12% of consumer spending. Bils (2009) uses CPI micro data to decompose the observed price increases of durable goods into quality changes and true inflation. Broda and Weinstein (2010) look at missing growth from entry and exit of products in the nondurable retail sector, using an AC Nielsen database. Byrne, Oliner and Sichel (2015) look at missing growth in the semiconductor sector.

missing growth for the entire private nonfarm sector from 1983–2013.⁷ Second, we focus on a neglected source of missing growth, namely imputation in the event of product exit. We isolate missing growth from creative destruction as opposed to the more familiar quality improvements by incumbents on their own products and expanding variety. The missing growth we identify is likely to be exacerbated when there is error in measuring quality improvements by incumbents on their own products.⁸

The rest of the paper is organized as follows. In Section 2 we lay out a growth model and derive the expression for missing growth and how it relates to creative destruction. In Section 3 we use the two alternative approaches, based on continuing plants' market shares and indirect inference, respectively, to estimate missing growth using U.S. Census data. Section 4 concludes.

2 A model of missing growth

In this section we develop a simple accounting framework that allows us to analyze the determinants of missing growth in the aggregate economy from imputation in the event of creative destruction.

⁷Broda and Weinstein (2010) used AC Nielsen data from 1994 and 1999–2003. This database is heavily weighted toward nondurables, particularly food. Bils and Klenow (2004) report a product exit rate of about 2.4% per month for nondurables (1.2% a month for food) versus about 6.2% per month for durable goods. Hence, it is important to analyze missing growth across many sectors of the economy, including durables.

⁸Unlike Broda and Weinstein (2010), we do not assume that the BLS makes no effort to quantify such quality improvements. Bils (2009) estimates that the BLS subtracted 0.7 percentage points per year from inflation for durables over 1988–2006 due to quality improvements. For the whole CPI, Moulton and Moses (1997) calculate that the BLS subtracted 1.8 percentage points in 1995.

2.1 Basic setup

2.1.1 Structure of the aggregate economy

Time is discrete and in each period output has a CES structure:

$$Y = \left(\int_0^N [q(j) y(j)]^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}. \quad (1)$$

We assume that Y is competitively produced from intermediate inputs $y(j)$ that come at quality $q(j)$.⁹ N is the number of intermediate varieties available, which can grow over time. An alternative interpretation of Y is that it denotes utility of a representative consumer. With either interpretation, $\sigma \geq 1$ denotes the constant elasticity of substitution between the different intermediate goods.¹⁰

2.1.2 Intermediate input production

Each intermediate input $y(j)$ is produced one-for-one with labor, i.e., we have

$$y(j) = l(j),$$

where $l(j)$ is the amount of labor used to produce intermediate good j .¹¹

2.1.3 Resource constraint and market structure

The final good sector is assumed to be competitive. Hence, each intermediate good is paid its marginal product in producing the final good. Intermediate

⁹We remove time subscripts when it does not risk confusion.

¹⁰The production function in (1) is not well defined for $\sigma = 1$. In this special case, we assume technology takes a Cobb-Douglas form

$$Y = N \exp \left(\frac{1}{N} \int_0^N \log [q(j) y(j)] dj \right).$$

We will use $\sigma > 1$ for our quantitative exercises based on estimates in the literature. But the special case with $\sigma = 1$ is helpful for highlighting channels of missing growth (see [Online Appendix C](#)).

¹¹Below we present evidence in favor of quality improvements over process innovations.

producers are monopolistic but potentially subject to a competitive fringe. A representative household supplies a fixed amount of labor every period. Labor is freely mobile across firms and the wage rate, W , is common across firms in equilibrium. Nominal expenditure by the representative household on the final good is given by M , which yields the budget constraint

$$M = PY.$$

P denotes the (quality-adjusted) price index that we will specify further below.

We will next analyze how, for a given path of M or W , innovations of different sorts affect the true (quality-adjusted) price index, P . This will allow us to decompose changes in nominal output, M , into inflation and growth in real output, Y . We will then model the statistical office's imputation procedure to measure price inflation in the economy and show how this imputation leads to a bias in estimated real output growth — i.e., to missing growth. Finally, we will highlight how this simple accounting framework can be used to quantify missing growth from data on market shares of surviving incumbent vs. newly entering producers.

2.1.4 Equilibrium prices

Suppose the existence of a competitive fringe which limits the markup, the producer of intermediate input j can charge at $\tilde{\mu} > 1$ over marginal cost $c(j) = W$. Profit maximization by the intermediate monopolist of each input j then implies that it is optimal to charge a markup factor of $\mu > 1$, where $\mu = \min \left\{ \tilde{\mu}, \frac{\sigma}{\sigma-1} \right\}$.¹² Hence, we obtain in equilibrium for the price of each intermediate good j

$$p(j) = \mu W, \forall j. \tag{2}$$

¹²In the absence of a binding competitive fringe on the intermediate producer's market, one can show that that it is optimal to set $\mu = \frac{\sigma}{\sigma-1}$.

2.1.5 Innovation

We model technical change as product innovation.¹³ At each point in time, and for each intermediate input j there is a common exogenous probability of creative destruction $\lambda_d \in [0, 1]$. I.e., with probability λ_d the incumbent firm of input j is replaced by a new producer. We assume that the new producer (who may be an entrant or an incumbent firm) improves upon the previous producer's quality by a factor $\gamma_d > 1$. The previous producer cannot profitably produce due to limit pricing by the new producer.¹⁴ If j is an existing variety where quality is improved upon by a new producer, we have

$$q_{t+1}(j) = \gamma_d q_t(j).$$

We refer to this innovation process as *creative destruction*.

In addition, for products j where the incumbent producer is not eclipsed by creative destruction, there is each period an exogenous arrival rate $\lambda_i \in [0, 1)$ of an innovation that improves their by factor $\gamma_i > 1$. Hence, if j is a variety where quality is improved upon by the incumbent producer, we have

$$q_{t+1}(j) = \gamma_i q_t(j).$$

We call this *incumbent own innovation*. The main difference from creative destruction is that the producer of j changes with creative destruction, whereas it stays the same with incumbent own innovation.

The arrival rates and step sizes of creative destruction and incumbent own

¹³This modeling choice matters in our empirical context since pure process innovation is arguably more likely to be captured by the statistical office. Yet, across firms and plants with price information in the Census of Manufacturing, we find that firm/plant revenues increase without a decline in unit prices. This suggests that innovations are rather of the product than of the process type. Hottman, Redding and Weinstein (2016) provide similar evidence for retail prices of consumer nondurable manufacturers.

¹⁴We assume Bertrand competition within each product market, which allows the new producer, who produces a better product at the same cost as the current producer, to drive the current producer out of the market.

innovation are constant over time and across varieties.

Finally, each period $t + 1$, a flow of $\lambda_n N_t$ new product varieties $\iota \in (N_t, N_{t+1}]$ are created and available to final goods producers from $t + 1$ onward. Consequently, the law of motion for the number of intermediate inputs is

$$N_{t+1} = (1 + \lambda_n)N_t.$$

The (relative) quality of new product varieties satisfies the following assumption:

Assumption 1 *A firm that introduces in period $t + 1$ a new variety ι starts with a quality that equals $\gamma_n > 0$ times the “average” quality of pre-existing varieties $j \in [0, N_t]$ in period t , or formally*

$$q_{t+1}(\iota) = \gamma_n \left(\frac{1}{N_t} \int_0^{N_t} q_t(j)^{\sigma-1} dj \right)^{\frac{1}{\sigma-1}}, \quad \forall \iota \in (N_t, N_{t+1}].$$

The average quality in Assumption 1 refers to the geometric average that depends on the elasticity of substitution. We do not put further restrictions on the value of γ_n so that new products may enter the market with above-average ($\gamma_n > 1$), average ($\gamma_n = 1$) or below-average quality ($\gamma_n < 1$).

To summarize, there are three sources of growth in this framework: First, the quality of some products increases due to *creative destruction*. Second, for some other products quality increases as a result of *incumbent own innovation*. Third, *new product varieties* are invented which affects aggregate output, because the production function (1) features love-for-variety.

2.2 Missing growth

2.2.1 Missing growth as mismeasured inflation

Because aggregate nominal output M is equal to the product of the price index, P , and real output, Y , gross real output growth between t and $t + 1$ can be

expressed as

$$\frac{Y_{t+1}}{Y_t} = \frac{M_{t+1}}{M_t} \cdot \frac{P_t}{P_{t+1}},$$

where $\frac{P_t}{P_{t+1}}$ is the inverse of the gross inflation rate. We assume that nominal output growth, $\frac{M_{t+1}}{M_t}$, is perfectly measured, in which case the mismeasurement in real output growth is entirely due to mismeasured inflation.

More formally, if $\widehat{\frac{P_t}{P_{t+1}}}$ denotes *measured* inverse inflation, then *measured* real output growth is equal to

$$\frac{\widehat{Y}_{t+1}}{Y_t} = \frac{M_{t+1}}{M_t} \cdot \frac{\widehat{P_t}}{P_{t+1}}.$$

Expressed in log first differences the rate of “missing” output growth is equal to

$$MG_{t+1} = \log\left(\frac{Y_{t+1}}{Y_t}\right) - \log\left(\frac{\widehat{Y}_{t+1}}{Y_t}\right) = \log\left(\frac{\widehat{P_{t+1}}}{P_t}\right) - \log\left(\frac{P_{t+1}}{P_t}\right). \quad (3)$$

Thus there will be missing growth if inflation is overstated.

2.2.2 True prices and inflation

The aggregate price index In the following we derive the exact welfare-based aggregate price index. The results follow from the final goods sector maximizing current final output, Y , with respect to $\{y(j)\}_{j=0}^N$ subject to $M = \int_0^N y(j)p(j)dj$.

Proposition 1 *In equilibrium: (i) the demand for an intermediate product $y(j)$ of quality $q(j)$ sold at price $p(j)$ is given by*

$$y(j) = q(j)^{\sigma-1} \left[\frac{P}{p(j)} \right]^{\sigma} \frac{M}{P}, \quad \forall j. \quad (4)$$

(ii) the equilibrium aggregate price index is given by

$$P = \left(\int_0^N \left[\frac{p(j)}{q(j)} \right]^{1-\sigma} dj \right)^{\frac{1}{1-\sigma}}. \quad (5)$$

Under optimal price setting of the firms we obtain

$$P = \mu W \left(\int_0^N q(j)^{\sigma-1} dj \right)^{\frac{1}{1-\sigma}}. \quad (6)$$

For the proof see [Online Appendix B](#).

The true inflation rate Using (6), we can compute the true inflation rate as a function of the arrival rates and the step sizes of the various types of innovations. We obtain the following proposition:

Proposition 2 *The true gross inflation rate in the economy is given by*

$$\frac{P_{t+1}}{P_t} = \frac{W_{t+1}}{W_t} [1 + \lambda_d (\gamma_d^{\sigma-1} - 1) + (1 - \lambda_d) \lambda_i (\gamma_i^{\sigma-1} - 1) + \lambda_n \gamma_n^{\sigma-1}]^{\frac{1}{1-\sigma}}. \quad (7)$$

For the proof see [Online Appendix B](#).

Proposition 2 shows how the arrival rates and step sizes of the different type of innovation affect the inflation rate (for a given change in wages). The term $\lambda_n \gamma_n^{\sigma-1}$ captures the effect of *variety expansion* on inflation, and the inflation rate is indeed falling in λ_n and γ_n . The term $(1 - \lambda_d) \lambda_i (\gamma_i^{\sigma-1} - 1)$ summarizes the effect of *incumbent own innovation* on price growth. The inflation rate is monotonically decreasing in λ_i and γ_i . The term $\lambda_d (\gamma_d^{\sigma-1} - 1)$ captures the effect from *creative destruction* on the inflation rate. The inflation rate is monotonically decreasing in γ_d . The only comparative static effect that is not immediately clear is the one with respect to λ_d . This ambiguity is due to the interaction between the arrival rates of creative destruction and incumbent own innovation. It is easy to show that, whenever creative destruction has a net positive effect on growth, the inflation rate is indeed decreasing in λ_d since then we must have

$$\gamma_d^{\sigma-1} - 1 > \lambda_i (\gamma_i^{\sigma-1} - 1). \quad (8)$$

This is a condition we will get back to further below.

2.2.3 Imputation and measured inflation

Throughout we assume the statistical office perfectly observes nominal output and wage growth. Hence, as highlighted in (3), missing growth arises if quality-adjusted price changes are overstated. We further assume that the statistical office has no problem measuring unit prices. Thus the difficulty in arriving at quality-adjusted price changes is in measuring quality changes (and in incorporating the benefits of new varieties).

There are well-known challenges to assessing quality changes when firms upgrade their own products, say from one model year to another Boskin et al. (1996). Quality improvements implemented by new producers (i.e., through creative destruction), pose an additional measurement challenge. When the item produced by a given seller has disappeared altogether, the standard procedure used by statistical offices is *imputation*.¹⁵ Imputation uses the rate of inflation for a set of continuing products that were not subject to creative destruction. This procedure is valid if the rate of inflation associated with creative destruction is the same as that for products with a surviving incumbent producer. But the vast majority of these surviving products are not being improved in a given period, according to the BLS.¹⁶ Instances of creative destruction are linked to innovative success, thus creative destruction is typically associated with more rapid quality improvements than surviving products taken as a whole.

We formally characterize imputation as follows:¹⁷

¹⁵Using statistics from Klenow and Kryvtsov (2008), [Online Appendix A](#) calculates that imputation was used 90% of the time in the years 1988–2004 when a producer ceased selling a product in the CPI.

¹⁶See Bils and Klenow (2004) and Nakamura and Steinsson (2008). In [Online Appendix A](#) we argue — based on CPI micro data — that at a monthly frequency λ_i is less than 1%.

¹⁷Erickson and Pakes (2011) likewise identify imputation as a source of upward bias in estimating the average price change for exiting products. In their application to consumer electronics, they argue that even the BLS's hedonic procedures — which estimate the missing price of exiting goods using regressions of prices on observable product characteristics — fail to eliminate the bias because they do not correct for time-varying unmeasured characteristics of the goods.

Assumption 2 *In the presence of product entry and exit the statistical office resorts to imputation, i.e., the set of products with a surviving incumbent producer is assumed to be representative and the economy-wide inflation rate is imputed from this subset of products.*

Products of continuing producers can be either subject to incumbent own innovation or no innovation at all. We denote the statistical office's estimates for the frequency and step size of quality-improving innovations on surviving products as $\widehat{\lambda}_i$ and $\widehat{\gamma}_i$.

Proposition 3 *Under Assumption 2, the measured inflation rate is given by*

$$\left(\frac{\widehat{P}_{t+1}}{P_t}\right) = \frac{W_{t+1}}{W_t} \left[1 + \widehat{\lambda}_i (\widehat{\gamma}_i^{\sigma-1} - 1)\right]^{\frac{1}{1-\sigma}}. \quad (9)$$

For a proof see [Online Appendix B](#).

Henceforth we assume that the statistical office perfectly observes the frequency and step size of incumbent own innovations, i.e., we have $\widehat{\lambda}_i = \lambda_i$ and $\widehat{\gamma}_i = \gamma_i$. We make this assumption to isolate missing growth due to imputation.¹⁸

2.2.4 Formula for missing growth

Recall that missing real output growth between period t and $t + 1$ is caused by upward bias in estimating inflation as in (3). Missing growth can be expressed in the following proposition:

Proposition 4 *Missing real output growth is given by*

$$MG = \frac{1}{\sigma - 1} \log \left[1 + \frac{\lambda_d [\gamma_d^{\sigma-1} - 1 - \lambda_i (\gamma_i^{\sigma-1} - 1)] + \lambda_n \gamma_n^{\sigma-1}}{1 + \lambda_i (\gamma_i^{\sigma-1} - 1)} \right]. \quad (10)$$

¹⁸In equation (31) below, we show how missing growth would change if the quality improvement of incumbents is not perfectly measured.

Proof. Equation (10) is obtained from combining (3), (7), and (9). ■

Because the statistical office imputes the inflation rate from products with surviving producers (Assumption 2), there are two sources of missing growth: (i) variety expansion; and (ii) creative destruction, which is imputed but not directly measured. The last term in the numerator in the square brackets in equation (4) captures new variety creation. As long as there is net entry, i.e., $\lambda_n > 0$, there is positive missing growth from this source. The first term in the numerator captures missing growth from not properly factoring in creative destruction.

The statistical office understates true output growth if the imputed growth from creative destruction (which is imputed from products with surviving producer) understates the true expected growth from creative destruction. This can happen in two ways: (i) creative destruction has a larger step size than incumbent own innovation ($\gamma_d > \gamma_i$) and/or (ii) not all of the surviving incumbents innovate ($\lambda_i < 1$). Note that missing growth from creative destruction remains positive even in the special case where creative destruction and incumbent innovation have the same step size, if not all of the surviving products are subject to incumbent own innovation (i.e., $\lambda_i < 1$). Missing growth is increasing in γ_d , and also in λ_d as long as $(\gamma_d^{\sigma-1} - 1) > \lambda_i(\gamma_i^{\sigma-1} - 1)$. This is the same condition which ensures that overall true growth is increasing in λ_d (see (8)). Missing growth is large if creative destruction is an important source of true growth.¹⁹

3 Estimating missing growth

In this section we explore two alternative approaches for quantifying missing growth in the U.S. data. Both approaches build on the model developed in the previous section. The first approach uses information on the market shares of entrants, exiters and survivors. We refer to it as the *market share approach*. The

¹⁹See [Online Appendix C](#) for an illustration with the Cobb-Douglas production function.

second approach uses the algorithm in Garcia-Macia, Hsieh and Klenow (2016) to infer arrival rates and step sizes of different type of innovations and compute missing growth. We refer to this as the *indirect inference method*.

3.1 The market share approach

Here we show how to use our model in the previous Section 2 to estimate missing growth using data on the market share of entering establishments (plants), surviving plants, and exiting plants. This approach does not allow us to differentiate between the different sources of missing growth (creative destruction vs. variety expansion), but it provides a simple and intuitive quantification.

3.1.1 Relating missing growth to market share dynamics

The idea behind the market share approach is that the nominal expenditure shares of different products contain information about their quality-adjusted prices. The imputation used by statistical offices assumes that the quality-adjusted price growth of products from surviving producers is representative for the economy-wide inflation rate. The CES framework of Section 2 suggests a simple test for the representativeness of this set of products: it is representative if and only if the market share of this set remains stable over time. Moreover, given an estimate for the elasticity of substitution between products, we will show how market shares can be used to quantify missing growth.

To construct the main argument, define the market share of a product j as: $s(j) \equiv \frac{p(j)y(j)}{M}$, where $M = PY$ denotes aggregate nominal output. Combining this definition with the demand curve in (4) gives

$$s(j) = \left(\frac{P}{p(j)/q(j)} \right)^{\sigma-1}.$$

Hence, the market share of product j is given by a power function of the quality-adjusted price of j relative to the aggregate price index, P . With a CES structure, a similar relationship also holds for the market share of a *subset* of products. We are particularly interested in the subset of products the statistical office bases its imputation on — those products with a surviving producer between the periods t and $t + 1$ — which we denote $\mathcal{I}_t = [0, N_t] \setminus \mathcal{D}_t$, where \mathcal{D}_t represents the set of products which underwent creative destruction between t and $t + 1$. In the following we refer to the set of products \mathcal{I}_t as the *continuers* between t and $t + 1$. In period t the aggregate market share of these continuers is given by

$$S_{I_t,t} = \int_{\mathcal{I}_t} \left(\frac{P_t}{p_t(j)/q_t(j)} \right)^{\sigma-1} dj.$$

A period later, the aggregate market share of the same continuers is

$$S_{I_t,t+1} = \int_{\mathcal{I}_t} \left(\frac{P_{t+1}}{p_{t+1}(j')/q_{t+1}(j')} \right)^{\sigma-1} dj'.$$

These two expressions together imply that the growth rate of the continuers' market share is given by

$$\frac{S_{I_t,t+1}}{S_{I_t,t}} = \left(\frac{P_{t+1}/P_t}{P_{I_t,t+1}/P_{I_t,t}} \right)^{\sigma-1}.$$

This highlights that, if the inflation rate of the continuers is representative of the aggregate inflation rate, then the market share of the continuers should remain stable over time. If, instead, the market share of continuers is systematically shrinking over time, we conclude that the inflation rate of the continuers must have been higher than the one of the aggregate economy (because $\sigma > 1$).

Using the price setting of firms (2) and the specified innovation processes, we can then express the growth rate of the market share of continuers as follows:

Proposition 5 *The market share growth of continuers from t to $t + 1$ is given by*

$$\frac{S_{I,t,t+1}}{S_{I,t,t}} = \left(\frac{P_{t+1}}{P_t} \right)^{\sigma-1} \left(\frac{W_{t+1}}{W_t} \right)^{1-\sigma} (1 - \lambda_i + \lambda_i \gamma_i^{\sigma-1}).$$

For a proof see [Online Appendix B](#).

As shown in Proposition 5, the price index of continuers grows at the rate of the nominal wage $\frac{W_t}{W_{t+1}}$ times $(1 - \lambda_i + \lambda_i \gamma_i^{\sigma-1})^{1/(1-\sigma)}$, which captures the impact of incumbent own innovation on prices. Under the assumption that the statistical office perfectly observes quality improvements due to incumbent own innovation — $\widehat{\gamma}_i = \gamma_i$ and $\widehat{\lambda}_i = \lambda_i$ — the market share growth of continuers can then be related to the measured inflation rate (see Proposition 3):

$$\frac{S_{I,t,t+1}}{S_{I,t,t}} = \left(\frac{P_{t+1}}{P_t} \right)^{\sigma-1} \left(\frac{\widehat{P}_{t+1}}{P_t} \right)^{-(\sigma-1)}. \quad (11)$$

The market share growth of continuers is a power function of the growth of their price index relative to the aggregate price index. These continuers are precisely the set of products that imputation is based on. Under the assumption that the statistical office perfectly measures the quality-adjusted price inflation of continuers, this gives us a simple test for the representativeness of these products. As in (11), the quality-adjusted inflation of continuers is representative if and only if their market share is stable over time. If instead the market share of continuers falls over time, then measured inflation $\frac{\widehat{P}_{t+1}}{P_t}$ is too high and there is missing growth (given $\sigma > 1$). For a given value of the elasticity of substitution, σ , (11) can be exploited to quantify missing growth. This is underscored in the next proposition.

Proposition 6 *If incumbent own innovation is captured perfectly, then missing growth from imputation can be calculated from the market share of continuers as*

$$MG_{t+1} = \log \left(\frac{\widehat{P}_{t+1}}{P_t} \right) - \log \left(\frac{P_{t+1}}{P_t} \right) = \frac{1}{\sigma - 1} \log \left(\frac{S_{I,t,t}}{S_{I,t,t+1}} \right). \quad (12)$$

Proof. Combining (3) and (11) directly proves the proposition. ■

Since $\sigma > 1$, missing growth is positive whenever the market share of continuers shrinks over time.²⁰ In the following we use equation (12) to quantify missing growth. Our next step is to explain how we try to measure continuer market shares in the data.

3.1.2 Measuring the continuers' market share

Our goal is to quantify missing growth in the aggregate economy over a time horizon of several decades. We therefore base our market share approach estimates of missing growth on the Longitudinal Business Database (LBD), which covers all nonfarm business sector plants with at least one employee. We use the employment/payroll information in this dataset to infer $S_{I,t}$.

Ideally we would have data at the *product* level for each firm. Unfortunately, such data does not exist for the aggregate U.S. economy outside of manufacturing (Census of Manufacturing) or consumer nondurable manufacturing (AC Nielsen scanner data). In lieu of such ideal data, we suppose that firms must add *plants* in order to produce new products. Such new plants could be at entering firms or at existing firms. And the products produced by new plants could be brand new varieties or the result of creative destruction. Moreover, we assume that all incumbent own innovation occurs at existing plants. Under these assumptions, we can use continuing *plants* as a proxy for continuers.

These assumptions are admittedly strong (we will relax them in our second approach, though indirect inference has limitations of its own). Our assumptions require that firms do not add products through existing plants. Bernard, Redding and Schott (2011) find that U.S. manufacturing plants do

²⁰Our market share approach is clearly related to Feenstra (1994) and Broda and Weinstein (2010). They also make use of the basic idea that, in a CES world, quality-adjusted prices are isoelastically related to nominal expenditure shares. But they assume that *all* quality growth is missed. We allow for the fact that statistical offices do attempt to measure quality changes, in particular for continuers. Our focus is therefore on the subtler question of whether imputation in the case of creative destruction leads to missing growth.

start up production in new industries. But perhaps our assumption is a better one outside manufacturing, such as in retail where location is a key form of product differentiation.

If existing plants do introduce new varieties or carry out creative destruction, then our market share approach is likely to understate missing growth. As we will explain below, our baseline specification will only assess market shares of new plants after a 5 year lag. Hence, the critical assumption is that plants do not add new products after the age of 5 years. We can offer two facts that provide some reassurance here. First, employment growth rates are much lower after 5 years than before 5 years of age (Haltiwanger, Jarmin and Miranda, 2013). Second, plant exit rates do fall with age, but not very sharply after age 5 (Garcia-Macia, Hsieh and Klenow, 2016). If plants added varieties as they aged, then one would expect exit rates to fall for surviving plants.

Note also that our approach uses employment (or payroll) data to measure market shares, rather than revenue data. In our model these variables are all proportional to each other across products. But in practice they could differ. Plant revenue data are only available at the firm level, and to us only every 5 years in the manufacturing sector. As a robustness check below we report on revenue results for manufacturing.

To be more exact, we implement the market share approach on plant data as follow. Let B denote the first year of operation and D denote the last year of operation of a plant. Then the continuing plants I_t are those plants who were born in t or before and who die in year $t + 1$ or later. That is, $B \leq t$ and $D > t + 1$. Define E_t as the group of plants entering in t ($B = t, D \geq t$) and X_t as the group of plants exiting between t and $t + 1$ ($B \leq t, D = t$). Let $L(t, \mathcal{M})$ denote the total employment or payroll in period t of plants belonging to group \mathcal{M} . We then measure the ratio $\frac{S_{I_t,t}}{S_{I_t,t+1}}$ on the right-hand side of (12) as

$$\frac{S_{I_t,t}}{S_{I_t,t+1}} = \frac{\frac{L(t, I_t)}{L(t, I_t) + L(t, X_t)}}{\frac{L(t+1, I_t)}{L(t+1, I_t) + L(t+1, E_{t+1})}}, \quad (13)$$

According to (13), missing growth is positive whenever the employment/payroll share of continuing plants shrinks between t and $t+1$. For our baseline results we will rely on employment data. In Section 3.1.4 we present results with payroll data as a robustness check.

Precisely how do we measure the first period of operation, B , and the last period of operation, D , in the data? We define a period t as a calendar year. We map D to the last year a plant is in LBD. We could map B to the first year the plant appears in the LBD. But it may take time for plants to accumulate customers and market share — even conditional on the price, quality, and variety of their products. And, as argued above, plants may be most likely to add products in their first 5 years, judging from their growth rates.

With these considerations in mind, we set B equal to $k \geq 0$ years after the plant first appears in the LBD. If B^d denotes the first year the plant appears in the database, we map B into $B^d + k$. We use $k = 5$ in our baseline specification. As mentioned, Haltiwanger, Jarmin and Miranda (2013) find that “the fastest-growing continuing firms are young firms under the age of 5” (see their figure 4B), and the same is true at the plant level. We check robustness to alternative lags in Section 3.1.4.

To reiterate, we use the market share after a lag to obtain a measure of market share that is arguably more tightly related to the quality-adjusted price of new plants. A new plant may take a number of years to adjust its inputs and accumulate customers. By abstracting from plants that enter and exit within the 5 year window, our approach may also be more robust to short run churning in the labor market that might be subject to cyclical taste shifts.

Finally, to quantify missing growth we need to parametrize the elasticity of substitution. As our baseline value we choose $\sigma = 4$ based on Redding and

Weinstein (2016) and Hottman, Redding and Weinstein (2016). We check robustness to higher and lower elasticities in Section 3.1.4.²¹

3.1.3 The market share of missing growth: results

Our baseline results calculate missing growth from 1983–2013 using LBD data. The LBD contains data on employment and payroll going back to 1976, but the payroll data features a number of implausible outliers before 1989. We therefore use employment data for our baseline estimates. Using the employment data, we identify entrants beginning in 1977. For our benchmark with 5 years lag, 1983 is the earliest year we can calculate missing growth (market share growth of survivors between 1982 and 1983) because we use plants that have been in the data for at least 5 years.

Table 1 presents our missing growth estimates with $\sigma = 4$ and $k = 5$, and compares them to official TFP growth. The entries are annual percentage points. We find on average 0.64 percentage points of missing growth per year from 1983–2013. BLS measured TFP growth over the same interval was 1.87 percentage points per year.²² If we add our missing growth to the BLS TFP series we arrive at “true” growth of 2.51% per year. Thus our baseline estimate is that about one-fourth of true growth is missed.

Table 1 also breaks the 30 year sample into three sub-periods: 1983–1995 (an initial period of average official growth), 1996–2005 (a middle period of rapid official growth), and 2006–2013 (a final period of low official TFP growth). Did missing growth contribute to the speedup or slowdown? Our estimates say yes, but modestly at most. Missing growth slowed down 11 basis points when official

²¹Our baseline is the median value of the elasticity of substitution within a product category across *different producers* from Hottman, Redding and Weinstein (2016). Our robustness checks cover the interquartile range of the elasticity of substitution estimates in Hottman, Redding and Weinstein (2016).

²²We put BLS TFP growth in labor-augmenting form, include the BLS estimates of the contribution of R&D and intellectual property to TFP growth. The BLS multifactor productivity series uses real output growth from the BEA. The vast majority of the price indices that go into constructing BEA real output growth come from the BLS (see U.S. Bureau of Economic Analysis (2014)), even if the BEA weights sectors differently than in the aggregate CPI or PPI.

Table 1: MEASURED VS. TRUE GROWTH, MARKET SHARE APPROACH

	Missing Growth	Measured Growth	“True” Growth	% of growth missed
1983–2013	0.64	1.87	2.51	25%
1983–1995	0.66	1.80	2.46	27%
1996–2005	0.55	2.68	3.23	17%
2006–2013	0.74	0.98	1.72	43%

Notes: Entries are percentage points per year. Missing growth is calculated using the market share approach in Proposition 6. The market share is measured as the employment share of plants taken from the Census’ Longitudinal Business Database (LBD) as described in (13). These baseline results assume a lag $k = 5$ and an elasticity of substitution $\sigma = 4$. Measured growth is calculated as the BLS MFP series + R&D contribution expressed in labor-augmenting terms. True growth is the sum of measured growth and missing growth.

growth accelerated by 88 basis points in the middle period. And missing growth sped up by 19 basis points when official growth dropped 170 basis points in the final period.

3.1.4 The market share of missing growth: robustness and discussion

In this section we discuss the robustness of our market share approach to quantifying missing growth. We show how our main results are affected by alternative assumptions about the lag used to defined new plants, the elasticity of substitution, and the data used to measure market shares (using payroll instead of employment data).

Using payroll to measure market shares Table 2 compares missing growth based on employment vs. payroll. The payroll data allows us to do this only from 1989 onward. The estimates are quite similar overall, but exhibit a bigger dip and bounce back across the three sub-periods.

Table 2: MISSING GROWTH WITH EMPLOYMENT VS. PAYROLL

	Employment	Payroll
1989–2013	0.70	0.72
1989–1995	0.86	0.97
1996–2005	0.55	0.43
2006–2013	0.74	0.83

Notes: Entries are percentage points per year.

Recall that, in manufacturing only, we also have access to revenue data. We find that the market share of continuing manufacturing shrinks more when based on revenue than employment or payroll, so that missing growth as defined by (12) is higher when market share is measured in terms of revenue.

Different lag k Our baseline results in Table 1 evaluate the market share of entrants after a lag of $k = 5$ years. Table 3 indicates how the results change if we instead look at market shares with no lag after entry. With $k = 0$ missing growth is much smaller, averaging 20 basis points per year rather than 64 basis points. Though not reported in Table 3, for $k = 3$ we obtain estimates of missing growth closer to $k = 5$. Increasing the lag beyond the baseline to $k = 7$ years increases missing growth only slightly compared to $k = 5$.

Elasticities of substitution As shown in (12), missing growth under the market share approach is proportional to $1/(\sigma - 1)$, where σ is the elasticity of substitution across varieties. Thus our missing growth estimates decline monotonically as we vary σ from 3 to 4 to 5 in Table 4. The lower is σ , the bigger must be the quality and variety improvements brought by new plants to explain a given decline in the observed market share of continuing plants.

In Online Appendix D we extend our framework to allow the elasticity of substitution to vary across sectors. Our market share approach applies easily

Table 3: MISSING GROWTH WITH DIFFERENT LAGS k

	$k = 5$	$k = 0$
1983–2013	0.64	0.20
1983–1995	0.66	0.28
1996–2005	0.55	0.20
2006–2013	0.74	0.07

Notes: Entries are percentage points per year.

to this more complicated nested structure. And below we show results under a two-tier specification with a lower elasticity of substitution within industries ($\sigma > 1$) than across industries.

Table 4: MISSING GROWTH WITH DIFFERENT ELASTICITIES σ

	Missing Growth		
	$\sigma = 3$	$\sigma = 4$	$\sigma = 5$
1983–2013	0.96	0.64	0.48
1983–1995	0.98	0.66	0.49
1996–2005	0.82	0.55	0.41
2006–2013	1.11	0.74	0.56

Notes: Entries are percentage points per year.

Sectoral missing growth The market share approach is well-suited to analyze missing growth at a more disaggregated level. To illustrate this point, we calculate missing growth using equation (12) and data on employment shares of incumbent, entering, and exiting plants within 2-digit NAICS sectors (2002 classification).

Table 5 shows missing growth in the five sectors with the highest shares of overall employment (with a combined share of 60% of all employment). Interestingly, we find very little missing growth in manufacturing or education. For health care, retail trade, and restaurants & hotels, in contrast, missing growth averages 80 basis points per year or more. The retail trade contribution may reflect the geographic spread of outlets by big box chains. Health care might capture how new facilities have more up-to-date equipment and treatment than dated facilities. Perhaps reassuringly, our aggregate missing growth is not driven by one particular sector.

Table 5: MISSING GROWTH BY SECTOR, 1983–2013 AVERAGES

Sector	Missing growth	Employment share
Manufacturing	0.04	16%
Health Care	0.80	13%
Retail Trade	0.91	12%
Education	0.06	11%
Restaurants & Hotels	1.64	7%

Notes: Entries are percentage points per year. The sectors shown are the top five 2-digit NAICS (2002 classification) in terms of employment.

The variation in missing growth across sectors in Table 5 suggests that official relative price trends across sectors may be biased considerably. Taken at face value, our estimates would alter the pattern of Baumol's cost disease and affect how one would estimate the elasticity of substitution across sectors.

We also calculated missing growth at a more detailed level and then aggregated the sectoral missing growth rates using Tornqvist employment shares of each sector as weights. This assumes a nested CES structure where products within a sector have an elasticity of substitution of 4 and sectors are aggregated in translog fashion. Table 6 compares our benchmark missing

growth estimates to those obtained by aggregating up missing growth in this way from the 2 to 5 digit sectoral level. The estimates are gently increasing in the level of disaggregation used (from 64 basis points to 72), and our broadly similar across sub-periods.²³

Table 6: MISSING GROWTH WITH DISAGGREGATED SECTORS

	1-sector	2-digit	3-digit	4-digit	5-digits
1983–2013	0.64	0.64	0.65	0.73	0.72
1983–1995	0.66	0.61	0.62	0.68	0.67
1996–2005	0.55	0.55	0.57	0.64	0.63
2006–2013	0.74	0.78	0.80	0.91	0.92

Notes: Entries are percentage points per year. The first column repeats our baseline results with a single sector. Other columns are employment-weighted averages of missing growth within different NAICS (2002) n-digit levels.

Declining dynamism and missing growth One may wonder why our missing growth estimates do not trend downward along with rates of entry, exit and job reallocation across firms — the “declining dynamism” documented by Decker et al. (2014). The answer is three-fold. First, we look at plants (establishments) not firms. Second, our market share equation for missing growth is tied to the *net* entry rate (weighted by employment), not the gross job creation rate due to entrants. Put differently, the growth of survivors’ market share is influenced by the *difference* between the job creation rate due to new plants and the job destruction rate due to exiting plants. Unlike gross flows at the firm level, we see no trend in the net job creation rate of plants over 1983–2013 in the LBD. Finally, we look at market shares five years after the plant appears in the LBD, rather than immediately upon entry.

²³We also aggregated sectoral missing growth rates using average employment shares over the entire 1983–2013 period to fully eliminate any trends due to changes in sectoral composition. The results were very similar to those in Table 6.

Table 7 illustrates these points of distinction between our market share approach and declining dynamism. Across the first two columns, missing growth drops dramatically when moving from plants to firms. Many new plants are at existing firms, and they are bigger on average than new plants in new firms. Also, unlike our plant-level estimates, missing growth declines over time when based on firm-level data.

Table 7: MISSING GROWTH VS. DECLINING DYNAMISM

	Plant level	Firm level	Net entry	Gross entry
1983–2013	0.64	0.22	0.37	0.19
1983–1995	0.66	0.33	0.54	0.70
1996–2005	0.55	0.17	0.40	0.06
2006–2013	0.74	0.09	0.06	-0.49

Notes: Entries are percentage points per year.

To demonstrate the importance of entrant/exiter vs. continuers growth, the third column of Table 7 shows how big missing growth would be using firm level data and assuming all firms were the same size. In this case, missing growth is larger than the firm-level estimate because entering firms tend to be smaller than the average firm. But missing growth continues to plummet over the sample.

Finally, the gross entry rate has declined more than the net entry rate, and our missing growth estimates focus on the net entry rate. To illustrate this point, the last column of Table 7 shows missing growth if all firms were equal-sized and the exit rate was fixed. Missing growth declines even more precipitously when measured in this counterfactual way. In the data, the exit rate fell along with the entry rate, dampening the decline in missing growth.

We base our missing growth estimates on plant dynamic, rather than firm dynamics, because we think it is much defensible to assume plants do not add

new products than to assume that firms do not do so.

In the [Online Appendix](#) we discuss further extensions of our analysis in this section: capital in production, varying markups, gains from variety, bias in measuring incumbent own innovation, importing and outsourcing.

3.2 The indirect inference method

In this section we provide an alternative, indirect inference approach to quantify missing growth. It allows us to, in principle, separate the two sources of missing growth — creative destruction and variety expansion. The method relies on first estimating the arrival rates and step sizes and then plugging these values into formula (10) in Proposition 4 to arrive at missing growth. To estimate the arrival rates and step sizes, we adapt the algorithm developed in Garcia-Macia, Hsieh and Klenow (2016), henceforth GHK.

The original GHK algorithm GHK's algorithm uses indirect inference to estimate the step size and arrival rate of three types of innovation: Own innovation (OI), Creative Destruction (CD), and New Varieties (NV). GHK estimate these parameters to fit aggregate TFP growth; the mean, minimum (one worker), and standard deviation of employment across firms; the share of employment in young firms (firms less than 5 years old); the overall job creation and destruction rates; the share of job creation from firms that grew by less 1 log point (three-fold) over a five year period; employment share by age; exit rate by size; and the growth rate in the number of firms (which is equal to the growth rate of employment in the model).²⁴ They calculate these moments in the LBD for 1983–1993, 1993–2003 and 2003–2013, respectively.²⁵

With their parameter estimates in hand, GHK decompose growth into

²⁴GHK assume each variety carries an overhead cost. Firms choose to retire a variety if the expected profits from that variety do not cover the overhead cost. GHK calibrate the overhead cost to match minimum employment in the data.

²⁵Their algorithm matches model steady state moments to data moments, and therefore does not produce annual estimates.

contributions from new varieties, incumbent innovation on their own products, creative destruction by incumbents, and creative destruction by entering firms.

How we differ from GHK Our model in Section 2 differs from the GHK model in several respects. GHK keep track of firms with multiple products, and estimate rates of creative destruction and new variety creation separately for entrants and incumbents. GHK also endogenize product exit: firms drop products whose quality relative to the average quality is below a certain cutoff. And rather than assuming a fixed step size for innovations, in GHK quality innovations are drawn from a Pareto distribution, where the same Pareto shape parameter (and hence the same average step size) is assumed for quality innovations from incumbents' own innovation and from innovations involving creative destruction. Finally, in GHK new varieties are drawn from a scaled version of the existing quality distribution rather than massing at a single point relative to the existing distribution.

As a result of these differences, GHK obtain an expression for true productivity growth which is somewhat different from that in our Section 2. Using our notation, true productivity growth $g \equiv \frac{Y_{t+1}}{Y_t} - 1$ in GHK is

$$1 + g = \left[(1 - \delta_o \psi) \left\{ [\lambda_i (1 - \lambda_{e,d} - \lambda_{i,d}) + (\lambda_{e,d} + \lambda_{i,d})] (\gamma_i^{\sigma-1} - 1) + 1 \right\} + (\lambda_{i,n} + \lambda_{e,n}) \gamma_n^{\sigma-1} \right]^{\frac{1}{\sigma-1}} \quad (14)$$

where δ_o denotes the share of products in the previous period whose quality falls below the obsolescence cutoff; ψ is the average quality of those below-cutoff products relative to the average quality;²⁶ λ_i is the share of products that are not obsolete, did not experience creative destruction, and did experience an innovation by the incumbent producer; $\lambda_{e,d}$ is the share of

²⁶ $\delta_o = \int_{q(j) < \bar{q}_t, q(j) \in \Omega_t} 1 \, dj$ and $\psi \delta_o = \frac{\int_{q < \bar{q}_t, q \in \Omega_t} q_t^{\sigma-1}(j) \, dj}{\int_{q \in \Omega_t} q_t^{\sigma-1}(j) \, dj}$. Ω_t denotes the set of products in t .

non-obsolete products with entrant creative destruction; $\lambda_{i,d}$ is the share of non-obsolete products with incumbent creative destruction; $\lambda_{i,n} + \lambda_{e,n}$ is the mass of new varieties from incumbents and entrants relative to the mass of products in the previous period; and γ_i and γ_d are the average step sizes of own innovation and creative destruction, respectively. As in GHK, we assume that the two step sizes are the same, which is why only γ_i appears in equation (14). Finally, γ_n is the average quality of a new variety relative to the average quality of varieties produced in the previous period. The term $1 - \delta_0\psi$ adjusts for the endogenous loss of varieties due to obsolescence.

The equation for measured growth \hat{g} in our modified GHK model is the same as in our Section 2 above:

$$1 + \hat{g} = (1 + \lambda_i(\gamma_i^{\sigma-1} - 1))^{\frac{1}{\sigma-1}}. \quad (15)$$

Recall that we assume the BLS accurately measures the arrival rate and the average step size of incumbents' own innovations. To adapt the GHK methodology to our model with missing growth, we make the following changes to the original GHK algorithm:²⁷

1. We choose parameters so that (15) matches the observed growth rates: 1.66% for 1983–1993, 2.29% for 1993–2003 and 1.32% for 2003–2013, according to the BLS.
2. We restrict the sum of the (unconditional) arrival rates of OI and CD to equal the cumulative rate of non-comparable substitutions from the CPI over 5 years.

Key advantages of this indirect inference method include: (i) we need not assume that creative destruction and new product varieties only come from new plants (the inference is on firm-level data and allows for multi-product firms); incumbent plants may also produce CD or NV innovations; (ii) we can

²⁷See our [Online Appendix H](#) for a more details description of the changes we made.

decompose missing growth into its CD and NV components using the arrival rates and step sizes of the various kinds of innovations; and (iii) we allow for the possibility of products disappearing because of obsolescence.

Results from indirect inference Table 8 defines the parameters and displays their estimated values for each the three samples: “1988” for 1983–1993, “1998” for 1993–2003, and “2008” for 2003–2013. The bottom panel of Table 8 reports the resulting estimates of measured, true, and missing growth. Missing growth is larger under this alternative approach than in the previous section using the market share approach for the sample periods 1988 and 1998: 1.25 percentage points per year (vs. 0.64 when using the market share approach) for the “1988” period, and 1.13 percentage points per year (vs. 0.55) for the “1998” period. For the last sample period, “2008”, the missing growth estimates from the indirect inference method are closer to those from the market share approach at (0.60 percent vs. 0.76). The fraction of total productivity growth that is missed is comparable across the three periods at around one-third of true growth.

As mentioned, an advantage of the indirect inference approach over the market share approach is that here we can decompose missing growth into its new varieties (NV) and creative destruction (CD) components. We calculate missing growth from creative destruction by taking the difference between measured productivity growth and the productivity growth that results when we set the total arrival rate for new varieties ($\lambda_{i,n} + \lambda_{e,n}$) equal to zero. We find that vast majority of the missing growth is due to creative destruction — around 80% in all three periods.

3.3 Comparison of the two quantification methods

As stressed, the market share approach uses plant-level data (assuming no added products per plant, and focusing attention on plants that are at least five years old), whereas the indirect inference approach uses firm-level data. The market share approach assumes that creative destruction only occurs through

Table 8: ESTIMATED PARAMETERS AND RESULTS WITH INDIRECT INFERENCE

Parameter	Definition	1988	1998	2008
λ_d	CD arrival rate	0.014	0.011	0.010
λ_i	OI arrival rate (if survive)	0.024	0.027	0.027
λ_n	NV arrival rate	0.004	0.002	0.002
γ_d, γ_i	Step size of CD, OI	1.106	1.125	1.074
γ_n	Step size of NV	0.328	0.482	0.366
	Measured growth per year (ppt)	1.66	2.29	1.32
	Missing growth (ppt)	1.25	1.13	0.60
	True growth per year (ppt)	2.91	3.42	1.92
	% of missing growth from CD	79.4%	79.7%	80.8%
	% of growth missed	43.0%	32.9%	31.2%

Notes: Results from running the algorithm from Garcia-Macia, Hsieh and Klenow (2016) on three samples: 1983–1993 (“1988”), 1993–2003 (“1998”) and 2003–2013 (“2008”). We adapt their algorithm to incorporate missing growth as explained in the text.

new plants. The indirect inference method allows for creative destruction by existing plants as well. This may be why we found larger average missing growth in this second quantification than in the market share approach.

The falling entry rate of new firms over the past three decades (“declining dynamism”) may explain why missing growth declines across the three periods in the indirect inference approach, which again uses firm-level data. The market share approach with plant-level exhibited no such decline. But when, as a robustness check, we applied the market share approach to firm-level data, we did obtain a sharp decline in missing growth across periods (Table 7).

As already mentioned above, indirect inference did not require that only entrant plants create new varieties or generate creative destruction; and this method allowed us to split overall missing growth into its NV and CD

components. On the other hand, the market share approach is simple, requires fewer model assumptions, and was not too data demanding.

4 Conclusion

In this paper we developed a Schumpeterian growth model with incumbent and entrant innovation to assess the unmeasured TFP growth resulting from creative destruction. Crucial to this missing growth was the use of imputation by statistical agencies when producers no longer sell a product line. Our model generated an explicit expression for missing TFP growth as a function of the frequency and size of creative destruction vs. other types of innovation.

Based on the model and U.S. Census data for all nonfarm businesses from 1983–2013, we explored two ways of estimating the magnitude of missing growth from creative destruction. The first approach used the employment shares of surviving, entering, and exiting plants. The second approach adapted the indirect inference algorithm of Garcia-Macia, Hsieh and Klenow (2016) to firm-level data. The former, “market share” approach is simple, intuitive and easy to replicate for different economies, but assumes all creative destruction occurs through new plants and cannot separate out missing growth from expanding variety. The latter “indirect inference” approach is more involved and less intuitive, but does not require creative destruction to occur through new plants. The second approach, moreover, allows us to decompose missing growth into that from creative destruction vs. expanding variety.

We found that: (i) missing growth from imputation was substantial — around one-fourth to one-third of total productivity growth; and (ii) it was mostly due to creative destruction.

We may be understating missing growth because we assumed there were no errors in measuring quality improvements by incumbents on their own products. We think missing growth from imputation is over and above (and

amplified by) the quality bias emphasized by the Boskin Commission.²⁸

Our analysis could be extended in several interesting directions. One would be to look at missing growth in countries other than the U.S. A second extension would be to revisit optimal innovation policy. Based on Atkeson and Burstein (2015), the optimal subsidy to R&D may be bigger if true growth is higher than measured growth. Conversely, our estimates give a more prominent role to creative destruction with its attendant business stealing.

A natural question is how statistical offices should alter their methodology in light of our results, presuming our estimates are sound. The market share approach would be hard to implement without a major expansion of BLS data collection to include market shares for entering, surviving, and exiting products in all sectors. The indirect inference approach is even less conducive to high frequency analysis. A feasible compromise might be for the BLS to impute quality growth for disappearing products based on its direct quality adjustments for those surviving products that have been innovated upon.²⁹

Our missing growth estimates have other implications which deserve to be explored further. First, ideas may be getting harder to find, but not as quickly as official statistics suggest if missing growth is sizable and relatively stable. This would have ramifications for the production of ideas and future growth (Gordon, 2012; Bloom et al., 2017). Second, the U.S. Federal Reserve might wish to raise its inflation target to come closer to achieving quality-adjusted price stability. Third, a higher fraction of children may enjoy a better quality of life than their parents (Chetty et al., 2017). Fourth, as stressed by the Boskin Commission, U.S. tax brackets and Social Security benefits may rise too steeply since they are indexed to measured inflation, the inverse of measured growth.

²⁸Leading economists at the BLS and BEA recently estimated that quality bias from health and ICT alone was about 40 basis points per year from 2000–2015 Groshen et al. (2017).

²⁹Erickson and Pakes (2011) suggest that, for those categories in which data are available to do hedonics, the BLS could improve upon the imputation method by using hedonic estimation that corrects for both the selection bias associated with exit and time-varying unmeasured characteristics.

References

- Abramovitz, Moses.** 1956. "Resource and Output Trends in the U.S. since 1870." *American Economic Review*, 46(2): 5–23.
- Acemoglu, Daron, Ufuk Akcigit, Nicholas Bloom, and William R. Kerr.** 2013. "Innovation, Reallocation and Growth." National Bureau of Economic Research NBER Working Papers 18993.
- Aghion, Philippe, and Peter Howitt.** 1992. "A Model of Growth through Creative Destruction." *Econometrica*, 60(2): 323–51.
- Aghion, Philippe, Ufuk Akcigit, and Peter Howitt.** 2014. "What Do We Learn From Schumpeterian Growth Theory?" In *Handbook of Economic Growth*. Vol. 2 of *Handbook of Economic Growth*, Chapter 0, 515–563. Elsevier.
- Akcigit, Ufuk, and William R. Kerr.** 2010. "Growth Through Heterogeneous Innovations." National Bureau of Economic Research NBER Working Papers 16443.
- Atkeson, Andrew, and Ariel Burstein.** 2015. "Aggregate implications of innovation policy." Society for Economic Dynamics 2015 Meeting Papers 640.
- Bartelsman, Eric J., and Mark Doms.** 2000. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature*, 38(3): 569–594.
- Bernard, Andrew B, Stephen J Redding, and Peter K Schott.** 2011. "Multiproduct firms and trade liberalization." *The Quarterly Journal of Economics*, 126(3): 1271–1318.
- Bils, Mark.** 2009. "Do Higher Prices for New Goods Reflect Quality Growth or Inflation?" *The Quarterly Journal of Economics*, 124(2): 637–675.

- Bils, Mark, and Peter J. Klenow.** 2001. “Quantifying Quality Growth.” *American Economic Review*, 91(4): 1006–1030.
- Bils, Mark, and Peter J. Klenow.** 2004. “Some Evidence on the Importance of Sticky Prices.” *Journal of Political Economy*, 112(5): 947–985.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb.** 2017. “Are Ideas Getting Harder to Find?” National Bureau of Economic Research NBER Working Papers 23782.
- Boskin, Michael, Robert J. Gordon Ellen Dullberger, Zvi Griliches, and Dale Jorgenson.** 1996. “Toward a more accurate measure of the cost of living: Final report of the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index.” *Advisory Commission to Study the Consumer Price Index*.
- Broda, Christian, and David Weinstein.** 2010. “Product Creation and Destruction: Evidence and Price Implications.” *American Economic Review*, 100(3): 691–723.
- Byrne, David M. John G. Fernald, and Marshall B. Reinsdorf.** 2016. “Does the United States have a productivity slowdown or a measurement problem?” Federal Reserve Bank of San Francisco Working Paper Series 2016-3.
- Byrne, David, Stephen Oliner, and Daniel Sichel.** 2015. “How Fast are Semiconductor Prices Falling?” National Bureau of Economic Research, Inc NBER Working Papers 21074.
- Chetty, Raj, David Grusky, Maximilian Hell, Nathaniel Hendren, Robert Manduca, and Jimmy Narang.** 2017. “The fading American dream: Trends in absolute income mobility since 1940.” *Science*, 356(6336): 398–406.
- Davis, Steven, John Haltiwanger, and Scott Schuh.** 1998. *Job Creation and Destruction*. Vol. 1. 1 ed., The MIT Press.

- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda.** 2014. "The Role of Entrepreneurship in U.S. Job Creation and Economic Dynamism." *Journal of Economic Perspectives*, 28(3): 3–24.
- Diewert, Erwin.** 2000. "The Challenge of Total Factor Productivity." *International Productivity Monitor*, 1: 45–52.
- Erickson, Tim, and Ariel Pakes.** 2011. "An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods." *American Economic Review*, 101(5): 1707–38.
- Feenstra, Robert C.** 1994. "New Product Varieties and the Measurement of International Prices." *American Economic Review*, 84(1): 157–77.
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow.** 2016. "How Destructive is Innovation?" National Bureau of Economic Research NBER Working Papers 22953.
- Gordon, Robert J.** 2012. "Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds." National Bureau of Economic Research NBER Working Papers 18315.
- Griliches, Zvi.** 1996. "The Discovery of the Residual: A Historical Note." *Journal of Economic Literature*, 34(3): 1324–1330.
- Groshen, Erica L. Brian C. Moyer, Ana M. Aizcorbe, Ralph Bradley, and David M. Friedman.** 2017. "How Government Statistics Adjust for Potential Biases from Quality Change and New Goods in an Age of Digital Technologies: A View from the Trenches." *Journal of Economic Perspectives*, 31(2): 187–210.
- Grossman, Gene M., and Elhanan Helpman.** 1991. "Quality Ladders in the Theory of Growth." *Review of Economic Studies*, 58(1): 43–61.

- Haltiwanger, John.** 2015. "Job Creation, Job Destruction, and Productivity Growth: The Role of Young Businesses." *Annual Review of Economics*, 7(1): 341–358.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda.** 2013. "Who Creates Jobs? Small versus Large versus Young." *The Review of Economics and Statistics*, 95(2): 347–361.
- Hottman, Colin J. Stephen J. Redding, and David E. Weinstein.** 2016. "Quantifying the Sources of Firm Heterogeneity." *The Quarterly Journal of Economics*, 131(3): 1291–1364.
- Hulten, Charles R.** 2000. "Total Factor Productivity: A Short Biography." National Bureau of Economic Research NBER Working Papers 7471.
- Jones, Charles I.** 1995. "Time Series Tests of Endogenous Growth Models." *The Quarterly Journal of Economics*, 110(2): 495–525.
- Jorgenson, Dale, and Zvi Griliches.** 1967. "The Explanation of Productivity Change." *Review of Economic Studies*, 34(3): 249–283.
- Klenow, Peter J. and Oleksiy Kryvtsov.** 2008. "State-dependent or time-dependent pricing: Does it matter for recent US inflation?" *The Quarterly Journal of Economics*, 123(3): 863–904.
- Klette, Tor Jakob, and Samuel S. Kortum.** 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy*, 112(5): 986–1018.
- Kortum, Samuel S.** 1997. "Research, Patenting, and Technological Change." *Econometrica*, 65(6): 1389–1420.
- Lipsey, Richard, and Kenneth Carlaw.** 2004. "Total factor productivity and the measurement of technological change." *Canadian Journal of Economics*, 37(4): 1118–1150.

- Lipsey, Richard G., and Kenneth Carlaw.** 2000. "What Does Total Factor Productivity Measure?" *International Productivity Monitor*, 1: 31–40.
- Moulton, Brent R. and Karin E. Moses.** 1997. "Addressing the Quality Change Issue in the Consumer Price Index." *Brookings Papers on Economic Activity*, 1997(1): 305–366.
- Nakamura, Emi, and Jon Steinsson.** 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *The Quarterly Journal of Economics*, 123(4): 1415–1464.
- Redding, Stephen J., and David E. Weinstein.** 2016. "A Unified Approach to Estimating Demand and Welfare." National Bureau of Economic Research, Inc NBER Working Papers 22479.
- Romer, Paul M.** 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98(5): 71–102.
- Syverson, Chad.** 2016. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." National Bureau of Economic Research NBER Working Papers 21974.
- U.S. Bureau of Economic Analysis.** 2014. "Principal Source of Data." In *Concepts and Methods of the U.S. National Income and Product Accounts*. Chapter 3. Washington DC:U.S. Department of Commerce.
- U.S. Bureau of Labor Statistics.** 2015. "Producer Prices." In *Handbook of Methods*. Chapter 14. Washington DC:U.S. Department of Labor.
- U.S. General Accounting Office.** 1999. "Consumer Price Index: Impact of Commodity Analysts' Decisionmaking Needs to be Assessed." GAO/GDD-99-84.