

FEDERAL RESERVE BANK OF SAN FRANCISCO

WORKING PAPER SERIES

Missing Growth from Creative Destruction

Philippe Aghion

College de France and London School of Economics

Antonin Bergeaud

London School of Economics

Timo Boppart

IIES, Stockholm University

Peter J. Klenow

Stanford University

Huiyu Li

Federal Reserve Bank of San Francisco

August 2018

Working Paper 2017-04

<http://www.frbsf.org/economic-research/publications/working-papers/2017/04/>

Suggested citation:

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, Huiyu Li. 2018. "Missing Growth from Creative Destruction," Federal Reserve Bank of San Francisco Working Paper 2017-04. <https://doi.org/10.24148/wp2017-04>

The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Board of Governors of the Federal Reserve System.

Missing Growth from Creative Destruction

Philippe Aghion, Antonin Bergeaud

Timo Boppart, Peter J. Klenow, Huiyu Li*

August 2018

Abstract

When products exit due to entry of better products from new producers, statistical agencies typically impute inflation from surviving products. This understates growth if creatively-destroyed products improve more than surviving products. Accordingly, the market share of surviving products should shrink. Using entering and exiting establishments to proxy for creative destruction, we estimate missing growth in U.S. Census data on non-farm businesses from 1983–2013. We find: (i) missing growth is substantial — around half a percentage point per year; but (ii) missing growth did not accelerate much after 2005, and therefore does not explain the sharp slowdown in growth since then.

*Aghion: Collège de France and London School of Economics; Bergeaud: Banque de France; Boppart: IIES, Stockholm University; **Klenow**: Stanford University; Li: Federal Reserve Bank of San Francisco. We thank Raouf Boucekkine, Pablo Fajgelbaum, Gita Gopinath, 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, four referees, and numerous seminar participants provided helpful comments. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the Federal Reserve System or the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is 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 (2018), 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 due to creative destruction could arise 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 estimate the magnitude of missing growth based on our model. We use micro data from the U.S. Census on employment at all private nonfarm businesses to estimate missing growth from 1983–2013. 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.

Our findings can be summarized as follows. First, missing growth from imputation is substantial: roughly one-half a percentage point per year, or around one-third of measured productivity growth. Second, we find only a modest acceleration of missing growth since 2005, an order of magnitude smaller than needed to explain the slowdown in measured growth.

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 ends our example, which hopefully clarifies the main mechanism by which imputation can miss growth from creative destruction.³

Our paper touches on the recent literature about secular stagnation and growth measurement. Gordon (2012) argues that innovation has run into diminishing returns, inexorably slowing TFP growth.⁴ Syverson (2016) and Byrne, Fernald and Reinsdorf (2016) conclude that understatement of growth in the ICT sector cannot account for the productivity slowdown since 2005. In contrast to these studies, we look at missing growth for the whole economy, not just from the ICT sector.

³This example is stylized. In practice, imputation by the BLS is carried out within the items category or category-region. See the U.S. General Accounting Office (1999).

⁴Related studies include Jones (1995), Kortum (1997), and Bloom et al. (2017).

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).⁵ We make two contributions relative to these important papers. First, we compute missing growth for the entire private nonfarm sector from 1983–2013.⁶ Second, we focus on a neglected source of missing growth, namely imputation in the event of creative destruction. The missing growth we identify is likely to be exacerbated when there is error in measuring quality improvements by incumbents on their own products.⁷

The rest of the paper is organized as follows. In Section 2 we lay out an environment relating missing growth to creative destruction. In Section 3 we estimate missing growth using U.S. Census data on plant market shares. Section 4 compares our estimates to those in three existing studies. Section 5 concludes.

2 A model of missing growth

In this section we relate our measure of missing growth to Feenstra (1994). We modify his approach to focus on bias from new *producers* of a given product

⁵Feenstra (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.

⁶Broda 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.

⁷Unlike 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.

line, as opposed to all new products. This allows us to highlight the role of creative destruction in missing growth, as opposed to quality bias from incumbent improvement of their own products.

2.1 Relating missing growth to market share dynamics

Time is discrete and in each period consumption has a CES structure

$$C_t = \left(\int_0^{N_t} [q_t(j) c_t(j)]^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}, \quad (1)$$

where $c_t(j)$ denotes quantity and $q_t(j)$ the quality of variety j . N_t is the number of varieties available, which can change over time. Here $\sigma > 1$ denotes the constant elasticity of substitution between varieties.

Let π_t denote true inflation, the log first difference between $t - 1$ and t in the minimum cost of acquiring one unit of composite C . Following Feenstra (1994) we can decompose true inflation as

$$\pi_t = \hat{\pi}_t - \frac{1}{\sigma - 1} \log \left(\frac{S_{I_t, t-1}}{S_{I_t, t}} \right), \quad (2)$$

where $\hat{\pi}_t$ is inflation for a subset of products denoted I_t , and $S_{I_t, \tau}$ is the share in nominal expenditure spent on this set of products at time τ . Here $\hat{\pi}_t$ can be constructed from the nominal expenditure shares $s_{I_t, \tau}(j) \equiv \frac{S_{j, \tau}}{S_{I_t, \tau}}$ and quality-adjusted prices of products $p_\tau(j)/q_\tau(j)$ within the subset I_t as

$$\hat{\pi}_t \equiv \int_{j \in I_t} f(j, \{s_{I_t, t-1}(i), s_{I_t, t}(i)\}_{i \in I_t}) \log \left(\frac{p_t(j)/q_t(j)}{p_{t-1}(j)/q_{t-1}(j)} \right) dj, \quad (3)$$

where $f(j, \{s_{I_t, t-1}(i), s_{I_t, t}(i)\}_{i \in I_t})$ is the Sato-Vartia weight.⁸

⁸Sato-Vartia weights are

$$f(j, \{s_{I_t, t-1}(i), s_{I_t, t}(i)\}_{i \in I_t}) = \frac{\frac{s_{I_t, t}(j) - s_{I_t, t-1}(j)}{\log s_{I_t, t}(j) - \log s_{I_t, t-1}(j)}}{\int_{i \in I_t} \frac{s_{I_t, t}(i) - s_{I_t, t-1}(i)}{\log s_{I_t, t}(i) - \log s_{I_t, t-1}(i)} di}.$$

Feenstra (1994) chose I_t to be the set of U.S. import category-country pairs that overlapped between $t - 1$ and t . If the U.S. import price index is captured by $\hat{\pi}_t$, then the second term in (2) reflects inflation bias due to the changing set of category-country pairs. Intuitively, if the market share of the subset of products used for measurement changes, then this subset is not representative. In particular, if the market share of the subset I_t is shrinking, then the relative price of the subset must be rising. For a given change in market shares, the bias is smaller when the σ elasticity is higher, because a smaller relative price change is needed to explain the observed change in market shares.

Broda and Weinstein (2010) chose I_t to be the set of UPC codes in AC Nielsen scanner data with positive sales in adjacent periods. They argue that quality is fixed over time for a product with a given UPC code. They therefore infer bias in measured inflation *à la* Feenstra from the change in market share of overlapping UPC codes. This assumes that all quality growth is missed in inflation measurement, including improvements by incumbents on their own products. In practice, statistical offices do attempt to measure quality changes and net them out of inflation, in particular for incumbent product upgrades.⁹

Our focus is on whether imputation leads to missing growth — in particular in cases of creative destruction. We therefore define the subset of products I_t as those produced by the same producer in overlapping periods. This is broader than the set used by Broda-Weinstein since it includes successive generations of products in a given line produced by the same firm. The omitted set of goods is exiting and entering producer-product pairs. As we argue in our [Online Appendix A](#), the BLS infers inflation for these entering and exit producer-product pairs using the inflation rate for continuing pairs. If the market share of overlapping pairs is shrinking over time, however, then the inflation rate of continuing pairs overstates overall inflation and understates real growth. Thus, when the market share of continuing pairs shrinks, we infer missing growth from creative destruction and/or brand new varieties.

⁹See Table 1 in Groshen et al. (2017) and Table B1 in Moulton (2018).

We make these links explicit in the next subsection by adding a supply to side to the model with growth coming from quality improvements by incumbent producers on their own products, creative destruction, and new variety creation. In the next section we describe how we implement this modified Feenstra approach on U.S. Census establishment data.

2.2 Framework of creative destruction and missing growth

There is a representative household supplying inelastically L units of labor each period and maximizes (1) subject to a budget constraint $\Pi_t + W_t L = \int_0^{N_t} p_t(j) c_t(j) dj$, where Π_t and $W_t L$ denotes profits and labor earnings.

On the supply side, each variety $c(j)$ is produced one-for-one with labor

$$c_t(j) = l_t(j),$$

by a monopolistically competitive producer, where $l_t(j)$ is the amount of labor used to produce good j in period t . Given the isoelastic demand that results from the household problem, it is optimal for each producer to set the price

$$p_t(j) = \frac{\sigma}{\sigma - 1} W_t, \tag{4}$$

where W_t is the nominal wage that equalizes across firms in the competitive labor market. We assume the nominal wage follows an exogenously given path over time which will imply that the other nominal variables like prices, $p(j)$, profits, Π , and nominal output will grow at the same rate over time.

We model technical change as product innovation.¹⁰ At each point in time,

¹⁰This 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.

and for each variety 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.¹¹ If j is an existing variety where quality is improved upon by a new producer, we have

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

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

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

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 innovation are constant over time and across varieties.

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

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

¹¹We assume $\gamma_d > \frac{\sigma}{\sigma-1}$ and Bertrand competition within each 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.

We allow the (relative) quality of new product varieties to differ from the “average” quality of pre-existing varieties by a factor γ_n .¹²

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 utility function (1) features love-for-variety.

True, measured and missing growth There are no investment goods so that aggregate nominal output equals aggregate nominal consumption expenditure.¹³ Nominal expenditure can be expressed as $P_t C_t$ where $P_t \equiv \left(\int_0^{N_t} (p_t(j)/q_t(j))^{1-\sigma} dj \right)^{1/(1-\sigma)}$ is the quality-adjusted ideal price index. In this economy the real (gross) output growth $g_t = \frac{C_t}{C_{t-1}}$ is given by nominal output growth divided by the inflation rate $\frac{P_t}{P_{t-1}}$.

As we show in [Online Appendix B](#), the true growth rate in the economy is

$$g = \left[1 + \lambda_d (\gamma_d^{\sigma-1} - 1) + (1 - \lambda_d) \lambda_i (\gamma_i^{\sigma-1} - 1) + \lambda_n \gamma_n^{\sigma-1} \right]^{\frac{1}{\sigma-1}}. \quad (5)$$

This equation shows how the arrival rates and step sizes affect the growth rate. The term $\lambda_n \gamma_n^{\sigma-1}$ captures the effect of variety expansion on growth, and the growth rate is increasing 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 growth. The term

¹²More formally, we assume that a firm that introduces in period t a new variety ι starts with a quality that equals $\gamma_n > 0$ times the “average” quality of pre-existing varieties $j \in [0, N_{t-1}]$ in period $t - 1$, that is

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

Average quality here is the geometric average, which 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 ($\gamma_n < 1$). As the relative quality of new varieties is a free parameter, expressing it this way is without loss of generality.

¹³In an [Online Appendix C](#) we allow for capital and show how this affects the interpretation of our missing growth estimates relative to measured growth.

$\lambda_d (\gamma_d^{\sigma-1} - 1)$ captures the effect from creative destruction on the growth rate.

Next we turn to measured growth. We posit that the statistical office resorts to imputation in response to producer-product entry and exit (we argue this is realistic in [Online Appendix A](#)). That is, they presume the set of continuing producer-product pairs is representative of the economy-wide inflation rate.¹⁴ For the products that are not subject to creative destruction we assume that statistical agency observes unit prices correctly. Their quality adjustments allow them to retrieve the true frequency and step size of quality improvements by incumbents on their own products. As we show in our [Online Appendix B](#), under these assumptions, the *measured* real growth rate is then

$$\hat{g}_t = [1 + \lambda_i (\gamma_i^{\sigma-1} - 1)]^{\frac{1}{\sigma-1}}. \quad (6)$$

We then define the log difference between true growth and measured growth as missing growth (MG). Combining (5), and (6), allows us to approximate missing growth as¹⁵

$$MG_t \approx \frac{\lambda_d (\gamma_d^{\sigma-1} - 1) + \lambda_n \gamma_n^{\sigma-1} - \lambda_d \lambda_i (\gamma_i^{\sigma-1} - 1)}{\sigma - 1}. \quad (7)$$

The first two terms are growth from creative destruction and new varieties, respectively. They are not missed entirely because growth is imputed based on incumbent own innovations in the event of creative destruction, which accounts for the last term. Growth is missed when the rate of innovation from creative destruction and new varieties exceeds that imputed from continuing producer-product pairs. Equation (7) allows us to theoretically decompose missing growth into its sources — creative destruction and variety expansion.

¹⁴BLS imputation is actually carried out within categories or category-regions. See the U.S. General Accounting Office (1999).

¹⁵The exact growth rate is

$$MG_t = \frac{\log [1 + \lambda_d (\gamma_d^{\sigma-1} - 1) + (1 - \lambda_d) \lambda_i (\gamma_i^{\sigma-1} - 1) + \lambda_n \gamma_n^{\sigma-1}] - \log (1 + \lambda_i (\gamma_i^{\sigma-1} - 1))}{\sigma - 1}.$$

Under the appropriate definition of the subset of products I_t , missing growth can again be expressed in terms of formula (2) above as

$$MG_t = \hat{\pi}_t - \pi_t = \frac{1}{\sigma - 1} \log \left(\frac{S_{I_t, t-1}}{S_{I_t, t}} \right), \quad (8)$$

where measured inflation $\hat{\pi}_t$ is as defined in equation (3). Importantly, as discussed, we define the set of overlapping products I_t in (3) as the set of continuing *producer-product pairs*. Growth is missed when the market share of these continuing pairs shrinks over time.

In our [Online Appendix D](#), we derive missing growth when the quality improvement of incumbents is not perfectly measured. In particular, we show that missing growth due to *creative destruction* would be larger if quality improvements by incumbents are understated, precisely due to imputation.

3 Estimates of missing growth

Here we present estimates of 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 which avoids having to estimate the size and frequency of the various types of innovations.

3.1 Measuring the market share of continuers

Our goal is to quantify missing growth in the aggregate economy over a time horizon of several decades. We therefore base our 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 information in this dataset to infer the market share of continuing plants, $S_{I_t, t}$.

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 the Census of Manufacturing, or consumer nondurables in the 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 continuing incumbent products.

These assumptions are admittedly strong. They require that firms do not add products through existing plants. Bernard, Redding and Schott (2011) find that U.S. manufacturing plants do start up production in new industries. Our assumption may be a better approximation outside manufacturing, such as in retail where location is a key form of product differentiation. Related, if creative destruction occurs through process innovations embodied in new establishments (e.g., new Walmart outlets), then the market share of new plants should in principle capture them.

If existing plants do introduce new varieties or carry out creative destruction, then our approach is likely to understate missing growth. As we explain below, our baseline specification will 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 is much lower after age 5 than for younger plants (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, 2018). If plants add varieties after age 5, then one would expect exit rates to fall rapidly beyond age 5.

Our baseline estimates use employment data to measure market shares, rather than revenue or even payroll data. In our model these variables are all proportional to each other across products. But in practice they differ. Annual revenue data are only available at the firm level. Plant level revenue data is

available to us only every 5 years in the manufacturing sector. As a robustness check we will report results with revenue for manufacturing. And we will show robustness to using payroll rather than employment for all sectors.

To be more exact, we calculate market shares using plant data as follows. 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 operated in both $t - 1$ and t , that is, all plants with $B \leq t - 1$ and $D \geq t$. Define E_t as the group of plants that first operated in period t ($B = t, D \geq t$) and X_t as the group of plants that last operated in period t ($B \leq t, D = t$). Let $L(t, \mathcal{M})$ denote the total employment in period t of plants belonging to group \mathcal{M} . We then measure the ratio $\frac{S_{I_t, t-1}}{S_{I_t, t}}$ on the right-hand side of (2) as

$$\frac{S_{I_t, t-1}}{S_{I_t, t}} = \frac{\frac{L(t-1, I_t)}{L(t-1, I_t) + L(t-1, X_{t-1})}}{\frac{L(t, I_t)}{L(t, I_t) + L(t, E_t)}}, \quad (9)$$

According to (9), missing growth is positive whenever the employment share of continuing plants shrinks between $t - 1$ and t .

We define a period t as a calendar year, and map D to the last year a plant is in LBD. We set B equal to $k \geq 0$ years after the plant first appears in the LBD. If τ is the first year the plant appears in the database, we set $B = \tau + k$. We use $k = 5$ in our baseline specification.

3.2 Elasticities of substitution

To quantify missing growth we need the elasticity of substitution across plants in all nonfarm business sectors. Hall (2018) estimates markups for 18 2-digit sectors using KLEMS data for 1988-2015.¹⁶ He regresses industry output growth on share-weighted input growth instrumented by military purchases and oil prices, in the hopes that these are orthogonal to residual productivity growth.

¹⁶Hall estimates that markups are modestly (but not precisely) trending up over time. His sample period is close to our 1983–2013 LBD sample time frame.

We convert Hall (2018)'s estimated markups to elasticities assuming the markups equal $\frac{\sigma}{\sigma-1}$. We display the implied elasticities for each 2-digit sector in the first few columns of Table 1. In a few of the industries, Hall (2018) reports markups that are less than one. In these industries (health care, construction and education), we assume $\sigma \rightarrow \infty$. As shown in the Table, Hall (2018)'s implied σ 's range from below 3 to over 26 in the industries with markups in excess of one.

3.3 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. 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 (1977 to 1982 at the beginning). We examine missing growth starting in 1983.¹⁷

In the fourth column of Table 1, we report our missing growth estimates for each 2-digit sector using Hall (2018)'s implied σ 's. We order the sectors by their contribution to missing growth on average over 1983–2013, which totals 54 basis points per year.¹⁸ The biggest contributors, by far, are NAICS 72 (hotels and

¹⁷More specifically, we calculate missing growth in each year t using data from years $t - 6$, $t - 5$, $t - 1$ and t . We use data in year $t - 6$ to identify plants that have been in the data for at least 5 years in $t - 1$ and calculate their total employment in year $t - 1$. Then we calculate the share of this total employment belonging to plants that survive to year t . Similarly, we use year $t - 5$ data to identify plants that are in the data for at least 5 years in t and calculate the t period employment share of plants surviving from period $t - 1$.

¹⁸We aggregate using Tornqvist averages of the current and previous year employment shares of the sector; this is a second order approximation to any smooth utility aggregator of 2-digit sectors. We calculate the contribution of a sector in a year by multiplying the missing growth in that sector with the Tornqvist employment share. Then we take averages over 1983 to 2013 to arrive at the average contribution reported in the table.

Table 1: AVERAGE MISSING GROWTH 1983–2013 BY 2-DIGIT SECTORS

NAICS	Sector name	Hall σ	Missing growth (ppt)	Contribution (ppt)	(%)
72	Hotels & Restaurants	2.82	2.70	0.18	33.9
44-45	Retail Trade	3.22	1.23	0.15	28.6
54	Professional Services	4.23	1.07	0.05	8.6
52	Finance and Insurance	3.17	0.60	0.03	5.4
81	Other Services	4.03	0.54	0.02	4.4
48-49	Transportation & Warehousing	4.23	0.58	0.02	3.7
71	Arts, Entertainment	2.92	1.51	0.02	3.6
51	Information	3.56	0.88	0.02	3.4
42	Wholesale Trade	4.70	0.29	0.01	2.5
31-33	Manufacturing	3.44	0.05	0.01	2.2
56	Admin & Support Services	26.0	0.21	0.01	1.7
55	Management of Companies	6.26	0.13	0.01	0.9
53	Real Estate	12.1	0.24	0.00	0.6
22	Utilities	9.33	0.16	0.00	0.3
21	Mining	4.85	0.37	0.00	0.3
62	Health Care	∞	0	0	0
23	Construction	∞	0	0	0
61	Education	∞	0	0	0
Total				0.54	

restaurants) and 44-45 (retail trade); they contribute 33 of the 54 basis points, or over 62% of the total. Their contribution may reflect the geographic spread of outlets by big national chains. No other sector contributes more than 5 basis points and we find very little missing growth in manufacturing.

The variation in missing growth across sectors in Table 1 suggests that official relative price trends across sectors may be biased considerably. Taken at face value, our estimates alter the pattern of Baumol’s cost disease cross sectors.

Table 2: MEASURED VS. TRUE GROWTH WITH HALL σ 'S

	Missing Growth	Measured Growth	“True” Growth	% of growth missed
1983–2013	0.54	1.87	2.41	22.4%
1983–1995	0.52	1.80	2.32	22.4%
1996–2005	0.48	2.68	3.16	15.2%
2006–2013	0.65	0.98	1.63	39.9%

Notes: Entries are percentage points per year. Missing growth is calculated using equation (2). The market share is measured as the employment share of plants in the Census Longitudinal Business Database (LBD) as in (9). These baseline results assume a lag $k = 5$ and use the 2-digit elasticities of substitution in Hall (2018). 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.

Table 2 compares our missing growth estimates to official TFP growth. The entries are annual percentage points. As mentioned, we find 0.54 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.¹⁹ If we add our

¹⁹We put BLS TFP growth in labor-augmenting form, and 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.

missing growth to the BLS TFP series we arrive at “true” growth of 2.41% per year. Thus our baseline estimate is that over one-fifth of true growth is missed.

Table 2 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 by 4 basis points when official growth accelerated by 88 basis points in the middle period. And missing growth sped up by 17 basis points when official growth dropped 170 basis points in the final period.

While missing growth did not accelerate significantly in the aggregate, it did so in the Information sector. In this sector, missing growth rose from 0.22 percent per year over 1983–2013 to 0.90 over 1996–2005 and to 1.90 over 2006–2013. Due to its small employment share, however, the Information sector contributed a small amount to overall missing growth: 0.00 over 1983–2013, 0.02 over 1996–2005 and 0.04 over 2006–2013.

3.4 Robustness and discussion

In this section we discuss how our estimates of missing growth are affected by the elasticity of substitution, the lag used to defined new plants, and the data used to measure market shares (e.g., using payroll instead of employment data).

Elasticities of substitution As a robustness check, we use the Visa data to estimate σ for the two sectors responsible for 57% of the total missing growth in Table 1: retail trade and restaurants.²⁰ Following the methodology of Dolfen et al. (2018), we estimate σ by using the location of Visa cardholders vs.

²⁰In Table 1, Retail Trade contributed 28.6% of the total missing growth whereas Hotels and Restaurants contributed 33.9%. We calculate the contribution of Retail Trade plus Restaurants by $28.6\% + 33.9\% \times 83.8\% = 57.0\%$ where 83.8% is the share of Restaurant employment (NAICS 722) in Hotels and Restaurants employment (NAICS 72) in the 2012 Census of Accommodation and Food Services.

physical stores and converting distance into effective price variation, based on the opportunity cost of time and the direct costs of travel. We add the differential travel costs to the average spending per visit at the store. Using all cards, we regress relative visits on relative prices across stores to estimate the elasticity of substitution in a given sector:

$$\log\left(\frac{c_{ij}}{c_{ik}}\right) = \log\left(\frac{q_j}{q_k}\right) - (\sigma - 1) \log\left(\frac{p_{jk} + \tau_{ij}}{p_{jk} + \tau_{ik}}\right). \quad (10)$$

Here i denotes a cardholder, j and k are competing merchants, c refers to the number of visits, q denotes residual quality (assumed to be orthogonal to customer distance to merchant j versus k), p_{jk} is the average card spending per visit at merchants j and k , and the τ 's are estimated costs of travel.²¹

We only compare stores of competing chains within 3-digit NAICS (e.g., general merchandisers Walmart vs. Target, or grocery stores Trader Joe's vs. Whole Foods). Across restaurants (NAICS 722) we estimate an elasticity of substitution of 2.92, not far from Hall (2018)'s estimate of 2.82. Across retail establishments (NAICS 44-45), we estimate an elasticity of 5.01, higher than Hall's estimate of 3.22.

In Table 3 we report missing growth for Retail Trade and Restaurants combined, where we aggregate these two industries using Tornqvist employment shares. Due to the higher σ estimate for Retail, the Visa data implies lower missing growth (1.26 percentage points per year) than when we use Hall (2018) σ values (1.68), and a correspondingly smaller contribution to aggregate missing growth of 23 vs. 30 basis points per year.

We next show the broader sensitivity of missing growth to the degree of substitution across plants. As shown in (8), missing growth is proportional to

²¹Dolfen et al. (2018) estimate 79 cents in direct costs (fuel, depreciation) and 80 cents in indirect costs (opportunity cost of time based on after-tax wages), for a total of \$3.18 per roundtrip mile. We follow them in using cardholder and store locations to calculate driving distance to stores, and multiplying the distance by this cost per mile. In using average spending per visit across the two merchants for p_{jk} , we are assuming that the price and bundle of items bought is the same at competing merchants, other than quality differences.

Table 3: Average missing growth for Retail Trade and Restaurants combined

	Using Hall σ 's	Using Visa σ 's
Missing growth in the sector	1.68	1.26
Contribution to overall missing growth	0.30	0.23

$1/(\sigma - 1)$. Thus missing growth declines monotonically as we raise σ in Table 4. The higher is σ , the smaller the quality and variety improvements by new plants to needed explain the observed decline in the observed market share of continuing plants.

Table 4: MISSING GROWTH WITH DIFFERENT ELASTICITIES σ

	Missing Growth		
	Higher elasticities	Benchmark	Lower elasticities
1983–2013	0.43	0.54	0.72
1983–1995	0.42	0.52	0.69
1996–2005	0.38	0.48	0.64
2006–2013	0.52	0.65	0.87

Note: The Benchmark columns uses σ values from Hall (2018). The Higher elasticities are 25% higher than the $\sigma - 1$ values estimated by Hall, and the Lower elasticities are 25% lower than Hall's.

Using payroll to measure market shares Table 5 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 5: MISSING GROWTH WITH EMPLOYMENT VS. PAYROLL

	Employment	Payroll
1989–2013	0.59	0.56
1989–1995	0.66	0.72
1996–2005	0.48	0.31
2006–2013	0.65	0.73

Note: Entries are percentage points per year.

Using revenue to measure market shares For manufacturing only, we have access to revenue data through the Census of Manufacturing (CMF). The CMF is available every 5 years. To calculate missing growth with lag $k = 5$ in a census year t , we define incumbents as plants that are in the census in year t and $t - 10$ and calculate the market share growth of incumbents between year $t - 5$ and t . We convert this growth over 5 years to annual growth by taking logs and dividing by 5. We then use $\sigma = 3.44$ for manufacturing from Hall (2018) to convert the resulting annualized growth into missing growth. Table 6 provides our missing growth estimates based on revenue versus employment for manufacturing. The time periods are altered slightly to coincide with Census years.

The market share of continuing manufacturing plants shrinks more in terms of revenue than employment, so that missing growth based on equation (8) is higher when market shares are measured in terms of revenue. As shown, there is little missing growth in manufacturing except for the end period, so on average this increase is relative to a small contribution. The last column of the table displays missing growth from LBD manufacturing plants for the same periods. The CMF and LBD yield similar missing growth for manufacturing, except for over 2007–2012.

Table 6: MISSING GROWTH WITH REVENUE VS. EMPLOYMENT

	Revenue (CMF)	Employment (CMF)	Employment (LBD)
1982–2012	0.04	-0.05	-0.06
1982–1997	0.09	0.05	0.14
1997–2007	0.06	-0.25	-0.37
2007–2012	0.48	0.33	0.06

Note: All entries are for manufacturing only. For CMF, entries are averages over quinquennial data. For LBD, entries are averages of annual data.

Different lag k Our baseline results in Table 2 evaluate the market share of entrants after a lag of $k = 5$ years. Table 7 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 23 basis points per year rather than 54 basis points. Though not reported in Table 7, 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 $k = 5$.

Table 7: MISSING GROWTH WITH DIFFERENT LAGS k

	$k = 5$	$k = 0$
1983–2013	0.54	0.23
1983–1995	0.52	0.24
1996–2005	0.48	0.27
2006–2013	0.65	0.14

Notes: Entries are percentage points per year.

Missing growth with more sectors Our baseline estimate of missing growth assumes a given elasticity of substitution across all plants within a 2-digit sector. We also calculated missing growth at a more detailed level and then aggregated the missing growth rates using Tornqvist employment shares within each subsector as weights (again, a second order approximation to any smooth utility aggregator). Within subsectors we set the elasticity of substitution equal to the Hall (2018) estimate for the 2-digit sector containing the subsector.

Table 8 compares our benchmark missing growth estimates to those obtained by aggregating up missing growth in this way up from the 3-, 4-, or 5-digit sectoral level. We present our baseline 2-digit estimates as well, for comparison. The estimates are gently increasing in the level of disaggregation (from 54 basis points to 65), and are broadly similar across sub-periods.²²

Table 8: MISSING GROWTH WITH DISAGGREGATED SECTORS

	2-digit	3-digit	4-digit	5-digits
1983–2013	0.54	0.56	0.57	0.65
1983–1995	0.52	0.53	0.46	0.58
1996–2005	0.48	0.54	0.63	0.68
2006–2013	0.65	0.64	0.70	0.72

Notes: Entries are percentage points per year. The first column repeats our baseline results. Other columns are Tornqvist 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

²²We 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 8.

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.

Table 9 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.²³

Table 9: MISSING GROWTH VS. DECLINING DYNAMISM

	Plant level	Firm level	Net entry	Gross entry
1983–2013	0.54	0.15	0.29	0.29
1983–1995	0.52	0.17	0.27	0.41
1996–2005	0.48	0.16	0.32	0.29
2006–2013	0.65	0.10	0.27	0.09

Notes: Entries are percentage points per year.

The third column of Table 9, labeled “Net Entry”, shows how big missing growth would be using firm-level data and assuming all firms were of the same size (i.e., had the same level of employment). In this case, missing growth is

²³We identify firms using the LBD firm id, which can change with mergers and acquisition events. We also calculated firm-level missing growth using an alternative firm id for each plant, where we set the firm id of a plant to its initial firm id throughout the lifetime of the plant. We find higher firm-level missing growth (0.31 percent per year versus 0.15 over 1983–2013) after netting out M&A activity in this way.

larger than the firm-level estimate because entering firms tend to be smaller than the average firm.

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 9 shows missing growth if all firms were equal-sized and the exit rate was fixed to the 1983–2013 average. 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.

To recap, declining dynamism is seen most strikingly in the gross entry rate, which is several steps removed from our missing growth calculation. We base our missing growth estimates on plant dynamics, 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.

4 External validation

Here we compare our sectoral missing growth estimates with those in two prominent papers that used more detailed data on prices and quantities, and to those obtained using a separate, indirect inference approach.

4.1 Bils (2009) and Broda and Weinstein (2010)

Bils (2009) uses scanner data to estimate quality bias in the CPI for consumer durables. He estimates a bias of 1.8% per year from 1988–2006. To obtain a comparable estimate, we first restrict our attention to Census retail NAICS for durable consumer goods, which include 441 (Motor Vehicle and Parts Dealers), 442 (Furniture and Home Furnishings Stores), and 443 (Electronics and Appliance Stores). Comparing stores of competing chains in these categories, we estimate $\sigma = 7.9$ using Visa data. With this Visa-based σ and the market share of continuers in the LBD in these industries, we arrive at missing growth

of 0.36 percentage points per year from 1988–2006. This compares to 1.8 percentage points per year in Bils (2009). Bils’s number is understandably higher because he includes understated improvements of incumbent products, whereas our focus is solely on new outlets.

Similarly, we can restrict attention to grocery and drug store retailers to facilitate comparison to Broda and Weinstein (2010). We map this to retailers selling nondurables other than gasoline: categories 445 (Food and Beverage Stores) and 446 (Health and Personal Care Stores). Comparing Visa spending at stores within these industries, we estimate $\sigma = 6.0$. We then combine this σ with data on continuer shares in the LBD for these industries over the period 1994–2003. We find annual missing growth of 0.43 percentage points per year, compared to the Broda and Weinstein (2010) estimate of 0.8 percent per year. Again, our estimate should be smaller in that we focus on new outlets, whereas their estimate would capture improvements in products at continuing retailers.

4.2 Indirect inference approach

Our market share approach is simple and intuitive. It does, however, require that a plant not add new product lines — all creative destruction must occur through new plants. Furthermore, it cannot separate out expanding variety from creative destruction. We therefore entertained an alternative “indirect inference” methodology which does not rely on such assumptions. This methodology starts from the decomposition of missing growth into its creative destruction and variety expansion components. We use the exact growth rate, but for intuition the approximate growth decomposition is useful:

$$MG_t \approx \frac{\lambda_d (\gamma_d^{\sigma-1} - 1) + \lambda_n \gamma_n^{\sigma-1} - \lambda_d \lambda_i (\gamma_i^{\sigma-1} - 1)}{\sigma - 1}.$$

where λ_d (λ_n) is the sum of the arrival rate of creative destruction (new variety) innovation by incumbents and entrants.

We estimate missing growth and its decomposition by first estimating the frequency and size of the various types of innovation, i.e., by estimating $(\lambda_i, \gamma_i, \lambda_d, \gamma_d, \lambda_n, \gamma_n)$. This in turn requires more data moments than when using the market share approach. This method is built on Garcia-Macia, Hsieh and Klenow (2018) (hereafter GHK), who back out arrival rates and step sizes $(\lambda_i, \gamma_i, \lambda_d, \gamma_d, \lambda_n, \gamma_n)$ by inferring parameter values to mimic moments on firm dynamics in the LBD. In the [Appendix](#) we describe the GHK algorithm in some detail. We modify it to incorporate how measured growth can differ from true growth; GHK assumed that growth was measured perfectly.

Table [A1](#) in the [Appendix](#) reports our indirect inference estimates — both parameter values and the missing growth they imply. We find more missing growth under this indirect inference approach (99 basis points per year on average) than under the market share approach (54 basis points per year). But, like the market share approach, the indirect inference approach yields no big acceleration in missing growth to account for the sharp slowdown in measured growth over 2003–2013. Finally, the indirect inference approach implies that the vast majority (over 80%) of missing growth is due to creative destruction. Expanding varieties play a limited role.

5 Conclusion

In this paper we lay out a model with incumbent and entrant innovation to assess the unmeasured TFP growth resulting from creative destruction. Crucial to this missing growth is the use of imputation by statistical agencies when producers no longer sell a product line. Our model generates 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, we estimated the magnitude of missing growth from creative destruction over the period 1983–2013. Our approach uses the market share of surviving, entering,

and exiting plants. We found that: (i) missing growth from imputation was substantial at around half a percentage point per year — or over one-fifth of measured productivity growth; and (ii) it has accelerated modestly since 2005, so that it accounts for only about one-tenth of the sharp growth slowdown since then.

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.²⁴

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 (2018), 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.²⁵

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.

²⁴Economists 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).

²⁵Erickson 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.

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.

A Appendix: 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 in the model of Section 2 and then calculating missing growth based on formulas derived from the model. To estimate the arrival rates and step sizes, we adapt the algorithm developed in Garcia-Macia, Hsieh and Klenow (2018), 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).²⁶ They calculate these

²⁶GHK 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.

moments in the LBD for 1983–1993, 1993–2003 and 2003–2013, respectively.²⁷

With their parameter estimates in hand, GHK decompose growth into 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 $g \equiv \frac{C_{t+1}}{C_t} - 1$ in GHK is

$$1 + g = [(1 - \delta_o \psi) \{[\lambda_i(1 - \lambda_{e,d} - \lambda_{i,d}) + (\lambda_{e,d} + \lambda_{i,d})] (\gamma_i^{\sigma-1} - 1) + 1\} + (\lambda_{i,n} + \lambda_{e,n}) \gamma_n^{\sigma-1}]^{\frac{1}{\sigma-1}}, \quad (A1)$$

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

²⁷Their algorithm matches model steady state moments to data moments, and therefore does not produce annual estimates.

below-cutoff products relative to the average quality;²⁸ 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 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 (A1). 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:

$$1 + \hat{g} = \left(1 + \lambda_i(\gamma_i^{\sigma-1} - 1)\right)^{\frac{1}{\sigma-1}}. \quad (\text{A2})$$

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:²⁹

1. We choose parameters so that (A2) 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.

²⁸ $\delta_o = \int_{q(j) < \bar{q}_t, q(j) \in \Omega_t} 1 dj$ 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 .

²⁹See our [Online Appendix E](#) for a more details description of the changes we made.

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 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 A1 defines the parameters and displays their estimated values for each the three samples: 1983–1993, 1993–2003, and 2003–2013. The bottom panel of Table A1 reports the resulting estimates of measured, true, and missing growth. Missing growth is larger under this alternative approach than under the market share approach for the first two sample periods: 1.25 percentage points per year (vs. 0.52 when using the market share approach) for the 1983–1993 period, and 1.13 percentage points per year (vs. 0.48) for the 1993–2003 period. For the last sample period, 2003–2013, the missing growth estimates from the indirect inference method are closer to those from the market share approach (0.60 percent vs. 0.65). Under the indirect inference approach, the fraction of total productivity growth that is missed is comparable across the three periods at around one-third of true growth. Just like in the market share approach, in the indirect inference approach we do *not* find that missing growth accelerated when measured growth fell sharply in the last interval.

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 —

Table A1: ESTIMATED PARAMETERS AND RESULTS WITH INDIRECT INFERENCE

Parameter	Definition	1983–	1993–	2003–
		1993	2003	2013
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 (2018) on three samples: 1983–1993, 1993–2003, and 2003–2013. We adapt their algorithm to incorporate missing growth as explained in the text.

around 80% in all three periods.

A.1 Comparison to the market share approach

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 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 gross entry of firms, we did obtain a sharp decline in missing growth across periods (Table 9).

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 is less data demanding.

References

- 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.** 2018. “Growth through Heterogeneous Innovations.” *Journal of Political Economy*, 126(4): 1374–1443.
- Atkeson, Andrew, and Ariel Burstein.** 2018. “Aggregate implications of innovation policy.” forthcoming.
- 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.

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.

Dolfen, Paul, Liran Einav, Peter J Klenow, Benjamin Klopach, Jonathan Levin, Larry Levin, and Wayne Best. 2018. “Assessing the gains from e-commerce.” Working paper, Stanford University.

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.” *The American Economic Review*, 157–177.

- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow.** 2018. “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.
- 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.
- Hall, Robert E.** 2018. “New Evidence on the Markup of Prices over Marginal Costs and the Role of Mega-Firms in the US Economy.” National Bureau of Economic Research Working Paper 24574.
- 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.
- Jones, Charles I.** 1995. “Time Series Tests of Endogenous Growth Models.” *The Quarterly Journal of Economics*, 110(2): 495–525.
- Kortum, Samuel S.** 1997. “Research, Patenting, and Technological Change.” *Econometrica*, 65(6): 1389–1420.
- Moulton, Brent.** 2018. “The Measurement of Output, Prices, and Productivity: Whats Changed Since the Boskin Commission?” The Hutchins Center on Fiscal and Monetary Policy at the Brookings Institution.

- 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.
- 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.