

Missing Growth from Creative Destruction*

Philippe Aghion Antonin Bergeaud Timo Boppart

Peter J. Klenow Huiyu Li

June 2017

Most recent version: <http://www.klenow.com/missing-growth.pdf>

Abstract

Statistical agencies typically impute inflation for disappearing products from the inflation rate for surviving products. As some products disappear precisely because they are displaced by better products, inflation may be lower at these points than for surviving products. As a result, creative destruction may result in overstated inflation and understated growth. We use a simple model to relate this “missing growth” to the frequency and size of various kinds of innovations. Using U.S. Census data, we then apply two ways of assessing the magnitude of missing growth for all private non-farm businesses for 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 — approximately 0.5 percentage points per year for both approaches; 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 Banque of France; Boppart: IIES, Stockholm University; Klenow: Stanford University (klenow@stanford.edu); Li: Federal Reserve Bank of San Francisco. We thank Raouf Boucekkine, Colin Hottman and Steve Redding for excellent discussions. Ufuk Akcigit, 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 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 and variety improvements 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 (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.

Creative destruction is known to be a key source of economic growth. See Aghion and Howitt (1992), Akcigit and Kerr (2010), and Aghion et al. (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 entire U.S. nonfarm business sector over the past three decades.

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, but also by competing incumbents and entrants (creative destruction). The model predicts missing growth due to creative destruction if the statistical office resorts to imputation.

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, 2015a, p.10) BLS explicit quality adjustments, such as hedonics, are used predominantly for goods that undergo periodic model changes by *incumbent* producers (Groshen et al., 2017).

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 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 signal their contribution to growth.

In the second approach, we extend the algorithm in [Garcia-Macia et al. \(2016\)](#) 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 new plants introduce new varieties and carry out creative destruction, but does rely on indirect inference.

Our findings from these two quantification exercises can be summarized as follows. First, missing growth from imputation is substantial. We estimate that missing growth averages around 0.5 percentage points per year over the past thirty years when using both the “market share” approach and the “indirect inference” method. Second, the primary source of missing growth has been creative destruction rather than new varieties.

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).² 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](#)).³

Our paper touches on the recent literature on 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.⁴ He argues that the innovation process has run into diminishing returns, leading to an irreversible slowing of TFP growth. [Syverson \(2016\)](#) and [Byrne et al. \(2016\)](#) conclude

²See also [Hulten \(2000\)](#) and [Lipsey and Carlaw \(2000, 2004\)](#).

³See also [Davis et al. \(1998\)](#) and [Bartelsman and Doms \(2000\)](#).

⁴Related studies include [Jones \(1995\)](#), [Kortum \(1997\)](#), and [Bloom et al. \(2016\)](#).

that understated growth in certain sectors, most notably the ICT sector, cannot account for the productivity slowdown since 2005 because, among many reasons, 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 declined over the past thirty years, and in fact seems to have risen modestly. A corollary is that missing growth appears to be a growing fraction of true productivity growth.

More closely related to our analysis are [Feenstra \(1994\)](#), [Bils and Klenow \(2001\)](#), [Bils \(2009\)](#), [Broda and Weinstein \(2010\)](#), [Erickson and Pakes \(2011\)](#), [Byrne et al. \(2015\)](#), 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 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.⁷

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, respectively based on survivor market shares and indirect inference, to compute missing growth estimates using U.S. Census data. [Section 4](#) concludes.

⁵[Feenstra \(1994\)](#) 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 non-durable retail sector, using the AC Nielsen Homescan database, which samples purchases by 55,000 households across 23 U.S. cities. [Byrne et al. \(2015\)](#) look at missing growth in the semiconductor sector.

⁶[Broda 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 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 1.76 percentage points on average for the year 1995.

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 biased measurement of quality improvement as well as expanding product variety.

2.1 Basic setup

2.1.1 Structure of the aggregate economy

Time is discrete and in each period output has a CES structure:⁸

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

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

2.1.2 Intermediate input production

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

$$y(j) = l(j), \quad (2)$$

where $l(j)$ is the amount of labor used to produce intermediate good j . In our [online appendix](#), we extend our core analysis and quantification of missing growth to the case where the production of intermediate inputs also uses capital, according

⁸We remove time subscripts for notational simplicity whenever it risks no confusion.

⁹The 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 (Section 2.3).

to the Cobb-Douglas technology

$$y(j) = (k(j)/\alpha)^\alpha (l(j)/(1 - \alpha))^{1-\alpha},$$

where $k(j)$ denotes physical capital used in production of input j . In the [online appendix](#) we argue that factoring in capital leads to either unchanged or increased missing growth as a fraction of measured TFP growth under reasonable assumptions as to how the growth in capital stock is measured by the statistical office.

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 productivity in producing the final good, whereas intermediate producers are monopolistic but potentially subject to a competitive fringe. There is a representative household supplying inelastically a fixed amount of labor every period. Labor is freely mobile across firms and the wage rate, W , consequently equalizes in equilibrium across all firms. Nominal expenditure by the representative household on the final output good are given by M which gives the budget constraint

$$M = PY.$$

Here, P denotes the price index that we will specify further below.

In the remaining part of the paper we will analyze how, for a given path of M or W (which is subject to normalization), innovations of different sorts affect the dynamics in the true (quality-adjusted) price index, P , of the economy. This will allow us to decompose changes in nominal output, M , into inflation and growth in real output, Y . Then, we will 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 firms. In the [online appendix](#) we show how the framework can be extended to the more general case where the elasticity of substitution and markups vary across sectors and/or over time.

2.1.4 Equilibrium prices

Suppose that all firms j can maximally charge a markup factor $\tilde{\mu} > 1$ over marginal cost $c(j) = W$. Profit maximization by each intermediate monopolist j then implies that it is optimal to charge a markup factor of $\mu > 1$, where we have $\mu = \min\{\tilde{\mu}, \frac{\sigma}{\sigma-1}\}$.¹⁰ Hence, we obtain in equilibrium for the price of each intermediate good j

$$p(j) = \mu W, \quad \forall j. \quad (3)$$

2.1.5 Innovation

We model technical change as product innovation.¹¹ At each point in time, and for each intermediate input j there is an exogenous probability of creative destruction $\lambda_d \in [0, 1)$, i.e., with probability λ_d a new entrant is replacing the incumbent firm. If a new entrant is entering a product market j the incumbent firm is pushed out of the market. We assume that the new entrant improves upon the incumbent's quality by a factor $\gamma_d > 1$. Then, formally, if j is an existing variety where quality is improved upon by a new entrant, we have

$$q_{t+1}(j) = \gamma_d q_t(j). \quad (4)$$

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

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

$$q_{t+1}(j) = \gamma_i q_t(j). \quad (5)$$

We call this type of innovation process *incumbent own innovation*.

Both the arrival rates of creative destruction and incumbent own innovation are constant and uncorrelated with the initial quality level.

¹⁰In the absence of a binding competitive fringe on the intermediate producer's market, one can show that that it is optimal to set $\mu = \frac{\sigma}{\sigma-1}$.

¹¹This is a modeling choice which does not affect the main result of our theoretical analysis. However, it matters in the empirical context since pure process innovation is arguably more likely to be captured by the statistical office. Yet, as it turns out, for products with price information in the Census of manufacturing firms and plants, we find that firm revenues increase without a decline in unit prices. This suggests that innovations are rather of the product than of the process type.

Finally, each period $t + 1$, $N_t \lambda_n$ firms are newly created and sell a new product variety $\iota \in (N_t, N_{t+1}]$ from $t + 1$ onward. Consequently, the law of motion of the number of intermediate inputs is given by

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

To complete our description of the innovation process, we need to state the (relative) quality of new product varieties. This is done through the following assumption.

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

$$q_{t+1}(\iota) = \gamma_n \left(\frac{1}{N_t} \int_0^{N_t} q_t(j)^{\sigma-1} dj \right)^{\frac{1}{\sigma-1}}, \quad \forall \iota \in (N_t, N_{t+1}]. \quad (6)$$

The “average quality” in Assumption 1 refers to the weighted 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 or below average market productivity.

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 growth, both because the production function (1) features gains from specialization and because new varieties may appear at an above-average quality.

2.2 Missing growth

2.2.1 Missing growth as mismeasured inflation

By definition, the aggregate nominal output, M , is equal to the product of the price index, P , and real output, Y . Hence, 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 (quality-adjusted) inflation.

More formally, if $\widehat{\frac{P_t}{P_{t+1}}}$ denotes *measured* inverse gross 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

$$MG_{t+1} = \log\left(\frac{Y_{t+1}}{Y_t}\right) - \log\left(\frac{\widehat{Y}_{t+1}}{Y_t}\right) = \log\left(\frac{\widehat{P}_{t+1}}{P_t}\right) - \log\left(\frac{P_{t+1}}{P_t}\right). \quad (7)$$

Thus there will be positive missing growth whenever inflation is overstated and vice versa.

2.2.2 True prices and inflation

The aggregate price index In the following we derive the “true” welfare based aggregate price index. The results are immediately obtained from the fact that, in each period, the final good sector maximizes current final output, Y , with respect to $\{y(j)\}_{j=0}^N$ subject to $M = \int_0^N y(j)p(j)dj$. We remove time subscripts here again for notational simplicity.

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

$$y(j) = q(j)^{\sigma-1} \left[\frac{P}{p(j)} \right]^\sigma \frac{M}{P}, \quad \forall j. \quad (8)$$

(ii) the equilibrium aggregate price index is given by

$$P = \left(\int_0^N \left[\frac{p(j)}{q(j)} \right]^{1-\sigma} dj \right)^{\frac{1}{1-\sigma}}. \quad (9)$$

Under optimal price setting of the firms we obtain

$$P = \mu W \left(\int_0^N q(j)^{\sigma-1} dj \right)^{\frac{1}{1-\sigma}}. \quad (10)$$

Proof. The first-order conditions when maximizing $Y = \left(\int_0^N [q(j)y(j)]^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}$ with respect to $\{y(j)\}_{j=0}^N$ subject to the (budget) constraint $M = \int_0^N y(j)p(j)dj$, can be written as

$$\xi p(j) = q(j)^{\frac{\sigma-1}{\sigma}} y(j)^{-\frac{1}{\sigma}} \left(\int_0^N [q(j')y(j')]^{\frac{\sigma-1}{\sigma}} d(j') \right)^{\frac{1}{\sigma-1}}, \quad \forall j \in [0, N],$$

where ξ is the Lagrange multiplier attached to the budget constraint. Integrating both sides of this equation over all j 's and combining it with (1) yields

$$\xi = \frac{Y}{M} = \frac{1}{P}.$$

Together with the above first-order conditions, this yields (8). Next, to derive expression (9) for P , note that (8) implies that

$$p(j)y(j) = \frac{M}{P} q(j)^{\sigma-1} P^\sigma p(j)^{1-\sigma}.$$

Integrating both side of this equation over all j 's then immediately yields (9). Finally, substituting for the equilibrium $p(j)$ using (3) in (9) yields equation (10). This establishes the proposition. ■

The true inflation rate Using the above expression for the aggregate price index, we can compute the true inflation rate as a function of the arrival rates and the quality 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*

$$\frac{P_{t+1}}{P_t} = \frac{W_{t+1}}{W_t} [1 + \lambda_d (\gamma_d^{\sigma-1} - 1) + (1 - \lambda_d) \lambda_i (\gamma_i^{\sigma-1} - 1) + \lambda_n \gamma_n^{\sigma-1}]^{\frac{1}{1-\sigma}}. \quad (11)$$

Proof. Taking gross growth factors of both sides of (10) gives

$$\frac{P_{t+1}}{P_t} = \frac{W_{t+1}}{W_t} \left(\int_0^{N_t} q_t(j')^{\sigma-1} dj' \right)^{\frac{1}{\sigma-1}} \left(\int_0^{N_{t+1}} q_{t+1}(j)^{\sigma-1} dj \right)^{\frac{1}{1-\sigma}}. \quad (12)$$

Next, note that the term, $\int_0^{N_{t+1}} q_{t+1}(j)^{\sigma-1} dj$, can be written as

$$\int_0^{N_{t+1}} q_{t+1}(j)^{\sigma-1} dj = \int_0^{N_t} q_{t+1}(j)^{\sigma-1} dj + \int_{N_t}^{N_{t+1}} q_{t+1}(\iota)^{\sigma-1} d\iota. \quad (13)$$

Furthermore, with Assumption 1 and $\frac{N_{t+1}-N_t}{N_t} = \lambda_n$, we obtain

$$\int_{N_t}^{N_{t+1}} q_{t+1}(\iota)^{\sigma-1} d\iota = \lambda_n \gamma_n^{\sigma-1} \int_0^{N_t} q_t(j)^{\sigma-1} dj. \quad (14)$$

The first term on the right-hand side of (13), $\int_0^{N_t} q_{t+1}(j)^{\sigma-1} dj$, can be rewritten as

$$\int_0^{N_t} q_{t+1}(j)^{\sigma-1} dj = \gamma_d^{\sigma-1} \int_{\iota \in \mathcal{N}_{d,t}} q_t(\iota)^{\sigma-1} d\iota + \gamma_i^{\sigma-1} \int_{j' \in \mathcal{N}_{i,t}} q_t(j')^{\sigma-1} dj' + \int_{\iota' \in \tilde{\mathcal{N}}_t} q_t(\iota')^{\sigma-1} d\iota'. \quad (15)$$

where $\mathcal{N}_{d,t}$ and $\mathcal{N}_{i,t}$ is the set of products with a successful creative destruction or incumbent own innovation and $\tilde{\mathcal{N}}_t = [0, N_t] \setminus \{\mathcal{N}_{d,t} \cup \mathcal{N}_{i,t}\}$ is the set of surviving incumbents that do not improve the quality of their product between t and $t+1$. We also know that $|\mathcal{N}_{d,t}| = \lambda_d N_t$ and $|\mathcal{N}_{i,t}| = (1 - \lambda_d) \lambda_i N_t$. Then, because the arrival rate of an innovation is independent of $q_t(j)$ (and there is a continuum of varieties) the distribution of productivity of the varieties with and without innovation coincide and then by the law of large numbers we have:

$$\begin{aligned} \int_{\iota \in \mathcal{N}_{d,t}} q_t(\iota)^{\sigma-1} d\iota &= \lambda_d \int_0^{N_t} q_t(j)^{\sigma-1} dj, \\ \int_{j' \in \mathcal{N}_{i,t}} q_t(j')^{\sigma-1} dj' &= (1 - \lambda_d) \lambda_i \int_0^{N_t} q_t(j)^{\sigma-1} dj, \\ \int_{\iota' \in \tilde{\mathcal{N}}_t} q_t(\iota')^{\sigma-1} d\iota' &= (1 - \lambda_d - (1 - \lambda_d) \lambda_i) \int_0^{N_t} q_t(j)^{\sigma-1} dj. \end{aligned}$$

This in turn implies that (15) can be expressed as

$$\int_0^{N_t} q_{t+1}(j)^{\sigma-1} dj = [1 + \lambda_d (\gamma_d^{\sigma-1} - 1) + (1 - \lambda_d) \lambda_i (\gamma_i^{\sigma-1} - 1)] \int_0^{N_t} q_t(j)^{\sigma-1} dj. \quad (16)$$

Putting equations (12), (14), and (16) together establishes the proposition. ■

Proposition 2 shows how the arrival rates and step sizes of the different type of

innovation affect the inflation rate in the economy (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. Again, it can directly be seen from (11) that the inflation rate in the economy is monotonically decreasing in λ_i and γ_i . Finally, the term $\lambda_d (\gamma_d^{\sigma-1} - 1)$ captures the effect from *creative destruction* on the overall quality-adjusted inflation rate. For a given path of nominal wages, 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 . The economy wide inflation rate is decreasing in λ_d if

$$\gamma_d^{\sigma-1} - 1 > \lambda_i (\gamma_i^{\sigma-1} - 1), \tag{17}$$

which is a condition we will get back to further below. This ambiguity is due to the non-trivial interaction between the arrival rates of creative destruction and incumbent own innovation. An increase in the rate of creative destruction raises growth for instance if surviving products are expected to improve less than creatively destroyed products.

2.2.3 Imputation and measured inflation

Throughout our analysis we assume that the statistical office perfectly observes nominal values such as nominal output and wage growth. Hence, as highlighted in (7), 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.

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 some form of *imputation*.¹² Imputation uses the rate of quality-adjusted price growth for a set of surviving products that were not sub-

¹²Based on the 1999 Report of the General Accounting Office (GAO) of the BLS (which is itself based on data from 1997), the Appendix calculates that imputation was used 92% of the time in 1997 when a seller ceased producing a product in the CPI.

ject to creative destruction. This procedure is valid if the rate of quality change associated with creative destruction is the same as that for surviving products. But the vast majority of surviving products are not being improved in a given period, according to BLS estimates (Bils and Klenow, 2004; Nakamura and Steins-son, 2008). Instances of creative destruction are linked to innovative success, so they may associated with more rapid quality improvements than surviving prod-ucts taken as a whole.

In line with the imputation used in practice, let us formally characterize the procedure in the following way:¹³

Assumption 2 *In the presence of new products the statistical office resorts to imputation, i.e., the set of surviving products is assumed to be representative and the economy wide inflation rate is imputed from this subset of products.*

Surviving products are either products with incumbent own innovation or no innovation at all. We denote statistical office 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*

$$\left(\frac{\widehat{P}_{t+1}}{P_t}\right) = \frac{W_{t+1}}{W_t} \left[1 + \widehat{\lambda}_i (\widehat{\gamma}_i^{\sigma-1} - 1)\right]^{\frac{1}{1-\sigma}}. \quad (18)$$

Proof. Under Assumption 2 we have

$$\left(\frac{\widehat{P}_{t+1}}{P_t}\right) = \frac{W_{t+1}}{W_t} \left(\int_{\mathcal{N}_{I,t}} q_t(j')^{\sigma-1} dj'\right)^{\frac{1}{\sigma-1}} \left(\int_{\mathcal{N}_{I,t}} q_{t+1}(j)^{\sigma-1} dj\right)^{\frac{1}{1-\sigma}}, \quad (19)$$

where $\mathcal{N}_{I,t} = [0, N_t] \setminus \mathcal{N}_{d,t}$ is the set of products that survive between period t and $t + 1$. Note that a fraction λ_i of these surviving products experiences incumbent own innovation (and the quality improves by a factor of γ_i) whereas for the remaining fraction, $1 - \lambda_i$, quality remains unchanged. Hence, we have $\int_{\mathcal{N}_{I,t}} q_{t+1}(j)^{\sigma-1} dj = \left(\int_{\mathcal{N}_{I,t}} q_t(j')^{\sigma-1} dj'\right) [1 - \lambda_i + \lambda_i \gamma_i^{\sigma-1}]$. Using this equation in (19) and replacing γ_i and λ_i by their estimates yields (18). ■

¹³As we do, Erickson and Pakes (2011) identify imputation exit as source of upward bias in estimating the average price change for exiting products. These authors explain why 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 bias because they do not correct for time-varying unmeasured characteristics of the goods.

Henceforth we assume that the statistical office perfectly observes frequency and step size of incumbent own innovation, i.e., we have $\widehat{\lambda}_i = \lambda_i$ and $\widehat{\gamma}_i = \gamma_i$.¹⁴

2.2.4 Missing growth

Recall that missing real output growth between period t and $t + 1$ is given by the inverse of the bias in the estimated inflation rate (see (7)). Under the above assumptions about innovation processes and about the procedure of the statistical office, missing growth can be expressed as in the following proposition.

Proposition 4 *Missing real output growth is given by*

$$MG = \frac{1}{\sigma - 1} \log \left(1 + \frac{\lambda_d [\gamma_d^{\sigma-1} - 1 - \lambda_i (\gamma_i^{\sigma-1} - 1)] + \lambda_n \gamma_n^{\sigma-1}}{1 + \lambda_i (\gamma_i^{\sigma-1} - 1)} \right). \quad (20)$$

Proof. The expression for missing growth is directly obtained from combining (7), (11), and (18). ■

Since the statistical office imputes the inflation rate from surviving products (see Assumption 2) there are two sources of missing growth: (i) variety expansion and (ii) creative destruction innovation are imputed but not directly measured. The last term in equation (20) captures the growth mismeasurement from missing out on new variety creation. The first term in (20) measures the missing growth from not properly factoring in creative destruction. We see that the statistical office understates true output growth if the imputed growth from creative destruction (which is imputed from surviving products) understates the true expected growth from creative destruction. This can happen in two ways: (i) creative destruction has larger step size than incumbent own innovation ($\gamma_d > \gamma_i$) and (ii) not all of the surviving incumbents innovate ($\lambda_i < 1$). Hence, 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. We see that missing growth is monotonically increasing in γ_d and is also increasing in λ_d as long as $(\gamma_d^{\sigma-1} - 1) > \lambda_i (\gamma_i^{\sigma-1} - 1)$. This is the same condition that ensures that overall true growth is increasing in λ_d (see (17)). Missing growth is large if creative destruction is an important source of true growth.

¹⁴In equation (27), we show how our main results would be affected if the quality improvement of incumbents is not perfectly measured.

2.3 An illustrative example: the Cobb-Douglas case

Even though this may not be the most realistic case, we use the special case where the production technology for the final good is Cobb-Douglas to illustrate how creative destruction can lead to missing growth. Hence, let us consider the limit case where final output is produced according to the Cobb-Douglas technology

$$Y = N \exp \left[\frac{1}{N} \int_0^N \log [q(j)y(j)] dj \right]. \quad (21)$$

We assume the number of varieties N is fixed because there is no love of variety under Cobb-Douglas aggregation.

Aggregate price index Since the final sectoral output producer produces competitively we get as a demand for product $y(j)$ that is sold at price $p(j)$

$$y(j) = \frac{PY}{Np(j)},$$

where P is the price index defined as

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

Under the optimal price setting rule we get the following expression for the aggregate equilibrium price index

$$P = \mu W \exp \left(-\frac{1}{N} \int_0^N \log (q(j)) dj \right). \quad (22)$$

The true inflation rate can then be expressed as

$$\frac{P_{t+1}}{P_t} = \frac{W_{t+1}}{W_t} \gamma_i^{-(1-\lambda_d)\lambda_i} \gamma_d^{-\lambda_d}.$$

Measured inflation and missing growth Under Assumption 2 measured inflation becomes

$$\left(\frac{\widehat{P}_{t+1}}{P_t} \right) = \frac{W_{t+1}}{W_t} \gamma_i^{-\lambda_i}.$$

Consequently, we obtain for missing growth as:

$$MG = \lambda_d \cdot (\log \gamma_d - \lambda_i \log \gamma_i). \quad (23)$$

This missing growth from creative destruction can be decomposed as

$$\lambda_d (\log \gamma_d - \lambda_i \log \gamma_i) = \lambda_d (1 - \lambda_i) \log \gamma_i + \lambda_d (\log \gamma_d - \log \gamma_i).$$

The first term in this decomposition captures the fact that not all incumbents innovate, whereas the second term captures the step size differential between creative destruction and incumbent own innovation.

Numerical example Here we perform an illustrative calculation based on (23). We assume: (i) no variety expansion; (ii) the same step size for incumbent own innovation (OI) and for creative destruction (CD), i.e., $\gamma_i = \gamma_d = \gamma$, and (iii) annualized arrival rates λ_i and λ_d of OI and CD by new entrants that are both equal to 10% which is close to the exit rate of plants/firms. Finally, we assume that *measured* annual real output growth is equal to 1%, which implies $\lambda_i \log \gamma_i = 1\%$ so that $\log \gamma_i = 10\%$. Then, the annual rate of missing growth from creative destruction is equal to

$$MG = 10\% \cdot (1 - 10\%) \cdot 10\% = 0.9\%.$$

Although this is just an illustrative exercise, we will see in the next sections that this simple example is not far off from what we obtain using plant- and firm-level data on employment dynamics to determine the step sizes and frequencies of the various types of innovations.

3 Estimating missing growth

In this section we explore two alternative approaches for quantifying missing growth in the data. Both approaches build on the model developed in the previous section, although differently and using different data sets. The first approach uses information on the market shares of entrants, exiters and survivors: we refer to it as the *market share approach*. The second approach uses the algorithm in [Garcia-Macia et al. \(2016\)](#) to infer arrival rates and step sizes of different type of

innovations and compute missing growth: we refer to it 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 shares throughout the time period we consider of entrant establishments (plants), of survivor plants that stay in the market, and of exiters. This approach does not allow us to differentiate between the different sources of missing growth 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 group of products are potentially observable and do contain information about the quality-adjusted prices. The imputation used by the statistical office implies that the quality-adjusted price growth of the surviving products is taken as representative for the economy wide inflation rate. The CES framework of Section 2 suggests a simple test for the representativeness of the set of surviving products; the quality-adjusted price dynamics of the survivors is representative if and only if their market share remains stable over time. Furthermore, given an estimate for the elasticity of substitution, we show how the dynamics of market shares can be used to quantify missing growth. To construct the main argument more formally, let us define the market share of a product j as follows: $s(j) \equiv \frac{p(j)y(j)}{M}$, where $M = \int_0^N p(j)y(j)dj = PY$. Combining this definition with the demand (8) gives

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

Hence, the market share of product j is given by an power function of the quality-adjusted price of j relative to the aggregate price index, P .

The statistical office's imputation is based on surviving products between two period t and $t + 1$, i.e., on the set of products $\mathcal{N}_{I,t} = [0, N_t] \setminus \mathcal{N}_{d,t}$. In the following we call these surviving products between t and $t + 1$ *continuers*. In period t the aggregate market share of these *continuers* is given by

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

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

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

Under the price setting of firms, (3), and the specified innovational processes we can then express the growth rate of the market share of continuers in the next proposition.

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

$$\frac{S_{I_t,t+1}}{S_{I_t,t}} = \left(\frac{P_{t+1}}{P_t} \right)^{\sigma-1} \left(\frac{W_{t+1}}{W_t} \right)^{1-\sigma} (1 - \lambda_i + \lambda_i \gamma_i^{\sigma-1}). \quad (24)$$

Proof. Using the price setting behavior of the firms, (3), yields for the mar-

ket share growth $\frac{S_{I_t,t+1}}{S_{I_t,t}} = \left(\frac{W_{t+1}/P_{t+1}}{W_t/P_t} \right)^{1-\sigma} \frac{\int_{\mathcal{N}_{I,t}} q_{t+1}(j')^{\sigma-1} dj'}{\int_{\mathcal{N}_{I,t}} q_t(j)^{\sigma-1} dj}$. Now note that

a fraction λ_i of continuers experience incumbent own innovation whereas for the remaining fraction, $1 - \lambda_i$, quality remains unchanged. Hence, we have $\int_{\mathcal{N}_{I,t}} q_{t+1}(j')^{\sigma-1} dj' = \int_{\mathcal{N}_{I,t}} q_t(j)^{\sigma-1} dj [1 - \lambda_i + \lambda_i \gamma_i^{\sigma-1}]$, which establishes the proposition. ■

Since we have $\sigma > 1$, the market share of continuers decreases if and only if the price index of these products $\left(\int_{\mathcal{N}_{I,t+1}} (p_{t+1}(j')/q_{t+1}(j'))^{1-\sigma} dj' \right)^{\frac{1}{1-\sigma}}$ grows faster than the aggregate price index in the economy, P_t . Proposition 5 additionally exploits that the price index of continuers grows at the rate of the nominal wage growth $\frac{W_t}{W_{t+1}}$ times $(1 - \lambda_i + \lambda_i \gamma_i^{\sigma-1})^{1/(1-\sigma)}$, which captures the impact of incumbent own innovation on prices.

Important to note is that in view of Proposition 3 (and under the assumption of $\widehat{\gamma}_i = \gamma_i$ and $\widehat{\lambda}_i = \lambda_i$) the market share growth of continuers can be related to the measured inflation rate as follows:

$$\frac{S_{I_t,t+1}}{S_{I_t,t}} = \left(\frac{P_{t+1}}{P_t} \right)^{\sigma-1} \left(\frac{\widehat{P}_{t+1}}{P_t} \right)^{-(\sigma-1)}. \quad (25)$$

The intuition behind this equation is that, as argued before, the market share

growth of continuers is a power function of the gross growth factor of the price index of these surviving firms relative to the aggregate price index. The surviving products however constitutes precisely the set of products the 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 this subset of products: As (25) highlights, the quality-adjusted price dynamics of the continuers is representative if and only if their market share is stable over time. However, if the market share of the continuers decreases over time, measured inflation $\frac{\widehat{P}_{t+1}}{P_t}$ is too high and there is missing growth (note that we have $\sigma > 1$). In addition, for a given value of the elasticity of substitution, σ , this can be exploited to quantify missing growth. This is highlighted in the next proposition.

Proposition 6 *Missing growth can be calculated from the market share dynamics of continuers as*

$$MG_{t+1} = \log \left(\frac{\widehat{P}_{t+1}}{P_t} \right) - \log \left(\frac{P_{t+1}}{P_t} \right) = \frac{1}{\sigma - 1} \log \left(\frac{S_{I_t,t}}{S_{I_t,t+1}} \right). \quad (26)$$

Proof. Combining (7) and (25) directly proves the proposition. ■

Since $\sigma > 1$, this proposition highlights again that missing growth is positive whenever the market share of continuers shrinks over time. In the following we use equation (26) in Proposition 6 to quantify missing growth.¹⁵

3.1.2 Measuring the continuers' market share

Let B denote the first period of operation and D denote the last year of operation for a plant. Then let $L(t, B \leq b, D \geq d)$ denote the total employment or payroll in period t of plants who were born before or in period b and die in period d or

¹⁵Proposition 6 assumes quality improvement by incumbent's own innovation is correctly measured, i.e., $\hat{\gamma}_i = \gamma_i$ and $\hat{\lambda}_i = \lambda_i$. Without this assumption, missing growth in this model is given by

$$MG_{t+1} = \frac{1}{\sigma - 1} \left[\log \left(\frac{1 + \lambda_i(\gamma_i^{\sigma-1} - 1)}{1 + \hat{\lambda}_i(\hat{\gamma}_i^{\sigma-1} - 1)} \right) + \log \left(\frac{S_{I_t,t}}{S_{I_t,t+1}} \right) \right]. \quad (27)$$

Missing growth consists of quality measurement error of incumbent own innovation plus imputation error. Our market share approach captures the imputation error.

after. The ratio $\frac{S_{I_t,t}}{S_{I_t,t+1}}$ on the right-hand side of (26) is then measured as

$$\left(\frac{L(t, B \leq t, D \geq t+1)}{L(t, B \leq t, D \geq t+1) + L(t, B \leq t, D = t)} \right) / \left(\frac{L(t+1, B \leq t, D \geq t+1)}{L(t+1, B \leq t, D \geq t+1) + L(t+1, B = t+1, D \geq t+1)} \right). \quad (28)$$

This implies that 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 whereas Section 3.1.4 shows the results with payroll data as a robustness check.

Note that our approach here uses information on entering and exiting *plants* to measure growth from increased product variety. Hence implicitly our approach assumes that the number of products per plant is constant over the lifecycle of a plant. If anything, this assumption is likely to lead to an underestimation of missing growth since the number of products per plant is rather growing than shrinking over time. An underestimation of missing growth then results because our approach assumes that growth from adding new product lines *within a given establishment* is perfectly observed by the statistical office although it is of the creative destruction or variety expansion type. However, as we will explain below, our baseline specification will only use the information of market size of new plants after a 5-year lag. Hence the critical assumption is that the number of products per plant is constant after the age of 5 years, which seems to be a reasonable approximation of reality.¹⁶

Note also that our approach uses employment (or payroll) data to measure market shares. Unfortunately, plants' revenue data are only available for manufacturing firms through the Census of Manufactures (CMF), and only every five years. Hence these data only allows us to calculate the cumulative market share growth of surviving plants over the last five years. Yet for these survivors, market share by revenue shrunk faster than market share by employment, so that missing growth as defined by (26) is higher when market share is measured by revenue share.¹⁷

¹⁶Note in particular, that if surviving plants of age greater than five were on average increasing their number of products over time, the elasticity of plant exit with plant age would steeply decline with age, but in fact one can show that it does not.

¹⁷Note also that the CMF and LBD data yield similar figures for missing growth when using employment shares to measure plants' market shares, which suggests that the CMF-based estimates of missing growth are comparable.

How is the first period of operation, B , and the last period of operation, D , measured in the data? In the following we define a period t as a calendar year. A natural way is then to map B and D to the first and last year the plant appears in the dataset. This would implicitly assume that entry and exit in our data correspond to entry and exit in the market. However, in practice entering the LBD database does not necessarily mean fully entering the market. Some establishments may appear in the database even during the development phase of their products with substantially less employment than while production takes place. Hence, the mapping between the model and the data is likely to be more accurate if we use the market share of an establishment a few years after the firm has appeared in the database.

Furthermore, note that the market share used in our approach is supposed to reflect true quality-adjusted prices. However in reality plants may take time to accumulate customers and market share—even conditional on the price, quality, and variety of their products. Finally, as argued above, the number of product per plant is more likely to stay constant some years after its birth.

All of these reasons lead us to map B into a year in the dataset plus $k \geq 0$ years of lag. More formally, if B^d denotes the first year the plant appears in the database, we map B into $B^d + k$, where we use $k = 5$ in our baseline specification. [Haltiwanger et al. \(2013\)](#) 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. Hence, our baseline specification of a 5-year lag looks reasonable, though we check the robustness of this assumption in Section 3.1.4

It is important to note that, although we use the market share with a lag, this is done to obtain a measure of market share that is more tightly related to the quality-adjusted price of new plants. If a new plant produces and sells a new product but sales and employment are low in the beginning and reach its true potential only after some time, we use in our baseline specification the information of employment after 5 years to assess the quality of the product. By abstracting from plants that enter and exit within the 5 year window the applied lag makes our approach also 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 et al. \(2016\)](#) and the robustness of our results

with respect to this choice is documented in Section 3.1.4.

3.1.3 The market share of missing growth: results

It is our explicit goal to quantify missing growth in the *aggregate economy* and over a time horizon of several decades. For these reasons we 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 this employment/payroll information to infer the dynamics in $S_{I,t}$. Our data comes from LBD data for the period 1983–2013. LBD contains data on employment and payroll going back to 1976, although the payroll data is only clean from 1989. Also, although the LBD starts in 1976, we find missing growth estimates that use the 1976 to 1977 longitudinal link are 3 to 4 times larger than subsequent years. Our conjecture is that this is due to longitudinal linking issues in 1976 and 1977 years of the LBD (see Table 2 of Jarmin and Miranda, 2002). Hence we use the LBD employment data from 1977 onward. 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.

To match our LBD sample, we use BLS measured TFP growth for the nonfarm private business sector from 1983–2013.¹⁸ We put TFP in labor-augmenting form and make sure to not net out the BLS estimates of the contribution of R&D and intellectual property to TFP growth.

Table 1 shows the results for missing growth of our market share approach in annualized percentage points. The numbers in this table are based on our baseline specification with $\sigma = 4$ and $k = 5$. We find on average 0.56 percentage points of missing growth per year over the overall period 1983–2013. When we decompose the period into sub-periods we see that missing growth has not declined. In particular, compared to the previous sub-periods, missing growth is on average the highest during the period 2006–2013.¹⁹

Our calculation of missing growth does not use data on measured growth. To give a sense of the magnitude of our results, we compare our results to the BLS

¹⁸The BLS multifactor productivity series uses real output growth from the BEA. Most of BEA's prices comes from the BLS (see U.S. Bureau of Economic Analysis, 2014). So the measurement error we identify applies to the BEA series even if the BEA does not directly use the CPI to construct real output growth.

¹⁹See our [online appendix](#) for missing growth in manufacturing and non-manufacturing sectors.

Table 1: MARKET SHARE APPROACH

	Missing Growth in ppt. per year
1983–2013	0.56
1983–1995	0.60
1996–2005	0.41
2006–2013	0.69

Notes: This table presents the baseline missing growth estimates for the period 1983–2013 (and different sub-periods) using the market share approach in Proposition 6. The growth numbers are expressed in (average) percentage points per year. The market share is measured as the employment share of plants taken from the Census’ Longitudinal Business Database (LBD) as described in (28). These baseline results assume a lag $k = 5$ and an elasticity of substitution $\sigma = 4$.

TFP growth series which, as mentioned earlier, we express in labor-augmenting terms. Table 2 shows the implied true productivity growth figures compared with their measured productivity growth counterparts. True productivity growth is constructed by adding our missing growth to the BLS TFP series. Since measured TFP is lower in 2006–2013 whereas missing growth is the largest in this sub-period, missing growth expressed as a fraction of true growth is remarkably high in the last sub-period. Hence the increase in missing growth can partially explain the decrease in measured TFP growth since the mid 2000s. Measured growth declined from 2.65% per year in the 1996–2005 period to 0.90% per year in the post 2005 period. Of this 1.75 percentage points decline, we find however that only 0.28 percentage points is due to an increase in the overstatement of inflation.

3.1.4 The market share of missing growth: robustness and discussion

In this section we show how our main results of missing growth are affected by alternative assumptions about the lag of measured market shares, the elasticity of substitution, as well as using payroll instead of employment data.

Table 2: MEASURED VS. TRUE GROWTH WITH THE MARKET SHARE APPROACH

	Measured Growth	“True” Growth	Missing Growth
1983–2013	1.93	2.49	0.56
1983–1995	2.01	2.61	0.60
1996–2005	2.65	3.06	0.41
2006–2013	0.90	1.59	0.69

Notes: This table presents measured growth, true growth, missing growth and the corresponding share of total growth that is missed by the statistical office for the whole period 1983–2013 as well as for different sub-periods. Measured, true and missing growth are expressed in (average) percentage points per year. The missing growth is the same as in Table 1. 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.

Payroll as market size Table 3 shows resulting missing growth when payroll instead of employment shares are used to measure the relative market shares. The payroll data allows us to do this only since 1989 onward. The estimates show slightly larger but overall comparable numbers of missing growth.

Different lag k Our baseline results in Table 1 uses a lag $k = 5$. While we believe that this is a reasonable assumption it is nevertheless interesting to see how the results are affected by this choice. Table 4 shows the main results for missing growth for alternative lag specification of $k = 0$. For $k = 0$ missing growth decreases significantly (in particular in the period 2006–2013). However, though not reported in Table 4, for $k = 3$ we obtain estimates of missing growth of a similar order of magnitude to $k = 5$. Also, increasing k to 7 years increases missing growth slightly compared to $k = 5$.

Different elasticities of substitution Table 5 shows the main results under different elasticities of substitution. The estimate for missing growth monotonically decreases in σ but remains in a similar order of magnitude for small variations in the elasticity of substitution.

Table 3: DIFFERENT MEASURES OF MARKET SHARE

	Employment	Payroll
1989–2013	0.60	0.72
1989–1995	0.77	0.97
1996–2005	0.41	0.43
2006–2013	0.69	0.83

Notes: This table presents missing growth estimates for the period 1989–2013 as well as different sub-periods using employment or payroll data to measure the market share. The growth numbers are expressed in (average) percentage points per year. The estimates based on employment data are identical to the baseline results in Table 1 and $\sigma = 4$ and $k = 5$ is assumed throughout the table.

Declining dynamism and missing growth One may wonder why we get missing growth estimates that remain high, even though creative destruction as measured by the rates of entry, exit and job-reallocation are known to have declined continuously over the whole period (declining dynamism as documented by [Decker et al. 2014](#)). The answer is two-fold. First, we look at plants (establishments) not firms, with the assumption of one (or a constant number) of product(s) per plant. Second, our market share equation for missing growth corresponds to the *net* job creation by entry: indeed, the growth of survivors’ market share when market share is measured by employment, is captured by the *difference* between the job creation rate by new plants and the job destruction rate by exiting plants; and [Figure 1](#) indeed shows no trend in the net job creation rate of plants over the period 1977–2014. Finally, we look at market shares five years after the plant appears in our dataset.

Table 6 compares the missing growth estimates based on plants’ market shares with those based on firms’ market shares: we see that missing growth estimates drop dramatically when moving from plants to firms, and become negative as of 1996. This is not surprising: the former stems from the fact that new plants in existing firms are typically bigger than new plants in new firms, and we are not picking up this size difference when looking at firms instead of plants; the latter is

Table 4: DIFFERENT LAG LENGTHS k

	$k = 5$	$k = 0$
1983–2013	0.56	0.20
1983–1995	0.60	0.28
1996–2005	0.41	0.20
2006–2013	0.69	0.07

Notes: This table presents missing growth estimates for the whole 1983–2013 period (as well as different sub-periods) using different assumptions about the lag k . The growth numbers are expressed in (average) percentage points per year. The results with $k = 5$ are identical to the results in Table 1. The elasticity of substitution, σ , is taken equal to 4 throughout the table.

Figure 1: MISSING GROWTH AND DECLINING DYNAMISM




Table 5: DIFFERENT ELASTICITIES σ

	Missing Growth		
	$\sigma = 3$	$\sigma = 4$	$\sigma = 5$
1983–2013	0.84	0.56	0.42
1983–1995	0.90	0.60	0.45
1996–2005	0.62	0.41	0.31
2006–2013	1.04	0.69	0.52

Notes: This table presents missing growth estimates for the whole 1983–2013 period (as well as different sub-periods) using different assumptions about the elasticity of substitution, σ . The growth numbers are expressed in (average) percentage points per year. The results with $\sigma = 4$ are identical to the results in Table 1. The lag, k , is taken equal to 5 years throughout the table.

simply due to the fact that the net entry rate for firms has fallen more over time than the net entry rate for plants.

3.2 The indirect inference method

In this subsection we rely on the algorithm in [Garcia-Macia et al. \(2016\)](#), henceforth GHK, to estimate the arrival rates and step sizes of the various types of innovation. This affords another way to estimate missing growth. Key advantages of this indirect inference method are that: (i) we need not assume that creative destruction (CD) and new product varieties (NV) only come from new plants: 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 the possibility of products disappearing because of obsolescence.

GHK use indirect inference to estimate the step size and arrival rate of various types of innovation. They assume that own innovation (OI) and creative destruction (CD) have the same step size. They fit aggregate TFP growth and mean employment per firm exactly. They put equal weight on fitting other moments, in particular: the standard deviation of log employment across firms in

Table 6: MISSING GROWTH USING PLANTS VS. FIRMS

	Plant level	Firm level
1983–2013	0.56	0.08
1983–1995	0.60	0.29
1996–2005	0.41	-0.03
2006–2013	0.69	-0.14

Notes: The entries are missing growth in percentage points per year. The results at the plant level are identical to the results in Table 1. $\sigma = 4$ and $k = 5$ are assumed throughout.

the cross-section, minimum employment of a firm (one worker), the percent of employment at young firms (firms less than 5 years old), the overall job creation rate, the overall job destruction rate, the percent of firms with job creation less than 1 (which corresponds to firms who triple in size over a five year period), and the growth rate of the number of firms (which is equal to the growth rate of employment in the model). These data moments are averages calculated from the LBD for two time periods, 1976–1986 and 2003–2013. With their estimated parameters in hand, GHK decompose growth into respective contributions from new varieties, incumbent innovation on their own products, creative destruction by incumbents, and creative destruction by entering firms.

The original GHK algorithm assumes that measured growth equals true growth. It chooses parameters so that true growth, as given by equation (29) below, equals measured growth in the data. We modify the GHK algorithm to allow measured growth to differ from true growth. Specifically, we choose parameters such that measured growth, according to equation (30) below, matches the growth rate observed in the data.

The original GHK codes estimates 5-year arrival rates and step sizes. Since BLS substitutions and imputation happen at the monthly or bimonthly frequency (depending on the item), we run a version of the GHK model that estimates bimonthly arrival rates and step sizes. For a given set of parameter estimates, we check if measured growth in the model matches our empirical growth targets

(1.03% for 1976–1986 and 1.44% for 2003–2013). We iterate on parameter values until the measured growth from the model matches the data growth targets. The measured growth target is the same series as that used in the market share section.

In addition to the distinction between measured and true growth, our model in Section 2 differs from the GHK model in a few details. They keep track of firms with multiple products, so they estimate rates creative destruction and new variety creation separately for entrants and incumbents. They also allow for endogenous retirement of products due to an overhead cost denominated in labor: firms retire products whose quality relative to the average quality is below a cutoff. Rather than a fixed step size, quality innovations are drawn from a Pareto distribution. They assume the same Pareto shape parameter (and hence average step size) for quality innovations from own innovation as from creative destruction. New varieties are drawn from a scaled version of the existing quality distribution.

Using the notation from our model, in GHK true growth g is given by

$$(1 + g)^{\sigma-1} = (1 - \delta_o \psi) \{ [\lambda_i(1 - \lambda_{e,d} - \lambda_{i,d}) + (\lambda_{e,d} + \lambda_{i,d})] (\gamma_i^{\sigma-1} - 1) + 1 \} + (\lambda_{i,n} + \lambda_{e,n}) \gamma_n^{\sigma-1} \quad (29)$$

δ_o is the share of products in the previous period who quality falls below the cutoff for obsolescence, and ψ is the average quality of such below-cutoff products relative to the average quality.²⁰ λ_i is the share of products that are not obsolete and did not experience creative destruction who did experience an innovation by the incumbent producer. $\lambda_{e,d}$ is the share of non-obsolete products with entrant CD and $\lambda_{i,d}$ is the share of non-obsolete products with incumbent CD. $\lambda_{i,n} + \lambda_{e,n}$ is the mass of new varieties from incumbents and entrants relative to the the mass of products in the previous period. γ_i and γ_d are the average step size of own innovation and creative destruction, respectively. As in GHK, we assume the two step sizes are the same, which is why only γ_i appears in equation (29). γ_n is the average quality of a new variety relative to the average quality of varieties produced in the previous period. The term $1 - \delta\psi$ adjusts for endogenous loss of varieties due to obsolescence.

The equation for measured growth \hat{g} is the same in the modified GHK model

²⁰Formally, $\delta_o = \int_{q(j) < \bar{q}_t, q(j) \in \Omega_t} 1 dj$ and $\psi \delta_o = \frac{\int_{q < \bar{q}_t, q \in \Omega_t} q_t^{\sigma-1}(j) dj}{\int_{q \in \Omega_t} q_t^{\sigma-1}(j) dj}$ where Ω_t denotes the set of products produced in period t .

as in our Section 2 model, as before assuming the BLS accurately measures the arrival rate and average step size of own innovations:

$$(1 + \hat{g})^{\sigma-1} = 1 + \lambda_i(\gamma_i^{\sigma-1} - 1) \tag{30}$$

The top panel of Table 7 displays the parameter definitions and their estimated values for each of 1976–1986 and 2003–2013. The bottom panel gives rates of measured, true, and missing growth, respectively. Our missing growth estimates with the indirect inference approach here are similar to what we found using the market share approach: 0.52 percent (versus 0.46) over 1983–1986, and 0.42 percent (versus 0.68 percent) over 2003–2013.²¹

With the indirect inference approach here, we can also decompose missing growth into that from new varieties vs. creative destruction. We calculate missing growth due to creative destruction by taking the difference between measured growth and growth that results when we set the arrival rates for new varieties ($\lambda_{i,n} + \lambda_{e,n}$) to zero. We find that almost all of the missing growth is due to creative destruction: 0.41 percentage points from CD (vs. 0.52 total) over 1976–1986, and 0.33 percentage points from CD (vs. 0.42 total) over 2003–2013.

²¹We compare the market share results from 1983–1986 to GHK’s results from 1976–1986 because of data availability. Our LBD series starts in 1978 and we use plants that have been in the data for at least 5 years. Hence 1983 is the earliest year we can calculate missing growth using the market share approach.

Table 7: PARAMETERS AND MISSING GROWTH WITH INDIRECT INFERENCE

Parameter	Definition	1976-1986	2003-2013
δ_o	Share of products in t going obsolete in $t + 1$	0.000	0.000
$\lambda_{i,d}$	Probability of a product having incumbent CD, conditional on not becoming obsolete	0.003	0.002
$\lambda_{e,d}$	Probability of a product having entrant CD conditional on not becoming obsolete	0.007	0.004
λ_i	Probability of incumbent OI conditional on not becoming obsolete and creatively destroyed	0.024	0.027
$\lambda_{i,n} + \lambda_{e,n}$	Mass of new varieties relative to number of varieties	0.004	0.002
γ_i, γ_d	Step size of CD and OI innovation	1.014	1.017
γ_n	Step size of NV innovation	0.289	0.376
ψ	Step size of obsolescent product	0.079	0.056
σ	Elasticity of substitution	4	4
A. measured growth per year (data)		1.03%	1.44%
B. true growth per year (model)		1.55%	1.86%
C. missing growth per year (B-A)		0.52%	0.42%
D. growth with no new varieties (model)		1.44%	1.77%
E. missing growth due to CD per year (D-A)		0.41%	0.33%

Notes: Estimated value of parameters from running the algorithm from [Garcia-Macia et al. \(2016\)](#) on two sample: 1976-1986 and 2003-2013. The algorithm has been extended to fit our hypothesis as explained in the text.

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 and the use of imputation in calculating inflation rates. Our model generated explicit expressions for missing TFP growth as a function of the frequency and size of creative destruction vs. other types of innovation.

Based on the model, using U.S. Census data on all nonfarm businesses, we explored two alternative approaches to estimate the magnitude of missing growth from creative destruction for the U.S. from 1983 to 2013. The first approach used the employment shares of surviving, entering and exiting plants. The second approach applied the indirect inference of [Garcia-Macia et al. \(2016\)](#) 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. Moreover the second approach allows us to decompose missing growth into that from creative destruction vs. expanding variety.

We found that: (i) missing growth from imputation is substantial — about 0.5 percentage points per year on average when using both approaches; and (ii) it is mostly due to creative destruction. According to our estimates, missing growth has not declined over the past thirty years; since measured growth has declined recently, the fraction of true growth missed may have risen.

The missing growth we identify is likely to be exacerbated when there is error in measuring quality improvements by incumbents on their own products. That is, we think missing growth from imputation is over and above (and amplified by) the quality bias emphasized by the Boskin Commission.²²

Our analysis could be extended in several interesting directions. One would be to look at missing growth for particular sectors of the economy, such as goods vs. services. Another would be to look at missing growth in countries other than the U.S. A third extension would be to revisit optimal innovation policy. Based on [Atkeson and Burstein \(2015\)](#), the optimal subsidy to R&D may be higher

²²The BLS and BEA recently estimated quality bias from health and ICT alone to be about 0.4 percentage points per year from 2000–2015 ([Groschen et al., 2017](#)).

because true growth is higher, or lower because more growth comes from 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 inflation for disappearing products based on inflation only for those surviving products that have been innovated upon. This would assume that the step size is the same for creative destruction as for own innovation by incumbents, and that there is no quality bias in the BLS estimate of the latter.²³

Finally, we think it is useful to accurately measure growth, even if it only raises growth by the same percentage points over time. We illustrate this with four examples. First, it means that ideas are not getting as hard to find as official statistics suggest, with implications for the production of ideas and future growth (Gordon, 2012; Bloom et al., 2016). 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., 2016). Fourth, as stressed by the Boskin Commission, U.S. tax brackets and Social Security benefits are indexed to measured inflation, the flip side of measured growth.

²³Erickson 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.

Appendix on imputation in the CPI

One of the most difficult problems when compiling a price index is to accurately adjust for quality changes. Let v denote an item produced at date t and which is replaced by a new item $v + 1$ at date $t + 1$. To integrate the corresponding item change in the overall price index, the statistical office needs to infer a value for either price $P(v + 1, t)$ or price $P(v, t + 1)$ when it has information only about $P(v, t)$ and $P(v + 1, t + 1)$. According to the [U.S. General Accounting Office \(1999\)](#) and to the Handbook of Methods from [U.S. Bureau of Labor Statistics \(2015b\)](#), the BLS largely chooses among four possible courses of action to handle these *item substitutions*.²⁴

The first course of action simply involves setting

$$P(v + 1, t) = P(v, t).$$

This no-adjustment strategy is pursued by the BLS when it deems the new and old item as *comparable*, by which the BLS means that the old and new items are essentially the same, so that no quality difference exists between the two items.

The interesting case is when the BLS judges the new and old items to be *non-comparable*. Then the BLS typically chooses between three remaining strategies. First is *direct quality adjustment*. This is when the BLS can perform hedonic regressions or has information on manufacturers' production costs. Direct quality adjustment involves the BLS setting

$$P(v + 1, t) = P(v, t) \cdot QA(t).$$

Viewed through the lens of our model, BLS quality adjustments are an estimate of the step size of innovations.

For those noncomparable substitutions where the BLS lacks the information to make direct quality adjustments, it resorts to *class-mean imputation* or *linking*. Class-mean imputation is based on the rate of price changes experienced by other item substitutions — those which the BLS considers comparable or can directly adjust. Linking, meanwhile, uses the average rate of price change among items without substitution, items with comparable substitutions, and items with non-comparable substitutions subject to direct quality adjustments. Both imputations

²⁴We use italics to highlight terminology used by the BLS.

are usually carried out within the item's category or category-region.

The [U.S. General Accounting Office \(1999\)](#) usefully reports statistics on the frequency with which each method was used by the BLS for item substitutions in 1997. In that year, the BLS judged 58% of item substitutions to be comparable; the prices for these items entered the CPI without adjustment. The remaining 42% (the noncomparable substitutions) broke down as follows:²⁵

- 31.1% direct quality adjustments
- 33.4% class-mean imputations
- 35.5% linking.

To estimate the fraction of creative destruction innovations that were effectively subject to imputation based on all surviving items (those not creatively destroyed), we make the following three assumptions:

1. Comparable item substitutions do not involve any innovation.
2. Direct adjustments are implemented when incumbents improve their own products (OI).
3. Creative destruction (CD) results in imputation by class-mean or linking in the proportions stated above.

Under these assumptions, we estimate that creative destruction (CD) innovations were treated with the equivalent of all-surviving-items imputation 92% of the time in 1997. To see why, let D , C , and L denote the numbers of item substitutions subject to direct adjustment, class-mean imputation, and linking, respectively. Let N denote the number of comparable item substitutions.

The number of item substitutions for which some form of imputation is done is $L + C$. The imputation in the two strategies, however, is based on different sets of products. Whereas linking imputes from all surviving products (as in our theoretical model), class-mean imputation is based on other substitutions. We are looking for the fraction E of the products $L + C$ for which imputation is effectively based on all surviving products, as opposed to just those surviving products with

²⁵From conversations with BLS personnel, our understanding is that class-mean imputations and linking have been deployed with similar frequency in the years since 1997.

incumbent own innovations (fraction $1 - E$). These include all cases of linking plus a fraction (call it x) of class-mean imputations:

$$E = \frac{L + x \cdot C}{L + C}. \quad (31)$$

How do we determine x ? Class-mean imputations C use a weighted average for inflation from item substitutions for which there was either no adjustment (fraction $N/(D + N)$) or a direct adjustment (fraction $D/(D + N)$). Since 42% of all substitutions in 1997 were noncomparable (31.1% of which were direct adjustments) and 58% of all substitutions were comparable, we get:

$$\frac{D}{D + N} = \frac{0.311 \cdot 0.42}{0.311 \cdot 0.42 + 0.58} \approx 0.184.$$

Using the [Garcia-Macia et al. \(2016\)](#) algorithm we estimated the fraction of incumbent own-innovations (OI) among surviving products (those not creatively destroyed) to be $\lambda_i \approx 2.4\%$ at a bimonthly rate. If the fraction of direct quality adjustments in class-mean imputations was also 2.4%, we would say class-mean imputation is just like linking (imputation based on all products not creatively destroyed). Because the fraction of direct quality adjustments in class-mean imputations (at 18.4%) was higher than 2.4%, we infer that class-mean imputation puts extra weight on OI:

$$\frac{D}{D + N} = x \cdot \lambda_i + (1 - x) \cdot 1, \quad (32)$$

where x is the weight on all surviving items (only fraction λ_i of which were innovations) and $1 - x$ is the weight on those surviving products which did experience incumbent innovations. Rearranging (32) and using the above percentages we get

$$x = \frac{N/(D + N)}{1 - \lambda_i} \approx \frac{0.816}{1 - 0.024} \approx 83.6\%.$$

Thus, class-mean imputation effectively puts 83.6% weight on all surviving items and 16.4% weight on innovating survivors. Given that class-mean imputation and linking were about equally common, we estimate that the BLS used imputation based on all surviving items the equivalent of 92% of the time in 1997. More

exactly, we substitute the numerical values for x , L and C into (31) to get

$$E = \frac{L + x \cdot C}{L + C} \approx \frac{0.355 + 0.836 \cdot 0.334}{0.355 + 0.334} \approx 0.920.$$

References

- Abramovitz, Moses**, “Resource and Output Trends in the U.S. since 1870,” *American Economic Review*, 1956, *46* (2), 5–23.
- Acemoglu, Daron, Ufuk Akcigit, Nicholas Bloom, and William R. Kerr**, “Innovation, Reallocation and Growth,” NBER Working Papers 18993, National Bureau of Economic Research, Inc April 2013.
- Aghion, Philippe and Peter Howitt**, “A Model of Growth through Creative Destruction,” *Econometrica*, March 1992, *60* (2), 323–51.
- , **Ufuk Akcigit, and Peter Howitt**, “What Do We Learn From Schumpeterian Growth Theory?,” in “Handbook of Economic Growth,” Vol. 2 of *Handbook of Economic Growth*, Elsevier, December 2014, chapter 0, pp. 515–563.
- Akcigit, Ufuk and William R. Kerr**, “Growth Through Heterogeneous Innovations,” NBER Working Papers 16443, National Bureau of Economic Research, Inc November 2010.
- Atkeson, Andrew and Ariel Burstein**, “Aggregate implications of innovation policy,” 2015 Meeting Papers 640, Society for Economic Dynamics 2015.
- Bartelsman, Eric J. and Mark Doms**, “Understanding Productivity: Lessons from Longitudinal Microdata,” *Journal of Economic Literature*, September 2000, *38* (3), 569–594.
- Bils, Mark**, “Do Higher Prices for New Goods Reflect Quality Growth or Inflation?,” *The Quarterly Journal of Economics*, 2009, *124* (2), 637–675.
- and **Peter J. Klenow**, “Quantifying Quality Growth,” *American Economic Review*, September 2001, *91* (4), 1006–1030.
- and – , “Some Evidence on the Importance of Sticky Prices,” *Journal of Political Economy*, October 2004, *112* (5), 947–985.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb**, “Are Ideas Getting Harder to Find?,” Unpublished paper, Stanford University 2016.

Boskin, Michael, Robert J. Gordon Ellen Dullberger, Zvi Griliches, and Daleale Jorgenson, “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,” 1996.

Broda, Christian and David Weinstein, “Product Creation and Destruction: Evidence and Price Implications,” *American Economic Review*, 2010, *100* (3), 691–723.

Byrne, David M., John G. Fernald, and Marshall B. Reinsdorf, “Does the United States have a productivity slowdown or a measurement problem?,” Working Paper Series 2016-3, Federal Reserve Bank of San Francisco March 2016.

Byrne, David, Stephen Oliner, and Daniel Sichel, “How Fast are Semiconductor Prices Falling?,” NBER Working Papers 21074, National Bureau of Economic Research, Inc 2015.

Chetty, Raj, David Grusky, Maximilian Hell, Nathaniel Hendren, Robert Manduca, and Jimmy Narang, “The Fading American Dream: Trends in Absolute Income Mobility Since 1940,” NBER Working Papers 22910, National Bureau of Economic Research, Inc 2016.

Davis, Steven, John Haltiwanger, and Scott Schuh, *Job Creation and Destruction*, 1 ed., Vol. 1, The MIT Press, 1998.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, “The Role of Entrepreneurship in U.S. Job Creation and Economic Dynamism,” *Journal of Economic Perspectives*, Summer 2014, *28* (3), 3–24.

Diewert, Erwin, “The Challenge of Total Factor Productivity,” *International Productivity Monitor*, Fall 2000, *1*, 45–52.

Erickson, Tim and Ariel Pakes, “An Experimental Component Index for the CPI: From Annual Computer Data to Monthly Data on Other Goods,” *American Economic Review*, August 2011, *101* (5), 1707–38.

Feenstra, Robert C, “New Product Varieties and the Measurement of International Prices,” *American Economic Review*, March 1994, *84* (1), 157–77.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow, “How Destructive is Innovation?,” NBER Working Papers 22953, National Bureau of Economic Research, Inc December 2016.

Gordon, Robert J., “Is U.S. Economic Growth Over? Faltering Innovation Confronts the Six Headwinds,” NBER Working Papers 18315, National Bureau of Economic Research, Inc August 2012.

Griliches, Zvi, “The Discovery of the Residual: A Historical Note,” *Journal of Economic Literature*, 1996, *34* (3), 1324–1330.

Groshen, Erica L, Brian C Moyer, Ana M Aizcorbe, Ralph Bradley, and David M Friedman, “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*, 2017, *31* (2), 187–210.

Grossman, Gene M. and Elhanan Helpman, “Quality Ladders in the Theory of Growth,” *Review of Economic Studies*, 1991, *58* (1), 43–61.

Haltiwanger, John, “Job Creation, Job Destruction, and Productivity Growth: The Role of Young Businesses,” *Annual Review of Economics*, 08 2015, *7* (1), 341–358.

– , **Ron S. Jarmin, and Javier Miranda**, “Who Creates Jobs? Small versus Large versus Young,” *The Review of Economics and Statistics*, May 2013, *95* (2), 347–361.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein, “Quantifying the Sources of Firm Heterogeneity,” *The Quarterly Journal of Economics*, 2016, *131* (3), 1291–1364.

Hulten, Charles R., “Total Factor Productivity: A Short Biography,” NBER Working Papers 7471, National Bureau of Economic Research, Inc 2000.

Jarmin, Ron and Javier Miranda, “The Longitudinal Business Database,” Working Papers, U.S. Census Bureau, Center for Economic Studies 2002.

Jones, Charles I., “Time Series Tests of Endogenous Growth Models,” *The Quarterly Journal of Economics*, 1995, *110* (2), 495–525.

Jorgenson, Daleale and Zvi Griliches, “The Explanation of Productivity Change,” *Review of Economic Studies*, 1967, 34 (3), 249–283.

Klette, Tor Jakob and Samuel S. Kortum, “Innovating Firms and Aggregate Innovation,” *Journal of Political Economy*, October 2004, 112 (5), 986–1018.

Kortum, Samuel S., “Research, Patenting, and Technological Change,” *Econometrica*, November 1997, 65 (6), 1389–1420.

Lipsey, Richard and Kenneth Carlaw, “Total factor productivity and the measurement of technological change,” *Canadian Journal of Economics*, November 2004, 37 (4), 1118–1150.

Lipsey, Richard G. and Kenneth Carlaw, “What Does Total Factor Productivity Measure?,” *International Productivity Monitor*, Fall 2000, 1, 31–40.

Moulton, Brent R and Karin E Moses, “Addressing the Quality Change Issue in the Consumer Price Index,” *Brookings Papers on Economic Activity*, 1997, 1997 (1), 305–366.

Nakamura, Emi and Jon Steinsson, “Five Facts about Prices: A Reevaluation of Menu Cost Models,” *The Quarterly Journal of Economics*, 2008, 123 (4), 1415–1464.

Redding, Stephen J. and David E. Weinstein, “A Unified Approach to Estimating Demand and Welfare,” NBER Working Papers 22479, National Bureau of Economic Research, Inc August 2016.

Romer, Paul M., “Endogenous Technological Change,” *Journal of Political Economy*, 1990, 98 (5).

Syverson, Chad, “Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown,” NBER Working Papers 21974, National Bureau of Economic Research 2016.

U.S. Bureau of Economic Analysis, “Principal Source of Data,” in “Concepts and Methods of the U.S. National Income and Product Accounts,” Washington DC: U.S. Department of Commerce, 2014, chapter 3.

U.S. Bureau of Labor Statistics, “Producer Prices,” in “Handbook of Methods,” Washington DC: U.S. Department of Labor, 2015, chapter 14.

– , “The Consumer Price Index,” in “Handbook of Methods,” Washington DC: U.S. Department of Labor, 2015, chapter 17.

U.S. General Accounting Office, “Consumer Price Index: Impact of Commodity Analysts’ Decisionmaking Needs to be Assessed,” GAO/GDD-99-84 1999.