NBER WORKING PAPER SERIES

OLDER AND SLOWER:
THE STARTUP DEFICIT'S LASTING EFFECTS ON AGGREGATE PRODUCTIVITY GROWTH

Titan Alon
David Berger
Robert Dent
Benjamin Pugsley

Working Paper 23875
http://www.nber.org/papers/w23875

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2017

Prepared for the Carnegie-Rochester-NYU Conference, NYU, April 2017. We are grateful for the very helpful comments from an anonymous referee as well from Steve Davis, Ryan Decker, John Haltiwanger, Greg Kaplan, Pete Klenow, Geoffrey Tate, and participants in the NBER Entrepreneurship and Economic Growth conferences, Northwestern University Macro Lunch, Carnegie-Rochester-NYU Conference on Public Policy, 2017 SED meetings, and the NBER Macro Productivity group. We also thank Javier Miranda and Ryan Decker for their generous assistance with the Business Register, Jim Spletzer for his documentation, Sara Moreira for her work merging the revenue data to the LBD and Harry Wheeler for his excellent research assistance. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau or the National Bureau of Economic Research. All results have been reviewed to ensure that no confidential information is disclosed.

Older and Slower: The Startup Deficit’s Lasting Effects on Aggregate Productivity Growth
Titan Alon, David Berger, Robert Dent, and Benjamin Pugsley
NBER Working Paper No. 23875
September 2017
JEL No. E01,E24

ABSTRACT

We investigate the link between declining firm entry, aging incumbent firms and sluggish U.S. productivity growth. We provide a dynamic decomposition framework to characterize the contributions to industry productivity growth across the firm age distribution and apply this framework to the newly developed Revenue-enhanced Longitudinal Business Database (ReLBD). Overall, several key findings emerge:(i) the relationship between firm age and productivity growth is downward sloping and convex; (ii) the magnitudes are substantial and significant but fade quickly, with nearly 2/3 of the effect disappearing after five years and nearly the entire effect disappearing after ten; (iii) the higher productivity growth of young firms is driven nearly exclusively by the forces of selection and reallocation. Our results suggest a cumulative drag on aggregate productivity of 3.1% since 1980. Using an instrumental variables strategy we find a consistent pattern across states/MSAs in the U.S. The patterns are broadly consistent with a standard model of firm dynamics with monopolistic competition.

Titan Alon
Northwestern University
Department of Economics
2001 Sheridan Road
Evanston, IL 60208
titanalon@u.northwestern.edu

David Berger
Department of Economics
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
and NBER
david.berger@northwestern.edu

Robert Dent
Nomura Securities
robcdent@gmail.com

Benjamin Pugsley
University of Notre Dame
Department of Economics
Notre Dame, IN 46556
bpugsley@nd.edu

1 Introduction

Over the last three decades, the U.S. business sector has experienced a collapse in the rate of new startups alongside an enormous reallocation of economic activity from entrants and young firms to older incumbents. Figures 1a and 1b reproduced from Pugsley and Sahin (2014) illustrate the trends.[1] The magnitude of the reallocation across firm age groups is quite startling and even surpasses those flows documented in the structural transformation literature from manufacturing to services over the same period.[2] These patterns are also widespread across industries and geographic markets, suggesting they are independent of any compositional variation in economic activity.[3] In what follows, we refer to the persistent and widespread collapse in startup rates and the subsequent aging of U.S. businesses as the startup deficit.

Recently, economists and policy makers have begun questioning whether the startup deficit may impact the health of the aggregate economy through a variety of channels. For example, growth theory often associates new firms with the introduction of new innovations and new products; in trade, these firms are typically responsible for opening new markets; in industrial organization, entrants play a critical role in maintaining market competition; and in the firm dynamics literature, new and young firms typically drive productivity gains from selection and reallocation. Work by Evans (1987) and Dunne, Roberts, and Samuelson (1989) in the manufacturing sector, and more recently Haltiwanger, Jarmin, and Miranda (2013) for the entire nonfarm business sector, highlight the key role of firm age, above and beyond firm size for firm dynamics.

Building on this earlier empirical work, a new vein of research has emerged attempting to better understand the economic significance of firm age and, more broadly, the life-cycle of the firm.[4] Together, these lines of research suggest that recent slowdowns in the rate of business creation and the shift in economic activity toward older incumbents could exercise significant drags on aggregate growth and employment dynamics. In this paper, we aim to build on these recent findings by examining explicitly the systematic relationship between firm age and labor productivity growth that is bound together with the traditional measures of reallocation. We also stress that while the focus of our analysis is the effects of the compositional change of firm age, recent work by Decker, Haltiwanger, Jarmin, and Miranda (2014) and Decker, Haltiwanger, Jarmin, and Miranda (2017) have highlighted that the majority of changes in employment and productivity dynamics occur *within* firm age groups. We view our analysis as complementary to theirs.

We conduct our investigation using U.S. Census Bureau data encompassing the entire nonfarm business sector from 1996-2012. As our aim is mainly descriptive, we apply a methodology that

[1] Decker, Haltiwanger, Jarmin, and Miranda (2014) and Hathaway and Litan (2014) also document the same change in age composition. Reedy and Strom (2012) are the first, to our knowledge, to document a decline in the aggregate startup rate that long pre-dates the Great Recession.

[2] See Dent, Karahan, Pugsley, and Şahin (2016). Over the 1987-2012 period, the mature employment share increases by roughly 17 percentage points. For comparison, over the same period the manufacturing employment share declines by 11 percentage points and the services employment share increases by 14 percentage points.

[3] See, again, Decker, Haltiwanger, Jarmin, and Miranda (2014), Hathaway and Litan (2014) and Pugsley and Sahin (2014) for details.

[4] Arkolakis, Papageorgiou, and Timoshenko (2014) and Pugsley and Sahin (2014) are recent examples.

remains agnostic about the underlying mechanism at work. Instead, we exploit the rich industry and geographic variation in the Census data to nonparametrically identify any common, underlying links between productivity growth and firm age. For shorthand, we refer to this relationship as the *age-productivity profile*. We find a robust and mostly stable relationship between our measure of productivity growth and firm age. We submit our main results to a large battery of robustness checks controlling for price effects, organizational structure of firms, industrial and geographic composition, and the pattern is little changed.

Given a robust set of estimates, we then use our results to assess the impact of the startup deficit has had on aggregate productivity. Our first approach applies the age-productivity profile results directly and shows how, under some empirically plausible assumptions, the age-productivity profiles we estimate can be linked directly to average labor productivity growth. While this exercise allows us to transparently quantify the significance of the shape of our estimated age-productivity profiles in light of the decline in firm entry, it does not admit any causal interpretation. We therefore complement these results with a series of cross-sectional regressions, which exploit plausibly exogenous variation in startup activity across geographic and industry markets. Using differences in startup rates across markets generated either by lagged demographics or shocks to the value of collateral we find markets with higher entry rates show significantly faster productivity growth.

Our decomposition-based results show that the age composition matters. The estimation procedure establishes a statistically significant and robust empirical link between the distribution of firm age and productivity growth which is independent of pricing, compositional, organizational, or cyclical variation. Our age-productivity profiles suggest that the relationship between firm age and productivity growth is downward sloping and convex, mirroring similar patterns uncovered in other work between firm age and employment growth. The differential in growth rates are substantial but converge quickly; while the youngest firms grow very quickly relative to older incumbents, nearly two-thirds of the effect is gone after five years and the total effect is nearly gone after ten. The pattern seems to be remarkably stable across a wide variety of robustness checks for all uncensored age groups. This stability does not seem to be shared, however, by the oldest age censored firms for whom we have more limited information. For these firms we document marked declines in labor productivity growth over our sample, driven primarily by falls in what we will label allocative efficiency. Quantitatively, our results suggest from 1980 to 2014–even holding constant the productivity level of the oldest censored firms– the startup deficit and subsequent aging of U.S. business sector has reduced aggregate productivity by a little more than 4 percent, or roughly 0.12 percentage points per year. Declines in allocative efficiency among the oldest firms, may have further contributed up to an additional 2 percentage points.

In order to better understand which economic mechanisms drive the age-productivity profile, we adapt the Dynamic Olley-Pakes (DOP) decomposition of Melitz and Polanec (2015) and apply it to the age-productivity profile[5]. Using changes in the cross-sectional distribution of productivity, the

[5]Note that here we are decomposing productivity growth over the life-cycle of a cohort of firms. This in contrast to the normal manner in which the DOP methodology is employed to decompose the productivity growth of an industry or economy.

DOP framework decomposes the change in average productivity into into (1) a contribution from exit of less productive firms, (2) a widespread shift in the average productivity of incumbents, and (3) the reallocation of market share to more productive incumbents, where the last two encompass the traditional Olley-Pakes decomposition. We find that the exit and reallocation channels are primarily responsible for the shape of the age-productivity profile. In appendix C, we demonstrate that a standard model of firm dynamics with monopolistic competition is broadly consistent with these patterns. Interestingly, among cohorts of firms founded since the late 1970s, the relative contributions of each component of productivity growth have been stable. This stability includes even the contribution from the reallocation channel, which contrasts with other recent work such as Hyatt and Spletzer (2013) and Decker, Haltiwanger, Jarmin, and Miranda (2016b) that show declines in measures of job and worker reallocation across all age groups and young firms. Among the oldest firms (formed before the late 1970s), however, we too observe declines in the contribution of changes in reallocation to productivity growth–consistent with these earlier findings in the literature.

Complimenting the counterfactual analysis, our cross-sectional regressions also uncover significant drags on productivity growth from declining entry. Across detailed geographic and industrial market we find that a one percentage point decline in the entry rate implies a roughly 1 to 2 percentage point drag on local annual productivity growth. The sign and approximate magnitude of the effect are preserved even after we employ two distinct instrumental variable approaches exploiting cross-sectional variation in demographic trends and new business financing conditions. Importantly, our cross sectional analysis admits both direct and indirect effects of the declines in entry, with the direct effect following from the compositional change and the indirect effect following from further within age group changes.

Our paper contributes to multiple literatures. First, several influential papers have studied the dynamics of productivity using similar data. Foster, Haltiwanger, and Krizan (2001) is the first paper, to our knowledge, to use administrative data on receipts to study the dynamics of productivity outside of the manufacturing sector. They show that receipts per worker is an informative proxy for labor productivity and provide comparisons with the standard approach of constructing productivity measures from survey data on inputs and outputs. Bartelsman, Haltiwanger, and Scarpetta (2013) show the importance of the reallocation channel that we capture in our dynamic Olley-Pakes decomposition, and provide a more structural interpretation in a model that supports within industry dispersion in labor productivity.

Second, we contribute to the large and growing literature on the decline of business dynamism. Decker, Haltiwanger, Jarmin, and Miranda (2016a) and Decker, Haltiwanger, Jarmin, and Miranda (2017) (using the same Revenue-enhanced Longitudinal Business Database developed by Haltiwanger, Jarmin, Kulick, and Miranda (2016a)) have identified a decline in reallocation as key factor in the sluggish growth of labor productivity over the last decade.

Of course, the aggregate productivity implications of shifts in the age distribution of domestic businesses represents only a fraction of the broad structural changes recently underway in the economy. The secular increases in market concentration and declines in job flows accompanying the start-up deficit recently documented in the literature leave ample room for aggregate effects above and beyond those we uncover here. For example Autor, Dorn, Katz, Patterson, and Van Reenen (2017) document increasing concentration of domestic sales, and Decker, Haltiwanger, Jarmin, and Miranda (2017) find that firm's job creation responses to profitability shocks has declined over time. Furthermore, while our work focuses on the aggregate implications of between age-group variation in labor productivity growth, the literature has also highlighted the importance of within age-group variation as an important determinant of recent declines in business dynamism[6]. The shifts in aggregate productivity reflect both compositional and within age group changes. Although our focus is primarily on the direct effects of the changes in the age distribution, we view it as complementary to efforts investigating within age group changes. Moreover, our cross sectional evidence, linking plausibly exogenous declines in the entry rate to declines in labor productivity growth, is compatible with both compositional and within-group channels.

The structure of the rest of the paper is as follows: Section 2 reviews our estimation methodology and discusses the construction of our main dataset, as well as our decomposition framework linking aggregate productivity growth to our profile estimates. Section 3 presents the empirical results and quantifies the macroeconomic significance of our findings. Section 4 presents the battery of robustness tests we conduct on our final estimates. Section 5 presents the cross-sectional evidence and our IV approach. Section 6 concludes the paper.

2 Data and Methodology

To begin our analysis we aim to estimate the empirical relationship between firm age and labor productivity growth. Our starting point is an identifying conjecture that their exists some stationary relationship between firm age and productivity growth which is common across industries up to scale. In this section, we estimate this relationship non-parametrically and evaluate its stability using rich Census data covering the entire nonfarm business sector.

2.1 Data

We use firm-level measures of labor revenue productivity encompassing the entire U.S. nonfarm business sector from 1996-2012. Our main data source is the Census Bureau's Revenue-enhanced Longitudinal Business Database (ReLBD), first constructed by Haltiwanger, Jarmin, Kulick, and Miranda (2016b). The ReLBD merges the Census Bureau's Longitudinal Business Database (LBD) with corresponding administrative records in the Census Bureau's Business Register (BR) containing revenues reported to the IRS from business tax filings.

The LBD provides high quality measures of employment, location, and industry with nearly universal coverage of the nonfarm business sector which are carefully linked over time at the es-

[6]For example, Decker, Haltiwanger, Jarmin, and Miranda (2016b) find that as much as 70-75 percent of recent declines in the job reallocation rate can be explained by within age-group variation over time

tablishment level. These longitudinal records may be used to calculate measures of employment growth, entry, exit and establishment age. Through data gathered in the Cenus Bureau annual Company Organization Survey and quinquennial Economic Census, the LBD also provides a firm identifier for each year that groups establishments at the highest level of operational control.[7] We assign an establishment age 0 in the year it hires its first employee. Following the approach of Davis, Haltiwanger, Jarmin, and Miranda (2006) and more recently Haltiwanger, Jarmin, and Miranda (2013), we then assign each firm that is comprised of more than one establishment the age of its oldest establishment The advantage of this approach is that age 0 firms are de novo firms, composed entirely of new establishments.[8] We then merge revenue records from the Census Bureau BR following the methodology of Haltiwanger, Jarmin, Kulick, and Miranda (2016b).[9] These records are merged at the level of the tax reporting unit (EIN) and aggregated to the firm level. Unfortunately, matching revenue records are not available for all firms, and the matching is non random. Revenue records may be incomplete for very large and very small firms. Following Haltiwanger, Jarmin, Kulick, and Miranda (2016b), after applying a set of filters to remove outliers, we estimate propensity scores using a set of observable firm characteristics and re-weight observations by the inverse of the predicted propensity scores to adjust for the non random matching. We then merge, where available, price indices at the NAICS 4-digit level from the Bureau of Labor Statistics Labor Productivity and Costs database. Finally, to construct real revenues, we deflate nominal revenues by the GDP implicit price deflator from the BEA.[10] Further details on the construction of the dataset and cleaning process are provided in appendix B. With the merged and cleaned dataset we construct firm-level measures of real revenues per employee, which may later be aggregated by firm age, industry and location.

2.2 Age-productivity profiles

Using our firm-level measure of real revenue R_{it} per worker E_{it}, we define firm i log labor productivity $\phi_{it} = \log(R_{it}/E_{it})$ and aggregate (industry j) log labor productivity as

$$\Phi_{jt} \equiv \sum_i s_{it}\phi_{it},$$

where s_{it} is the market share for firm i measured here as employment share. Throughout, labor productivity always refers to its log level, Φ. To describe the age-productivity profile, we measure

[7] See Jarmin and Miranda (2002) for additional details on the LBD and its construction.

[8] There are, however, alternative measures of firm age that may also be relevant for firm productivity. For example, we might want to classify a firm by the tenure of its ownership and management, since if a firm is acquired by a private equity company that replaces its management team, its dynamics may more closely resemble a startup than a well established company. Considering alternative age concepts is an exciting area for future research. We thank Steve Davis for this insight.

[9] The details of the merge are nontrivial and described in an internal Census Bureau technical memo. We thank Javier Miranda and Jim Spletzer for assisting us with this merge.

[10] Our baseline specification also includes industry and time fixed effects to capture some of the variation that occurs across industries and time. In the robustness section, we explicitly deflate our series using the most comprehensive set of price indices that are available.

productivity by age group. We also adapt the Dynamic Olley-Pakes (DOP) decomposition of Melitz and Polanec (2015) to measure the changing sources of productivity growth over firms' life-cycles.[11] Instead of applying the decomposition directly to the aggregate economy (or an industry) as those authors do, however, we use the methodology to study how the productivity of a cohort of firms evolves as they age. Specifically, We let $\Phi_{ajt} = \sum_{i \in a}(s_{it}/s_{at})\phi_{it}$ be the productivity of a *cohort* of age a firms in period t in industry j. The productivity of this group of firms in period $t-1$ can be expressed as the employment share weighted average productivity of those firms that will survive (s) to period t and those firms that will exit (x) before then such that

$$\begin{aligned}\Phi_{a,t-1} &= s_{sa,t-1}\Phi_{sa,t-1} + s_{xa,t-1}\Phi_{xa,t-1} \\ &= \Phi_{sa,t-1} + s_{xa,t-1}(\Phi_{xa,t-1} - \Phi_{sa,t-1})\end{aligned}$$

where s_{iat} for $i \in \{s, x\}$ represents the share of firms i *within* cohort a (in industry j) at time t. For simplicity, we suppress the industry subscript j when obvious. Noting then that the cohort's productivity in period t will be constituted by survivors alone we can express the productivity growth rate of the cohort as

$$\Delta\Phi_{a,t} = \Delta\Phi_{sa,t} - s_{xa,t-1}(\Phi_{xa,t-1} - \Phi_{sa,t-1})$$

which captures growth from survivors and a contribution from selective exit. Applying the Olley-Pakes decomposition to the first component yields

$$\Delta_{at}\Phi_{a,t} = \underbrace{\Delta\bar{\phi}_{sa,t}}_{\text{Within}} + \underbrace{\Delta\widetilde{\text{Cov}}_{sa}(s_{it}, \varphi_{it})}_{\text{Allocative Efficiency}} - \underbrace{s_{xa,t-1}(\Phi_{xa,t-1} - \Phi_{sa,t-1})}_{\text{Selection}} \tag{1}$$

which decomposes the productivity growth of a cohort into its firm dynamic components. The component $\Delta\bar{\phi}_{sa,t}$ is the change in the (unweighted) mean productivity across surviving firms and captures any broad based changes in productivity *within* firms as they age, such as those emerging from learning or process innovations. The component $\Delta\hat{Cov}_{sa}(s_{it}, \varphi_{it})$ measures the change in covariances between a firm's market share (s_{it}) and its productivity (φ_{it}) which captures the *allocative efficiency* of the cohort insofar as it increases as higher productivity firms in the cohort capture larger fractions of the market share.[12] The Olley-Pakes distinction between within and allocation attributes shocks to all firms as within, and shocks either to the market shares or the productivity of specific firms to allocation. If within firm changes are concentrated rather than broad based, they will also appear in this covariance term, i.e., the covariance increases if highly productive firms gain market share *or* if highly productive become even more productive. The final term represents the contribution of *selection* and contributes positively to cohort productivity

[11] See also Diewert and Fox (2005) for an alternative dynamic decomposition that accounts for entry and exit.

[12] We use the $\tilde{\cdot}$ notation to denote that the term is technically a quasi-covariance term between market shares and productivities since, as in Melitz and Polanec's decomposition, the $1/N$ term is embeded in the shares. Stated more precisely, the term is the inner product of the deviations from the cohort mean of firm market shares and firm productivity.

growth provided exiting firms are on average less productive than surviving incumbents within the same cohort.

2.3 Estimation

To estimate the age-productivity growth profile we pool annual cross-section of productivity growth by firm age and 4-digit NAICS industries measured in the Census data. Our main identifying assumption is that we can express an age group's productivity growth as

$$\Delta_{at}\Phi_{ajt} = \nu_j + \mu_t + \sum_{l=1}^{a} \delta_l + \varepsilon_{ajt} \qquad a = 1 \ldots, A, \tag{2}$$

where $\mathbb{E}[\varepsilon_{ajt}|a, j, t] = 0$. We also apply the same decomposition of average productivity growth into time, industry and age effects to each component of equation (1). This specification imposes some strong assumptions: first, industry, time, and age effects are additively separable; second, there are no cohort effects in the *growth* rate of labor productivity, which is consistent with recent work by Moreira (2016) who finds persistent cohort effects in the level but not the growth rate of productivity. Because our specification is in differences, it already removes any *level* fixed effects across industries or time. The inclusion of fixed effects in the differenced specification then allows us to also control for differences in time and industry *trends* in our estimation procedure.

Additionally, we are assuming that revenue labor productivity, $\frac{P_i Y_i}{L_i}$, which is what we can measure since we lack firm specific price information, is informative about physical labor productivity, $\frac{Y_i}{L_i}$. Previous empirical evidence suggests that this assumption holds in the data (Bartelsman, Haltiwanger, and Scarpetta, 2013).[13]

Given the identifying assumptions implied by our specification and discussed above, we are able to semi-parametrically estimate the age-group profiles by projecting age group productivity growth, $\Delta\Phi_{ajt}$, and its components in equation (1), on a full set of industry, time, and age group fixed effects. Pooling samples from 1996 to 2012, we estimate equation (2) by OLS and WLS where we weight by an industry average employment share, in order to hold industry composition constant.

The full set of estimated age group coefficients $\hat{\delta}_a$ provide semi-parametric estimates of average productivity growth across firm ages, which we refer to as age-productivity profiles. In our analysis of aggregate productivity growth, our foremost focus is the role of firm age on average productivity growth. Recent works, however, by Haltiwanger, Jarmin, Kulick, and Miranda (2016b) and Decker, Haltiwanger, Jarmin, and Miranda (2016a) have also shed light on how the dispersion in both productivity and employment growth rates across firms has changed over time, even within age groups. Although it would not alter our results, recent evidence suggests the stability we find in the average profile does not extend to higher order moments.

[13] This paper also points out that most models do not imply that there is any firm-level dispersion in revenue labor productivity because optimizing firms will set the marginal revenue product to a constant (across firms marginal cost. The paper also shows that is possible to induce a positive correlation in these labor productivity measures as well as dispersion by adding overhead labor, adjustment costs or non-CES demand.

The upper-case A represents the fact that, due to data limitations, our ability to observe productivity growth by age is right-censored. However, since we are estimating a profile that is stationary across time, we are able to ameliorate the censoring by conducting estimation on a triangular panel of firm ages that grows as we gain more year observations that reveal the behavior of older age groups in later years. This approach allows us to estimate the profile for firms through age 30, rather than being forced to curtail estimate at age 15 had we worked with a balanced panel. As a robustness check, we verify that the triangular panel approach does not significantly impact the precision of early age estimates that we would have gotten with the balanced panel alone.

The differenced specification above will not be suitable for new entrants. While this is not crucial for understanding the dynamics of firm age and productivity growth, identifying trends in the productivity of entering cohorts will be crucial in linking our results to aggregate productivity in section 2.4. For this purpose, we run an auxiliary regression on the sub-sample of new entrants, controlling as best as possible for well known issues with industry heterogeneity through a specification in differences with industry fixed effects. Specifically, we estimate:

$$\Phi_{E,it} = \eta t + \nu_i + \epsilon_{E,it} \tag{3}$$

where we interpret η as a common deterministic trend in the productivity of new entering cohorts. In appendix figure A2 we explicitly estimate cohort effects for non-censored age groups and verify that the assumption of a linear trend is representative of long-term trajectory patterns in the data.

2.4 Aggregating Firm-Level Findings

Given a robust set of estimates of the relationship between firm age and productivity growth, this section establishes a framework with which to interpret the results and link them to aggregate productivity. As our ultimate aim is to quantify the aggregate productivity implications of the startup deficit and subsequent aging, we need to establish a framework linking the latter to the former. To isolate the aging effects, we derive an aggregation in the absence of any aggregate time or industry composition contributions. In this case, we can use our estimates in equations (2) and (3) which indicate that conditional on no time or industry effects we have

$$\mathbb{E}\left[\Delta\Phi_{ait}|\mu = 0, \nu = 0\right] = \delta_a \qquad \mathbb{E}\left[\Delta\Phi_{E,it}|\mu = 0, \nu = 0\right] = \eta$$

This in turn allows us to rewrite a cohort's productivity in any given time period (dropping the conditional expectation for brevity) as:

$$\Phi_{a,t} = \sum_{l=1}^{a} \delta_l + \Phi_{E,t-a} \qquad a = 1, \dots, A, \tag{4}$$

where $\Phi_{E,t-a}$ is the initial cohort productivity of firms created in period $t - a$ (corresponding to the cohort of age a firms in period t). The expression clarifies that once we condition away aggregate and industry effects, the difference between firms of a given age across time can be pinned down by differences in initial cohort productivity fed through the life-cycle profiles. Given this observation, we can isolate the effects of aging on aggregate productivity by decomposing the latter into contributions across cohorts and within-cohorts as the age distribution shifts. To see this, consider rewriting aggregate productivity growth in period t, conditional on aggregate and industry effects, so that:

$$\begin{aligned}
\Delta\Phi_t &= \sum_a^A s_{a,t}\Phi_{a,t} - \sum_a^A s_{a,t-1}\Phi_{a,t-1} \\
&= \sum_a^A (s_{a,t} - s_{a,t-1})\Phi_{a,t} + \sum_a^A s_{a,t-1}(\Phi_{a,t} - \Phi_{a,t-1}) \\
&= \sum_a^A (s_{a,t} - s_{a,t-1})\Phi_{a,t} + \sum_a^A s_{a,t-1}(\Phi_{E,t-a} - \Phi_{E,t-a-1}) \\
&= \sum_a^A \Delta s_{a,t} \sum_{j=1}^a \delta_j - \eta \sum_a^A \Delta s_{a,t} a + \eta
\end{aligned} \tag{5}$$

where the third and fourth equalities follow from plugging in equation (4), using the trend in entering cohorts, and the fact that the age shares are exhaustive. The expression makes clear that, when one isolates the contribution of changes in the age distribution over time, there are both aging effects (first term) and cohort-composition effects (the second term). Specifically, the expression $(\Phi_{E,t-a} - \Phi_{E,t-a-1})$ captures the differences in the initial productivity of entering cohorts created in period $t-a$ and $t-a-1$.

Given the framework above, we are now in a position to construct empirical counter-factuals of the lasting impact of the startup deficit on aggregate productivity. Let $s_{a,t}$ represent the historical distribution of economic activity across age groups. We can represent the net impact of a counter-factual path $s^{cf}_{a,t}$, holding constant time and industry effects, by:

$$\mathbb{E}[\Delta\Phi|s^{cf}_{a,t}] - \mathbb{E}[\Delta\Phi|s_{a,t}] = \underbrace{\sum_a^A (\Delta s^{cf}_{a,t} - \Delta s_{a,t}) \sum_{j=1}^a \delta_j}_{\text{life-cycle}} - \underbrace{\eta \sum_a^A (\Delta s^{cf}_{a,t} - \Delta s_{a,t}) a}_{\text{Cohort}} \tag{6}$$

where the expression follows from differencing equation (5) for two different paths of employment shares. The first term captures the contribution of life-cycle effects captured by our profiles, the second term captures the cohort-composition effects due to the trend in entering cohorts productivity. Note that the two effects work in opposite directions here. As economic activity shifts away from young firms, aggregate productivity growth will rise as more activity is concentrated at older firms that have already gone through the crucible of selection and are higher on the age-productivity profile – this is captured by the first term. At the same time, the shift will push down aggregate productivity growth as economic activity moves away from the younger firms who are entering with better vintages of techniques and technology. The shape of the transition and its net effect therefore depends on the rate of the shifts and the relative size of the trend in new entrants versus the steepness in the age profile. Finally, as the counter-factual is defined relative to the historical evolution of shares, the result is to be interpreted as the net effect on productivity growth had shares followed the counter-factual path rather than its historical evolution.

3 Results

3.1 Age-productivity profile

Productivity growth varies significantly over the firm life-cycle. Using the ReLBD pooled across all years 1996-2012 for firms age 1 to 15, we estimate the common age-productivity profile across 4-digit industries for each of the components of the DOP decomposition described above in section 2.2. Figure 2 plots the incumbent firm age-productivity profile for each components of the DOP decomposition along with its 95 percent confidence set. The upper left panel is net productivity growth by firm age $\Delta\Phi_{at}$, which is the sum of the selection, within and reallocation terms. The net productivity growth profile is convex and downward sloping. Expected revenue productivity growth at young firms is approximately 15 percent in the first year and falls quickly towards 0 within the first 5 years of a firm's life. After 5 years, expected productivity growth is statistically indistinguishable from zero. Looking across the components of net productivity growth, reallocation and selection account for roughly two-thirds and one-third of net productivity growth, respectively. Interestingly, almost none of the expected growth is captured by the within term. The significant productivity gains at young firms stem entirely from the high exit rates of less productive young firms and the accumulation of additional market share of the already more productive firms. Our benchmark results, because we control for year and industry fixed effects, express the age-productivity profile relative to firms age 11 to 15. In practice, the productivity growth for this group is close to zero.[14]

Despite the prominence of within firm productivity gains in the economic literature, there are a number of reasons why we shouldn't be overly surprised to find such a small contribution in our empirical result. First, our estimation procedure already allows for differential trend growth rates in productivity across detailed 4-digit industries[15]. If within firm productivity gains are highly correlated across firms within detailed industries then these will have already been removed in the first stage of the estimation process and so do not appear prominently in the age profile decomposition. Furthermore, there are also mechanical reasons related to the DOP decomposition which might dampen the contributions of the within piece. For instance, if firms that register large within productivity gains also capture larger market shares then it is possible that the DOP methodology ascribes these gains to increases in allocative efficiency rather than within firm gains. Nevertheless, given the popularity of the DOP approach we continue to rely on it in order to facilitate comparability with results elsewhere in the literature.

[14] In appendix figure A1 we plot the estimated profile in levels (rather than relative to the 11-15 group) from a specification without time and industry fixed effects against our benchmark. Year and industry fixed effects do little to change the shape of the age-productivity profile.

[15] Captured by the ν_i in our estimation framework

Overall, three key findings emerge from our estimate: (i) the age-productivity profile is downward sloping and convex and mirrors patterns estimated between employment growth and age; (ii) the magnitudes are significant but fade quickly, with nearly 2/3 of the effect disappearing after five years and nearly the entire effect disappearing after ten; (iii) the higher productivity growth of young firms is driven nearly exclusively by the forces of selection and reallocation.

3.2 Stability of age-productivity profile

Our counterfactual relies on the shape of the age-productivity profile changing little over time. Although we include a more exhaustive set of robustness checks in section 4, before constructing a counterfactual, we verify that the shape of the age-productivity profile we estimate captures fundamental dynamics of firm age and productivity and is uncorrelated with any time varying factors such as business cycles or other lower frequency changes. Our key identifying assumption presupposes that to the extent that there are cyclical or non-stationary effects on the age-productivity profile, these forces enter as level effects and so don't drive variation across age groups. We test this supposition by dividing our sample into a high-growth productivity period (1996-2004) and the more recent period of sluggish productivity growth (2005-2012) and re-run our estimation procedure including interaction terms that allow the profile to shift between the two periods. We then test to see if the interaction terms capture statistically significant shifts anywhere in the profile, or jointly, between the high-growth and low-growth periods.

Figure 3 plots the results of splitting our estimation across a high and low growth time period. The figure shows our baseline estimates and the changes captured by the interaction terms for the low growth period along with a 95 percent confidence set. All the interacted terms suggest small movements between the high growth and low growth period and are statistically insignificant. The results suggest the profiles are stationary across time.[16]

3.3 A no startup deficit counterfactual

To quantify the effects of the startup deficit and subsequent aging of U.S. business on aggregate productivity we implement the decomposition in equation (6). We use the expression to assess the net effect on aggregate productivity if, ceteris paribus, the startup deficit had never occurred and instead the entry rate and age distribution had remained at 1980 levels so that $\Delta s_{at}^{cf} = 0$. To evaluate the expression we plug in our estimates of the age-productivity profile, $\hat{\delta}_a$, and trend in entrant productivity, $\hat{\eta}$, from equations (2) and (3), respectively.[17] We then feed through the histor-

[16] See section 4.1 for evidence on the stability of our estimates across local markets.

[17] Specifically, we use estimates of the profile from our triangular panel approach in Figure 6. We assign to our censored 16+ age group a profile estimate equal to 0.005, which is the average of all age groups older than 16 which we can observe in the triangular approach. For η we use an estimate of 0.9, corresponding to the employment weighted estimate for the whole sample and very close to the fitted trend line in the non-parametric cohort effects we estimate in figure A2. See table A0 for various estimates of η over the entire sample as well as with a break in 2005 and under alternative identifying assumptions. In figure A3 we plot η_t estimated nonparametrically relative to 1997 using time effects. From equation (6) it should be clear that the effect of using different estimates of η is to scale the contributions of the cohort effects. To illustrate the point, figure 5 plots the total cumulative effect from our main

ical evolution of employment shares by age derived from the Census Bureau's Business Dynamics Statistics (BDS) database versus a counter-factual path where the distribution stays constant at its 1980 levels.

Figure 4 displays the cumulative effect on aggregate productivity from 1980-2014 had the startup deficit and accompanying shifts in the age distribution never occurred. The empirical results show a 3.1% cumulative reduction in aggregate productivity by 2014 relative to a world where the distribution of activity remained constant at its 1980 level. While the per annum effect of the transition is small, roughly 10 basis a year from 1980-2014, the cumulative effect of the process is large in economic terms. To put things in perspective, the results imply that, in 2014 alone, real median household income would have been roughly $1,600 higher had the startup deficit never occurred.[18] The cumulative lost income over the 35 year period since 1980 would clearly be magnitudes larger.

The decomposition also allows us to assess the contribution of the life-cycle and cohort components separately, which are plotted in the same figure. The first channel captures the fact that firms of different ages are at different points in their life-cycle. Our age-profile results indicate that, due to the forces of selection and reallocation, surviving firms register large increases in productivity in the early years of life. Hence, as economic activity shifts toward older firms it is also moving toward firms that, as a group, are more productive, raising aggregate productivity. The results in figure 4 suggest that the cumulative effect of these effects is roughly -1% by 2014. In other words, the startup deficit actually *raised* aggregate productivity by nearly 1% by reallocating activity from less productive entrants to more productive, older incumbent firms.

In addition to differences over the life-cycle of a firm, compositional changes can also induce changes through differences across cohorts of entrants. This effect is captured by the second channel in our decomposition which accounts for the fact that entering cohorts start with different techniques and vintages of capital. This channel is perhaps that which comes to mind most readily when people think about the importance of entrants in the aggregate economy; namely, that the productivity of entrants improves over time as subsequent generations of entrants adopt the latest vintages of techniques and equipment. This channel creates a negative drag on productivity as the startup deficit reduces the number of firms entering with the latest techniques and technologies. Our estimates put the cumulative effect of these cohort effects at nearly 4% by 2014, more than offsetting the negative effect from the life-cycle channel and accounting for the lion's share of the effects we capture with our approach.

3.4 Older firms

At older ages the expected productivity growth may become negative. If expected net productivity growth were always weakly positive and entrants were initially less productive than incumbents, an increase in entry and its dynamic effects on the share of young businesses would reduce aggregate

results in figure 4 for various η within the range of our estimates.

[18] Calculated using the Census Bureau's estimates of real median household income in 2014 of $53,718

productivity growth by reducing the relative market share of older more productive businesses. Our main profile results only consider firms up to age 15 because we can observe this age group for all years in the ReLBD.[19] By relaxing the requirement that we observe each age group in each year of our data, we can construct an "unbalanced" triangular panel, where older age groups are included only in years when that age group is not censored. Figure 6 plots the same components estimated on the triangular panel. Although for the more advanced ages, standard errors are slightly larger because of fewer years of data, the point estimates begin to push below zero because of a further decline in the reallocation and selection terms.[20] A decline in productivity for older firms should not be surprising. In a relatively standard model of firm dynamics with diminishing marginal revenue product and endogenous exit, older firms are more likely to be close their optimal scale and thus there are no longer gains from reallocation. Additionally, older firms are more likely to be sufficiently profitable that idiosyncratic reasons uncorrelated with productivity account for a greater share of exits than shocks to profitability.

Our use of the triangular panel leans heavily on the assumption that, conditional on age group, the productivity growth terms are approximately stationary. As we show above in section 3.2, there is no change in the average DOP profiles between the early 1996-2004 high growth period to the later 2005-2012 low growth period.[21]

However, when looking at the left censored group, who entered before 1979, we do see declines over time in the allocation DOP terms. We apply the DOP decomposition from (1) to firms age 16 or more. This includes the left censored group of the very oldest firms. Then for each term in the DOP decomposition for just the oldest age group by industry and year we project on to industry fixed effects and a dummy variable for the 2005 to 2012 period. Table 1 reports the estimated coefficient on this dummy variable, which should be interpreted as the average change, from 1996-2004 to 2005-2012, in the components of productivity growth for firms at least 16 years old. Average annual productivity growth for this age group declines by 2.3 percentage points from the early period to the late period, primarily from a decline in the allocation term. This finding is consistent with Decker, Haltiwanger, Jarmin, and Miranda (2017) who apply the DOP decomposition to *all* firms and identify a decline in the allocation term over time. Given the stable DOP profile we find for firms age 1 to 15 in section 3.2, table 1 reveals that the decline they describe is driven by the oldest firms, which constitute more than 70 percent of total employment.

This decline among the oldest firms could be a consequence of the startup deficit. One possibility is that it reflects compositional changes among the oldest firms from the declining inflows into this age group: the age-productivity profile may be stable, but employment is shifting further along

[19]Because age is assigned when a firm hires its first employee, firm age will always be right censored (birth year is left censored), where the uncensored maximum age increases with each additional year of data. We treat any birth year before 1979 as left censored.

[20]Strictly speaking, less than the 11-15 group, because we condition on year and industry fixed effects. However, note from figure A1 without controls that the growth of the 11-15 group is almost exactly zero.

[21]We also find no statistically significant differences for older uncensored firms using the triangular panel, however because older age groups feature fewer observations over time and are observed only for a subset of the early period, this is necessarily a lower powered test

to a negative part of the profile. An alternative possibility is that, for these firms, the profile is nonstationary with a worsening of the allocation term over time. This latter possibility may also be a consequence of the a startup deficit if increasing concentration within industries from declining entry impedes gains from reallocation. Lacking measures of age for this group we cannot distinguish these two potential explanations. Nevertheless, if the declining productivity growth among the oldest firms is a consequence of the declines in entry, then the effects of the startup deficit on productivity growth we describe above are significantly understated. We leave exploration of this interesting topic to future work.

4 Robustness

We have shown that firm age matters and that the profile of labor productivity growth is to a first approximation stable across time. In this section, we review potential threats to the validity and generality of these findings and discuss how we address them. Specifically, our estimates aim to pin down variation in growth rates of labor productivity over the life-cycle of a firm. Our identifying conjecture is that such a stable relationship exists independently of compositional and time effects so that we can identify it up to some level scale. Even though the results of our estimation procedure suggest that this is a good characterization of empirical regularities in the data, there are good a priori reasons to believe this structure is too strong and that alternative approaches to the data may cause our results to disappear. Below, we address these concerns by testing whether our results are robust to pricing effects, organizational status, industrial composition, and other sources of non-stationarity.

Overall, we find that the convex pattern we uncover in the baseline estimate is robust to a number of tests. While alternative specifications do induce some twisting of the profile, these effects are almost always statistically insignificant or otherwise minor. Across all tests the shape of the profile is always preserved and most meaningful differences emerge only as level effects.

4.1 Stability: across markets

Given the length of our time period, it is be hard to test for cyclical effects in the time-series at any finer level than the split above. However, we can test our assumption that the profiles are constant up to a level effect by exploiting variation in business cycles across regions. We define local markets at the CBSA and CBSA-by-Industry level and then divide local market data into terciles depending on (i) average startup rates, and (ii) average change in startup rates over our sample period. We then re-run our estimates with interaction terms to capture changes across high startup activity locations and low startup activity locations. The idea is that if the shape of our profile estimates are sensitive to cyclical fluctuations in startup activity we would see meaningful differences across local markets with high-activity and those with low-activity in the cross-section. To the extent that patterns are preserved up to levels, we can conclude that our results are robust to this dimension.

Figures 12 and 13 show how the estimated profiles change when we restrict attention to high activity and low activity geographic markets grouped in terms of the level of startup activity and average growth rates in startup activity, respectively. The changes we find here are the most substantial in all our robustness tests but are still consistent with their being an underlying convex relationship between age and productivity growth, as in our baseline model. Moving from the low tercile to high tercile, we see that the magnitude of growth rates can nearly double at a given age, which should be expected from the definition of the groups. The more remarkable outcome, however, is that the curvature of the profile which dictates the differences across age groups is nearly entirely preserved, suggesting that and consistent with our main identification assumption, the heterogeneity across market states itself enters mainly as a level shift.

4.2 Price effects: nominal versus real

One of the most well known empirical issues with existing large-scale firm level datasets is that they often lack reliable firm-level price information. In certain settings, failing to account for this pricing heterogeneity, even in narrowly defined industries, can result in misleading conclusions about productivity growth rates (see Foster, Haltiwanger, and Syverson (2008)). While we cannot directly control for firm level pricing heterogeneity, we can control for pricing heterogeneity across narrowly defined 4-digit NAICS industries both indirectly, through the use of fixed effects, and directly, by using publicly available price indices from the BEA. This cross-industry variation in prices is likely the biggest source of potential bias as our identification strategy relies on exploiting detailed cross-industry variation over time. To tackle this directly, we re-run our estimation procedure using industry measures of output-per-worker calculated by deflating our revenue data using the most comprehensive set of BLS price indices available. The results are presented in figure 7. Adjusting our data with BLS price series had almost no noticeable impact on our estimates, suggesting that variation in prices across industries was already well controlled for by working in log growth rates and including industry and time fixed effects.

While in principle it is still possible that their exist differences in pricing strategies between young and old firms *within* industries that could be driving our profile estimates, we view this fact as an interpretation rather than a threat to validity. The extent to which these systematic age-pricing differences exist across all detailed industries in the nonfarm business sector is an indication of an important economic mechanism at work causing the distribution of firm age to matter for aggregate outcomes. As stated above, we remain agnostic as to the underlying mechanism driving our results and focus instead on establishing a characterization and quantification of the role of firm age and its distribution in aggregate outcomes. We therefore leave open the possibility that our results are driven by the evolution of pricing strategies over the life cycle of firms as a potential mechanisms to rationalize the data. We hope future researchers will bring better data to bear on this important question.

4.3 Organizational: single-unit versus multi-unit status

Another potential threat to validity is the failure to account for heterogeneity in the organizational structure of firms. Our concern arises from the fact that there is significant age-bias in the distribution of organizational status across firm types: most entrants and young firms are single-units and multi-unit firms are mostly concentrated in the older part of the age distribution. Figure 9 highlights the extent to which this occurs in terms of employment and number of firms. Any systematic differences between these organizational types then might pollute our estimates of age-effects if not properly controlled for. To address this, we split our sample based on organizational status and re-estimate the profiles over a sub-sample of single-unit firms only and one of multi-unit firms only. Comparing the results allows us to assess to what extent organizational status may be confounding the patterns we uncover between age and productivity growth. Interestingly, the profile estimates hardly change when we restrict ourselves only to single-unit firms, confirming that heterogeneity in organizational structure is not driving our results either.

4.4 Compositional: industry representativeness

It is also well known that there is substantial variation in firm dynamics across industries which makes comparing levels of productivities across industries potentially problematic. In our baseline analysis, we address this critique by conducting our estimation in growth rates and accounting for different industry trends through the use of industry fixed effects. Nevertheless, one remaining concern is that there exists a wide variation in the age-productivity profile across industries and by exploiting this variation for estimation we generate results for a "representative industry" that displays patterns not present in any given industrial group. To ensure that our findings are representative *within* industry groups, and not just across them, we divide our data across 2-digit NAICS sectors are re-run out estimation procedures across detailed industries within each group. Doing so allows us to assess the extent to which our aggregate profile represent trends common across and within industries groups in the nonfarm business sector.

Figures 10 and 11 summarize the robustness tests for industry composition. Figure 10 super imposes our baseline estimates over those estimates derived *within* 11 different 2-digit NAICS sectors. In this exercise, we include all sectors except those corresponding to raw materials (i.e. agriculture, mining) or utilities. What the figure makes clear is that the overall profile pattern is present within each of the industry groups and does not deviate significantly in curvature or magnitude. Figure 11 plots the profiles separately for each industry group to better identify where the deviations come from. It is clear that the life-cycle pattern of labor productivity growth is remarkably constant across industries.

Despite the robustness of the underlying age-productivity profile's shape, differences do exist, particularly at the young tail of the estimates. Moreover, declines in entry rates, and the subsequent shifts in the age distributions, have been uneven across sectors. Combined, these between sector effects continue to suggest the possibility of significant divergences between our representative sector approach and the outcomes of individual industries. To examine this in more detail, we compute the counterfactual from section 2.4 using detailed industry variation within NAICS supersectors. This involves calculating age-productivity profiles, trends in entering cohorts, and shifts in the age distribution for each sector separately. We plot the new counterfactual in figure A4. Each line represents the cumulative difference in a sector's average productivity holding the sector's employment by age group constant relative to the actual evolution of those within sector employment shares. Figures A5 and A6 plot the life-cycle and cohort components, respectively, of the cumulative difference. While on average these are consistent with our aggregate counterfactual presented in section 3.3, there are some significant differences between sectors: sectors like the finance, insurance and real estate sector and the general services sector, where the latter contains many small scale service businesses such as auto mechanics or nail salons, were little changed by the startup deficit, whereas our counterfactual identifies sectors such as construction and professional and business services as changed sharply. These differences coincide somewhat, but not perfectly, with the pace of further declines in each sector's young firm share. Nevertheless, the qualitative results of our analysis remain unchanged.

5 Cross-sectional Evidence

As a complement to our decomposition-based counterfactual, we now explore a different source of variation to reveal the relationship between startups, aging firms, and productivity growth. Here we relax the strong identifying assumption of a stable age-productivity growth profile, and instead ask whether areas with relatively higher startup rates and younger firms also exhibit faster productivity growth. On its own, this exercise would raise significant concerns of reverse causality: startup rates could be elevated because of local innovations to productivity rather than the reverse. To address this possibility, we adopt two different instrumenting strategies to generate plausibly exogenous shifts across local markets in the level of startup activity. In both cases, we find an economically and statistically significant effect of changes in the startup rate on labor productivity growth. Importantly, these effects could follow from the direct compositional changes we characterize in our counterfactual or from spillovers that change the dynamics of productivity growth even within age group. This last channel is especially important among older firms, which comprise the bulk of business sector employment and for which in section 3.4 we observed declines in the contribution of reallocation to productivity growth.[22]

We begin by exploring the reduced form relationship between startup activity and within-industry labor productivity growth by exploiting the rich geographic variation across local markets: both states and CBSAs. Our dependent variable is the annual labor productivity growth in year t for industry j and area k. Pooling all years, we project industry x area productivity growth on the area startup rate SR_{kt} as well as year, industry and area fixed effects. That is we estimate the following model:

[22] See also Decker, Haltiwanger, Jarmin, and Miranda (2017) who find similar declines in reallocation.

$$\Delta\Phi_{jkt} = \mu_t + \nu_j + \gamma_k + \beta SR_{kt} + \varepsilon_{jkt}. \tag{7}$$

Column (1) of table 2 reports the estimated $\hat{\beta}$ from estimating this specification by OLS. The model exploits cross state and industry variation in the startup rate (equally weighted) and standard errors are clustered at the state level, allowing for serial correlation within state.[23] There is a clear and strong correlation between states with relatively high startup rates and productivity growth within industries in those states. These OLS estimates confirm that states with relative increases in the entry rate are also ones with relative increases in industry productivity growth. The magnitudes are sizable: a 1 percentage point decline in the startup rate predicts that industry productivity growth within the area would decline by 0.796%.

We can also measure the correlation of within state changes in entry and gross output per worker growth within the state by using real gross state products published by the Bureau of Economic Analysis. This measure is closer in spirit to the aggregate nonfarm business sector productivity growth. Column (3) of table 2 reports the estimated $\hat{\beta}$ using a state's GSP/worker growth as the dependent variable. Here we also observe a strong correlation with the startup rate, although economically smaller. Gross state product grows relatively faster in states with increasing startup rates.

These reduced-form results cannot say whether increasing startup rates lead to higher productivity or the reverse. Faster productivity growth could also lead to increasing entry as businesses form to take advantage of the opportunities created by the gains in productivity. To learn about whether shifts in entry may *cause* shifts in productivity growth, we look to two different instrumenting strategies.

Demographic instrument Our first IV relies on the relationship between slow-moving demographic shifts and the entry rate. Our approach draws on the recent literature studying the determinants of entrepreneurship and startup activity. Karahan, Pugsley, and Şahin (2016) show that changing demographics play a significant role in the equilibrium startup rate both theoretically and empirically. Karahan, Pugsley, and Şahin (2016) develop an a demographic instrument, based on long lags of a state's fertility rate, to generate shifts in that state's contemporary growth in the working age population or labor force. With imperfect mobility, increases in a state's births will lead to an increase in the growth rate of the labor supply when that birth cohort enters the working age population. Arguably, conditional on state and future year fixed effects, forecasts of future businesses conditions are unrelated to a fertility decision many years in advance. Using this instrument they find that states with larger declines in the growth rate of their labor force, predictable only by lagged demographics, also have larger declines in their startup rates. With this mechanism in mind, we adopt the same demographic instrument to generate plausibly exogenous shifts in a states startup rate vis-a-vis the demographic channel. The exclusion restriction, that lagged demographics may change productivity growth only through their effects on labor supply

[23] Results using employment weights were similar.

growth, is more challenging to defend in this context. Lagged demographics, by changing the age composition of the workforce, may also change observed labor productivity through its effects on average worker quality. To the extent this caveat is applicable, it should weigh against any effect since increases in labor supply growth from fertility would shift the composition by increasing the share of younger and less experienced workers.

Table 2 column (2) reports the estimates of equation (7), where SR_{kt} is instrumented with 20 year lags of the state's fertility rate.[24] When a state's startup rate is increasing because of demographics and not current business conditions, industry productivity growth increases. This semi-elasticity is larger than the one estimated in the reduced form results and suggests that a one percentage point decrease in the startup rate (e.g. from 6% to 5%) lowers labor productivity by 1.46%. Again, standard errors are clustered at the state level. Column (4) presents results when the outcome is instead state output per work. Again, the IV results are stronger than the OLS results and highly significant.

Collateral value instrument The second IV approach draws on a growing literature studying how financing opportunities for new and young firms differ from those of established incumbents. In particular, we appeal to the finding that a large number of new startups are financed through the home equity of entrepreneurs in the early years of their operation (See Adelino, Schoar, and Severino (2015), Robb and Robinson (2014) and references, therein). As a result, exogenous increases in local housing prices could loosen financing constraints faced by would-be entrepreneurs and young firms and encourage startup activity through a *collateral channel*. To identify such variations, we use the housing price booms caused by speculative activity in the run up to the great recession identified by Charles, Hurst, and Notowidigdo (2016).

The idea behind their identification is that some of the variation in housing prices over the boom and bust was a result of non-rational changes in house price expectations (speculative bubbles) and not from changes in income, construction costs or population. This "bubble" component of house prices can be isolated by identifying structural breaks in local house prices series. The main identification assumption is that standard drivers of house prices (income, population growth and housing costs) are smoothly incorporated into house prices while the bubble component induces breaks in the price series.Charles, Hurst, and Notowidigdo (2016) provide evidence that these these "sharp" breaks in house prices are not systematically related to pre-period levels and changes of income, population and housing costs. Furthermore, they show that this instrument has power: these break explains a significant portion of MSA level house price variation over the period 2000-2006.

One potential challenge with this approach is that the increase in home equity likely also stimulates local demand and might induce local businesses to undertake other productivity enhancing investments that confound our measurements. To address this issue, we focus these IV regressions

[24] For additional power, we also include 10 year lags of the shares of pre-working age and pre-retirement population. These serve to generate additional variation in the growth rate of labor supply by explicitly changing the inflows and outflows of the working age population across states.

on the growth of young firms in the tradeable sector. By grouping young firms with new entrants we seek to identify the population of firms that would have benefited most from the collateral channel[25]. By excluding both the construction and non-tradedable sector, we seek to identify firms for whom the local demand effect is relatively minor and so avoid the confounding effects of increased local demand.

Our goal is to estimate the causal relationship between the change in startup activity and labor productivity growth across local markets. Our dependent variable is the change in log labor productivity growth in CBSA k over the period 2000-2006. This is the same time horizon over which the Charles, Hurst, and Notowidigdo (2016) housing demand instrument is defined. In particular, we estimate the following equation:

$$\Delta\Phi_k = \gamma_k + \beta\Delta SR_k + \varepsilon_k. \tag{8}$$

Before estimating equation (8), we establish that our housing demand instrument predicts variation in the startup rate. Columns (5)-(6) of table 2 below report the F-stats from these first stage results. These show that our housing demand instrument has strong predictive power for changes in startup activity. We find an F-stat of over 100 when we construct productivity using all industries (column 5), and an F-stat over 20 even when we exclude the construction and non-tradeable sectors[26]. Thus in all cases the F-statistic is well above 10 indicating that our instrument is strongly correlated with changes in the startup rate over this time period.

Table 2 column (5) reports the estimates of equation (8), where ΔSR_k is instrumented with Charles, Hurst, and Notowidigdo (2016) housing demand instrument. When a state's startup rate is increasing because collateral constraints are relaxed, we see a strong increase in MSA level productivity. In our baseline semi-elasticity specification, we find that a percentage point decline in the startup rate leads to 2.8% decline in MSA labor productivity growth over the same period that is strongly statistically significant even when standard error clustered at the state level. Interestingly, despite using a very different identification scheme, we find results similar to our demographic instruments. Column (6) shows this result is robust to excluding both the construction and non-tradeable goods sectors. This helps alleviate concerns that our results are being driven by unobserved changes in local demand.[27]

Overall, we uncover robust evidence that increases in entry, ceteris paribus, leads to increases in productivity growth. Because our cross sectional approach does not rely on the strong assumptions of the counterfactual, it captures both the direct effects of the decline in entry embodied in the counterfactual as well as any indirect effects from within age group changes induced by the decline in entry. These indirect effects are compatible with both our observed decline in the reallocation

[25] Our empirical results below suggest that the age-productivity dynamics we identity are most pronounced from entry through the first five years of a firm's life. Therefore, we view grouping young firms together with startups as still consistent with our underlying approach and an improvement in the identification of the collateral IV.

[26] We use the classification system of Main and Sufi (2014).

[27] We also find that our results are stronger when excluding highly density areas. This is consistent with our story that relaxation of credit constraints is driving our results since less populated areas typically have less access to credit.

component of productivity growth among the oldest firms as well as other recent evidence on the declines in within age group measures of job reallocation, such as in Decker, Haltiwanger, Jarmin, and Miranda (2016b). We should also stress that this does not mean on net productivity growth must increase. In fact, the robust productivity growth of the late 1990s and early 2000s occurred even as startup rates were continuing to slow. Instead, our estimates imply that the decline in entry and its effects on the age distribution restrained the effects of these gains in productivity.

6 Conclusion

This paper studies the link between declining firm entry, the aging of the firm distribution and productivity growth using U.S. Census data representative of the nonfarm business sector. Consistent with a growing body of research, we find that age composition plays a key role in shaping the dynamics of labor productivity growth.[28]

We show that the relationship between firm age and productivity is downward sloping and convex. The magnitudes of the differences are substantial but short lived. Conditional on surviving, new entrants register cumulative productivity growth of roughly 20% in the first 5 years of operation. After year 5, however, the productivity profile flattens dramatically and is statistically at or near zero for the remainder of the age distribution we observe.

Applying the dynamic Olley-Pakes decomposition to the profile, we find that the strong performance of young firms is driven nearly exclusively by the forces of selection and allocation. In other words, the fast gains in productivity of young firms is driven by the fact that inefficient entrants lose market share and exit quickly, rather than productivity growth which occurs *within* surviving firms. In the last section of our paper, we show how the driving forces of selection and allocation in shaping the age-productivity profile we uncover emerge easily from a variant of the workhorse Hopenhayn (1992) model of firm dynamics.

Our results suggest that the start-up deficit and subsequent aging of the U.S. business sector have had a considerable impact on aggregate productivity. Using a model-free aggregation technique, we show that our results suggest the start-up deficit and accompanying aging have reduced aggregate productivity by roughly 0.10 percentage points a year from 1980-2014. While the per annum rate is small, the cumulative effect over the whole period is substantial, reducing the level of aggregate productivity by 3.1% by 2014.

However, this counterfactual may understate the importance of firm ages. We document that since 2005 mature firms (age 20+) have become an even greater drag on productivity growth. We apply our DOP decomposition and identify the source of this drag as a slowdown in the allocation component. In other words, there was a decline in allocative efficiency for the most mature firms in our sample in the sense that market share was flowing more slowly to the most productive mature

[28]While we focus our analysis on the effects of the compositional change of firm age, we readily acknowledge that this is not the whole story. Recent work by Decker, Haltiwanger, Jarmin, and Miranda (2014) and Decker, Haltiwanger, Jarmin, and Miranda (2017) have highlighted large changes in employment and productivity dynamics even *within* firm age groups. Understanding these trends is an important area for future study.

firms.

Our main results are complemented by a series of cross-sectional IV regressions that, unlike our decompositions, admit a causal interpretation. By exploiting plausibly exogenous variation in start-up activity through demographic and collateral channels we are able to show that the local labor productivity growth does indeed to exhibit a causal link to start-up activity across geographic and industrial markets.

Given that our aim is mainly empirical, we hope that this paper provided many useful facts for applied modelers to explain and calibrate their models to. Going forward, we think there are many interesting follow up papers to be written. Chief among them is to develop a better understanding of the economic mechanisms behind the decline in allocation among mature firms since understanding this phenomenon will be useful for understanding what will happen to labor productivity growth over the next 5-10 years.

References

ADELINO, M., A. SCHOAR, AND F. SEVERINO (2015): "House prices, collateral, and self-employment," *Journal of Financial Economics*, 117(2), 288–306.

ARKOLAKIS, C., T. PAPAGEORGIOU, AND O. TIMOSHENKO (2014): "Firm learning and growth," Discussion paper, Yale, mimeo.

AUTOR, D., D. DORN, L. KATZ, C. PATTERSON, AND J. VAN REENEN (2017): "The fall of the labor share and the rise of superstar firms," .

BARTELSMAN, E., J. HALTIWANGER, AND S. SCARPETTA (2013): "Cross-country differences in productivity: The role of allocation and selection," *The American Economic Review*, 103(1), 305–334.

CHARLES, K. K., E. HURST, AND M. J. NOTOWIDIGDO (2016): "Housing Booms, Manufacturing Decline, and Labor Market Outcomes," *March. http://faculty. wcas. northwestern. edu/noto/research/CHN_manuf_decline_housing_ booms_mar2016. pdf.*

DAVIS, S. J., J. HALTIWANGER, R. JARMIN, AND J. MIRANDA (2006): "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," *NBER macroeconomics annual*, 21, 107–179.

DECKER, R., J. HALTIWANGER, R. JARMIN, AND J. MIRANDA (2014): "The role of entrepreneurship in US job creation and economic dynamism," *The Journal of Economic Perspectives*, 28(3), 3–24.

DECKER, R. A., J. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA (2016a): "Declining Business Dynamism: Implications for Productivity?," *Brookings Institution, Hutchins Center Working Paper.*

DECKER, R. A., J. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA (2016b): "Where has all the skewness gone? The decline in high-growth (young) firms in the U.S.," *European Economic Review*, 86, 4 – 23, The Economics of Entrepreneurship.

DECKER, R. A., J. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA (2017): "Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown," *American Economic Review, Papers and Proceedings.*

DENT, R. C., F. KARAHAN, B. PUGSLEY, AND A. ŞAHIN (2016): "The Role of Startups in Structural Transformation," *The American Economic Review*, 106(5), 219–223.

DIEWERT, W. E., AND K. J. FOX (2005): "On Measuring the Contribution of Entering and Exiting Firms to Aggregate Productivity Growth," *Vancouver School of Economics Discussion Paper*, 05-02.

Dunne, T., M. J. Roberts, and L. Samuelson (1989): "The growth and failure of US manufacturing plants," *The Quarterly Journal of Economics*, 104(4), 671–698.

Evans, D. S. (1987): "The relationship between firm growth, size, and age: Estimates for 100 manufacturing industries," *The journal of industrial economics*, pp. 567–581.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?," *The American economic review*, 98(1), 394–425.

Foster, L., J. C. Haltiwanger, and C. J. Krizan (2001): "Aggregate productivity growth: lessons from microeconomic evidence," in *New developments in productivity analysis*, pp. 303–372. University of Chicago Press.

Haltiwanger, J., R. S. Jarmin, R. Kulick, and J. Miranda (2016a): *High Growth Young Firms: Contribution to Job, Output, and Productivity Growth*University of Chicago Press.

Haltiwanger, J., R. S. Jarmin, R. Kulick, and J. Miranda (2016b): "High growth young firms: contribution to job, output, and productivity growth," in *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*. University of Chicago Press.

Haltiwanger, J., R. S. Jarmin, and J. Miranda (2013): "Who creates jobs? Small versus large versus young," *Review of Economics and Statistics*, 95(2), 347–361.

Hathaway, I., and R. E. Litan (2014): "Declining business dynamism in the United States: A look at states and metros," *Brookings Institution*.

Hopenhayn, H. A. (1992): "Entry, exit, and firm dynamics in long run equilibrium," *Econometrica: Journal of the Econometric Society*, pp. 1127–1150.

Hyatt, H., and J. Spletzer (2013): "The Recent Decline in Employment Dynamics," .

Jarmin, R. S., and J. Miranda (2002): "The longitudinal business database," .

Karahan, F., B. Pugsley, and A. Şahin (2016): "Demographic Origins of the Startup Deficit," .

Melitz, M. J., and S. Polanec (2015): "Dynamic Olley-Pakes productivity decomposition with entry and exit," *The Rand journal of economics*, 46(2), 362–375.

Moreira, S. (2016): "Firm dynamics, persistent effects of entry conditions, and business cycles," .

Pugsley, B. W., and A. Sahin (2014): "Grown-up business cycles," *FRB of New York Staff Report*, (707).

Reedy, E. J., and R. J. Strom (2012): *Starting Smaller; Staying Smaller: America's Slow Leak in Job Creation*pp. 71–85. Physica-Verlag HD, Heidelberg.

Robb, A. M., and D. T. Robinson (2014): "The Capital Structure Decisions of New Firms," *Review of Financial Studies*, 27(1), 153–179.

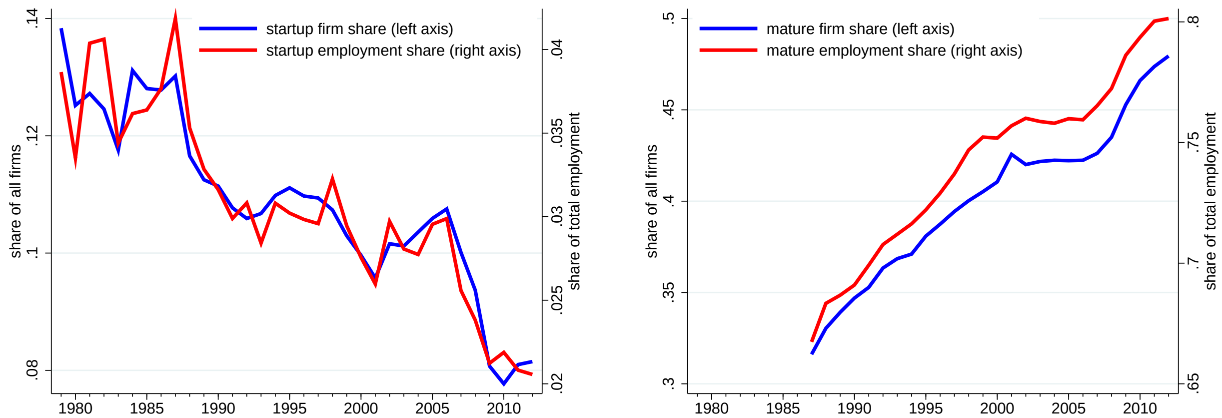

(a) Startups (Age 0) (b) Mature firms (Age 11+)

Figure 1: Employment and Firm Shares, 1979-2012

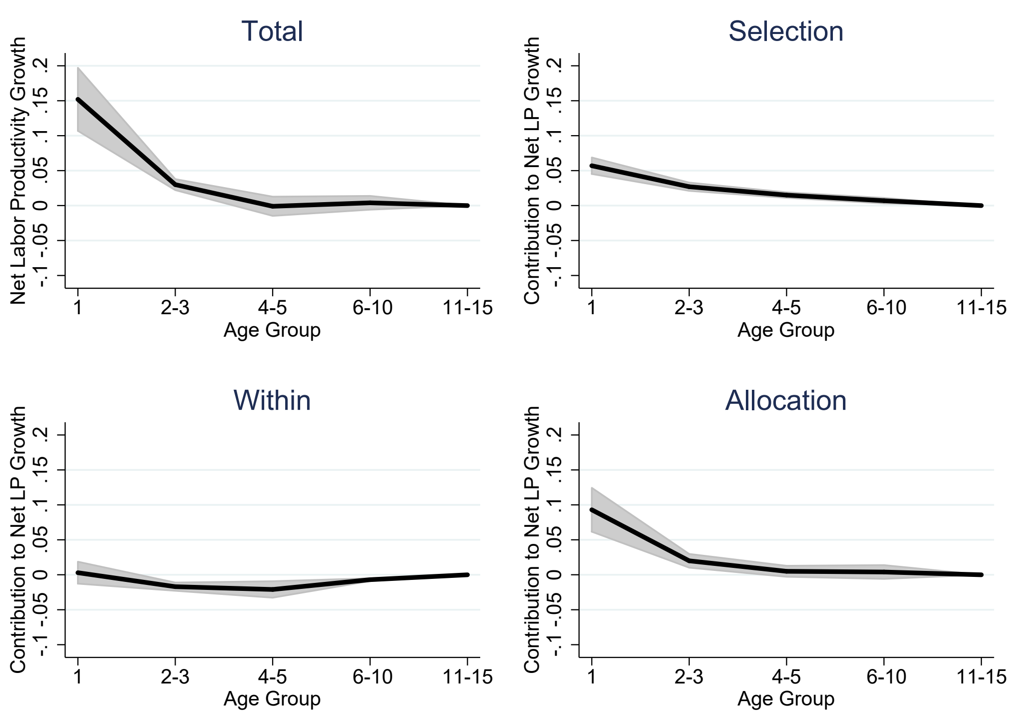

Figure 2: DOP Decomposition by Firm Age

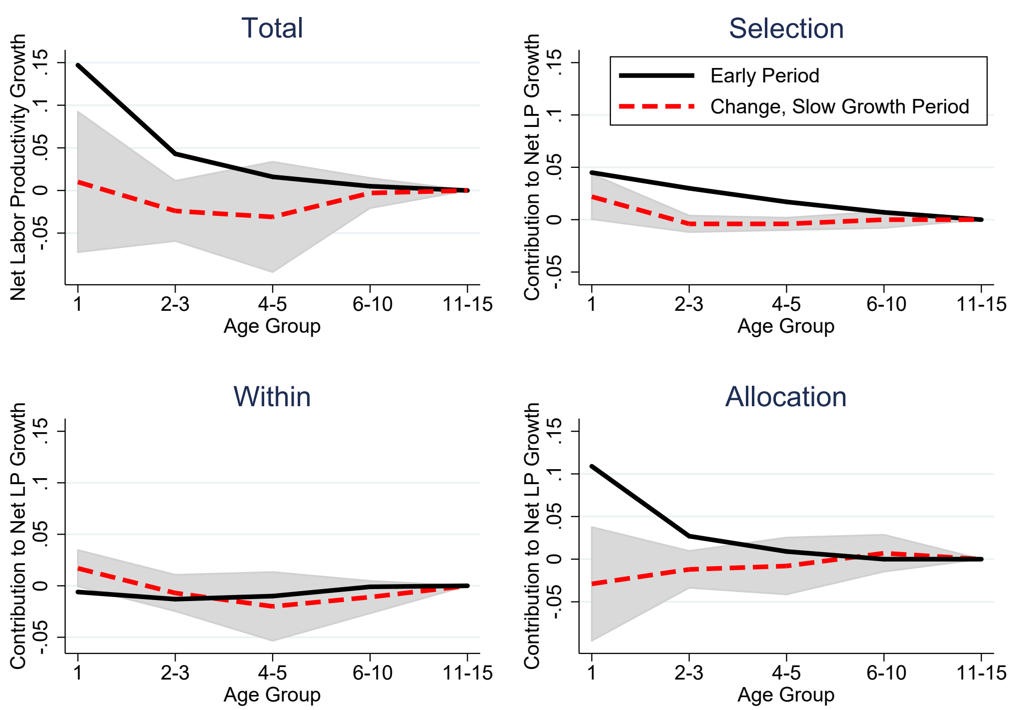

Figure 3: DOP Decomposition by Firm Age: Early ('96-'04) vs. Late ('05-'12)

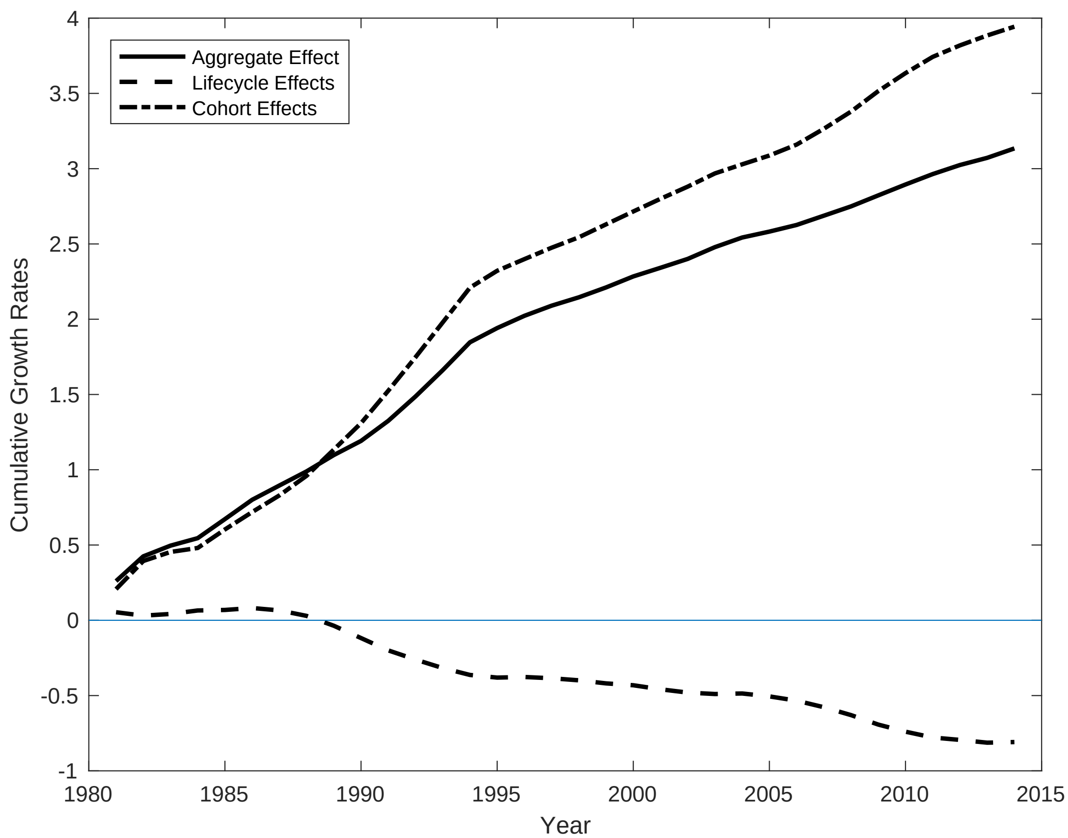

Figure 4: Empirical Counter-Factual and Components

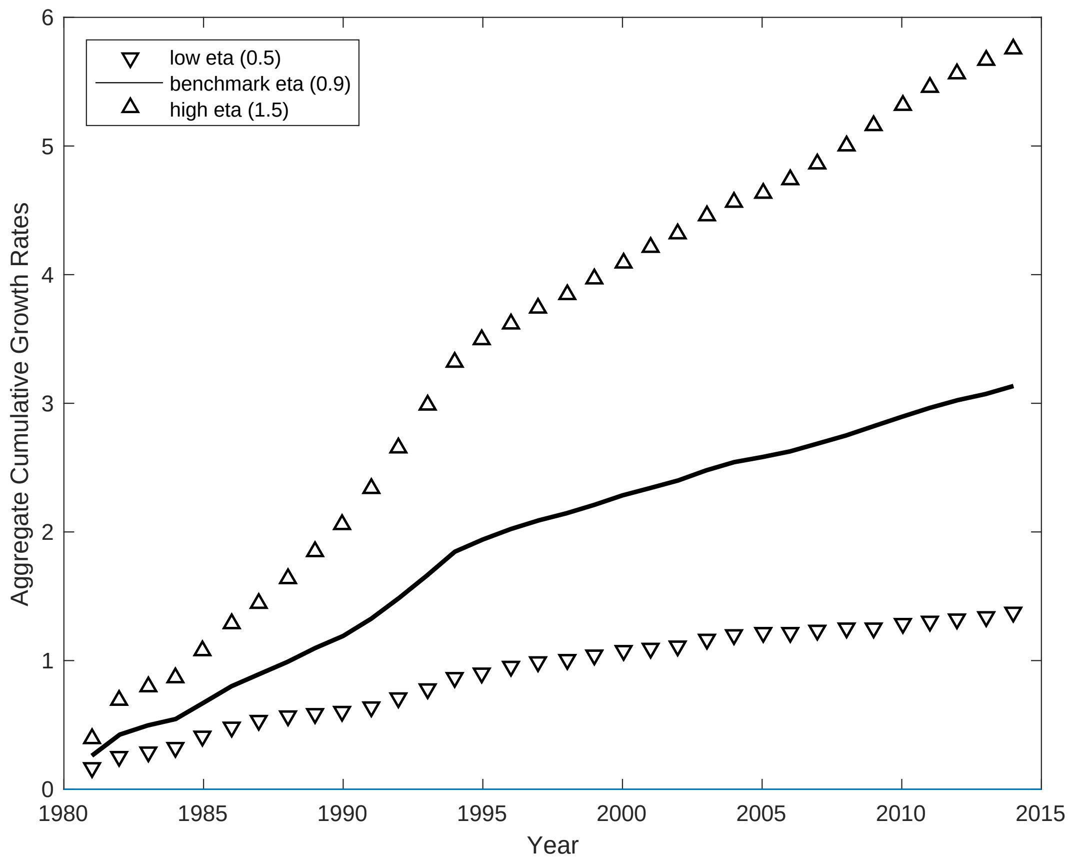

Figure 5: Total Counter-Factual with various η

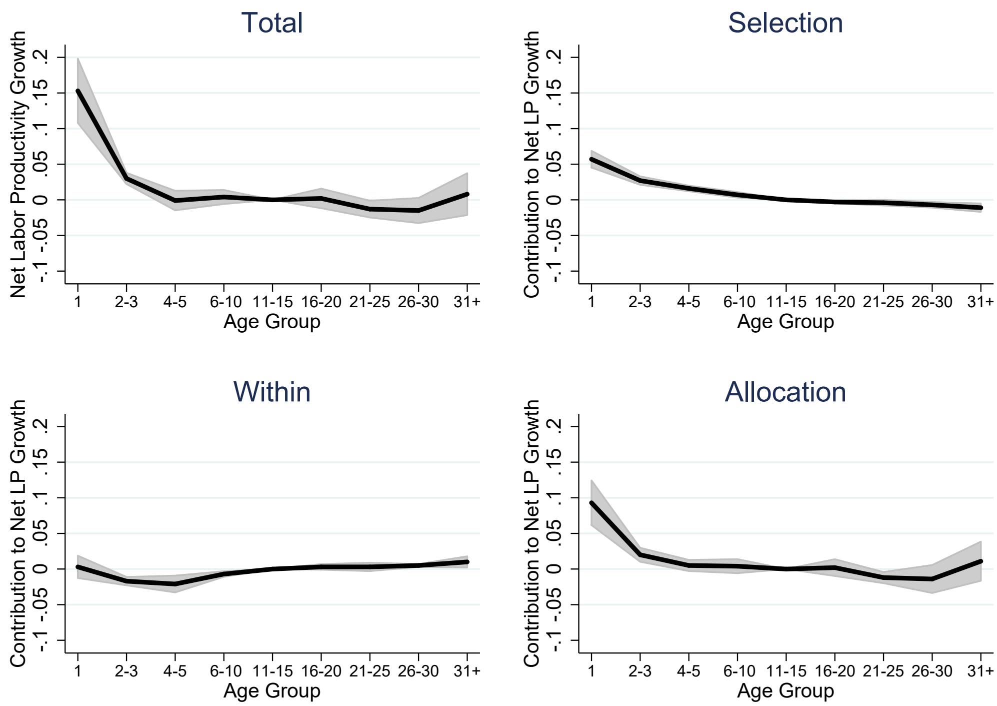

Figure 6: DOP Decomposition by Firm Age

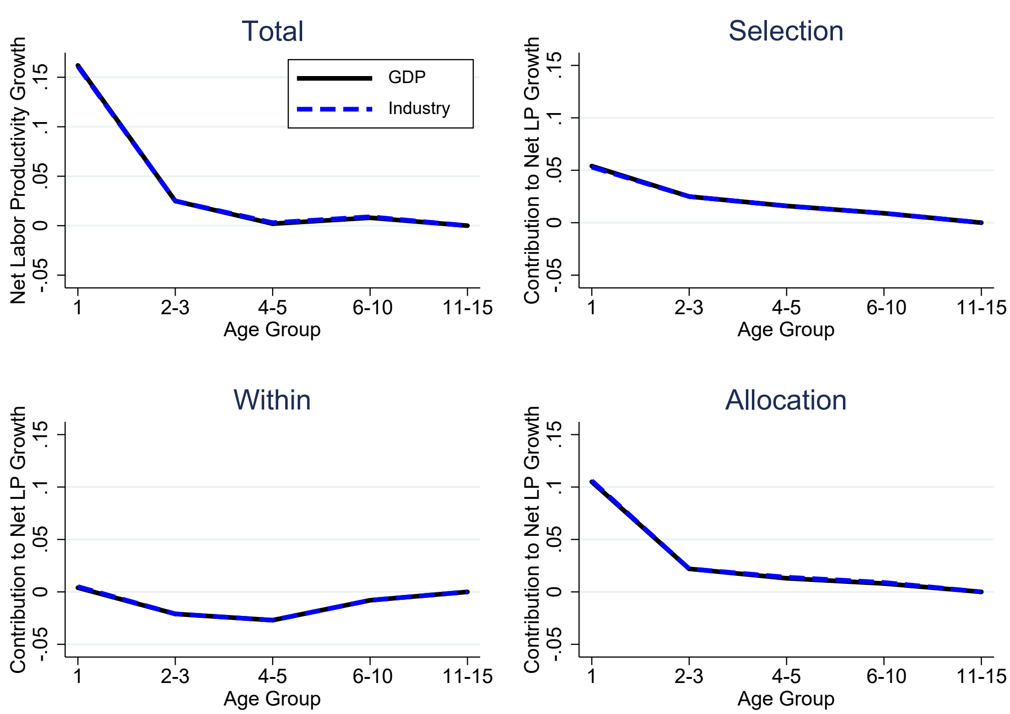

Figure 7: DOP decomposition by firm age: common GDP deflator and BLS industry-specific price deflators

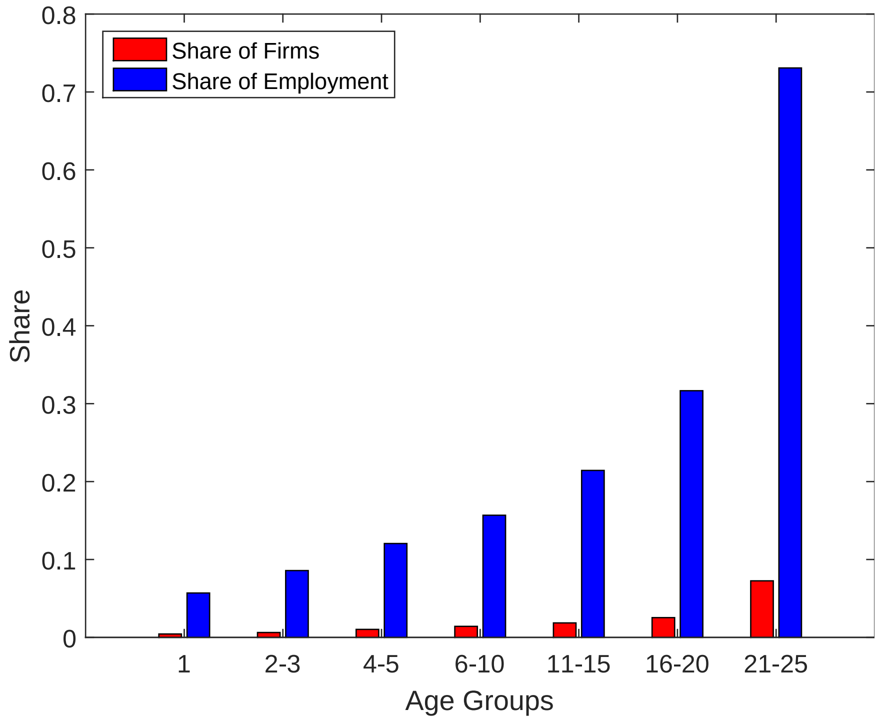

Figure 8: Share of multi establishment firms by firm age

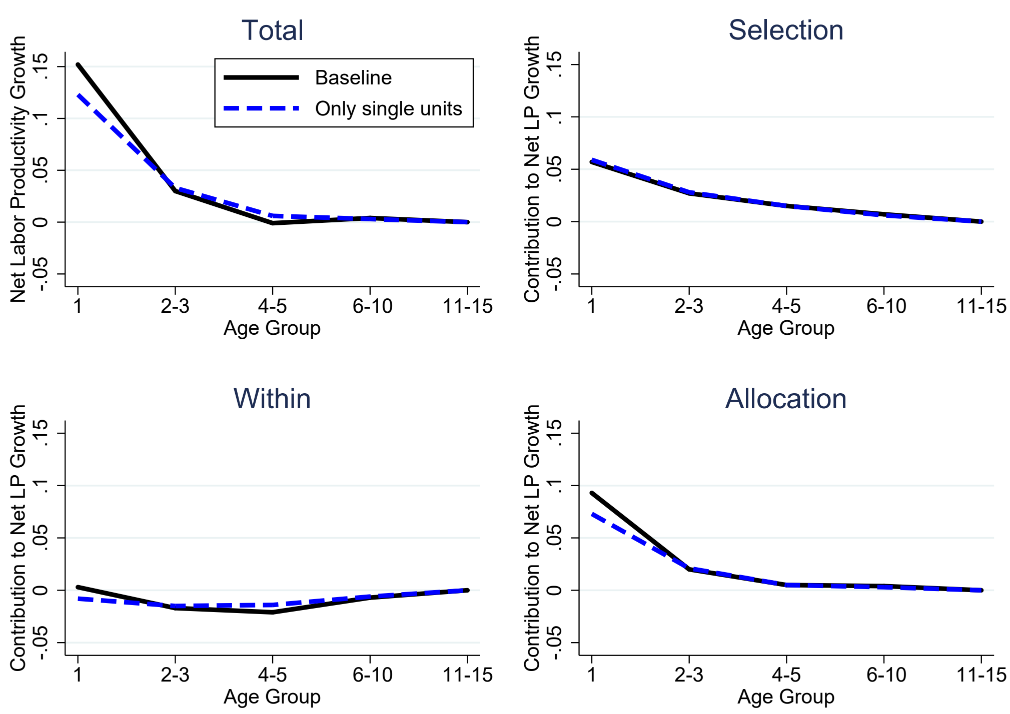

Figure 9: DOP decomposition by firm age: overall and only single establishment firms

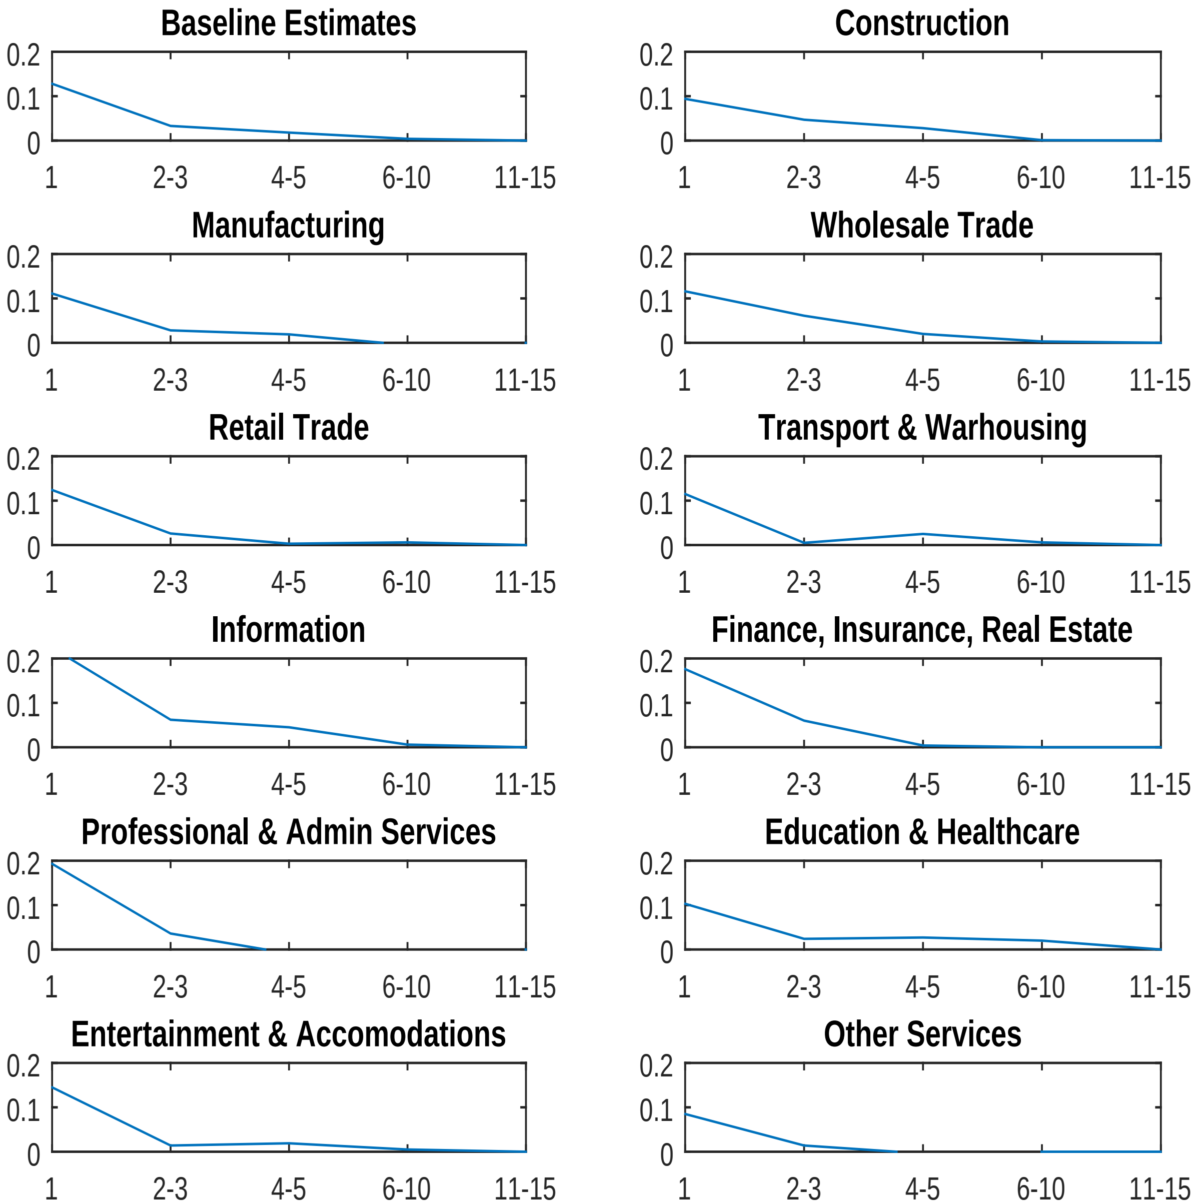

Figure 10: Net age-productivity growth profile by sector

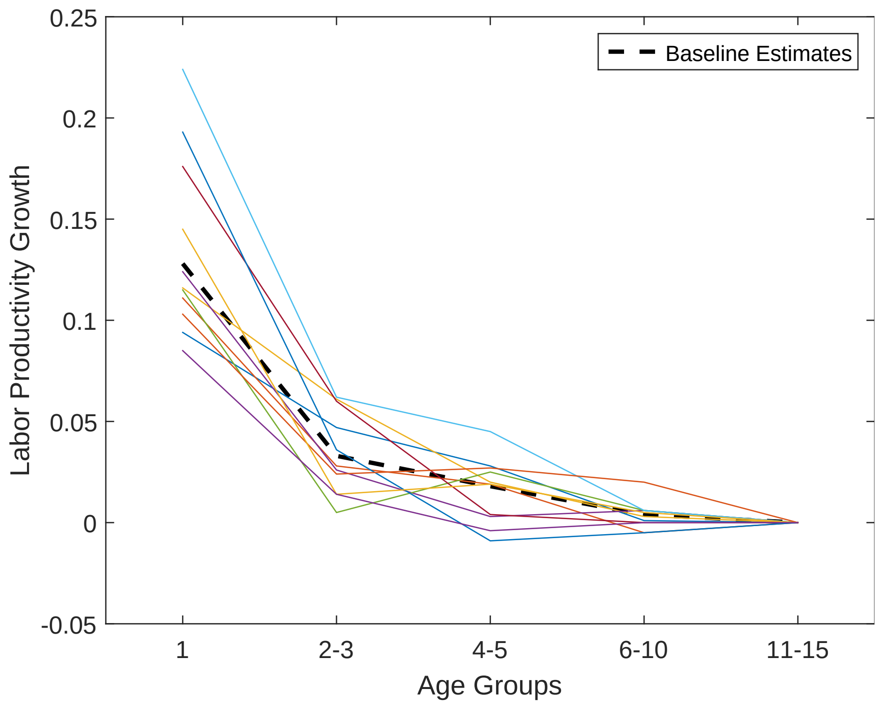

Figure 11: Combined net age-productivity growth profile by sectors

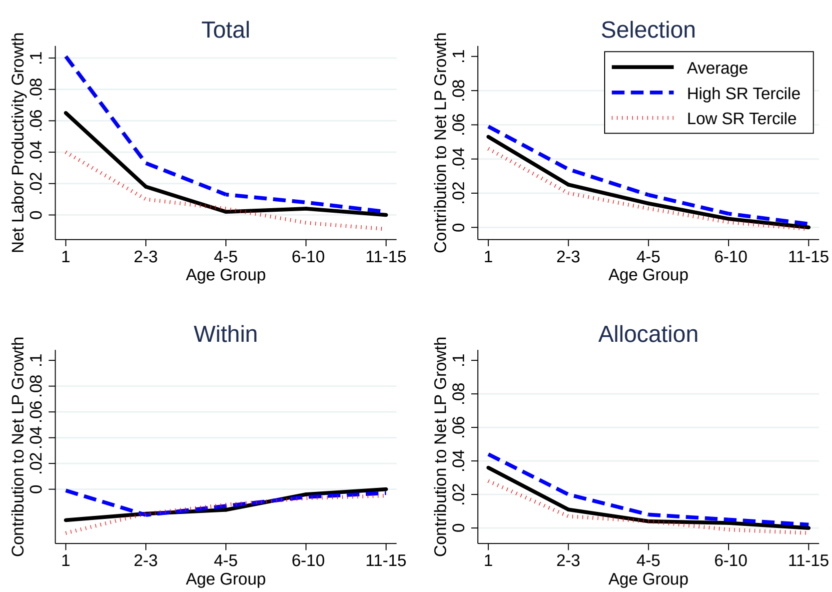

Figure 12: DOP decomposition by firm age: top versus bottom tercile of startup rates by CBSA

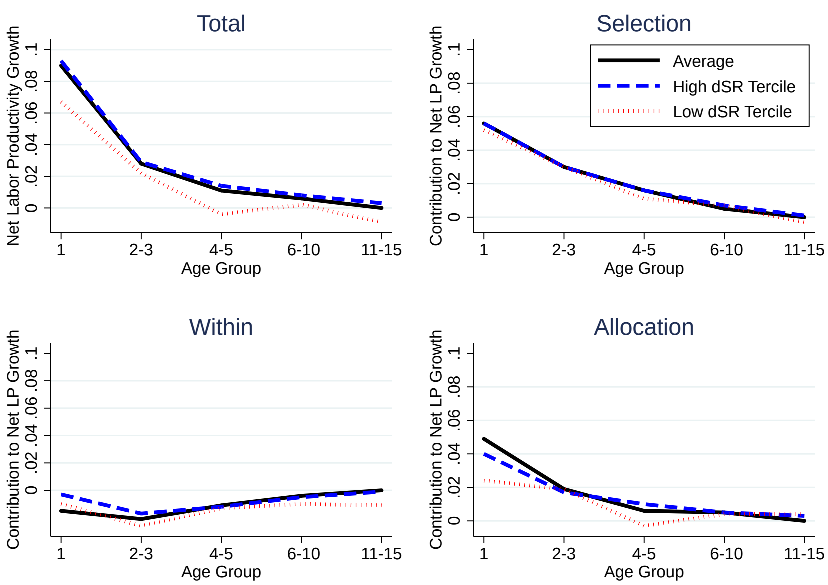

Figure 13: DOP decomposition by firm age: top versus bottom tercile of *change* in startup rates by CBSA

Table 1: Change in average productivity growth of mature (age 16+) firms from 1996-2004 to 2005-2012

	Change in Average Mature (Age 16+) Firm Industry Productivity Growth $\Delta\Phi_{16+jt}$			
	Total	Within	Allocation	Selection
Late Period	-0.023	-0.006	-0.019	0.002
(2005-2012)	(0.018)	(0.004)	(0.017)	(0.002)

Note: U.S. Census Bureau Revenue Enhanced Longitudinal Business Database, 1996-2000 and 2003-2012. OLS regression of average productivity growth of 16+ group on late period (2005-2012) dummy and NAICS4 industry fixed effects. Weighted by average industry employment for all years. Standard errors clustered by industry.

Table 2: Productivity growth and startup rates: regional and industry evidence

	Industry Productivity Growth $\Delta\Phi_{jkt}$		GSP/Worker Growth		MSA Productivity Growth	
	(1)	(2)	(3)	(4)	(5)	(6)
Startup Rate	0.796***	1.46**	0.22**	1.71***		
	(0.209)	(0.620)	(0.108)	(0.480)		
Δ Startup Rate					2.801***	6.089**
					(0.908)	(2.603)
Fertility IV	No	Yes	No	Yes		
Collateral IV					Yes	Yes
Year FE	Yes	Yes	Yes	Yes	No	No
NAICS4 FE	Yes	Yes	Yes	Yes	No	No
State FE	Yes	Yes	Yes	Yes	No	No
	OLS	IV	OLS	IV	IV	IV
	State x Ind	State x Ind	State x Ind	State x Ind	MSA (All Ind)	MSA (All Ind Ex-Con and NT)
R^2	0.097	0.107			0.042	0.553
N	20000	200000	1200	1200	300	300
F-test statistic		>10		10.35	105.7	21.7

Note: U.S. Census Bureau Revenue Enhanced Longitudinal Business Database, 1996-2000 and 2003-2012. Columns (1-2): OLS and IV regression of average change in log net receipts per worker by industry and State on startup rate by area with year, NAICS4 digit industry and area fixed effects. Columns (3-4) OLS and IV regression of BEA Annual GSP/worker growth for years 1980 to 2007. Columns (5-6) We regress the change in log net receipts per on the change in the startup rate by CBSA over the time period 2000-2006 using the Charles, Noto, and Hurst (2016) housing demand instrument. The time period was chosen for comparability. Column 5 includes all industries within an MSA; column 6 excludes construction and non-tradable industries. Standard errors clustered by area. Number of observations are rounded to the nearest thousand.

Table 3: Counterfactual cumulative percent difference in productivity growth by sector

	Lifecycle	Cohort	Total
Construction (23)	−2.42	10.96	8.55
Manufacturing (31-33)	−0.29	5.64	5.36
Wholesale Trade (42)	−1.24	8.75	7.51
Retail Trade (44-45)	−0.71	7.27	6.56
Transport and Warehousing (48-49)	−0.65	6.30	5.65
Information (51)	0.15	−0.85	−0.70
Finance, Insurance, Real Estate (52-53)	−0.31	1.85	1.54
Professional and Admin Services (54-56)	−1.10	12.06	10.96
Education and Healthcare (61-62)	0.02	0.20	0.22
Entertainment and Accommodations (71-72)	−0.48	5.80	5.32
Other Services (81)	−0.16	0.64	0.48

Note: U.S.C̃ensus Bureau Revenue Enhanced Longitudinal Business Database. We apply the decomposition described by equation (6) individually to each two digit sector. We report the cumulative difference in productivity growth for each component and the total from 1996 to 2012.

A Additional Tables and Figures

Table A0: Estimates of trend growth in entrant productivity Φ_{Et} 1997 to 2012

	$\Phi_{Et} = \nu_i + \eta t +_{it}$				$\Delta\Phi_{Et} = \eta +_{it}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Linear Trend	0.00650***		0.00932***					
	(0.00122)		(0.00114)					
1997-2004		0.0207***		0.0306***				
		(0.00521)		(0.00490)				
2005-2012		0.00567***		0.00810***				
		(0.00125)		(0.00117)				
Stochastic Trend					0.0131		0.0146	
					(0.0119)		(0.00991)	
1997-2004						0.0233		0.0518***
						(0.0227)		(0.0190)
2005-2012						0.00925		0.000603
						(0.0140)		(0.0116)
Weighted	No	No	Yes	Yes	No	No	Yes	Yes
N	3,594	3,594	3,594	3,594	3,032	3,032	3,032	3,032
R^2	0.008	0.010	0.018	0.024	0.000	0.000	0.000	0.002

Note: U.S. Census Bureau Revenue Enhanced Longitudinal Business Database, 1996-2000 and 2003-2012. Columns (1) to (4) are OLS/WLS regressions of industry entrant productivity (relative to its value in 1997) on a deterministic linear time trend or piecewise linear trend. Columns (2) and (4) are weighted across industry by the average industry employment computed over the entire sample 1997-2012. Columns (5) to (8) are OLS/WLS regressions of the change in industry entrant productivity on a constant, or time period dummy variable. Columns(6) and (8) are weighted by the industry average employment. Standard errors clustered by industry.

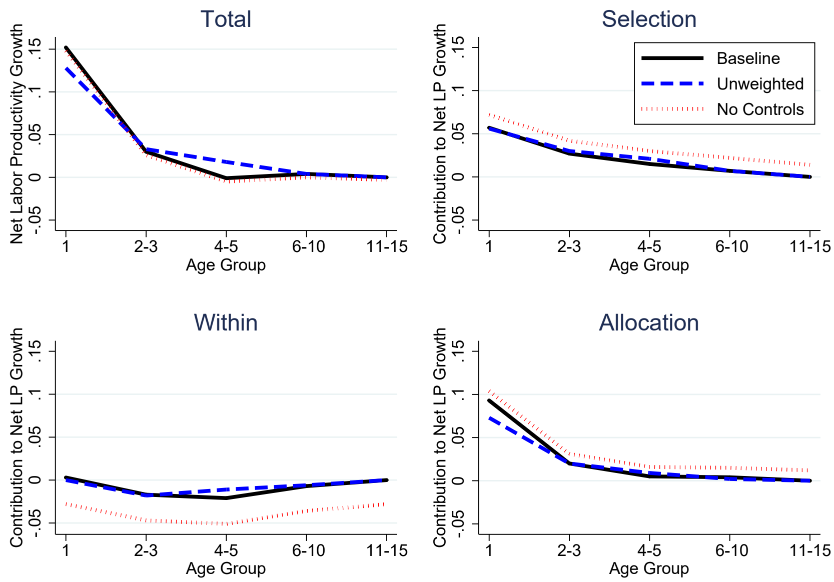

Figure A1: DOP Decomposition by Firm Age: Robustness

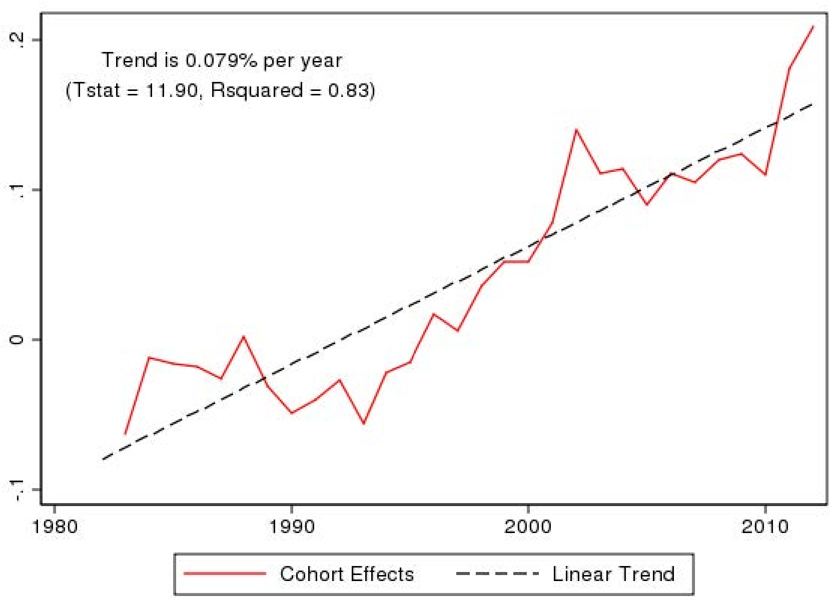

Figure A2: Estimated Cohort Effects

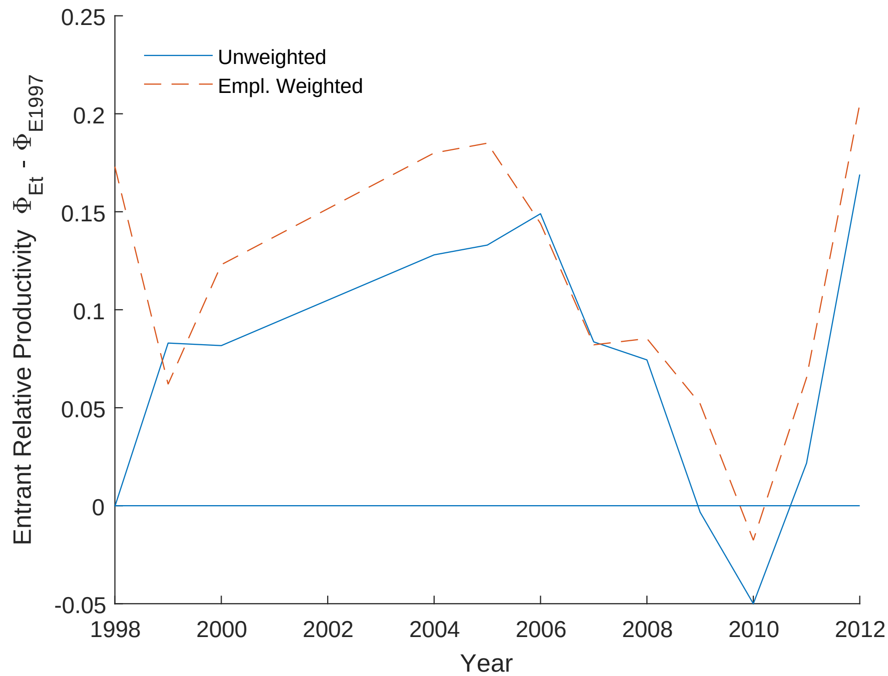

Figure A3: Estimated Cohort Effects

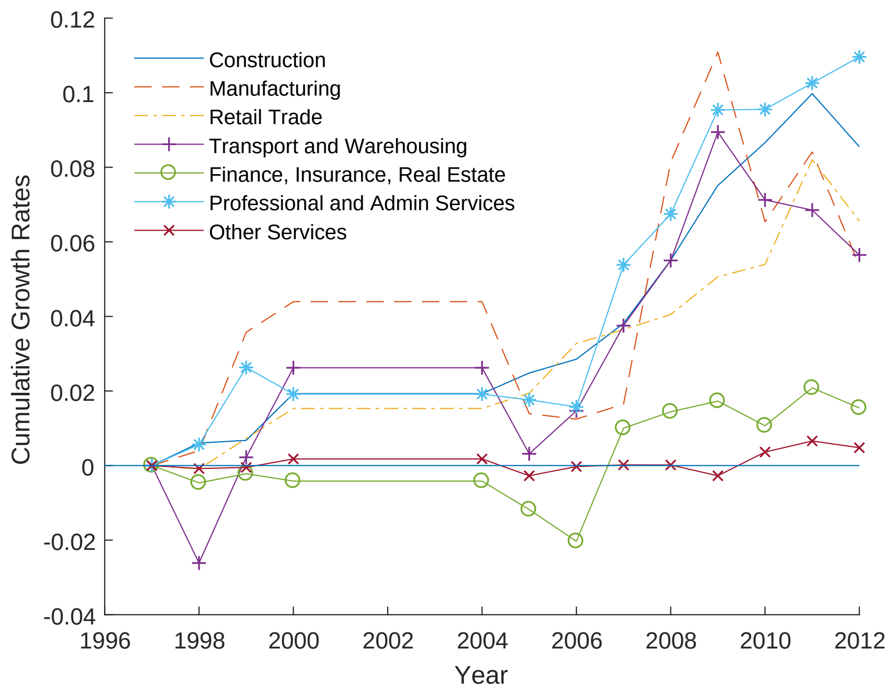

Figure A4: Difference in cumulative productivity growth counterfactual by sector

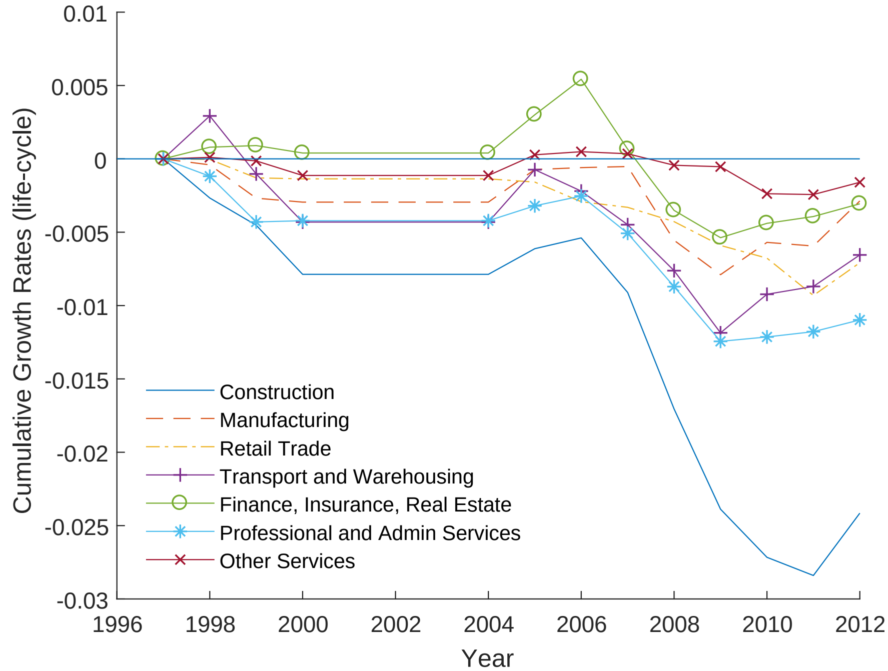

Figure A5: Difference in cumulative productivity growth counterfactual by sector from life-cycle component

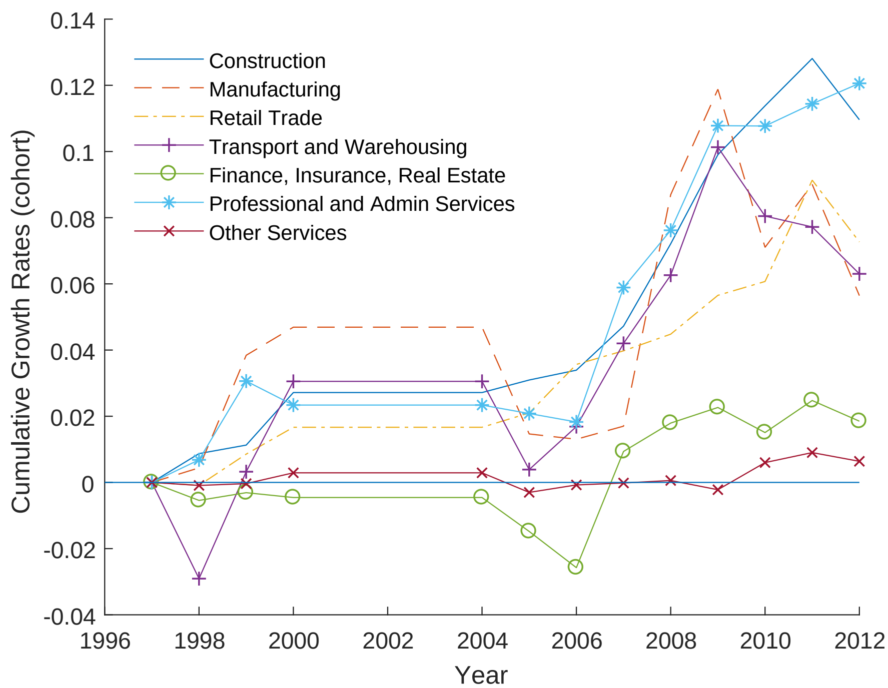

Figure A6: Difference in cumulative productivity growth counterfactual by sector from cohort component

B Data Appendix

To construct our empirical counter-factual we rely on public-use data from the BDS on employment by firm-age going back to 1980. We choose to define our age categories consistently with the BDS categorization, as is generally done in the literature. Our counter-factual age groups are therefore ages 1,2,3,4,5,6-10,11-15, and 16+. In evaluating the cohort trend contribution, we assign each binned group its midpoint (i.e. 8 for group 6-10). For the top censored age group, we assign an age of 23, which is roughly the midpoint for this category in 2014, the year where the censoring issue is least severe.

The major issue with the BDS data is that age is censored for firms born before 1977 and so in early years we do not observe the distribution of employment activity across older firms which is necessary to construct our counter-factual. In addition, age bins employment levels are often introduced once the youngest age in the bin is observed and so these numbers are partly censored in the first few years after they are reported, leading to large fluctuations in reported levels.

To address these issues, we appeal to the fact that employment re-allocations across the age distribution is a persistent and slow moving process. This allows us to represent changes in employment by age groups as a long tailed moving average process. To fill in censored observations, we begin in the first year where an age group is fully observed and calculate the average annual growth in employment 5-years forward. We then use the growth rate to roll back the process and fill in censored years. We iterate the process by allowing the 5-year window to roll back as we fill in censored observations. Specifically, for the employment level of age group a in year t we use

$$\hat{E}_{at} = \frac{E_{a,t+1}}{1 + \bar{g}_{at}}$$

where $1 + \bar{g}_{at}$ is the arithmetic mean of the growth rates in employment for age group a in periods $t + 1$ through $t + 5$. Furthermore, to ensure that aggregate employment levels stay the same, we define the top censored age category employment level as the residual employment after we've estimated all the younger age groups. We apply this process to the 11-15 age bin for years 1980-1991, for the 6-10 age bin for years 1980-1986, and for age five firms for 1980-1981 and for age 4 firms in 1980 only. Given that most movement in the age-productivity is concentrated in year 0-3 we choose not to roll back our estimates before 1980 so as to avoid imputing any values for these age groups. We then convert all employment numbers into shares to get the distributional changes necessary to evaluate the counter-factual equation 6. The imputations effect less then 7% of our data on shares and we conduct several robustness tests to confirm that the estimates are sensible and that our results are not unduly sensitive to the choice of window or imputation process.

The resulting employment shares used in the counter-factual are shown in Figure A7. The observations after the dashed vertical line correspond to raw BDS data and the observation preceeding are imputed contain MA estimates of the shares as described above. As a robustness, we also explore MA processes ranging from 3-6 year windows and find our final results are not sensitive to these changes.

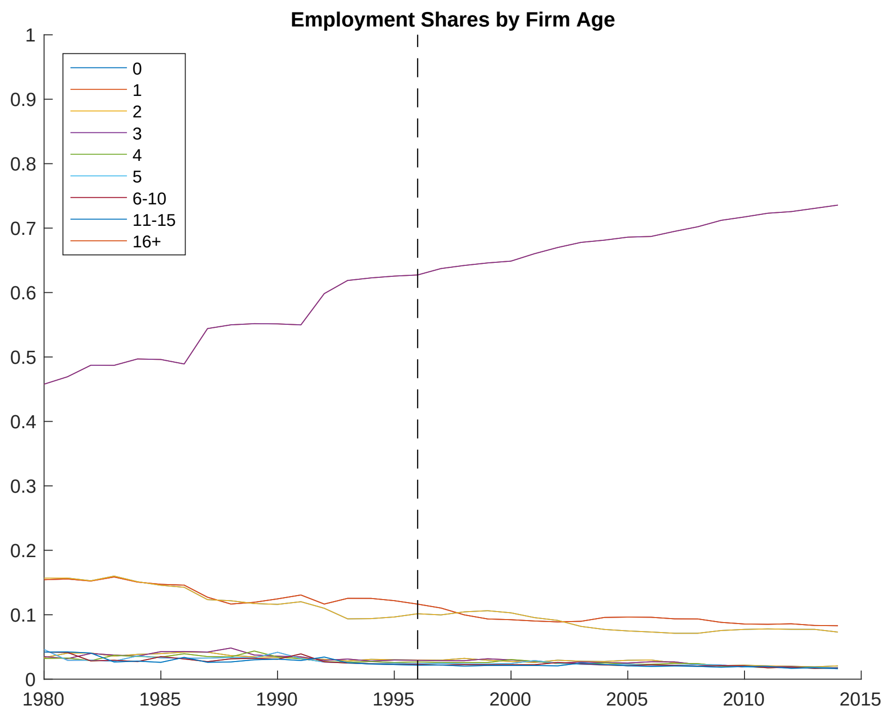

Figure A7: Employment Shares by Age Group

C An Illustrative Model

In this section we show how a small modification to the workhorse heterogenous firms model of Hopenhayn (1992) is able to generate life-cycle growth rates in labor productivity that are roughly consistent with our empirical findings. The issue with using the standard Hopenhayn (1992) set-up is that there exists no dispersion in labor productivity across operating firms due to the decreasing returns to scale production technology. To adapt the model to our purposes then, we replace the assumptions of decreasing returns to scale and competitive markets with a constant returns to scale production technology and monopolistic competition. Monopolistic competition is necessary when linear production technologies are used to guarantee non-trivial distributions of economic activity across firms in equilibrium. The change leaves a model that is still tractable and easy to compute while generating a non-degenerate distribution of labor productivity across firms. We keep the model intentionally simple to highlight the fact that a very standard model of firm dynamics is able to generate our main empirical findings.[29]

We assume there is a single consumption good produced by a competitive final goods sector which aggregates all varieties of intermediate inputs subject to a constant returns to scale technology:

$$Y = \left[\int_{j\in\Omega} Y_j^{\gamma}\right]^{\frac{1}{\gamma}}$$

where $\gamma = 1 - \frac{1}{\sigma}$ is function of elasticity of substitution, σ.

Each intermediate good j is produced by a monopolistically competitive firm which chooses production to maximize operating profits conditional on deciding to produce output:

$$\pi_j(A_j, L_j) = \underset{L_j}{max} \;\; P_j(Y_j)Y_j - wL_j - c_f$$

Here A_j denotes the firm specific stochastic productivity which we assume follows an $AR(1)$ process $log(A') = a_0 + \rho_A log(A) + \epsilon$ and L_j is hired labor. The production technology of all intermediate producers is linear so that $Y_j = A_j L_j$. To stay in operation, each firms must pay a flow fixed cost c_f each period it produces.

In contrast to standard competitive industry models where labor productivity is constant across firms[30], these assumptions imply that there will be dispersion in physical revenue per worker, $\frac{Y_j}{L_j}$ driven by variation in A_j. Because the model is frictionless, the model does not generate any variation in revenue productivity per worker, $\frac{P_j * Y_j}{L_j}$, because optimization implies marginal revenue product of labor is equal to the (common) marginal cost of labor. Our preferred interpretation

[29] This simplicity is not without loss of generality. Our model generates dispersion in physical productivity per worker $\frac{Y_j}{L_j}$ but not revenue per worker $\frac{P_j * Y_j}{L_j}$ thus there is some tension between our empirical results, which show dispersion in revenue per worker, and our model results. For our purposes we do not think this is a big concern because a) the existing empirical evidence suggests these measures all highly correlated at the firm/plant level and b) one can induce dispersion in both labor productivity measures by introducing additional frictions such as overhead labor or adjustment costs (see Bartelsman, Haltiwanger, and Scarpetta (2013) for a recent example), without affecting our main result. We leave these extensions to future work.

[30] It is proportional to the common output price or wage, whichever is not the numeraire.

is that our empirical age-productivity profile mostly reflects underlying variation in physical labor productivity across firms[31], nonetheless, we readily acknowledge that our controls are imperfect so clearly some of this variation is driven by within industry differences in prices.

That being said, we believe that revenue labor productivity is a useful proxy for physical labor productivity because empirically these two labor productivity concepts are highly correlated at the plant/firm level, and, moreover, they both predict plant/firm outcomes such as employment growth. Still, the fact that our model cannot generate dispersion in both labor productivity measures is a limitation of our model. We think that a simple extension of it, such as adding overhead labor (Bartelsman, Haltiwanger, and Scarpetta (2013)) would generate this dispersion without changing our interpretation so one interpretation is that we are assuming that one of these mechanisms is implicitly. present. Due to our desire to highlight the essential mechanisms that drive the age-productivity profile (selection and reallocation), we leave the actual incorporation of these features to future work.

Consider the problem of a firm who is choosing whether to continue operating or whether to shut down. Letting $\pi^*(A_j)$ denote optimal operating profits at $L_j^*(A_j)$, we can write the firm's value function as:

$$V(A) = \pi^*(A) - c_f + \beta max\{\mathbb{E}_A V(A')\ ,\ 0\}$$

which captures the endogeneity of the exit decision. Exiting itself is an absorbing state for firms and so if a firm chooses to leave the market they receive a continuation value of 0 for all time thereafter.

Each period there is a mass J of potential entrants of which E actually decide to enter. Before entering, firms must pay an entry fee c_E to get an initial productivity draw from a stationary productivity distribution $\varphi(s) = be^{-bs}$ and then they can decide whether to produce or exit immediately in the first period. If they do not enter, they earn a zero payoff forever. If they enter, their problem is identical to the production decision of an incumbent firm which faces shock draws A and currently employs no workers. Because the entry decision is made before the idiosyncratic shock is drawn, some entrants choose to exit immediately after receiving their initial shock draw. Free entry implies that the expected value of entering is equal to the cost of entering:

$$\mathbb{E}V(A) - c_E \geq 0$$

To close the model, we assume that there is a single unit of labor supplied inelastically in competitive labor markets by a unit mass of households. We define a *stationary recursive competitive equilibrium* in our model as consisting of a (i) value function $V(A)$, (ii) policy functions $X(A)$ and $L(A)$, (iii) A wage w, incumbent measure μ, and entrant measure M such that

1. Optimality: $V(A)$, $L(A)$, and $X(A)$ solve incumbent's problem

[31] Our baseline empirical specification uses real revenue per worker, our estimates use log differences and includes time and industry fixed effects in the log difference specification designed to soak up much of the within industry price variation.

2. Labor Market Clearing

$$1 = \int L(A)d\mu + \int L(A)d\varphi$$

3. Measure of Actual Entrants: $\forall t \geq 0$,

$$M = J \int [1 - X(A)]\, d\varphi$$

4. Model Consistent Dynamics $T(\mu, J)$

$$\mu = T(\mu, J) = \int \int [1 - X(A)]\, d(A'|A)d\mu + J \int [1 - X(A)]\, d\varphi \tag{9}$$

C.1 Calibration and Simulation

Our model has six parameters for calibration: $a_0, \rho_a, \sigma_\epsilon, b, c_e, c_f$. To fit these, we follow the literature in setting calibration target to match the distribution of activity and size of firms in the BDS. Specifically, we choose parameters to match twenty-two moments: size distribution of incumbent firm (5 moments), distribution of incumbent employment shares (5 moments), size distribution of entrants (5 moments), distribution of entrant employment shares (5 moments), average size of new entrants, and the exit rate. The results are shown in A7. The best fit parameters that fit our model are $(a_0, \rho_A, \sigma_\varepsilon, b, c_e, c_f) = (0, 0.95, 0.40, 0.25, 15, 12.87)$ and provide a reasonably good match of moments in the data. footnoteWe also tried calibrated to match the productivity profiles directly. This approach gave similar qualitative results.

We now use our calibrated model to simulate life-cycle profiles for firms in the stationary equilibrium and to assess whether the dynamics of our model are consistent with our empirical findings. To do so, we calculate the stationary equilibrium and then simulate 500 paths of life-cycle labor productivity growth for a cohort of firms and average them. For each simulation we calculate the gains in physical labor productivity by cohort at each stage in life and then also calculate the associated DOP components contributing to life-cycle growth. Again, notice that we are implicitly assuming that our empirical age-productivity profile reflects underlying variation in true, physical labor productivity and does not sole reflect variation in prices across firms.[32]

Figure A8 contains the result of the exercise. What is clear from the result is that our modified version of Hopenhayn (1992) is able to qualitatively replicate most of the empirical patterns. In particular, it generates a sharply declining and convex patter for the net growth in labor productivity as a function of age as well as the empirically correct contribution patterns for the allocation and selection effect components.

The positive allocation effect comes from two facts: entrants are on average less productive than survivors and selection. By definition, a positive allocation effect means that on average firms that are higher productivity, gain market share. In our model firms that receive positive shocks are the

[32]As noted before, we think this is a reasonable assumption. Still, without detailed data on prices it is impossible to know for sure.

Estabs Size	Share of incumbents Data	Share of incumbents Model	Employment share Data	Employment share Model	Share of entrants Data	Share of entrants Model	Employment share Data	Employment share Model
1-19 emps	0.832	0.781	0.175	0.115	0.928	0.924	0.436	0.400
20-99 emps	0.134	0.168	0.232	0.256	0.065	0.070	0.323	0.396
100-499 emps	0.023	0.037	0.183	0.265	0.006	0.006	0.151	0.153
500-999 emps	0.005	0.008	0.123	0.148	0.001	0.000	0.054	0.033
1000+ emps	0.001	0.006	0.277	0.216	0.000	0.000	0.036	0.018
Avg size of entrants	7.40	7.53						
Exit rate	0.087	0.084						

Source: Business Dynamic Statistics. Average of annual, 1977-2014

Table A7: Calibration Targets

ones that grow and those that receive negative shocks are the firms that shrink or exit. For young cohorts, many firms are near the exit threshold, so conditional on surviving they more than likely received a positive shock. Thus, conditional on survival, the fraction of positive shocks is greater than the fraction of negative shocks leading to positive allocation. The reason the allocation effect dies off is that for older cohorts the distribution of shocks is more symmetric because the mean productivity level is far away from the exit threshold. The presence of a selection effect is easier to explain. Many firms start near the exit threshold and the least productive ones exit. Over time, fewer firms are near the threshold so the selection effect becomes weaker.

The one component we cannot match with our baseline model is the within effect; the model suggests this effect is convex and large while the empirical data suggests it is nearly linearly and mostly flat. The model has a hard time generating a small within effect because our calibrated shock process is not persistent enough. For young firms, the within effect is positive because firms that survive likely received a positive shock. For old firms, the opposite is true. The reason is that while the unconditional distribution shocks is symmetric, the distribution of shocks is not conditional on selection. If a firm survives long enough to become old, that means the firm is likely large. Given mean reversion in the shocks, this means that that more likely than not the firm will shrink next period leading to a negative within effect. Overall, the model performs quite well suggesting that our facts about the age profile of labor productivity growth can easily be generated with standard mechanisms.

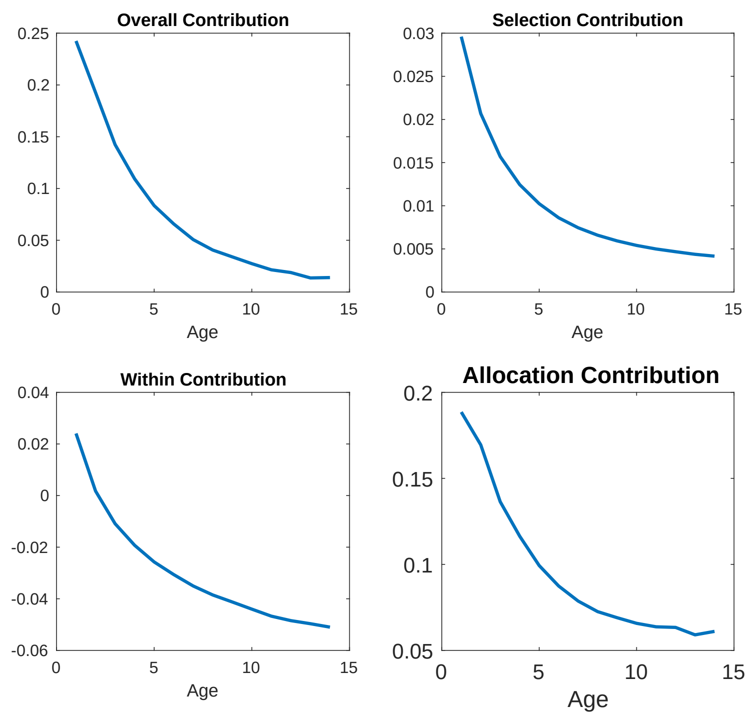

Figure A8: Age-Productivity Profile in the Model