Older and Slower: The Startup Deficit's Lasting Effects on Aggregate Productivity Growth[*]

Titan Alon
Northwestern
titanalon@gmail.com

David Berger
Northwestern and NBER
david.berger@northwestern.edu

Robert Dent
New York Fed
robcdent@gmail.com

Benjamin Pugsley
New York Fed
benjamin.pugsley@gmail.com

April 16, 2017

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. We apply this framework to the newly developed Revenue-enhanced Longitudinal Business Database, ReLBD, and document a convex age-productivity profile with the largest growth from new and young employers. For firms with uncensored measure of age, the profile is stable. The shape of the profile as well as a decline in the gains from reallocation among firms born prior to 1979 imply significant drag in annual productivity growth from the "startup deficit". Using an instrumental variables strategy we find a consistent pattern across states in the U.S. The patterns are consistent with a standard model of firm dynamics with monopolistic competition.

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 from Pugsley and Sahin (2014) illustrate the trends. The magnitude of the reallocation is quite startling and even surpasses those flows documented in the structural transformation literature from manufacturing to services over the same period.[1] These patterns are also widespread across industries and geographic markets, suggesting they are independent of any compositional variation in economic activity.[2] Following Pugsley and Sahin (2014), we refer to the persistent and widespread collapse in startup rates and the implied aging of U.S. businesses as the startup deficit.

Recently, economists and policy makers have begun questioning how structural changes like the startup deficit, may be impacting the health of the aggregate economy, and the economic literature suggests a number of channels through which it could have significant macroeconomic consequences. 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 the opening of 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 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 lifecycle of the firm (Arkolakis, Papageorgiou, and Timoshenko (2014), Pugsley and Sahin (2014), and Pugsley, Sedáček, and Sterk (2016)). 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.

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. Even more recently, Decker, Haltiwanger, Jarmin, and Miranda (2016) and Decker, Haltiwanger, Jarmin, and

---

[1]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.

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

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. 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 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 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 uses the results directly and shows how, under some empirically plausible assumptions, the age-productivity profiles we estimate can be linked directly to aggregate productivity growth. While this exercise allows us to transparently quantify the significance of our results in a macroeconomic sense, it does not admit any causal interpretation. We therefore compliment these results with a set of Cross-Sectional regressions that exploit plausibly exogenous variation in startup activity across detailed geographic and industrial markets.

Our results suggest 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. Furthermore, after roughly 15 years the average firms begin to exhibit a systematic tendency toward negative labor productivity growth. This latter observation is economically meaningful and underlies the ultimate drag on aggregate labor productivity growth associated with an aging business sector. Quantitatively, our results suggest from 1980 to 2014 that the startup deficit may have reduced aggregate productivity by a little more than 4 percent, or roughly 0.12 percentage points per year. If a consequence of the declines in entry, more recent declines in the productivity gains from reallocation among the oldest firms since 2004 could subtract another 2 percentage points.

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 4 presents the battery of robustness tests we conduct on our final estimates. Section 3 presents the empirical results and quantifies the macroeconomic significance of our findings. Section 5 presents the cross-sectional evidence and our IV approach. Section 7 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 the 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 seek to estimate this relationship non-parametrically using rich panel variation in U.S. Census data encompassing the entire 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) 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 establishment level. These longitudinal records may be used to calculate measures of employment growth, entry, exit and establishment age. The Census Bureau, through data gathered in its annual Company Organization Survey and quinquennial Economic Census also provides a firm identifier for each year that groups establishments at the highest level of operational control.[3] Following Davis, Haltiwanger, Jarmin, and Miranda (2006) and Haltiwanger, Jarmin, and Miranda (2013) we assign an establishment age 0 in the year it hires its first employee. We then assign each firm, which be be 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. We then merge revenue records from the Census Bureau BR following the methodology of Haltiwanger, Jarmin, Kulick, and Miranda (2016b).[4] 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

---

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

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

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.[5] Further details on the construction of the dataset and cleaning process will be available in the data appendix for this paper. 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 $\phi_{it} = R_{it/E_{it}}$ and aggregate (industry j) 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. To describe the age productivity profile, we measure productivity by age group. We extend the Dynamic Olley-Pakes (DOP) decomposition of Melitz and Polanec (2015) to measure the changing sources of productivity growth over firms' lifecycles. 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 (s) that will survive to period t and those firms that will exit (x) before then such that (suppressing the industry index j)

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

where s_{iat} for $i \in \{s,x\}$ represents the share of firms i *within* cohort a (in industry j) at time t. 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

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

---

[5]Our baseline specifications use nominal revenues deflated by the implicit price deflator for GDP, but also includes industry and time fixed effects capture some of the variation of industry prices. In the robustness section, we explicitly deflate our series using the most comprehensive set of price indices that are available.

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[6] 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. The final term represents the contribution of *selection* and contributes positively to cohort 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 exploit panel variation in the rich cross-section of 4-digit NAICS industries available in the Census data. Our main identifying assumption is that we can express an age group's productivity growth as

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

where $\mathbb{E}[\varepsilon_{ajt}|a, j, t] = 0$. We also apply the same decomposition into time, industry and age effects to each component of equation (1). Given this identifying assumption, we are able to semi-parametrically estimate the age-group profiles by projecting age group productivity growth, $\Delta\Phi_{ajt}$, as 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 a industry average employment share, in order to hold industry composition constant. The estimated coefficients $\hat{\delta}_a$ on a full set of age group dummies provide semi-parametric estimates of the age-productivity profiles. Note that the specification is in terms of labor productivity growth and so 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.

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.

Clearly, the differencing specification above will not be suitable for new entrants. While this

---

[6]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.

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:

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

where we interpret η as a common trend in the productivity of new entering cohorts.

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:

$$\Delta\Phi_t = \sum_a^A s_{a,t}\Phi_{a,t} - \sum_a^A s_{a,t-1}\Phi_{a,t-1}$$

$$= \sum_a^A (s_{a,t} - s_{a,t-1})\Phi_{a,t} + \sum_a^A s_{a,t-1}(\Phi_{a,t} - \Phi_{a,t-1}) \quad (5)$$

$$= \sum_a^A (s_{a,t} - s_{a,t-1})\Phi_{a,t} + \sum_a^A s_{a,t-1}(\Phi_{E,t-a} - \Phi_{E,t-a-1})$$

where the third equality follows from plugging in equation (4) and the fact that the set of age groups is exhaustive so shares sum to 1 within each period. 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 $t - a - 1$.

Estimating differences in the productivity of entrants poses several conceptual challenges, such as comparing productivity *levels* across industries, as well as identification challenges, in identifying the common levels independently of aggregate time effects. To circumvent these issues we take a simplified approach and assume that the differences in the productivity of entering cohorts is well represented by a time trend, η so that $\Phi_{E,t} - \Phi_{E,t-1} = \eta$ for all t. This is precisely the estimation approach we use, along with additional controls for industry heterogeneity, in equation (3). This simplification, while convenient, can be considered consistent with balanced growth path equilibrium concepts common in macroeconomic models.

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 $s_{a,t}^{cf}$, holding constant time and industry effects, by:

$$\mathbb{E}[\Delta\Phi|s_{a,t}^{cf}] - \mathbb{E}[\Delta\Phi|s_{a,t}] = \sum_a^A (\Delta s_{a,t}^{cf} - \Delta s_{a,t})\Phi_{a,t}$$

$$= \sum_a^A (\Delta s_{a,t}^{cf} - \Delta s_{a,t})\sum_{j=1}^a \delta_j + \sum_a^A (\Delta s_{a,t}^{cf} - \Delta s_{a,t})\Phi_{E,t-a} \quad (6)$$

$$= \underbrace{\sum_a^A (\Delta s_{a,t}^{cf} - \Delta s_{a,t})\sum_{j=1}^a \delta_j}_{\text{Lifecycle}} - \underbrace{\eta\sum_a^A (\Delta s_{a,t}^{cf} - \Delta s_{a,t})a}_{\text{Cohort}}$$

where the second equality uses equation 4 and the third equality uses the linear trend in entrants and the fact that age groups are exhaustive.[7] The first term captures the contribution of lifecycle

---

[7]To see more clearly where the third line comes from note that $\Phi_{E,t-a} = \eta(t-a) + \Phi_{E,0}$ so that $\sum_a^A (\Delta s_{a,t} - \Delta s_{a,t}^{cf})\Phi_{E,t-a} = \sum_a^A (\Delta s_{a,t} - \Delta s_{a,t}^{cf})\eta(t-a) + \sum_a^A (\Delta s_{a,t} - \Delta s_{a,t}^{cf})\Phi_{E,0} = -\eta\sum_a^A (\Delta s_{a,t} - \Delta s_{a,t}^{cf})a$. The last step

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 lifecycle. Using the ReLBD pooled across all years (1996-2001 and 2003-2012) 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.[8]

Overall, several 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 fives

---

follows from the fact that the age groups are exhaustive and so shares sum to one.

[8]In appendix figure A2 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.

years and nearly the entire effect disappearing after ten; and (iii) the profile suggests that while young firms register large and consistent positive rates of productivity growth at early stages of life, firms at the older end of the age distribution register marginally negative productivity growth. This pattern, combined with equation (5), suggests a hump-shaped response for aggregate productivity growth in response to the startup deficit and aging of the business sector. Initial declines raise aggregate labor productivity by shifting more economic activity to youngish firms that have already survived early years of competition and registered the steepest gains of the lifecycle profile. However, as time passes the deficit will push a greater mass of economic activity into the negative tail of the profile dragging down aggregate productivity growth. The net long-term effect then will depend on the relative magnitudes of these effects and how economic activity gets redistributed in the long run.

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 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.

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 occured and instead the entry rate and age distribution had remained at 1980 levels. 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. We then feed through the historical evolution of

employment shares by age derived from the Census Bureau's Business Dynamics Statistcs (BDS) database[9] 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 4.15% 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, just over 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 $2,200 higher had the startup deficit never occurred.[10] 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 1% by reallocating activity from less productive entrants to more productive, older incumbent firms.

In addition to differences over the lifecycle 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 just over 5% 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 productivity growth by reducing the relative market share of older more productive businesses.

---

[9]See the data appendix for details on the BDS data preparation and use.

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

Our main results only consider firms up to age 15 because we can observe this age group for all years in the ReLBD.[11] 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 5 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.[12] 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.[13]

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 2 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 2 reveals that the decline they describe is driven by the oldest firms, which constitute more than 70 percent of total employment.

[INSERT TABLE 2: DOP 16+ GROUP ABOUT HERE]

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

---

[11]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.

[12]Strictly speaking, less than the 11-15 group, because we condition on year and industry fixed effects. However, note from figure A2 without controls that the growth of the 11-15 group is almost exactly zero.

[13]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

age group: the age-productivity profile may be stable, but employment is shifting further along 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.

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 test 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 cycle 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 11 and 12 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 show in figure 6. 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 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.

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 8 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 9 and 10 summarize the robustness tests for industry composition. Figure 9 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 this industry groups and does not deviate significantly in curvature or magnitude. Figure 10 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.

5 Cross-sectional Evidence

As an alternative to our industry labor productivity growth decompositions, we explore a different source of variation to reveal the relationship between startups, aging firms and productivity growth.

Here we make no assumptions about a stable age-productivity growth profile, and instead ask whether areas with relatively higher startup rates and younger firms also exhibit faster productivity growth, exploiting rich geographic variation within the ReLBD. On its own, this exercise would raise significant concerns about reverse causality: startup rates could be elevated because of local innovations to productivity. To address this possibility, we adopt two different instrumenting strategies to generate plausibly exogenous shifts across areas in the level of startup activity, and we find a statistically and economically significant link between the pace of business creation and productivity growth.

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:

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

Table 3 below reports the estimated $\hat{\beta}$ for this specification from OLS regressions. Columns (1) and (2), which estimate the model off cross-state variation in the startup rate (equally weighted and employment weighted, respectively) show a strong correlation between states with relatively high startup rates and productivity growth within industries in those states. Columns (3) and (4) report similar estimates instead using cross-CBSA variation in startup rates. In both cases, standard errors are clustered at the unit of geography, allowing for serial correlation within state or CBSA. These OLS estimates confirm that states with relative increases in the entry rate are also ones with relative increases in industry productivity growth. A roughly 1 percentage point decline in the startup rate would predict a slightly smaller sized decline in industry productivity growth within the area.

[INSERT TABLE 3: OLS REGRESSIONS ABOUT HERE]

At the state level, using real gross state products published by the Bureau of Economic Analysis, we can also measure the correlation of within state changes in entry and gross output per worker growth within the state. This measure is closer in spirit to the aggregate nonfarm business sector productivity growth for the U.S. Column (6) of table 3 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. Standard models of firm dynamics with free entry imply that changes in the growth rate of the labor supply require equilibrium shifts in the market's startup rate. 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.

Table 4 column (4) reports the estimates of equation (7), where SR_{kt} is instrumented with 20 year lags of the state's fertility rate. When a state's startup rate is increasing because of demographics and not current business conditions, industry productivity growth increases. This elasticity is larger than the one estimated in the reduced form results. Standard errors are again clustered at the state level. Although the standard error only makes this elasticity significant at the 10 percent level, when we add additional demographic instruments based on the same mechanism, the results are again positive and statistically significant at a smaller level.[14] We can also apply the same instrument to the measures of a state's growth in gross state product per worker. Here we find a economically and statistically significant elasticity of roughly 1.5 to 2, similar to the elasticity for industry productivity growth, implying that exogenous shifts in a state's startup rate could lead to significant shifts in the state's growth in output per worker.

[INSERT TABLE 4: IV REGRESSIONS ABOUT HERE]

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

---

[14]These results in columns (5) and (6) have not yet been released by the Census Bureau and we report only their sign and significance.

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). 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 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[15]. By restricting attention to the tradeable 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.

COLLATERAL IV RESULTS IN PROGRESS

Overall, without appealing to the identifying assumptions embedded in our decomposition based account of the link between labor productivity growth and new business creation, we see robust evidence that increases in entry, all things equal lead to increases in productivity growth. It is important to note 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   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 lifecycle growth rates in labor productivity that are 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.

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 technol-

---

[15]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.

ogy:

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

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:

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

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.

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

2. Labor Market Clearing

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

19

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

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

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

$$\mu = T(\mu, J) = \int \int [1 - X(A)] \, d(A'|A) d\mu + J \int [1 - X(A)] \, d\varphi \tag{8}$$

6.1 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 1. 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 lifecycle 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 lifecycle labor productivity growth for a cohort of firms and average them. For each simulation we calculate the gains in labor productivity by cohort at each stage in life and then also calculate the associated DOP components contributing to lifecycle growth.

Figure A1 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 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,

Estabs Size	Share of incumbents		Employment share		Share of entrants		Employment share	
	Data	Model	Data	Model	Data	Model	Data	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 1: Calibration Targets

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.

7 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.

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 0.12 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 4.15% 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 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.

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 (2016): "Declining Business Dynamism: Implications for Productivity?," *Brookings Institution, Hutchins Center Working Paper.*

——— (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.

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.

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.

PUGSLEY, B., P. SEDÁČEK, AND V. STERK (2016): "The Nature of Firm Growth," Discussion paper, UCL, mimeo.

PUGSLEY, B. W., AND A. SAHIN (2014): "Grown-up business cycles," *FRB of New York Staff Report*, (707).

ROBB, A. M., AND D. T. ROBINSON (2014): "The Capital Structure Decisions of New Firms," *Review of Financial Studies*, 27(1), 153–179.
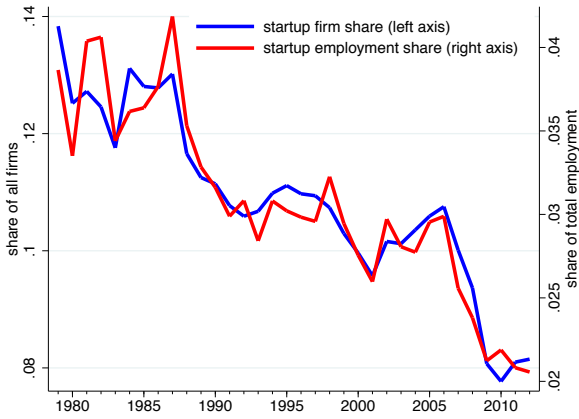
(a) Startups (Age 0)
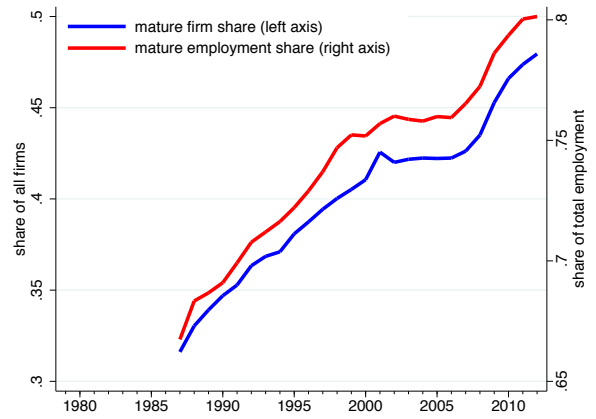
(b) Mature firms (Age 11+)

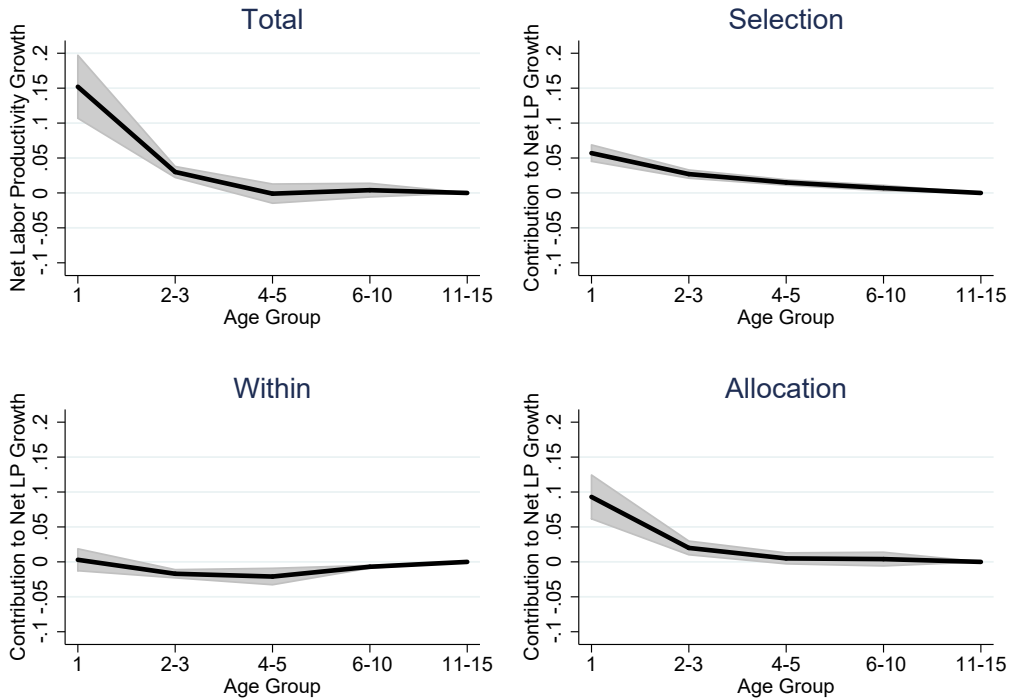Figure 1: Employment and Firm Shares, 1979-2012
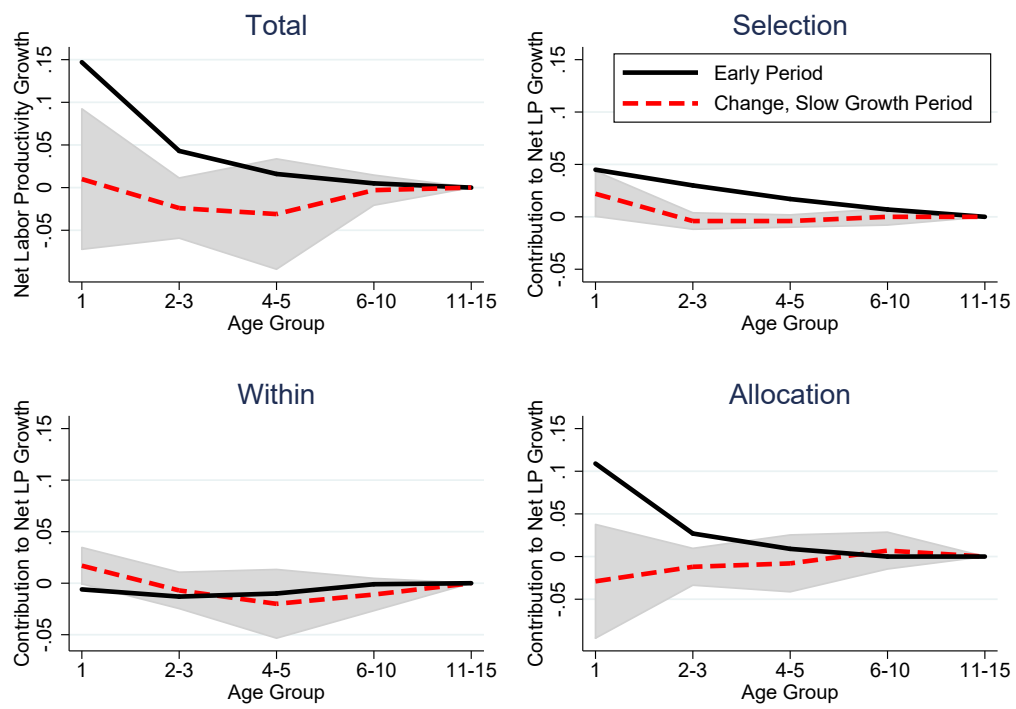

Figure 2: DOP Decomposition by Firm Age

Figure 3: DOP Decomposition by Firm Age: Early ('96-'04) vs. Late ('05-'12)
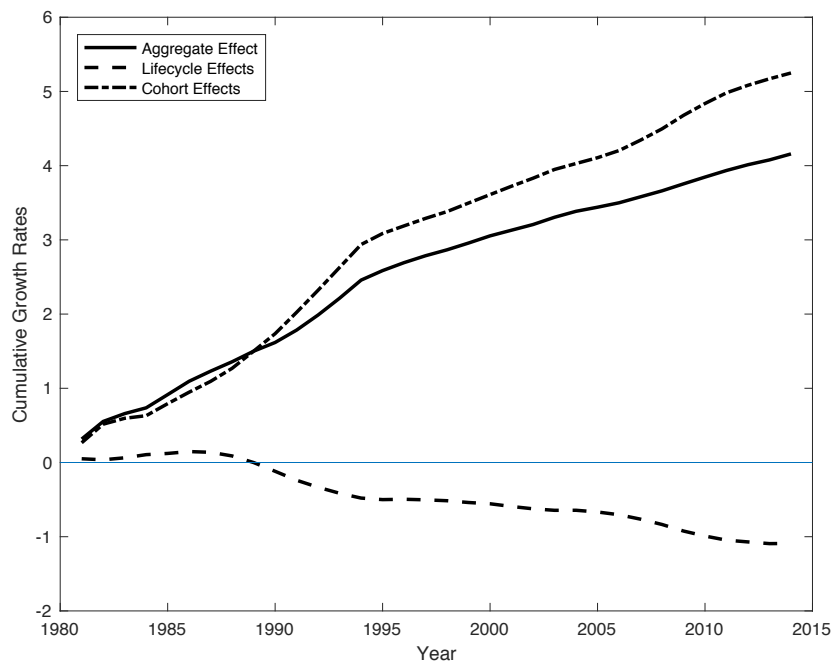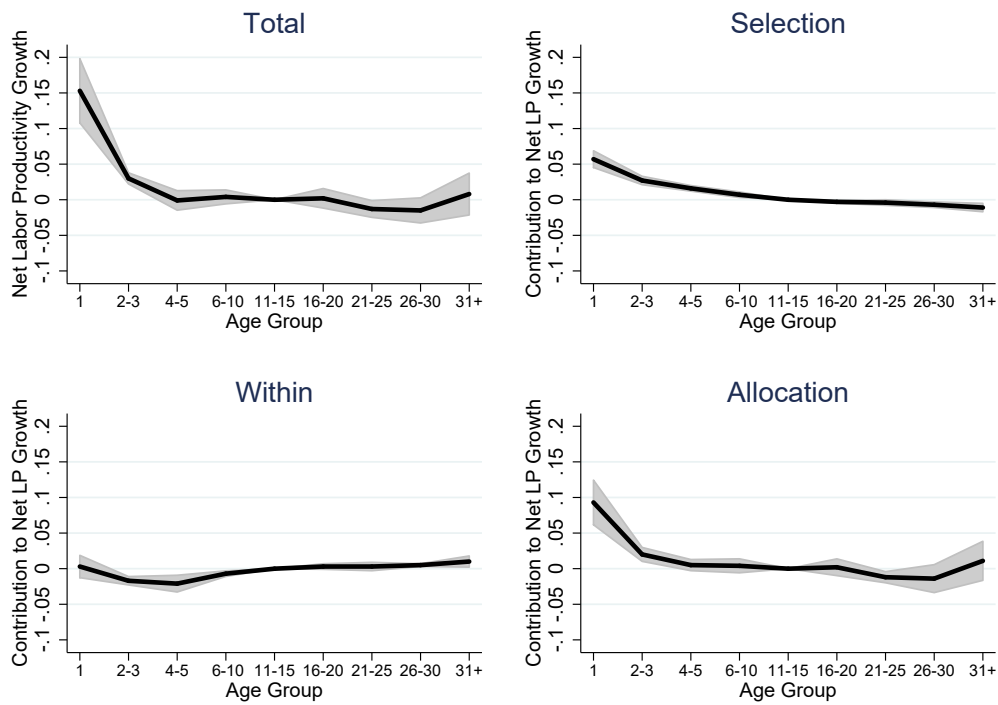
Figure 4: Empirical Counter-Factual and Components

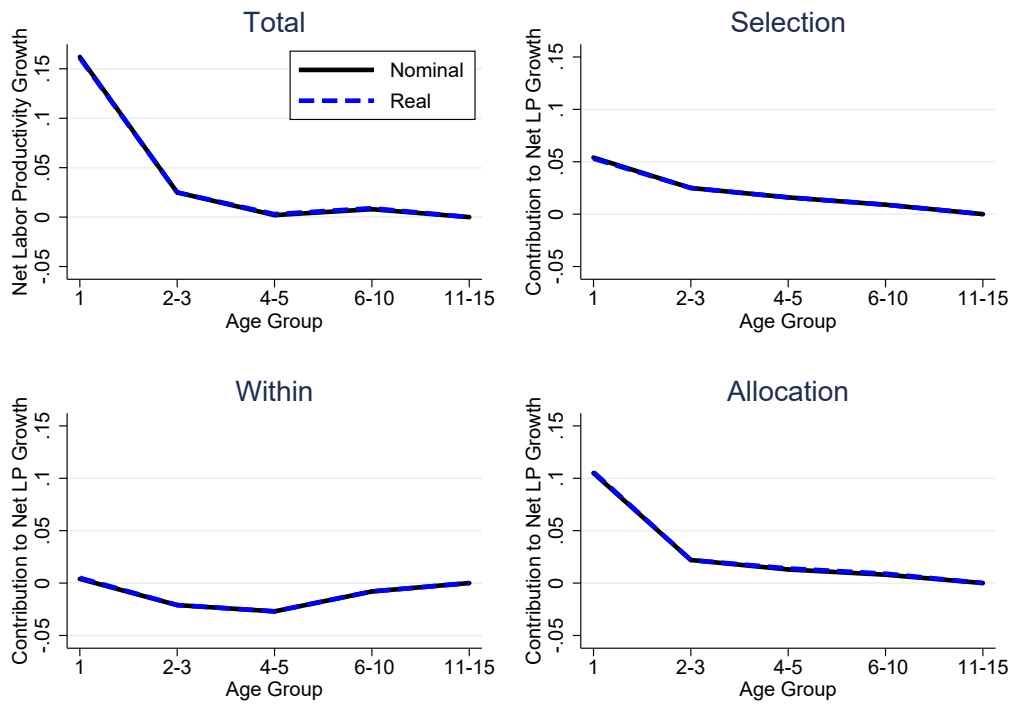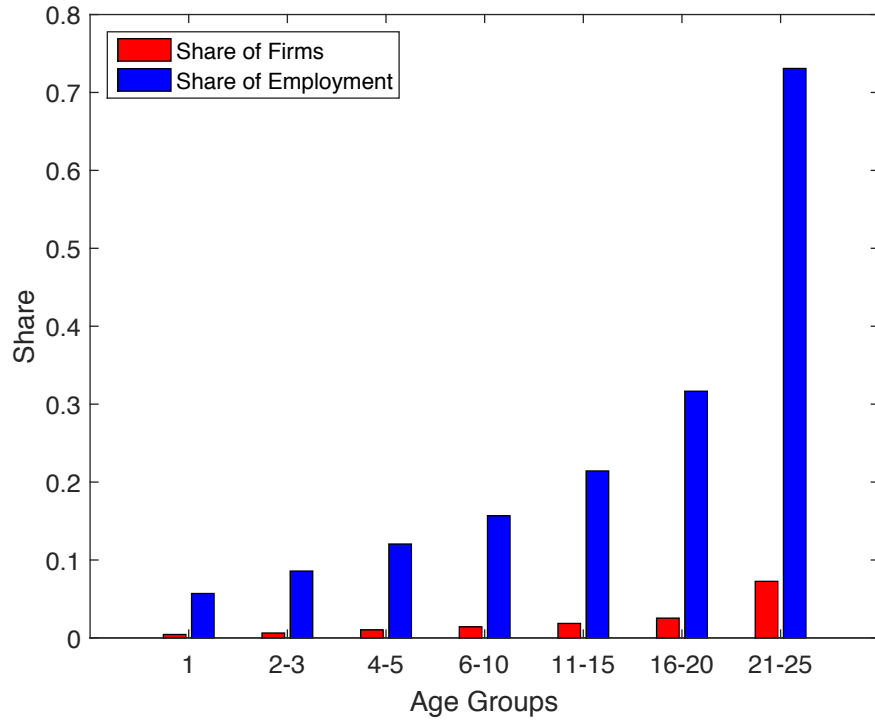Figure 5: DOP Decomposition by Firm Age
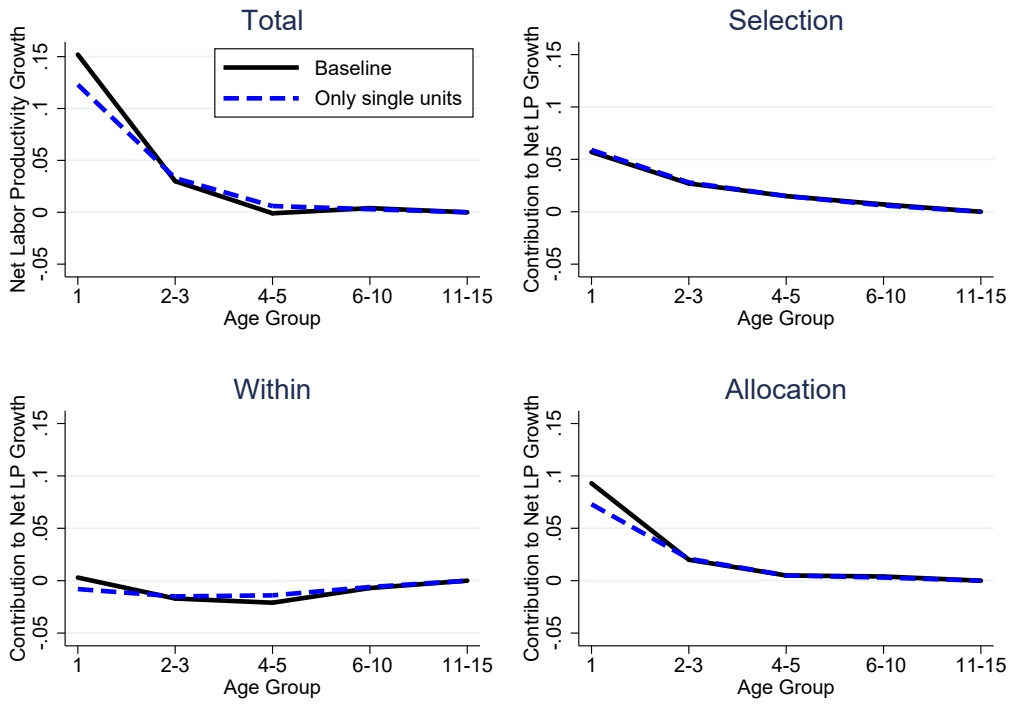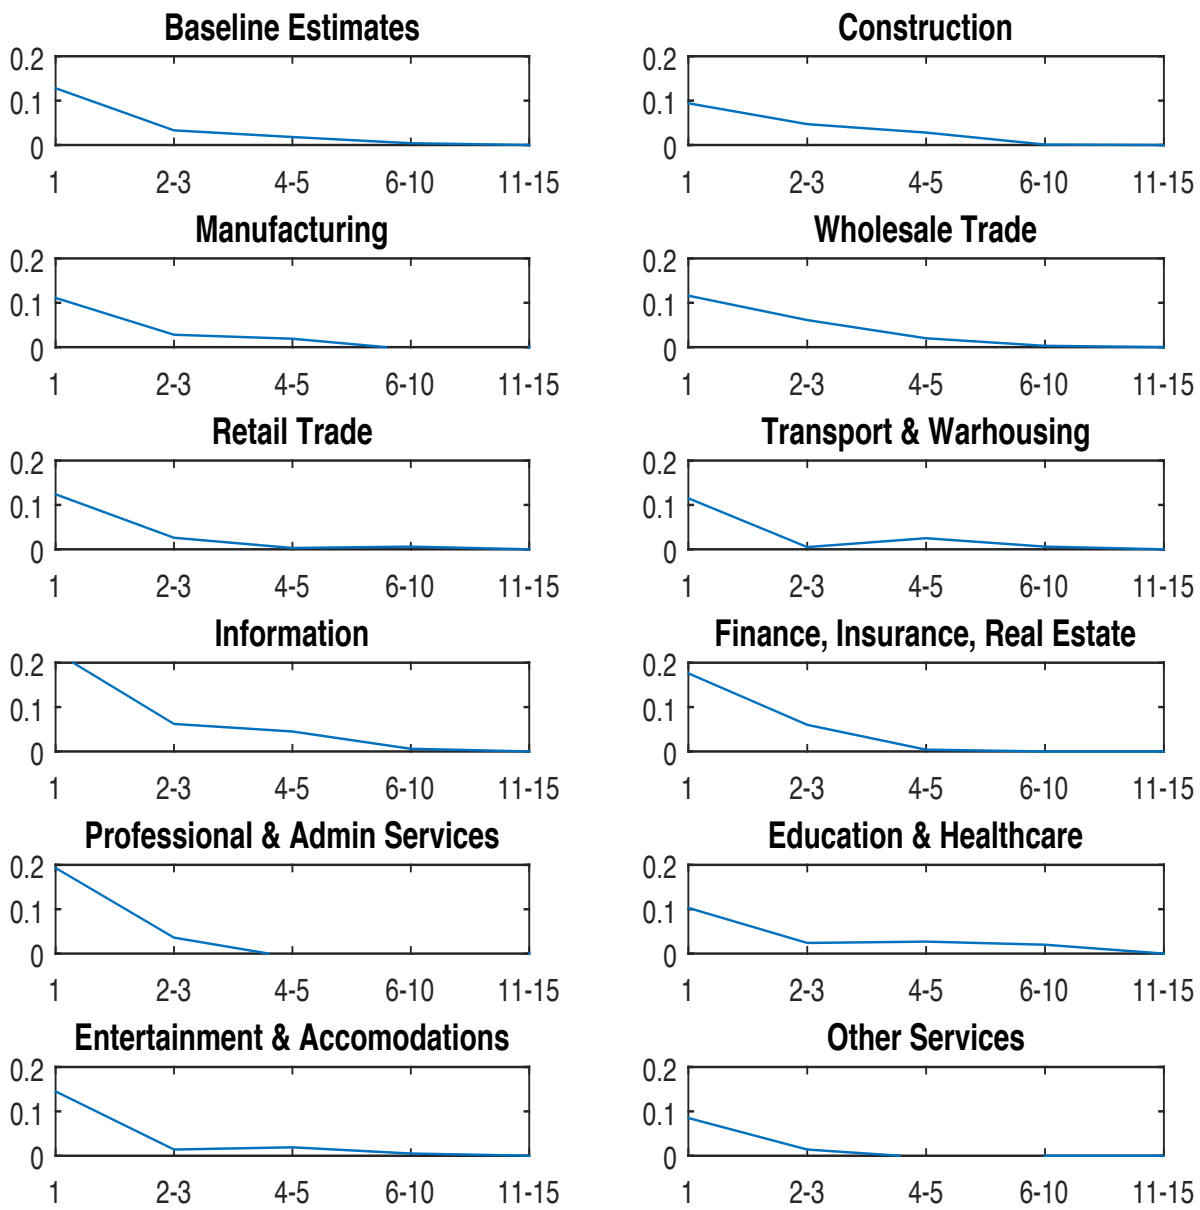

Figure 6

Figure 7

Figure 8

Figure 9: Net age-productivity growth profile by sector
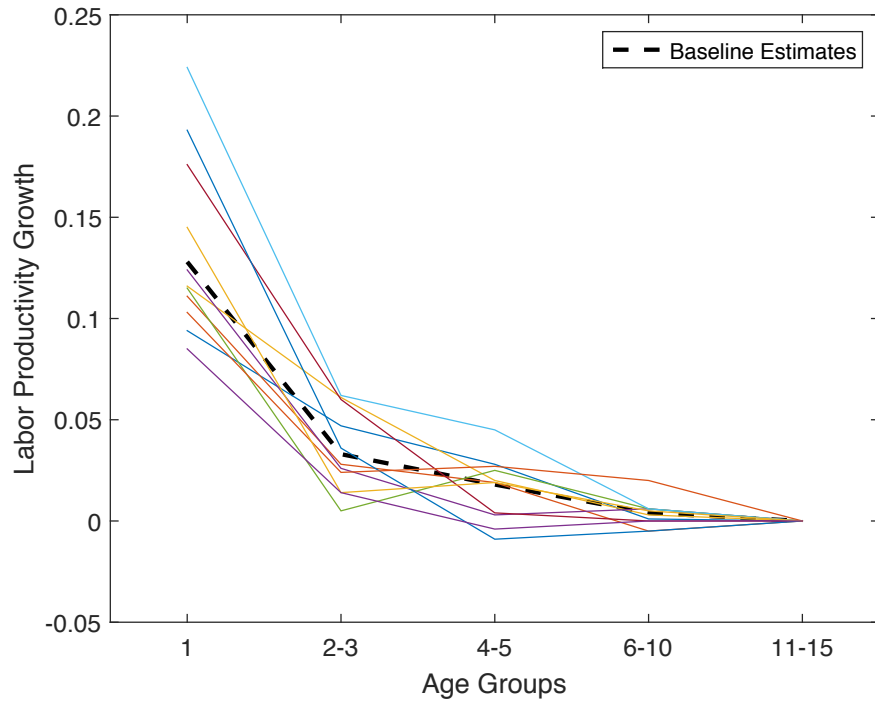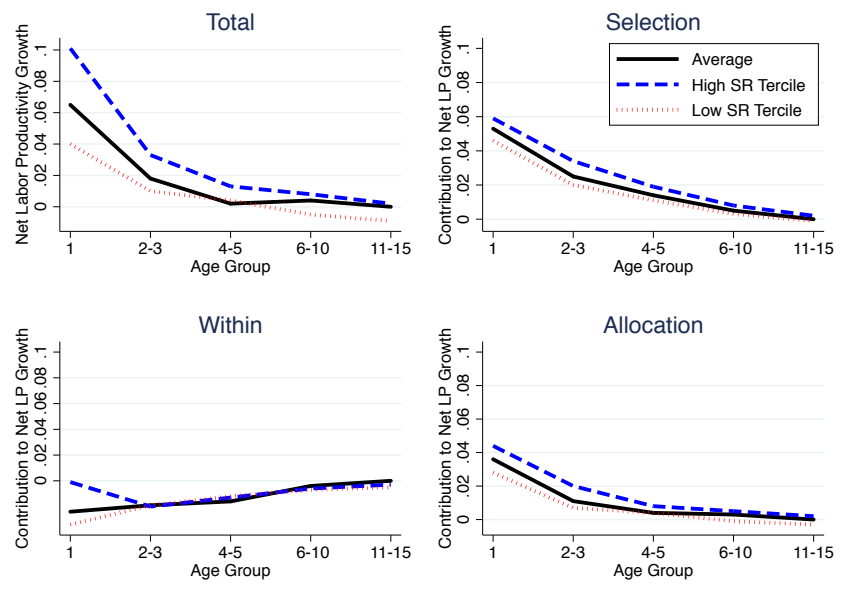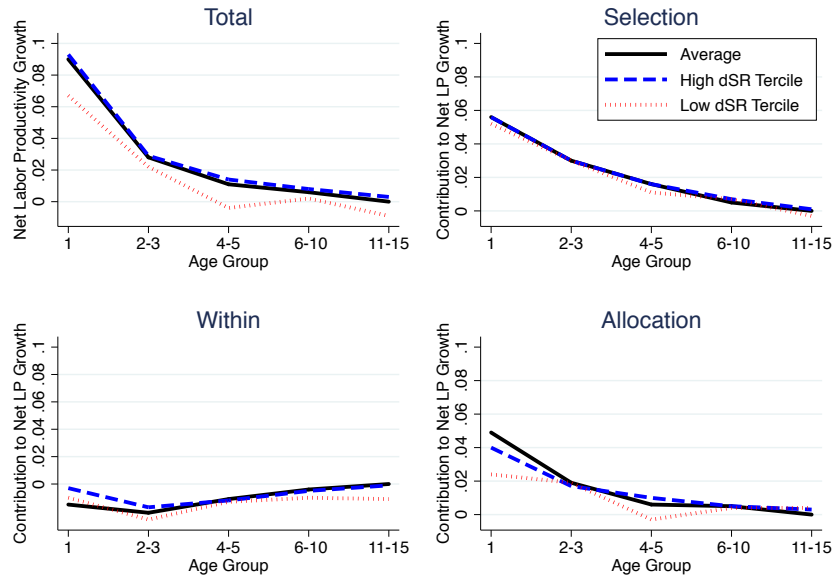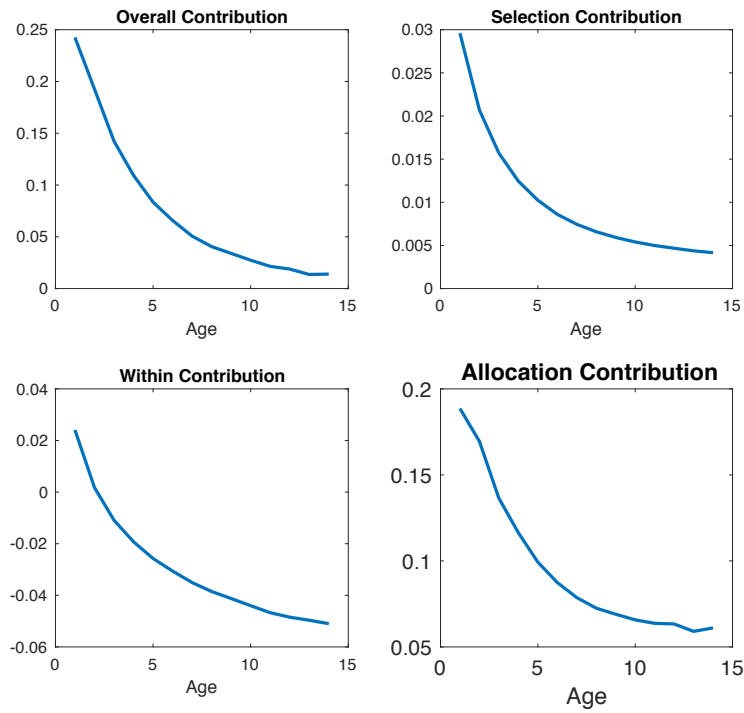
Figure 10

Figure 11

Figure 12

Figure 13: Age-Productivity Profile in the Model

Table 2: 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 3: Productivity growth and area startup rates: OLS regressions

	Industry Productivity Growth $\Delta\Phi_{jst}$				GSP/Worker Growth
	(1)	(2)	(3)	(4)	(5)
Startup Rate	0.796***	0.745*	1.796***	0.772***	0.22**
	(0.209)	(0.456)	(0.010)	(0.058)	(0.108)
Year FE	Yes	Yes	Yes	Yes	Yes
NAICS4 FE	Yes	Yes	Yes	Yes	—
State FE	Yes	Yes	—	—	Yes
CBSA FE	—	—	Yes	Yes	—
Weighted	—	Yes	—	Yes	—
R^2	0.097	0.07	0.368	0.058	0.289
N	200000	200000	1900000	1900000	1222

Note: U.S. Census Bureau Revenue Enhanced Longitudinal Business Database, 1996-2000 and 2003-2012. OLS regression of average change in log net receipts per worker by industry and area (State or CBSA) on startup rate by area with year, NAICS4 digit industry and area fixed effects. BEA Annual GSP/worker growth for years 1980 to 2007. Weighted regressions are weighted by employment. Standard errors clustered by area. Number of observations are rounded to the nearest thousand.

Table 4: Productivity growth and area startup rates: demographic IV

	Startup Rate First Stage			Industry Productivity Growth $\Delta\Phi_{jkt}$			GSP/Worker Growth
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Startup Rate				2.76* (1.68)	+XXXXX**	+XXXXX**	1.71*** (0.48)
Lagged Fertility	XXXXX**	+XXXXX**	+XXXXX**	Yes	Yes	Yes	Yes
Lagged 15-19 Share		+XXXXX**	+XXXXX**		Yes	Yes	
Lagged 60-64 Share		-XXXXX**	-XXXXX**		Yes	Yes	
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NAICS4 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	—	Yes	Yes
Weighted	Yes	—	Yes	Yes	Yes	Yes	No
R^2							
N	20000	200000	200000	200000	200000	200000	1222
J-test p-value	—	> 0.5	> 0.5				
F-test statistic	8.4	>10	>10				10.35

Note: U.S. Census Bureau Revenue Enhanced Longitudinal Business Database, 1996-2000 and 2003-2012. Industry productivity growth 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 state fixed effects. BEA Annual GSP/worker growth for years 1980 to 2007. Weighted regressions are weighted by employment. Standard errors clustered by state. Number of observations are rounded to the nearest thousand.
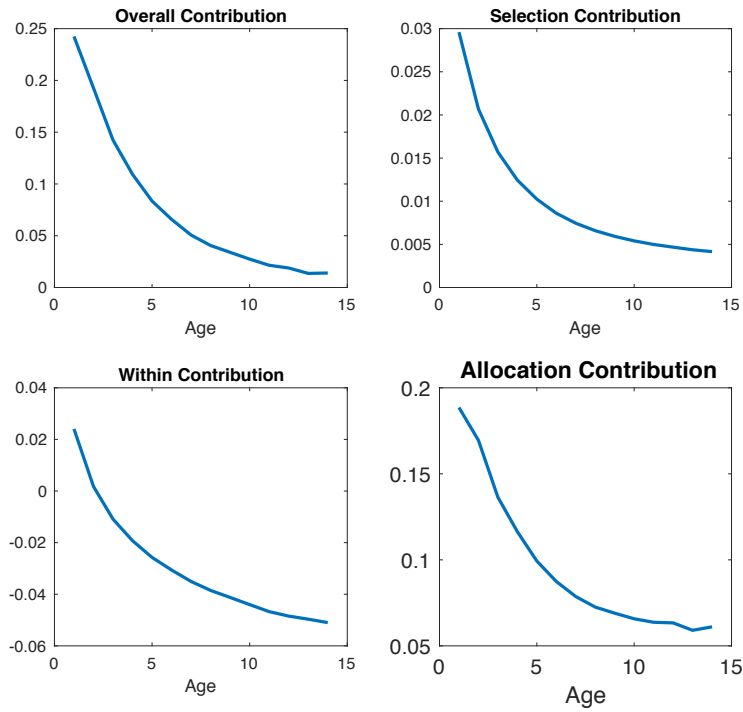
A Additional Tables and Figures
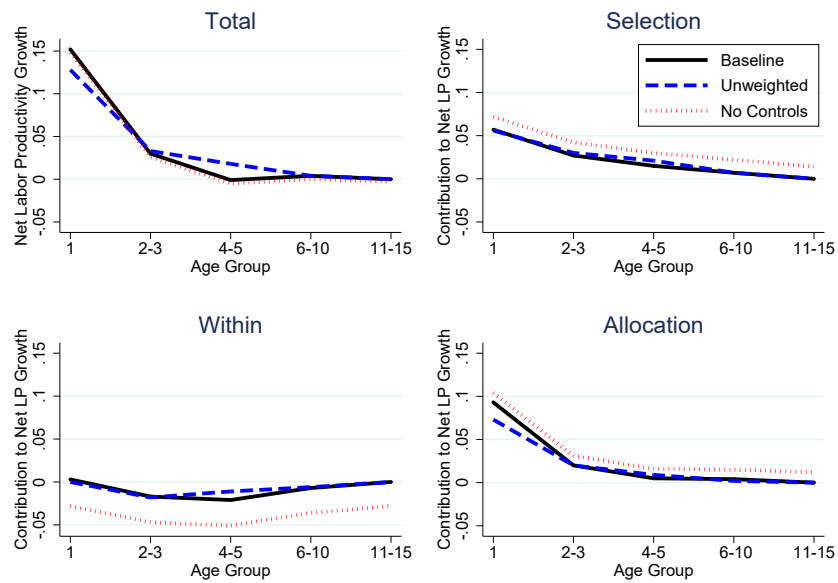
Figure A1: Age-Productivity Profile in the Model

Figure A2: DOP Decomposition by Firm Age: Robustness