Changing Business Dynamism and Productivity: Shocks vs. Responsiveness

Ryan A. Decker, John Haltiwanger, Ron S. Jarmin, and Javier Miranda*

February 2018

Abstract

Job reallocation has declined in all U.S. sectors since 2000. In standard models with adjustment costs, aggregate reallocation depends on (a) the dispersion of business-level productivity shocks and (b) the responsiveness of businesses to those shocks. Using novel business microdata, we infer that the pervasive post-2000 decline in reallocation reflects weaker responsiveness and not lower shock dispersion. Within-industry dispersion of TFP and output per worker has risen, while the marginal responsiveness of employment growth to business-level productivity has weakened substantially especially for young firms in the high-tech sector. The declining responsiveness yields a significant drag on aggregate productivity.

I.   Introduction and motivation

Business dynamics—the process of business birth, growth, decline and exit—is a critical driver of the productivity-enhancing reallocative process that characterizes market economies. But the pace of job reallocation in the U.S. has fallen in recent decades, and since 2000 this downward trend has accelerated and become pervasive across all sectors.[1] We argue that the post-2000 decline in job reallocation reflects declining responsiveness of U.S. businesses to idiosyncratic productivity shocks, implying a drag on aggregate productivity.

We present several novel facts from business-level data, and we study the changing pattern of reallocation by drawing inference from models of firm dynamics. In standard models with adjustment costs, job reallocation arises from the growth and survival responses of businesses to their idiosyncratic productivity draws, so declining reallocation arises from either of two forces: First, a decline in the cross-sectional dispersion of business-level productivity "shocks" reduces reallocation by reducing idiosyncratic incentives for businesses to create or destroy jobs. Second, declining dynamic "responsiveness" to business-level productivity shocks reduces reallocation when there is an increase in adjustment frictions. This "shocks" versus "responsiveness" framework guides our analysis.

We find support for the "responsiveness" hypothesis for the pervasive post-2000 reallocation decline by studying four key empirical moments implied by standard adjustment cost models.[2] First, within-industry dispersion of establishment TFP (which can be measured in manufacturing) has risen. Second, within-industry dispersion of firm labor productivity (which can be measured in all industries) has risen. Third, business growth and survival has become less responsive to idiosyncratic productivity (both TFP and labor productivity) with an especially large decline in responsiveness for young firms in the high-tech sector.   Fourth, counterfactuals reflecting the reduced responsiveness imply a drag on industry-level productivity.

---

[1] See Davis et al. (2007), Hyatt and Spletzer (2013), Davis and Haltiwanger (2014) and Decker et al. (2014). Job reallocation is defined as the sum of gross jobs created by expanding and entering establishments and gross jobs destroyed by downsizing and exiting establishments, expressed as a rate by dividing by average two-year employment as in Davis, Haltiwanger, and Schuh (1996).

[2] See Berger and Vavra (2017) for an application of the "shocks vs. responsiveness" approach in a different context; that paper and others cited therein likewise find an important role for the responsiveness factor in explaining aggregate outcomes.

We use a standard adjustment cost model, in the tradition of Hopenhayn and Rogerson (1993), to derive hypotheses for each of the multiple moments that we study empirically.[3] In this widely used model framework, rising labor adjustment costs can generate each of the moments we document; we therefore think rising adjustment costs is the most natural interpretation of our empirical findings. Other mechanisms may explain some of our empirical patterns but it is a challenge to explain all of them; for example, an increase in market power among firms can explain declining responsiveness and reallocation but does not itself generate rising dispersion of labor productivity. That said, we also discuss and explore other potential mechanisms that might be driving some of our results.

Consistent with our approach of focusing on multiple moments, we combine several data sources to shed light, from many angles, on the question of declining reallocation. We first study the manufacturing sector, with a particular focus on high-tech manufacturing given the role of high-tech in recent productivity dynamics. In a large sample of plant-level manufacturing data spanning 1980-2010, we can closely pair theory and empirics due to our ability to measure TFP (which we construct using multiple methods). We then extend our analysis to the entire private, non-farm sector using recently developed data on output per worker for all U.S. firms starting in the late 1990s. Our economy-wide exercises confirm that our manufacturing findings are not artifacts of sampling errors and that our substantive results are generalizable across industries.

At first glance, recent fluctuations in U.S. productivity growth do not match up with patterns of reallocation: productivity surged in the 1990s through the early 2000s before slowing after 2003, while aggregate job reallocation (and startup activity) fell throughout the 1980-2013 period. A careful review of the evidence resolves the inconsistency: prior to 2000, the decline in reallocation was dominated by the productivity-boosting consolidation of retail trade.[4] Job reallocation (and business formation) in high-tech actually rose during the 1980s and 1990s before declining after 2000 along with high-growth firm activity more generally (Haltiwanger,

---

[3] While not their explicit purpose, Hopenhayn and Rogerson (1993) indeed note that, in their model, rising adjustment costs result in not only lower reallocation and productivity but also higher labor productivity dispersion. The reduction in productivity in their model is in levels which implies reduced productivity growth if adjustment frictions increase over time and/or if the transition dynamics from an increase in adjustment frictions takes time.

[4] Foster et al. (2006), Jarmin, Klimek, and Miranda (2009), and others document the shift away from single unit establishment firms ("mom and pop shops") to national chains. Foster et al. (2006) and Foster et al. (2015) show that establishments of national chains are more productive and more stable. We discuss this evidence further below.

Hathaway, and Miranda (2014)), roughly coinciding with the ICT-driven surge in productivity from the late 1980s to early 2000s and subsequent decline after 2003 (Fernald (2014)).

Our approach of using multiple moments to draw inference is strengthened by this cross-sector variation in patterns of reallocation. In high-tech manufacturing, a hump-shaped pattern of reallocation is matched by a hump-shaped pattern of productivity responsiveness, consistent with theory linking responsiveness and reallocation. Importantly, however, TFP dispersion in high-tech manufacturing rose gradually or was flat throughout recent decades. These patterns are consistent with the strong role of "responsiveness" instead of "shocks" in accounting for the patterns of reallocation. We also find that changing responsiveness—not changing shocks—accounts for the patterns of reallocation in non-tech manufacturing.

Our results have implications for aggregate productivity: Counterfactual exercises suggest that the increased responsiveness of the 1980s and 1990s yielded as much as half a log point annual boost in industry-level TFP in high-tech manufacturing by the second half of the 1990s, then the declining responsiveness of the 2000s yielded as much as a two-log-point drag on annual industry-level TFP by 2010. Moreover, evidence based on labor productivity suggests that the finding of declining responsiveness since 2000 generalizes beyond high-tech manufacturing to other high-tech businesses as well as other areas of the economy. The pre-2000 rise and post-2000 fall of productivity responsiveness in the high-tech sector coincides with the ICT-driven rise and fall of aggregate productivity growth in the U.S. Our findings suggest the declining responsiveness in the post-2000 period has yielded about a 5 percentage point reduction in productivity within the private sector in 2013 relative to what it would have been in the absence of the decline in responsiveness.

An alternative potential cause of changes in reallocation and the responsiveness of firms is the recent decline in startup rates (Decker et al. (2014)). If young firms are more responsive to productivity shocks, changes in the average age of U.S. firms would mechanically reduce overall responsiveness, and the question of declining responsiveness and reallocation would boil down to the question of why startup activity has declined. Consistent with earlier findings (Davis et al. (2007), Decker et al. (2014)), we show that the changing firm age structure induced by declining startup rates accounts for just one quarter of the overall decline in the job reallocation rate. We focus on changes in reallocation and responsiveness within firm age groups to mitigate the

identification challenge posed by declining startup rates (while recognizing that the decline in startup rates is itself an important research area).[5]

We briefly explore other elements of changing responsiveness and productivity. The responsiveness decline in high-tech manufacturing is also evident when we measure responsiveness in terms of equipment investment instead of employment growth. We find mixed evidence linking falling responsiveness with import competition. Our finding of rising within-industry labor productivity dispersion is not consistent with slowing innovation in a Gort and Klepper (1982) framework.

Given the importance of the post-2000 period, an obvious question whether the Great Recession drives our findings. First, we find evidence of declines in the pace of dynamism in all sectors prior to the Great Recession. Second, in our econometric analysis quantifying the decline in responsiveness after 2000 we include rich controls for the business cycle, including year effects and state-level cyclical variables fully interacted with micro-level productivity measures. The latter capture changing responsiveness over the cycle and are interesting in their own right (see Foster, Haltiwanger and Grim (2016)). However, our focus is on lower-frequency trends, so inclusion of these variables is to abstract from this variation. With the inclusion of these controls, we view our results on the post-2000 period as not being driven by the Great Recession.

Section II describes key facts about the declining pace of business dynamism. Section III describes the datasets we employ. In section IV, we use establishment-level data for manufacturing, with a particular focus on high-tech, to study whether the evidence implies "changing shocks" or "changing responsiveness," and we analyze the implications of our findings for aggregate productivity growth. Section V looks beyond manufacturing and investigates the same questions using firm-level labor productivity and employment data for all U.S. sectors. Concluding remarks are in section VI.

II. Business Dynamics: Basic facts
A. *Sectoral Patterns of Reallocation and Young Firm Activity*

Starting with Davis et al. (2007), many studies have documented a decline in the pace of aggregate job reallocation and other indicators of business dynamism in the last few decades. Decker et al. (2016) report wide cross-sector variation in patterns of reallocation: retail trade

---

[5] For example, see Alon et al. (2017), who quantify the productivity implications of declining entry.

exhibits the strongest decline during the 1980s and 1990s, while information and high-tech saw rising reallocation over that period before falling sharply after 2000. These patterns are depicted on Figure 1 (using HP trends) for selected NAICS sectors as well as high-tech (as defined by Hecker (2005)).[6] Figure 2 illustrates similar patterns in the share of employment accounted for by young firms: retail trade saw declining startup activity throughout the 1980s-2010s, while information and high-tech saw rising startup activity prior to 2000.[7] Figures 1 and 2 also single out the high-tech component of manufacturing. Reallocation and startup activity behave similarly in high-tech manufacturing to the high-tech sector more generally. Information, which includes a heavy contingent of tech industries, behaves similarly to high-tech broadly.

The changing prevalence of young firms—which have high reallocation rates—accounts for some of the reallocation patterns in Figure 1.[8] Figure 3 reports annualized changes in reallocation rates for select sectors (and economy wide) for three periods: the late 1980s to late 1990s (1987-1989 to 1997-1999), the late 1990s to mid-2000s (1997-1999 to 2004-2006), and the mid-2000s to early 2010s (2004-2006 to 2011-2013). We use three-year averages at business cycle peaks to abstract from cyclical concerns.[9] Solid bars indicate the actual annualized change in reallocation rates over the period. Patterned (non-solid) bars indicate annualized changes resulting from a shift-share exercise freezing the age composition of businesses at its initial state; that is, the non-solid bars describe within-age group changes in reallocation rates. We use seven age groups (firm age 0, 1, 2, 3, 4, 5, and 6+).

During the 1990s (i.e., 1987-1989 to 1997-1999), the sharp decline in reallocation in the retail trade sector and the increase in the information sector are evident. The patterned bars show that falling young firm activity accounts for about a third of the reallocation decline in retail trade (i.e., two thirds of the decline occurred within age groups), and rising young firm activity accounts for about a tenth of the reallocation increase in information. Services saw a modest

---

[6] As noted above, high-tech is a particular focus of this paper due to its role in aggregate productivity dynamics (Fernald (2014)). Hecker (2005) defines industries as high-tech based on the 14 four-digit NAICS industries with the largest share of STEM workers. The high-tech sector thus defined includes industries in manufacturing (NAICS 3254, 3341, 3342, 3344, 3345, and 3364), information (5112, 5161, 5179, 5181, and 5182), and services (5413, 5415, and 5417). Notably, certain industries in the information sector are not high tech (e.g., book publishing).
[7] Guzman and Stern (2016) focus on extreme high-potential startups and also find 2000 to be an important turning point, with fewer high-growth outcomes for startups identified as having high potential.
[8] Young firms may be more volatile for a variety of reasons, such as the learning dynamics of Jovanovic (1982).
[9] 2011-13 is not a cyclical peak but our sample ends in 2013.

reallocation decline during the 1990s which is entirely accounted for by falling young firm activity.

The pace of decline in several sectors accelerated after the late 1990s. This can be seen in services, which had a more modest decline during the early 1990s. More notably, though, reallocation rates in information fell markedly during the early 2000s after rising during the 1990s, with about a fifth of the early-2000s decline accounted for by falling young firm activity. Each sector continued declining during the late-2000s, and in each case the change in reallocation can be partially but not completely explained by falling young firm activity. This is the main inference we draw from Figure 3: while changing startup rates can account for a nontrivial portion of the overall change in job reallocation rates since the 1980s, most of the variation occurred within firm age groups. This finding encourages us to focus on changing patterns of responsiveness within firm age classes; focusing on patterns within firm age classes also permits us to abstract from factors that may underlie both the changing pace of startups and the change in the age structure of firms.

B. Possible Sources of Changes in Young Firm Activity

Though it is not our focus, it is worth noting that a number of competing hypotheses may account for the variation in startup rates and young firm activity seen in recent decades. Changes in startup rates may endogenously reflect changes in the pace of innovation in an industry for reasons hypothesized by Gort and Klepper (1982): a period of rapid innovation leads to a surge in entry, reallocation and subsequent productivity growth.[10] Moreover, Gordon (2016) has argued that most of the 1980s-1990s high-tech innovations had already been implemented by the early 2000s, and the productivity slowdown since that time is due to slowed innovation and implementation. Taken together, these hypotheses suggest that the changing pace of both startup activity and reallocation in the high-tech sector in recent decades could have been caused by an exogenously changing pace of innovation.[11]

In retail trade, the share of sales and employment accounted for by single unit establishment firms fell from half to a third from 1977 to 2007 (see Foster et al. (2006), Jarmin et al. (2009), and Foster et al. (2015)). This dramatic transition is almost entirely accounted for by

---

[10] Foster et al. (2017) provide supportive empirical evidence for these dynamics for the 1990s U.S. high-tech sector.
[11] In the Gort and Klepper (1982) framework, declining innovation should be accompanied by declining dispersion of productivity within industries. While certainly not dispositive for the hypothesis of slowing innovation, we find the opposite below.

the rise of large, national "big box" retailers, which are more productive (by about 30 log points) and have lower entry and exit rates (by a factor of 15) than single-unit operations. Retail consolidations were likely facilitated by globalization and advances in information technology that permitted the development of large and efficient supply chains and distribution networks. Retail trade is an example of a sector in which declining reallocation and entrepreneurship has been productivity enhancing, a change that is reflected in the age structure of firms (from which we abstract in our analyses).

Yet another factor contributing to variation in the pace of startups is demographics-driven changes in labor force growth. Karahan, Pugsley, and Sahin (2015) show that variation in labor force growth driven by exogenous changes in population growth is positively associated with startup activity, an insight consistent with Hopenhayn (1992)-type models in which labor force growth is accommodated by adjustment in the number of firms.

Each of these factors underlying the changing share of young firm activity likely has some merit, and we seek to abstract from them to focus on the shocks vs. responsiveness hypotheses. We therefore study changes in the pace of reallocation within firm age groups. By abstracting from the changing age distribution, we may be understating the contribution of declining responsiveness to the post-2000 decline in productivity growth since the hypothesis that is the main focus of this paper—rising frictions inducing lower responsiveness of businesses—may be contributing to the pervasive, all-sector decline in startup rates since 2000. An increase in adjustment frictions raises the cost of business activity and reduces the expected discounted value of profits for entrants, a key quantity governing entry in standard models.

III.   Data and Measurement

The backbone dataset for our analysis is the Longitudinal Business Database (LBD), to which we attach other data as detailed below. The LBD includes annual location, employment, and industry for the universe of private non-farm establishments, with firm identifiers based on operational control (not an arbitrary tax identifier).[12] We use the LBD for 1979-2013 (during which consistent establishment NAICS codes are available from Fort and Klimek (2016)). As in previous literature, we construct firm age as the age of the firm's oldest establishment when the firm identifier first appears in the data, after which the firm ages naturally.

---

[12] See Jarmin and Miranda (2002) for a full description of the LBD.

A. Manufacturing and TFP

We construct TFP measures for 2 million plant-year observations (1981-2010) using data from Foster, Grim, and Haltiwanger (2016) (hereafter FGH) combining the Annual Survey of Manufacturers (ASM) with the quinquennial Census of Manufacturers (CM). The ASM-CM is representative of the manufacturing sector in any given year, but it is based on a rotating sample and thus lacks the complete longitudinal coverage of the LBD. To compensate, we integrate the ASM/CM TFP data into the LBD to obtain establishment-level employment growth.[13]

We construct two alternative empirical measures of TFP for our analysis. The first, which has been commonly used in the literature (see, e.g., Baily, Hulten and Campbell (2001), Foster, Haltiwanger and Krizan (2001), Syverson (2011), Ilut, Kehrig and Schneider forthcoming), is a cost share-based index given by:

$$lnTFP_{et} = lnQ_{et}^R - \alpha_K lnK_{et} - \alpha_L lnL_{et} - \alpha_M lnM_{et} - \alpha_E lnE_{et} \qquad (1)$$

where Q^R is real output, K is real capital, L is labor, M is materials, E is energy, α denotes factor elasticities, e denotes individual establishments, and t denotes time. Output is total value of shipments plus total change in the value of inventories, deflated by industry deflators from the NBER-CES Manufacturing Industry Database. Capital is measured separately for structures and equipment using a perpetual inventory method. Labor is total hours of production and non-production workers. Materials are measured separately for physical materials and energy (each is deflated by an industry-level deflator). Outputs and inputs are measured in constant 1997 dollars. Factor elasticities are estimated using industry cost shares (of total factor costs) with a Divisia index that allows cost shares to vary over time.[14] More details are in FGH.

This measure of TFP is a revenue-based measure and is increasingly referred to as TFPR. TFPR is defined by Foster, Haltiwanger, and Syverson (2008) as P*TFPQ, where P is the plant-level price and TFPQ is the typical measure of plant-level technical efficiency in economic models such as the model we consider below. If plants are price takers, within-industry variation

---

[13] We use propensity score weights (based on a logit model on industry, firm size, and firm age) to adjust the ASM/CM/LBD sample to represent the LBD (in the cross section) in each year (see FGH for details). These weights are cross-sectionally representative in any given year but are not ideal for using samples of ASM/CM that are present in both t and t+1. We discuss this further below.

[14] Cost shares yield factor elasticities under the assumptions of cost minimization and full adjustment of factors. We are not assuming full adjustment for each plant at each unit of time but rather that this holds approximately when pooling across all plants in the same industry over time.

in TFPR only reflects TFPQ and any exogenous (to the plant) variation in prices.[15] If plant-level prices are endogenous, TFPR still will be highly correlated with TFPQ in the adjustment cost framework we specify below. Moreover, as we show below, TFPR-based inferences about changing responsiveness are still valid in such a framework. However, with endogenous prices, variation in dispersion of TFPR will reflect not only shocks to fundamentals such as TFPQ but also adjustment costs.

Given possibly endogenous plant-level prices, we also consider an alternative measure of TFP that has been increasingly used in recent literature (e.g., Gopinath et al. (2017) and Foster et al. (2017)). Consider a plant-level isoelastic demand function $P_{et} = D_{et} Q_{et}^{\varphi-1}$ (where D_{et} is an idiosyncratic demand shock and $\varphi - 1$ is the inverse demand elasticity), a plant production function that is Cobb-Douglas with factor elasticity for factor i equal to α_i, and TFPQ equal to A_{et}. Then plant revenue is given by (lower case variables are in logs):

$$p_{et} + q_{et} = \beta_k k_{et} + \beta_l l_{et} + \beta_m m_{et} + \beta_e e_{et} + \varphi a_{et} + d_{et} \tag{2}$$

where $\beta_i = \varphi \alpha_i$ for factor i. That is, the β_i coefficients are the revenue elasticities that reflect both demand parameters and the production function factor elasticities. Given revenue elasticity estimates, the "revenue productivity residual" (RPR) is given by:

$$RPR_{et} = \varphi a_{et} + d_{et}, \tag{3}$$

that is, RPR is solely a function of idiosyncratic TFPQ and demand shocks. This implies that (as discussed in detail in Foster et al. (2017)) RPR exhibits positive dispersion regardless of frictions and distortions. We estimate RPR by estimating the revenue function in (2) using the GMM approach in Wooldrige (2009) (see Appendix C for more details).[16]

For each of these measures of productivity (which we denote as TFP for convenience), we take the log of TFP and deviate it from its detailed industry-by-year mean. These alternative measures are therefore within-industry measures that abstract from aggregate and industry-specific shocks and are unaffected by mismeasurement of industry-level prices (Byrne and Corrado (2015, 2016)). We model TFP as an AR(1) process. The current-period realization of

---

[15] Assuming price taking behavior is not equivalent to assuming homogenous goods and a single price in an industry. If plants in an industry have different product segments but are price takers within product segments then TFPR still only reflects fundamentals reflecting the quality differentials accounted for by price heterogeneity within an industry. TFPR is a referable measure to TFPQ in this case since it captures quality differentials.

[16] Gandhi et al. (2016) argue that if some factors are completely variable then the Wooldridge (2009) method may not be identified. As noted in Appendix C, our results are robust to an alternative estimation method that addresses this identification concern.

the idiosyncratic component of TFP is the shock, and we also consider innovations to these shocks by estimating the AR(1) process below.

In practice, we find that TFPR and RPR are highly correlated (about 0.8), consistent with the findings in Foster et al. (2017). Moreover, Foster, Haltiwanger, and Syverson (2008, 2016) find that TFPR and TFPQ are highly correlated (about 0.75) for the selected set of products where P and Q data are available to construct direct measures of TFPQ. Unsurprisingly, then, the main results of our empirical analysis on changing responsiveness and changing shocks are robust to using TFPR or RPR. For the sake of brevity, we focus on the TFPR results in the main text, but we discuss the results for RPR throughout. In addition, the details of the results for RPR are provided in Appendix C.

B. Economywide Labor Productivity

In Section IV we extend our analysis to nearly the entire economy by constructing measures of firm-level labor productivity. Combining LBD employment (collapsed from the establishment to the firm level) with revenue measures in the Census Bureau's Business Register (BR) (aggregated across EIN reporting units to the firm level) yields an enhanced LBD that we refer to below as the RE-LBD. Revenue data are available from 1996 to 2013; see Haltiwanger et al. (2017) for more details.[17] Consistent with previous literature, we construct annual firm employment growth rates on an "organic" basis to represent changes in establishment-level employment rather than artificial growth caused by mergers and acquisitions.

Similar to our TFP construction, we use (log) revenue per worker deviated from detailed (6-digit NAICS) industry-by-year means as a measure of firm labor productivity. We thereby control for price differences across industries such that our labor productivity measure is a within-industry relative gross output per worker measure; Foster, Haltiwanger, and Krizan (2001, 2006) show that within-industry relative gross output per worker is highly correlated with within-industry relative value added per worker and strongly correlated with within-industry relative TFP (suggesting that materials and capital shares are similar across firms within industries). We omit firms in the Finance, Insurance and Real Estate sectors (NAICS 52-53)

---

[17] About 20-percent of LBD firm-year observations cannot be matched to BR revenue data because firms can report income under EINs that may fall outside of the set of EINs that the Census considers part of that firm for employment purposes**.** We address potential match-driven selection bias by constructing inverse propensity score weights (separately for births, deaths, and continuers) such that the RE-LBD is representative of the LBD universe in terms of the size, age, employment growth rate, broad industry, and single/multi-unit structure of firms.

from all analysis due to the difficulty of measuring output and productivity in those sectors. As we show below, in our adjustment cost framework inferences regarding changing responsiveness can also be made using revenue per worker.

IV.   Change in shocks vs. change in responsiveness

A. *Theoretical motivation*

Models of firm[18] dynamics with adjustment frictions imply that a within-sector change in the pace of reallocation is due to either a change in the dispersion of shocks faced by firms or a change in firms' responses to those shocks. We analyze a standard model with adjustment frictions in the tradition of Hopenhayn and Rogerson (1993). Technical details are in Appendix B, but we provide an overview of the model and key predictions in this section. Firms face idiosyncratic productivity shocks, where the realization of productivity in the current period, A_{et}, is drawn from a persistent AR1 process. Net hiring and downsizing are subject to non-convex (kinked) adjustment costs.[19] The resulting decision rule for firms' net hiring rates reflects adjustment costs and is given by $g_{e,t} = f_t(A_{et}, E_{et-1})$, where the state variables are the productivity realization A_{et} and initial employment E_{et-1}, both of which are observed prior to the net hiring decision.[20] We do not model entry or exit but discuss these margins below. For purposes of discussion in this section and appendix B, A_{et} is referred to as TFP or TFPQ. If there are demand shocks, this measure should be interpreted as a composite shock measure reflecting both TFPQ and demand shocks.

We calibrate the model and report numerical analysis to motivate the empirical specifications and moments we consider below (see appendix B for calibration details). Our benchmark model and calibration allow for endogenous plant-level prices using a CES monopolistic demand structure. Strictly speaking, this model implies that the more relevant empirical measure is RPR. However, importantly for our empirical approach, in the model revenue productivity measures with endogenous plant-level prices —TFPR or revenue labor productivity—are highly correlated with TFPQ (pairwise correlations of about 0.90 in our

---

[18] We use the term "firms" loosely in this subsection for expositional ease.

[19] Qualitative predictions on the multiple moments of interest are robust to convex adjustment costs.

[20] A similar rule would exist for investment in a model with capital. For net hiring rate dynamics, see, e.g., Cooper, Haltiwanger, and Willis (2007) and Elsby and Michaels (2013). For investment dynamics, see, e.g. Cooper and Haltiwanger (2006).

benchmark calibration). As discussed in detail below, this implies that the core prediction that a rise in adjustment frictions yields declining responsiveness to TFPQ carries over to revenue productivity measures (either TFPR or revenue labor productivity). The high correlation between TFPQ and TFPR in the model is not a targeted moment but is consistent with our empirical finding that TFPR and RPR are highly correlated (about 0.8).

While we leave model details to the appendix, here we summarize the model's rich empirical predictions for two experiments: a decline in the dispersion of TFP and an increase in labor adjustment frictions. We can easily generate changes in the rate of job reallocation with these experiments; our interest is in observing implications for other moments that would allow us empirically to distinguish between the "shocks" hypothesis and the "responsiveness" hypothesis (with the latter corresponding to changing adjustment frictions). We study the experiments' effect on "responsiveness" by using a regression estimate (on simulated data) of the net hiring—or, equivalently, employment growth—policy function $g_{et} = f(A_{et}, E_{et-1})$; we simply regress employment growth on lagged productivity and employment then observe the coefficient on productivity.[21] The other key moment of study is the standard deviation of labor productivity.

First, we hold adjustment frictions constant at the baseline calibration and vary the dispersion of TFPQ (see Figure 4a).[22] A decline in TFPQ dispersion yields: (1) lower job reallocation; (2) weaker responses of firm-level growth (t to $t + 1$) to the realization of both TFP and labor productivity in t (conditional on employment in t);[23] and (3) lower standard deviation of labor productivity (labor productivity in the model is measured as revenue per worker, consistent with our empirical approach).

Second, we hold TFPQ dispersion constant at its baseline calibration and vary the magnitude of adjustment frictions (see Figure 4b). An increase in adjustment frictions yields: (1) lower job reallocation; (2) weaker responses of firm-level growth (t to $t + 1$) to the realizations

---

[21] That is, we focus on β from $g_{et+1} = \eta + \beta \ln A_{et} + \delta \ln E_{et} + \varepsilon_{et+1}$. This timing convention is designed to match our empirical work; the qualitative theoretical results are robust to other timing conventions.

[22] Holding constant the persistence of TFP so that this implies we are reducing the variance of innovations to TFP.

[23] Kinked adjustment costs give rise to inaction ranges. As TFP dispersion falls there is a decrease in the fraction of firms that make zero adjustment (i.e., the "real options" effect). But declining TFP dispersion also implies smaller adjustments among those firms that do adjust (i.e., the "volatility" effect). Vavra (2014) argues that the volatility effect dominants the real options effect in the steady state, a general result extending back to Barro (1972). Bloom et al. (2016) and others use a similar model to study the effects of uncertainty on business cycles; even in their model, the volatility effect dominates at annual frequency.

of both TFP and labor productivity in t (conditional on employment in t); and (3) higher standard deviation of labor productivity.

These empirical predictions are sufficient to distinguish between the "shocks" and "responsiveness" hypotheses even in the absence of observable TFP dispersion. The key distinguishing moment is labor productivity dispersion. For example, in the presence of declining reallocation and responsiveness, rising labor productivity dispersion would imply that adjustment costs have increased while falling labor productivity dispersion would imply that TFP dispersion is declining. Conveniently, we also observe TFP dispersion (in manufacturing), enhancing our ability to draw empirical inference. This focus on multiple empirical moments is a key strength of our approach.

Our model neglects firm entry and exit, but its key predictions are robust to consideration of these margins. Specifically, Hopenhayn and Rogerson (1993) find that a rise in adjustment frictions reduces entry and exit. In their model, the lower bound of productivity necessary for survival declines as frictions increase. The empirical prediction, then, is that not only will firm growth for continuers become less responsive to productivity when adjustment frictions rise, but so will exit. We explore this prediction in the empirical analysis below. Our model also neglects any learning dynamics of young firms. We seek to control for the latter by permitting the firm dynamics of young firms to differ from those of mature firms in our empirical analysis. This approach is consistent with our emphasis on studying empirical moments within firm age groups to abstract from the contribution of changing startup rates.

Our theoretical model in Appendix B also demonstrates that a decline in responsiveness from an increase in adjustment frictions yields a drag on aggregate productivity (see Figure B4). In our theoretical analysis, we develop a diff-in-diff counterfactual methodology that quantifies the contribution of the decline in responsiveness to aggregate productivity. Details of this counterfactual methodology are in Appendix B, and we discuss the empirical implementation below.

Our primary objective is to quantify the role of shocks vs. responsiveness empirically, so we do not identify empirically a structural model of adjustment frictions. However, we think this is a rich area for future research. One potential use of our empirical findings would be as moments to discipline such analysis.[24]

---

[24] See Cooper and Haltiwanger (2006).

B. *Empirical Analysis of U.S. Manufacturing*

In this section, we investigate these issues with establishment-level data for U.S. manufacturing with a particular focus on high-tech manufacturing.[25] We first study the "shocks" hypothesis by directly exploring the evolution of TFP dispersion (i.e., the dispersion of establishment productivity draws), quantified as the standard deviation of (log) within-industry TFP (see Section III for TFP measurement details). Figure 5 reports TFP dispersion separately for plants of young and mature firms, in high-tech and non-tech manufacturing.[26] We focus on low-frequency variation by reporting HP trends.

Figure 5 shows that TFP dispersion has risen gradually or been flat in high-tech manufacturing since the early 1980s and in non-tech manufacturing since the early 1990s.[27] Within-industry TFP dispersion is large (consistent with, *e.g.*, Syverson (2004, 2011)); for example, a level of 0.4 (40 log points) on Figure 5 implies that a plant one standard deviation above the mean for its industry is about $e^{0.4} \approx 1.5$ times as productive as the mean. Within-industry TFP dispersion is about the same for plants of young and mature firms. Figure C1 in appendix C shows very similar results for the alternative RPR productivity measure based on Wooldridge (2009). Bils, Klenow and Ruane (2017) suggest that the observed rising within-industry TFP dispersion may be due to rising survey-based measurement error in the ASM, but as we discuss below (and show on Figure A6 in appendix A) we find rising revenue productivity dispersion even in administrative data.

Employment dynamics at the plant level depend not only on dispersion but also on persistence of idiosyncratic TFP: plants facing adjustment costs are more likely to respond to TFP shocks if TFP is more persistent (Cooper and Haltiwanger (2006); Cooper, Haltiwanger, and Willis (2007)). Our data are not ideally suited for estimating TFP persistence and innovations due to ASM panel rotation issues, but Figure A4 in appendix A suggests that persistence is reasonably stable with an estimated AR(1) coefficient of about 0.6 to 0.7, and

---

[25] These include NAICS codes 3341 (computer and peripheral equipment), 3342 (communications equipment), 3344 (semiconductor and other electronic components), 3345 (navigational, measuring, electromedical, and control instruments), 3254 (pharmaceutical and medicine), and 3364 (aerospace product and parts).

[26] Our unit of analysis in this section is the establishment (plant), but the LBD permits us to classify plants based on the age of the firm to which they belong.

[27] Bloom et al. (2016) report dispersion of a different measure of productivity shocks; while we study the within-industry dispersion of TFP draws for the manufacturing sector generally, those authors study overall dispersion of innovations to TFP among a selected subset of plants that appear in manufacturing samples for 25 years or more.

trends of TFP innovation dispersion (Figure A5) mimic trends of TFP dispersion.[28]

Figures 5, C1, A4 and A5 suggest that changing reallocation is not driven by changing TFP dispersion or persistence. Consider high-tech: Figure 1 shows reallocation rising during the 1990s then falling after 2000. For dispersion and persistence of TFP to account for the reallocation trend we would expect dispersion and/or persistence to mimic these patterns; or, conversely, given the patterns of TFP dispersion and persistence, we should see flat or rising reallocation in the manufacturing sector in the post-2000 period. That we see the opposite is evidence against the "shocks" hypothesis for declining reallocation during that period. Moreover, as we note below in our investigation of firm labor productivity, labor productivity dispersion has also risen in manufacturing (see Figure A6 in appendix A).

Consistent with our multiple moments approach, we next estimate the relationship between growth (and survival) and TFP realizations at the establishment level. Our main dependent variable of interest is establishment employment growth from year t to $t+1$ using the Davis, Haltiwanger, and Schuh (1996) (hereafter DHS) concept that accommodates exit (by using the two-year average of employment as the denominator). We estimate the following:

$$
g_{e,t+1} = \sum_{age=y,m} [\beta_{age}TFP_{et} + \delta_{1age}TFP_{et} * Trend_t
$$
$$
+ \delta_{2age}TFP_{et} * Trend_t^2] * I_{age,et} + X'_{et}\Theta + \varepsilon_{e,t+1}
$$

(4)

where $g_{e,t+1}$ is the DHS employment growth rate for establishment e from time t to $t+1$, TFP_{et} is (log) industry-deviated TFP for establishment e at time t, and $Trend_t$ is a simple linear time trend.[29] The responsiveness to TFP in terms of the main and trend effects can vary by firm age with $I_{age,et}$ an indicator for young (age<5, subscript y) and mature (subscript m) plants (these dummy variables are denoted Young and Mature in the discussion below). X_{et} includes year effects, establishment size, firm size, state effects and a state-level business cycle measure (the change in state unemployment rates). We also interact the state cyclical measure with TFP and the young and mature dummies; our liberal inclusion of cyclical indicators is intended in part

---

[28] When we estimate the AR1 structure we exclude first panel years of the ASM since the panel of pairwise continuers is not representative in those years.

[29] Ilut, Kehrig and Schneider (forthcoming) estimate broadly similar reduced-form policy functions with a focus on asymmetric responsiveness, finding that businesses respond more strongly to negative than to positive shocks. These authors do not study changes in this asymmetry over time, a potentially interesting question given our findings.

to avoid result contamination from the Great Recession. We estimate equation (4) for 1981-2010 using propensity score weights relating the ASM/CM to the LBD. Since the measure of micro TFP is deviated from detailed industry-by-year effects, we are estimating the responsiveness of establishment growth to the idiosyncratic component of productivity.[30]

This reduced form specification is broadly consistent with the specifications of selection and growth dynamics from the literature discussed above, and it is consistent with our model-based exercises (which estimate the equivalent of equation (4) on simulated data). Moreover, by using DHS growth rates we can incorporate both the intensive margin and the extensive margin (exit) of plant-level growth; as noted above, adjustment cost models of employment growth predict that growth (and exit) from t to $t + 1$ is related to the realization of TFP in period t and state variables, and standard empirical specifications of exit (e.g., Syverson (2011)) likewise find that exit is related to TFP realizations. In this sense, equation (4) produces a reduced-form yet direct estimate of policy functions generated by standard models.[31]

Our question is whether the response to idiosyncratic productivity shocks has changed over time. The inclusion of the $Trend_t$ variable allows us to estimate a time-varying relationship between productivity and growth. In unreported results we have considered alternative ways to capture a changing trend (*e.g.*, interacting a linear trend with decade dummies), and results are robust to considering such alternatives.

On Table 1, we report the main effects for TFP by firm age group and the interactions with the trend terms. Columns 1 and 2 show the growth regressions of equation (4); columns 3 and 4 show a linear probability model with exit as the dependent variable (but otherwise identical to equation (4)). The estimates for β_y and β_m are given by the "TFP*Young" and "TFP*Mature" rows. These positive (negative) coefficients show that, consistent with previous literature, productivity and growth (exit) are positively (negatively) related.[32] The growth coefficients are stronger for establishments of young firms, consistent with intense selection

---

[30] Results are robust to also including industry-by-year effects as controls.

[31] Results are robust to estimating (4) with the innovation to TFP in period t. We focus on the estimation of (4) using the realization of TFP in period t to avoid the complications of panel rotation issues in the ASM (see Foster et al. (2016) for details). TFP in period t is measured for calendar year t while establishment growth is measured from March of t to March of t+1. Thus, the empirical timing of the data is closer to the timing in the theoretical specifications in Appendix B than might first appear. In Appendix B, we show declining responsiveness of firm-growth to current or lagged realizations in productivity from an increase in adjustment frictions.

[32] The coefficients relating productivity with growth or exit are statistically significant at the 1 percent level in all but one case: the coefficient for exit among young high-tech establishments is significant at the 10 percent level.

working on recently started businesses; the exit coefficients follow the same pattern in non-tech manufacturing, though interestingly this is not the case in high-tech.

The estimates of δ_{1y} and δ_{1m} are given by the "TFP*Young*Trend" and "TFP*Mature*Trend" rows of Table 1, respectively. These coefficients show how the marginal responsiveness of establishments to their idiosyncratic productivity has changed with time. Notably, in high-tech manufacturing δ_{1y} and δ_{1m} are positive (negative) and significant for the growth (exit) of plants of both young and mature firms, with the exception of the exit coefficient for mature firms, suggesting that productivity responsiveness generally strengthened in the early years of the sample (which begins in 1980), while the coefficients are close to zero among non-tech establishments. Both inside and outside of high-tech, however, the growth (exit) coefficients on the quadratic term (δ_{2y} and δ_{2m}) are negative (positive).

We next graphically illustrate the implications of the combined linear and quadratic trend terms. Since TFP is measured relative to industry-year means, we can calculate the growth differential between a "productive" plant—the plant with a TFP draw one standard deviation above its industry mean—and the average plant in an industry by multiplying the total regression coefficient (including trend effects) by the within-industry TFP standard deviation.[33] To abstract from changing TFP dispersion, we fix the standard deviation at 0.40 for high-tech and 0.37 for non-tech (roughly the respective averages across time). Figure 6 shows the resulting growth differentials averaged by decade.

First note that young firm plants are more responsive to productivity than are mature firm plants, especially in high-tech. In the 1980s (black bars), the growth differential among young high-tech plants was 13 percentage points: the plant with productivity one standard deviation above its industry mean grew 13 percentage points faster, over a one-year period, than the plant with industry mean productivity, compared with 6 percentage points among mature high-tech plants. This and Figure 5 imply that the high pace of reallocation of young-firm plants is not driven by a high variance of TFP but rather by a high responsiveness to TFP differences consistent with, for example, a learning model. The difference in responsiveness between plants in young and mature firms implies that overall responsiveness depends in part on the age distribution—hence our within-age group approach.

---

[33] We set the cyclical indicator (state change in unemployment) to zero to evaluate effects at a neutral cyclical state.

Our main focus is the variation in responsiveness over time. First, consider high-tech manufacturing. For plants in young firms, the growth differential rises from 13 to 16 percentage points from the 1980s to the 1990s then declines to 9 percentage points in the 2000s. For plants in mature firms, responsiveness initially declines modestly from the 1980s to the 1990s then accelerates into the 2000s, with the growth differential stepping down from 6 to 5 percentage points then dropping to 3 percentage points. These declines in responsiveness are large in magnitude. High-tech plants responsiveness in the post-2000 period is only about half that in the 1990s for plants of both young and mature firms.

Next, consider the non-tech results. Again, plants in younger firms are more responsive to TFP draws. Among young firms, the growth rate differential was about 10 percentage points in the 1980s, 9 percentage points in the 1990s, and 6 percentage points in the 2000s. Among mature firms, the growth differential was just above 5 percentage points in the 1980s and 1990s and fell by about half a percentage point in the 2000s.

On Figure C3 of appendix C, we report exercises using the Wooldridge (2009) RPR as our TFP measure.[34] In high-tech manufacturing, young-firm RPR responsiveness rises then falls as with TFPR, and mature-firm RPR responsiveness falls sharply in the post-2000 period in a manner similar to the TFPR results. For non-tech plants, the overall drop in responsiveness from the 1980s to the 2000s is similar with RPR compared to TFPR. The RPR results tell broadly the same story as the TFPR results, with reasonable similarity both quantitatively and qualitatively.

As can be seen from column 3 of Table 1, part of the growth responsiveness pattern is driven by selection dynamics associated with changing exit responsiveness (Figure A1 in appendix A shows exit charts analogous to Figure 6). Among young firm high-tech establishments, exit selection intensified from the 1980s to the 1990s then weakened in the 2000s; young non-tech establishments saw steadily weakening selection throughout the period. Among mature firm plants, selection intensity did not vary notably until it weakened somewhat in the 2000s. The findings on exit are interesting in their own right as they imply that in the post-2000 period low-productivity plants are more likely to survive, constraining aggregate productivity (and potentially raising TFP dispersion).

An alternative story is that rising dispersion of TFP (and its innovations) in the post-2000

---

[34] The decline from the 1990s to the 2000s among young high-tech businesses is not as notable in the RPR-based regressions as it is for TFPR, but it is still substantial as we show below in aggregate productivity counterfactuals.

period not only could be partially endogenous to changing selection but also could independently contribute to weakening growth responsiveness: in the presence of non-convex adjustment costs, higher TFP dispersion widens inaction bands and reduces the frequency of adjustment, a mechanism that has inspired a large literature on uncertainty and business cycles.[35] But these concerns are not likely to be playing a dominant role: during the 1990s, we find increased responsiveness of exit in high-tech despite mild increases in TFP dispersion, a finding that also holds for RPR (Figure C1 in appendix C). More broadly, the combined 30-year patterns of dispersion and responsiveness, across age and industry groups, cannot tell the alternative story coherently. As noted above, our model and a broader literature theorize that the "frequency" effect of widening inaction bands is not likely to dominate the "volatility" effect of larger adjustment-conditional changes in employment due to higher TFP dispersion.[36]

Taken together, these results have important implications for the evolution of firm dynamics in recent decades. The way in which individual businesses respond to their idiosyncratic realizations of productivity has changed. The positive relationship between realized productivity and subsequent employment growth remains robust, but it has weakened, particularly since 2000. Through the lens of firm dynamics models, our results are evidence that establishment-level policy functions have changed over time, particularly for young firms but also for older ones.[37] In the post-2000 period, these changes are consistent with an increase in adjustment costs or other frictions that reduce marginal responsiveness to productivity in these models. The changes are most striking among high-tech businesses, where we observe a pattern of rising and falling productivity responsiveness that coincides with the ICT-driven acceleration and deceleration of aggregate productivity growth documented by Fernald (2014) and others.

C. Implications for aggregate (industry-level) productivity

How important are the changes in responsiveness for aggregate fluctuations in productivity? For this purpose, we compute the following diff-in-diff counterfactual that is inspired by, but distinct from, the well-known Olley and Pakes (1996) decomposition:

---

[35] See Bloom (2009), Bachmann and Bayer (2013).

[36] Bloom et al. (2016) construct a model similar to ours in which the real options effect dominates the volatility effect in the short term (less than four quarters) at high frequency. We study annual responses in a steady state setting, consistent with our long-term (rather than cyclical) focus.

[37] Karahan, Pugsley, and Sahin (2016) argue that the dynamics of incumbent firms have not changed over this time period based on average growth rates for various age classes. We differ from their approach by directly estimating policy functions at the establishment level. Viewed through their framework, our results suggest that factors in addition to changes in the growth of the labor force are likely relevant for understanding the decline in startup rates.

$$\Delta_t^{t+1} = \sum_e (\theta_{e,t+1}^T - \theta_{e,t+1}^{NT}) a_{et} \qquad (5)$$

where a_{et} is log(TFP), θ_{et+1}^T is the predicted employment share for establishment e in period $t + 1$ based upon the full empirical model that includes trend patterns in responsiveness (the T superscript refers to "trend"), and θ_{et+1}^{NT} is the predicted employment share for establishment e in period $t + 1$ predicted by the estimated model with parameters reflecting responsiveness at the beginning of the sample period (that is, we set the trend terms δ_{ij} equal to zero, so NT means "no trend").[38] The employment share prediction for an establishment in a given period ($\theta_{e,t+1}^T$) is based on the actual realizations of productivity and initial employment for that establishment in the previous period, fed through the estimated growth rate model.[39]

This diff-in-diff object is distinct from changes in the Olley-Pakes (OP) covariance because it reflects the changing marginal responsiveness from our estimated empirical model. The standard OP decomposition approach is based on the weighted average of establishment- (or firm-) level productivity. As shown in Appendix B, this weighted index is equivalent to standard industry measures of productivity, defined as industry output per unit input, only under constant returns to scale and perfect competition. Under these assumptions, the marginal revenue product of the composite input at the establishment level does not change with the level of inputs. An implication is that, in the absence of frictions, all inputs should be allocated to the most productive establishment. In contrast, under revenue function curvature from decreasing returns to scale and/or imperfect competition, the weighted average of establishment-level productivity is not equivalent to industry productivity. A corollary is that under revenue function curvature it is generally not optimal to allocate all inputs to the most productive establishment, even in the absence of frictions. Instead, inputs should be reallocated to establishments with higher marginal productivity.

These properties imply that, for any increase in frictions, the standard OP covariance declines more quickly than true aggregate productivity in the presence of revenue curvature. We

---

[38] We set the cyclical effects to zero by setting the state-level change in unemployment to zero.

[39] This approach is related to other accounting productivity decompositions in the literature (see, e.g., Foster, Haltiwanger, and Krizan (2001) for a review). Our present approach focuses only on model-driven reallocation arising from variation in productivity across businesses, holding constant the productivity distribution. Decker et al. (2017) use the Dynamic Olley-Pakes (DOP) decomposition developed by Melitz and Polanec (2015) to show that these accounting decompositions also imply a decline in the contribution of the change in the covariance terms after 2000. Alon et al. (2017) use the DOP decomposition to study the cumulative contribution of changes in entry rates.

show in Appendix B (Figure B5) that the diff-in-diff counterfactual in (5), on the other hand, closely tracks the true aggregate productivity effects of changing adjustment costs and responsiveness in our benchmark model. The reason for the superior performance of our diff-in-diff counterfactual versus the standard OP covariance is intuitive: unlike the standard OP covariance, our diff-in-diff approach uses (estimated) optimal policy functions that reflect the impact of revenue function curvature on marginal revenue products. It is easily shown in the model that the responsiveness of establishment-level growth to realizations of productivity declines with greater curvature in the revenue function.

Another attractive feature of this diff-in-diff counterfactual is that it only captures the effect of time-varying responsiveness within firm age groups. Differences in responsiveness between young and mature firms will be present in both the counterfactual with and the counterfactual without the trend, as will the changing age structure of firms overall. Moreover, this diff-in-diff design mechanically abstracts from potential effects of changing TFP dispersion.

We report Δ_t^{t+1} for each year on Figure 7. For example, the observation for $t + 1 = 1981$ has $\Delta_{1980}^{1981} = 0$ because the trend variable begins then, and for high-tech the year 2001 again gives $\Delta_{2000}^{2001} = 0$. But the 2004 observation for high-tech shows that, given the productivity and size distributions of 2003, if responsiveness from 2003 to 2004 had been at the 1981 pace instead of the actual pace (as estimated by our model) then the productivity index in 2004 would have been about half a log point higher ($\Delta_{2003}^{2004} = -0.005$). For high-tech manufacturing plants, the increasing responsiveness over the 1980s and 1990s yields an implied counterfactual increase in the index that peaks at about half a log point per year in the 1990s. The sharp decline in responsiveness during the post-2000 period implies a decline in the productivity index of as much as 2 log points per year by 2010. Some caution needs to be used in interpreting the magnitude at the end points—and certainly extrapolating out of sample—since the pattern in Figure 6 is driven by fitting a quadratic trend. But the drag on this index of industry level productivity due to the decline in responsiveness may be quite substantial.

Figure C4 in appendix C reports the same exercise but using the RPR productivity concept; in high-tech, the RPR results are quite similar—qualitatively and quantitatively—to the TFPR results from Figure 7, while outside of high-tech the productivity drag implied by the RPR regressions starts somewhat sooner. The basic message of the TFPR and RPR results is the same, however, particularly in high-tech: changing responsiveness has quantitatively large

implications for aggregate productivity. In high-tech, changing responsiveness starts to be a drag on productivity around 2003, about the time that Fernald (2014) finds a trend break in productivity growth in the IT sector. Outside of high-tech, both the TFPR and RPR results show a decline in aggregate productivity from the 1980s to the 2000s from declining responsiveness.

Some caution should be used in interpreting our counterfactual results as yielding patterns that mimic actual aggregate (industry-level) productivity growth since there may be changes in the within-plant productivity components of aggregate (industry-level) growth that we have not estimated in this context. Fernald (2014), Byrne et al. (2016) and Gordon (2016) highlight many factors that are likely contributing to within-plant (and within-firm) declines in productivity growth in the post-2000 period. In addition to the factors they emphasize, there may be a role for declining entrepreneurship in declining within-firm productivity growth given the contribution of young firms to innovative activity (Acemoglu et al. (2013) and Alon et al. (2017)). We examine the within-firm productivity growth patterns later in the paper.

D. Changing Business Models

It is possible that changing responsiveness reflects changes in business models that are benign for productivity. For example, perhaps businesses increasingly respond to shocks by adjusting their capital stock instead of labor (a form of capital/labor substitution). We repeat the regressions from equation (4), replacing the employment growth rate with the investment rate (investment divided by initial capital) and adding initial capital as an additional control.[40] This regression therefore includes the key state variables: productivity, employment, and capital.

Table 2 reports the regression results for high-tech manufacturing, and Figure 8 shows results analogously to Figure 6. As with employment, young firms' investment is more responsive than is mature firms'. Investment responsiveness in high-tech manufacturing displays a qualitatively similar pattern to employment responsiveness, with a significant decline in the 2000s among young firms: In the 1990s, a young-firm plant with a TFP draw one standard deviation above its industry-year mean had an equipment investment rate 8 percentage points higher than the plant at the mean; this differential is about 3 percentage points in the post-2000 period. The decline in employment responsiveness was not accompanied by stronger investment responsiveness in high-tech manufacturing. Among non-tech manufacturing businesses,

---

[40] See appendix D for more detail. Note that the theory linking adjustment costs to employment growth applies equally to investment in models of firm dynamics (Cooper and Haltiwanger (2006)).

however, there is rising investment responsiveness from the 1980s to the 1990s, with responsiveness remaining elevated in the 2000s, suggesting that capital-labor substitution may play some role outside of the high-tech sector. More broadly, we cannot rule out other forms of capital investment—like intangibles—as substitute adjustment mechanisms.

In appendix D, we describe two other exercises exploring changes in business model. First, we find mixed evidence that industries facing increased import competition saw bigger declines in responsiveness, suggesting that globalization may be an interesting avenue for future work. Second, we find no evidence that industry composition shifts within high-tech manufacturing explain declining responsiveness.

V. Beyond Manufacturing

Thus far we have focused on the manufacturing sector for which we have high-quality TFP data. An important question is whether the patterns of productivity dispersion and responsiveness we have described are present outside manufacturing. For example, the information sector has been a key contributor to U.S. innovation in recent years. Moreover, changes in startup rates and in the dispersion and skewness of firm growth rates are even more dramatic in non-manufacturing components of the high-tech sector (Decker et al. (2016)).

We next conduct the same exercises as in section IV, but with firm-level gross output per worker as our productivity concept and with the full private nonfarm sector as our sample.[41] Output per worker cannot be used to directly track the pattern of shocks, but our adjustment cost framework in appendix B shows that moments based on output per worker move systematically with changes in adjustment frictions and shocks. We employ RE-LBD data (described in Section III), which permit the measurement of revenue per worker at the firm level for the entire U.S. private, non-farm sector from the mid-1990s to 2013. A contribution of this section is new evidence on the relationship between productivity and reallocation dynamics outside manufacturing.[42]

---

[41] We omit finance, insurance, and real estate (NAICS 52-53) from our sample. This exercise requires that we assign each firm an industry code; we do this by choosing the industry that accounts for the largest share of the firm's employment. An alternative approach is to construct firm-level within-industry labor productivity as the deviation of firm output per worker from a full set of relevant industry fixed effects. Our results are robust to this alternative approach to controlling for firm industry activity.

[42] Relatively little is known about these issues outside manufacturing. Exceptions include several retail trade studies (Foster, Haltiwanger, and Krizan (2006), Jarmin, Klimek, and Miranda (2009) and Foster et al. (2015)).

The inferences we draw in this section recognize that output per worker endogenously reflects not only TFP but also changes in adjustment frictions. However, as our benchmark adjustment cost model illustrates, several empirical moments based on output per worker are informative for changing adjustment frictions. First, recall from section IV that an increase in adjustment frictions implies an increase in the within-industry dispersion of labor productivity: adjustment frictions dampen the tendency for marginal revenue products to be equalized, implying higher labor productivity dispersion. Second, an increase in adjustment frictions also reduces the responsiveness of firm-level employment growth from t to $t + 1$ to the realization of revenue labor productivity in t (controlling for employment in t). Finally, an increase in adjustment frictions reduces the diff-in-diff counterfactual using labor productivity in a manner that tracks the implications for aggregate productivity. Our multiple moments approach, then, is still well suited to a study of labor productivity.

A. Productivity and growth at the firm level

Figure 9 reports the standard deviation of within-industry labor productivity for young and mature firms, in and out of high-tech; labor productivity dispersion has risen for each of these groups (note that our definition of high-tech now includes certain industry groups in services and information as well as manufacturing). Notably, unlike TFP, labor productivity is more dispersed among young than mature firms; younger firms likely face greater learning or other frictions and may also be more heterogeneous in capital intensity.

To provide perspective on the relationship between the findings in this section using firm-level data for the private, non-farm sector and the earlier analysis using establishment-level manufacturing data, Figure A6 in Appendix A reports within-industry revenue labor productivity dispersion for manufacturing in both the RE-LBD and the ASM. Labor productivity dispersion has risen in both the firm-level administrative data and the establishment-level survey data, contrary to the Bils et al. (2017) hypothesis of rising measurement error in the ASM.[43]

Our finding of rising within-industry productivity dispersion is consistent with other work documenting increased differences between firms. For example, Andrews, Criscuolo, and Gal (2015) find a widening productivity gap between "frontier firms" and others, arguing that the

---

[43] Figure A6 shows that these findings are robust to employment weighting industries. Interestingly, rising revenue labor productivity dispersion is more apparent in the administrative data than the survey data, raising more doubts about the measurement error hypothesis for rising productivity dispersion from the survey data.

pace of technological diffusion has slowed. While the diffusion hypothesis could play a role, our estimates of TFP persistence (Appendix A Figure A4) suggest that the group of "frontier firms" is sufficiently fluid to somewhat limit the diffusion story's explanatory power. Increased adjustment frictions is an alternative, but not mutually exclusive, explanation. Both explanations allow for a decoupling of technological progress and aggregate productivity growth.[44]

Rising labor productivity dispersion is evidence against the "shocks" hypothesis for falling reallocation in various U.S. sectors. We next estimate equation (4)—the regression we used to measure changing TFP responsiveness in manufacturing—except that we now use firm-level data (vs. establishment), labor productivity in place of TFP, all U.S. sectors (except finance, insurance and real estate), and only the years 1997-2013.[45]

Table 3 reports results of these regressions. The first two columns report regressions using the DHS growth rate denominator inclusive of exit; the last two columns report results using only a binary exit outcome as the dependent variable. Figure 10 graphically shows the time series pattern of the growth coefficients; as with TFP results, we report the growth rate differential between the firm one standard deviation above its industry mean and the mean. Growth is indeed related to revenue labor productivity, as theorized; that is, firms with higher output per worker are more likely to grow.[46] Figure A2 in Appendix A shows a strong relationship between labor productivity and exit as well. Young firms are particularly sensitive to labor productivity, including on the exit margin, indicating that labor productivity is correlated with selection mechanisms. Moreover, the relationship of labor productivity with growth and survival has weakened over time, particularly among young high-tech firms (where the growth differential has fallen by 10 percentage points), consistent with the TFP-based evidence from Section IV. This decline implies that responsiveness of young high-tech firms in 2013 is only about two thirds of what it was in the late 1990s. Broadly speaking, the evidence suggests that the survival and growth differential between high- and low-productivity firms is declining over time, particularly in high-tech.

---

[44] Andrews, Criscuolo, and Gal (2015) (ACG) provide evidence of rising productivity dispersion within broad sectors using ORBIS data on both labor productivity (similar to our approach here) and multifactor productivity (similar to our analysis in Section IV). ACG measure the difference between "frontier firms" and average firms, where the frontier firms are usually defined as the top 50 or 100 firms within a broad (2-digit) sector, and in the case of the U.S. their unit of analysis is actually the establishment (Pinto Ribeiro, Menghinello and De Backer (2010)).

[45] We also apply propensity score weights; see Section III for RE-LBD details.

[46] Growth differentials for labor productivity may seem large compared with TFP-based differentials from the previous section; this is partly because labor productivity dispersion is higher than TFP dispersion.

The data on both labor productivity dispersion and the relationship linking labor productivity with growth and survival indicate that the TFP-based patterns we found in manufacturing are likely to hold in other sectors. Again, the framework of our model, applied to multiple moments of evidence, suggests slowing reallocation is a symptom of increased frictions rather than changes in the distribution of idiosyncratic productivity shocks.

B. Reallocation and aggregate labor productivity

Following the approach from Section IV, we quantify the labor productivity regression results by relating them to aggregate productivity growth using the diff-in-diff counterfactual approach from equation (5). As shown in Figure B5 in Appendix B, in the calibrated model the diff-in-diff counterfactual using labor productivity tracks the impact of increased adjustment frictions on true aggregate (industry-level) productivity quite well.[47]

The diff-in-diff counterfactual for the high-tech (not just manufacturing) and non-tech sectors is presented in Figure 11. By 2013, the weakening responsiveness of growth and survival to productivity accounts for more than 5 log points in the diff-in-diff counterfactual. This implies that if responsiveness during 2012-2013 had been as strong as in 1996, aggregate productivity in 2013 would be 5 log points higher (given the firm distribution of 2012). In contrast to the TFP-based results from manufacturing, our labor productivity-based calculations for the entire economy show a similar pattern for firms inside and outside high-tech. In unreported results we find that this is driven by particularly strong declines in the sensitivity of exit to productivity among firms outside high-tech; moreover, within manufacturing specifically, we do find stronger results in high-tech than outside of it in the labor productivity counterfactuals, as in the TFP counterfactuals.

C. Changing Patterns of Within-Firm Productivity Growth

Has a falling productivity contribution from reallocation been offset by stronger within-firm productivity growth? We use the firm labor productivity database to construct two related but distinct measures of within-firm productivity growth. The first measure is the simple unweighted mean of annual within-firm productivity growth. The second is the employment-weighted mean of annual within-firm productivity growth using time-t employment weights for

---

[47] Though we are using the diff-in-diff counterfactual, it is useful to note that for gross output per worker the weighted mean of micro productivity tracks gross output per worker at the industry level quite well (see Figure A3 in Appendix A)

productivity growth from t to $t + 1$. We compute these measures at the 6-digit industry level then aggregate using time-invariant employment weights for each industry.[48]

Figure 12 shows trends in within-firm productivity growth, both weighted and unweighted, for the average industry, separately for high-tech and non-tech. For high-tech, within-firm productivity growth declines using both measures; for non-tech, weighted within-firm productivity growth declines, but the unweighted measure exhibits less systematic variation. The weighted measure is much larger than the unweighted measure for both tech and non-tech, and the unweighted measure is always negative for non-tech and turns negative for high-tech early in the sample; this might be surprising since it implies negative productivity growth for the average firm. However, as discussed by Decker et al. (2017), the unweighted measure overwhelmingly describes very small firms (more than 90 percent of firms have fewer than 20 employees).[49] In sum, within-firm improvements (e.g., innovation by incumbents) have not quickened to compensate for weaker reallocation.[50]

VI. Conclusion

Reallocation has declined in all sectors—particularly the high-tech sector—since the early 2000s. In the 1980s and 1990s, the declining overall pace of reallocation was dominated by sectors such as retail trade, while innovative high-tech sectors (including high-tech manufacturing) exhibited rising reallocation. Within-industry TFP dispersion has risen gradually in recent decades, both in and out of high-tech manufacturing, as has within-industry labor productivity dispersion throughout U.S. industries. The marginal employment growth response of businesses to idiosyncratic productivity draws has mimicked the pattern of aggregate reallocation over time, particularly among young firms, as has the relationship between productivity and exit. The decline in responsiveness is especially large in the high-tech sector, with the responsiveness of young firms in the post-2000 period only about half (manufacturing)

---

[48] Importantly, these measures of within-firm productivity growth exploit the RE-LBD's longitudinal links and are therefore distinct from exercises that measure average productivity growth among specific groups of firms.

[49] Decker et al. (2017) further show that the positive difference between the weighted and the unweighted means reflects a positive relationship between within-firm productivity growth and initial shares (i.e., larger firms have higher within-firm productivity growth).

[50] Alon et al. (2017) also use the Dynamic Olley Pakes decomposition method described by Melitz and Polanec (2015) to study the productivity slowdown. The authors show that declining entry has had a significant cumulative negative effect on aggregate productivity growth, consistent with our emphasis on the contribution of changing firm dynamics.

to two thirds (economy-wide) of the peak responsiveness in the 1990s. Counterfactual exercises imply that the decline in responsiveness yields a significant drag on aggregate (industry-level) productivity, as much as 2 log points in high-tech manufacturing and more than 5 log points economy-wide in recent years.

These novel facts, taken together and studied through the lens of standard models of firm dynamics, imply that changing reallocation is the result not of changes in the dispersion or intensity of idiosyncratic shocks but rather of changes in the dynamic responsiveness to those shocks arising from an increase in adjustment frictions. Moreover, the timing of reallocation and responsiveness patterns in high-tech is consistent with the timing of the productivity slowdown, which evidence indicates was driven by ICT-producing and using industries. Importantly, our evidence abstracts from the confounding effect of declining startup rates since we study responsiveness within firm age groups. Our main results are based on plant-level TFP in high-tech manufacturing, but the results extend to manufacturing broadly and, when focusing on firm-level labor productivity, to other U.S. industries more generally.

In addition to shedding light on the drivers of declining business dynamism, these findings comprise a novel contribution to the literature on the U.S. productivity slowdown in the post-2000 period. The cross-sector and time series dimensions of reallocation and responsiveness are consistent with the timing of aggregate productivity fluctuations, and our counterfactuals demonstrate the quantitative importance of this relationship. Notably, productivity dispersion within industries has risen in the post-2000 period, while a slowing pace of innovation would produce falling dispersion in a Gort and Klepper (1982) framework. Slowing reallocation and business-level responsiveness is, at least, complementary to innovation-based explanations for the productivity slowdown such as Gordon (2016) (and is consistent with Byrne, Fernald and Reinsdorf (2016) and Syverson (2016), who find that the productivity slowdown is not an artifact of mismeasurement).

We document several other interesting patterns. The responsiveness of equipment investment to TFP in high-tech manufacturing follows a similar pattern to employment responsiveness, rising during the 1980s and 1990s then falling sharply after 2000, while investment responsiveness in non-tech manufacturing was flat throughout the 1990s and 2000s after rising in the 1980s. The strong relationship between growth and productivity that has previously been documented for TFP in manufacturing also holds for labor productivity in other

28

to two thirds (economy-wide) of the peak responsiveness in the 1990s. Counterfactual exercises imply that the decline in responsiveness yields a significant drag on aggregate (industry-level) productivity, as much as 2 log points in high-tech manufacturing and more than 5 log points economy-wide in recent years.

These novel facts, taken together and studied through the lens of standard models of firm dynamics, imply that changing reallocation is the result not of changes in the dispersion or intensity of idiosyncratic shocks but rather of changes in the dynamic responsiveness to those shocks arising from an increase in adjustment frictions. Moreover, the timing of reallocation and responsiveness patterns in high-tech is consistent with the timing of the productivity slowdown, which evidence indicates was driven by ICT-producing and using industries. Importantly, our evidence abstracts from the confounding effect of declining startup rates since we study responsiveness within firm age groups. Our main results are based on plant-level TFP in high-tech manufacturing, but the results extend to manufacturing broadly and, when focusing on firm-level labor productivity, to other U.S. industries more generally.

In addition to shedding light on the drivers of declining business dynamism, these findings comprise a novel contribution to the literature on the U.S. productivity slowdown in the post-2000 period. The cross-sector and time series dimensions of reallocation and responsiveness are consistent with the timing of aggregate productivity fluctuations, and our counterfactuals demonstrate the quantitative importance of this relationship. Notably, productivity dispersion within industries has risen in the post-2000 period, while a slowing pace of innovation would produce falling dispersion in a Gort and Klepper (1982) framework. Slowing reallocation and business-level responsiveness is, at least, complementary to innovation-based explanations for the productivity slowdown such as Gordon (2016) (and is consistent with Byrne, Fernald and Reinsdorf (2016) and Syverson (2016), who find that the productivity slowdown is not an artifact of mismeasurement).

We document several other interesting patterns. The responsiveness of equipment investment to TFP in high-tech manufacturing follows a similar pattern to employment responsiveness, rising during the 1980s and 1990s then falling sharply after 2000, while investment responsiveness in non-tech manufacturing was flat throughout the 1990s and 2000s after rising in the 1980s. The strong relationship between growth and productivity that has previously been documented for TFP in manufacturing also holds for labor productivity in other

sectors. The decline in productivity-enhancing reallocation has not been offset by stronger within-firm labor productivity growth.

We do not study specific policy or other factors that may be contributing to declining responsiveness, a large task that we leave for future work. Our theoretical framework focuses on increased adjustment costs, but broader interpretations may be plausible. Globalization may have played a role in subdued business-level growth responsiveness by facilitating cross-border factor adjustment (see Appendix D). Other explanations could include any forces that raise the cost of, or reduce the incentive for, factor adjustment; possibilities include unlawful discharge regulations, occupational licensing rules, scope of intellectual property protection, land use regulations, rules or norms that increase job match specificity, or various other state or federal regulations.[51] Declining intensity of competition or increased prevalence of winner-take-all economics could also produce some of the empirical patterns we document (as argued by De Loecker and Eeckhout (2017)). Given the implications of declining responsiveness for productivity growth, this is an important area for future research.

We conclude by reemphasizing the strength of our multiple moments approach to studying changing business dynamism. For example, rising markups are a potential mechanism that could alternatively account for declining responsiveness and reallocation. However, a rise in average markups does not inherently predict rising within-industry revenue labor productivity dispersion.[52] While this does not rule out the potential role of changing competition and markups, it is a reminder that competing explanations need to confront the changing patterns of multiple moments that are accompanying the changing patterns of business dynamism.

---

[51] Using industry variation, Goldschlag and Tabarrok (2014) find no evidence that federal regulation counts relate with changes in the pace of gross flows, but both state and federal regulation may present further scope for research. Davis and Haltiwanger (2014) find evidence relating employment protection policies to lower rates of reallocation, consistent with earlier work by Martin and Scarpetta (2012) and Autor, Kerr and Kugler (2007). Kleiner and Krueger (2013) review occupational licensing data and research. Molloy et al. (2016) find that states with tighter land use restrictions as of the early 2000s did not see larger declines in labor flows in recent decades, but the effect of changes in land use regulations is unknown; Hsieh and Moretti (2017) estimate significant static misallocation from land use regulations.

[52] This argument is most transparent in Appendix B where we consider a frictionless model. In that case, a rise in markups leads to declining responsiveness and reallocation but no change in revenue labor productivity dispersion (since, in the absence of frictions and distortions, labor productivity is equal across firms). We include some analysis and discussion of markups in Appendix B. An accompanying increase in dispersion in markups can cause an increase in revenue labor productivity dispersion.

References

Ackerberg, Daniel, Kevin Caves, and Garth Fraser. 2015. "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83(6): 2411–2451.

Alon, Titan M., David Berger, Rob Dent, and Benjamin Pugsley. 2017. "Older and Slower: The Startup Deficit's Lasting Effects on Aggregate Productivity Growth." Forthcoming, *Journal of Monetary Economics.*

Acemoglu, Daron, Ufuk Akcigit, Nicholas Bloom, and William R. Kerr. 2013. "Innovation, Reallocation, and Growth." NBER Working Paper No. 18993.

Andrews, Dan, Chiara Criscuolo, and Peter N. Gal. 2015. "Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries." OECD: *The Future of Productivity: Main Background Papers.*

Autor, David H., William R. Kerr, and Adriana D. Kugler. 2007. "Does Employment Protection Reduce Productivity? Evidence from US States." *The Economic Journal* 117 no. 521:F189-F217.

Bachmann, Rüdiger and Christian Bayer. 2013. "Wait-and-See Business Cycles." *Journal of Monetary Economics* 60(6): 704-19.

Baily, Martin Neil, Charles Hulten, and David Campbell. 1992. "Productivity Dynamics in Manufacturing Plants." In *Brookings Papers on Economic Activity: Microeconomics*, ed. Clifford Winston and Martin Neil Baily. Washington, DC: Brookings Institution Press.

Barro, Robert J. 1972. "A Theory of Monopolistic Price Adjustment." *The Review of Economic Studies* 39(1): 17-26.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta. 2013. "Cross-Country Differences in Productivity: The Role of Allocation and Selection." *American Economic Review* 103 no. 1:305-34.

Berger, David and Joseph Vavra. 2017. "Shocks vs. Responsiveness: What Drives Time-Varying Dispersion," mimeo.

Bernard, Andrew B., J. Bradford Jensen, and Peter K. Schott. 2006. Survival of the Best Fit: Exposure to Low-Wage Countries and the (Uneven) Growth of US Manufacturing Plants. J*ournal of International Economics* 68(1): 219-237.

Bils, Mark, Peter J. Klenow, and Cian Ruane. 2017. "Misallocation or Mismeasurement?" Working paper.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 no. 3:623-85.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2016. "Really Uncertain Business Cycles." Working paper.

Byrne, David M. 2015. "Domestic Electronics Manufacturing: Medical, Military, and Aerospace Equipment and What We Don't Know about High-Tech Productivity," FEDS Notes 2015-06-02. Board of Governors of the Federal Reserve System.

Byrne, David M. and Carol A. Corrado. 2015. "Recent Trends in Communications Equipment Prices." *FEDS Notes*, no. 2015-06-02.

Byrne, David M. and Carol A. Corrado. 2016. "ICT Prices and ICT Services: What Do They Tell Us about Productivity and Technology?" Working paper.

Byrne, David M., John G. Fernald and Marshall B. Reinsdorf. 2016. "Does the United Have a Productivity Slowdown or a Measurement Problem?" *Brookings Papers on Economic Activity*, Spring 2016.

Cooper, Russell and John Haltiwanger. 2006. "On the Nature of Capital Adjustment Costs." *Review of Economic Studies* 73(3); 611-33.

Cooper, Russell, John Haltiwanger, and Jonathan Willis. 2007. "Search Frictions: Matching Aggregate and Establishment Observations." *Journal of Monetary Economics* 54 Supplement 1: 56-78.

Davis, Steven J. and John Haltiwanger. 2014. "Labor Market Fluidity and Economic Performance." NBER Working Paper 20479 (published in the 2014 Federal Reserve Bank of Kansas City Jackson Hole Conference Symposium).

Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda. 2007. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." Chap. 2 in *NBER Macroeconomics Annual 2006* edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, vol. 21, pp. 107-80. Cambridge, MA: MIT Press.

Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction.* Cambridge, MA: MIT Press.

De Loecker, Jan and Jan Eeckhout. 2017. "The Rise of Market Power and the
Macroeconomic Implications." NBER Working Paper no. 23687.

Decker, Ryan, John Haltiwanger, Ron Jarmin and Javier Miranda. 2014. "The Role of
Entrepreneurship in US Job Creation and Economic Dynamism" *Journal of Economic
Perspectives* 28(3) 3-24.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2016. "Where Has All
the Skewness Gone? The Decline in High-Growth (Young) Firms in the U.S.*" European
Economic Review* 86 (July) 4-23.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2017. "Declining
Dynamism, Allocative Efficiency, and the Productivity Slowdown*" American Economic
Review Papers and Proceedings* (May).

Elsby, Michael W. L. and Ryan Michaels. 2013. "Marginal Jobs, Heterogeneous Firms, and
Unemployment Flows." *American Economic Journal: Macroeconomics* 5 no. 1:1-48.

Fernald, John. 2014. "Productivity and Potential Output Before, During, and After the Great
Recession." Chap. 1 in *NBER Macroeconomics Annual 2014* edited by Jonathan A.
Parker and Michael Woodford, vol. 29, pp. 1-51. Cambridge, MA: MIT Press.

Fort, Teresa and Shawn Klimek. 2016. "The Effects of Industry Classification Changes
on US employment composition." Mimeo.

Foster, Lucia, Cheryl Grim, and John Haltiwanger. 2016. "Reallocation in the Great
Recession: Cleansing or Not?" *Journal of Labor Economics.*

Foster, Lucia, Cheryl Grim, John Haltiwanger, and Zoltan Wolf. 2017. "Micro and Macro
Dynamics of Productivity," NBER Working Paper No. 23666.

Foster, Lucia, John Haltiwanger, Shawn Klimek, C.J. Krizan and Scott Ohlmacher. 2015.
"The Evolution of National Chains: How we got here", Center for Economic
Studies Working Paper No. CES-WP-15-10.

Foster, Lucia, John Haltiwanger, and C.J. Krizan. 2001. "Aggregate Productivity
Growth: Lessons from Microeconomic Evidence." In *New Developments in
Productivity Analysis*, eds. Edward Dean, Michael Harper and Charles Hulten,
303-363. Chicago: University of Chicago Press.

Foster, Lucia, John Haltiwanger, and C.J. Krizan. 2006. "Market Selection, Reallocation and Restructuring in the U.S. Retail Trade Sector in the 1990s." *The Review of Economics and Statistics* 88(4): 748-58.

Foster, Lucia, John Haltiwanger, and Chad Syverson, 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1): 394-425.

Gandhi, Amit, Salvador Navarro, David Rivers, (2016), "On the Identification of Production Functions: How Heterogeneous is Productivity?", mimeo.

Goldschlag, Nathan and Alex Tabarrok. 2014. "Is Regulation to Blame for the Decline in American Entrepreneurship?" GMU Working Paper in Economics No. 15-11.

Gopinath, Gita, Sebnem Kalemli-Ozcan, Loukas Karabarbounis and Carolina Villegas-Sanchez 2017. "Capital Allocation and Productivity in South Europe" *Quarterly Journal of Economics*, (forthcoming).

Gordon, Robert J. 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living since the Civil War*. Princeton: Princeton University Press.

Gort, Michael and Steven Klepper. 1982. "Time Paths in the Diffusion of Product Innovations," *The Economic Journal,* 92(367): 630-653.

Guzman, Jorge and Scott Stern. 2016. "The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 15 US States, 1988-2014." *Mimeo*.

Haltiwanger, John, Ian Hathaway, and Javier Miranda. 2014. "Declining Business Dynamism in the U.S. High-Technology Sector." The Kauffman Foundation.

Haltiwanger, John, Ron S. Jarmin, Robert Kulick and Javier Miranda. 2017. "High Growth Young Firms: Contribution to Job Growth, Revenue Growth and Productivity" In *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges* edited by John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar. Chicago: University of Chicago Press.

Hecker, Daniel. 2005. "High-technology employment: a NAICS-based update," *Monthly Labor Review* (U.S. Dept. of Labor and U.S. Bureau of Labor Statistics), Volume 128, Number 7, July 2005: 58.

Hugo A. Hopenhayn. 1992. "Entry, Exit, and firm Dynamics in Long Run Equilibrium." *Econometrica*, Vol. 60(5): 1127-50.

Hopenhayn, Hugo and Richard Rogerson. 1993. "Job Turnover and Policy Evaluation: A General Equilibrium Analysis." *Journal of Political Economy* 101(5): 915-38.

Hsieh, Chang-Tai and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics*, 124(4): 1403-48.

Hsieh, Chang-Tai and Peter J. Klenow. 2017. "The Reallocation Myth." 2017 Jackson Hole paper.

Hsieh, Chang-Tai and Enrico Moretti. 2017. "Housing Constraints and Spatial Misallocation." Mimeo.

Hyatt, Henry R. and James Spletzer. 2013. "The Recent Decline in Employment Dynamics." *IZA Journal of Labor Economics* 2(3): 1-21.

Ilut, Cosmin L., Matthias Kehrig, and Martin Schneider. Forthcoming. "Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News." *Journal of Political Economy*.

Jarmin, Ron S., Shawn Klimek, and Javier Miranda. 2009. "The Role of Retail Chains: National, Regional and Industry Results." Chap. 6 in *Studies in Income and Wealth* edited by Timothy J. Dunne, J. Bradford Jensen, and Mark J. Roberts, vol. 68, *Producer Dynamics: New Evidence from Micro Data,* pp. 237-62. Chicago, IL: National Bureau of Economic Research, University of Chicago Press.

Jarmin, Ron S., and Javier Miranda. 2002. "The Longitudinal Business Database." CES Working Paper no. 02-17. Washington, D.C.: Center for Economic Studies, U.S. Census Bureau.

Jovanovic, Boyan. 1982. "Selection and the evolution of industry." *Econometrica*, 50(3): 649–70.

Karahan, Fatih, Benjamin Pugsley, and Aysegul Sahin. 2015. "Understanding the 30-Year Decline in the Startup Rate: A General Equilibrium Approach." Mimeo.

Kleiner, Morris and Alan B. Krueger. 2013. "Analyzing the Extent and Influence of Occupational Licensing on the Labor Market." *Journal of Labor Economics* 31 no. 2:S173-S202.

Levinsohn, James and Amil Petrin (2003). "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies* 70 no. 2:317-41.

Martin, John P. and Stefano Scarpetta. 2012. "Setting It Right: Employment Protection, Labour Reallocation and Productivity." *De Economist* 160 no. 2:89-116.

Melitz, M. J., & Polanec, S. (2015). Dynamic Olley-Pakes productivity decomposition with entry and exit. *The RAND Journal of Economics*, *46*(2), 362–375.

Molloy, Raven, Christopher L. Smith, Riccardo Trezzi, and Abigail Wozniak. 2016. "Understanding Declining Fluidity in the U.S. Labor Market." *Brookings Papers on Economic Activity*.

Olley, G. Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Industry." *Econometrica,* 64 no. 6:1263-98.

Petrin, Amil, T. Kirk White, and Jerome P. Reiter. 2011. "The Impact of Plant-Level Resource Reallocations and Technical Progress on U.S. Macroeconomic Growth." *Review of Economic Dynamics* 14 no. 1:13-26.

Pierce, Justin and Peter Schott. 2016. "The Surprisingly Swift Decline of U.S. Manufacturing Employment." *American Economic Review* 107 no. 7:1632-62.

Pinto Ribeiro, Samuel, Stefano Menghinello and Koen De Backer. 2010. "The OECD ORBIS Database: Responding to the Need for Firm-Level Micro-Data in the OECD." *OECD Statistics Working Papers* No. 2010/01, OECD Publishing, Paris.

Restuccia, Diego and Richard Rogerson. 2008. "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments." *Review of Economic Dynamics* 11 no. 4:707-20.

Schott, Peter. 2008. "The Relative Sophistication of Chinese Imports." *Economic Policy*, 5-49.

Syverson, Chad. 2004. "Product Substitutability and Productivity Dispersion." *Review of Economics and Statistics* 86(2): 534-50.

Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49(2): 326-65.

Syverson, Chad. 2016. "Challenges to Mismeasurement Explanations for the U.S. Productivity Slowdown." NBER Working Paper no. 21974.

Vavra, Joseph. 2014. "Inflation Dynamics and Time-Varying Volatility: New Evidence and SS Interpretation." *Quarterly Journal of Economics* 129(1): 215-58.

Wooldridge, Jeffrey M. 2009. "On Estimating Firm-Level Productivity Functions Using Proxy Variables to Control for Unobservables." *Economics Letters* 104 no. 3:112-4.

Figure 1: Job reallocation patterns vary by sector
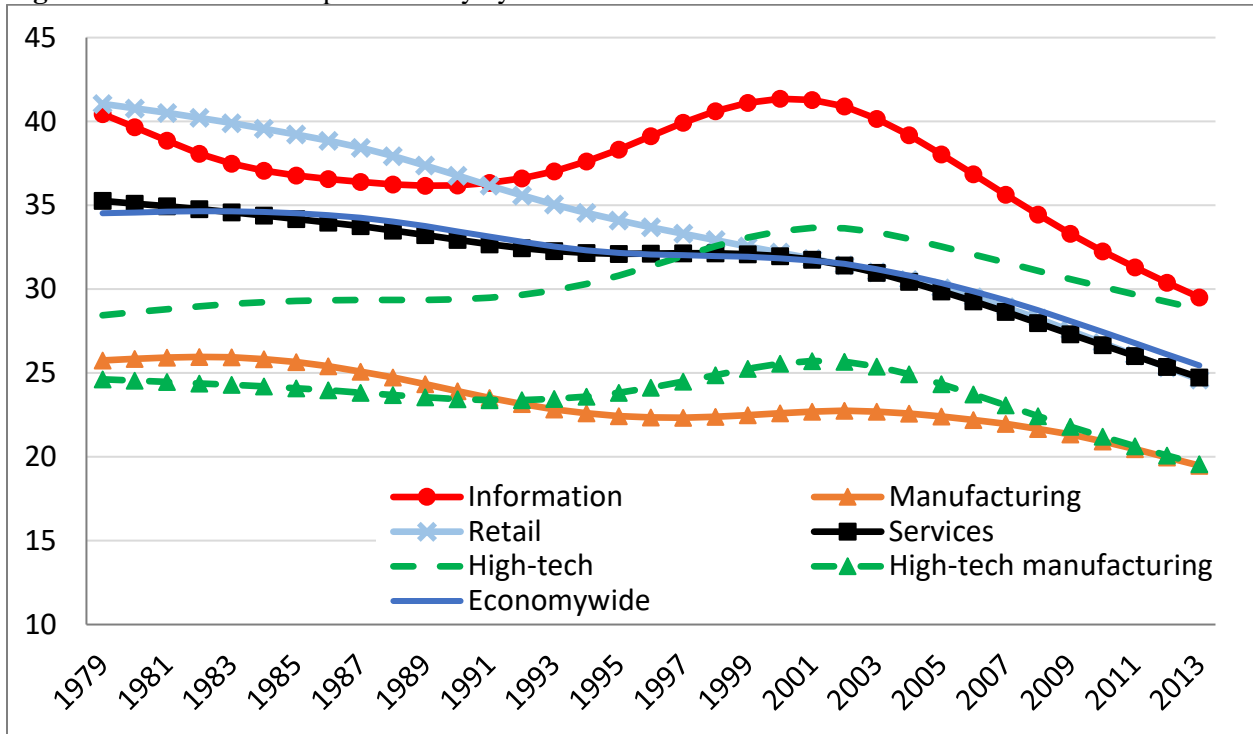
Note: Y axis does not start at zero. HP trends using parameter set to 100. Industries defined on a consistent NAICS basis; high-tech is defined as in Hecker (2005). Data include all firms (new entrants, continuers, and exiters). Author calculations from the Longitudinal Business Database (LBD).

Figure 2: Young firm share patterns vary by sector
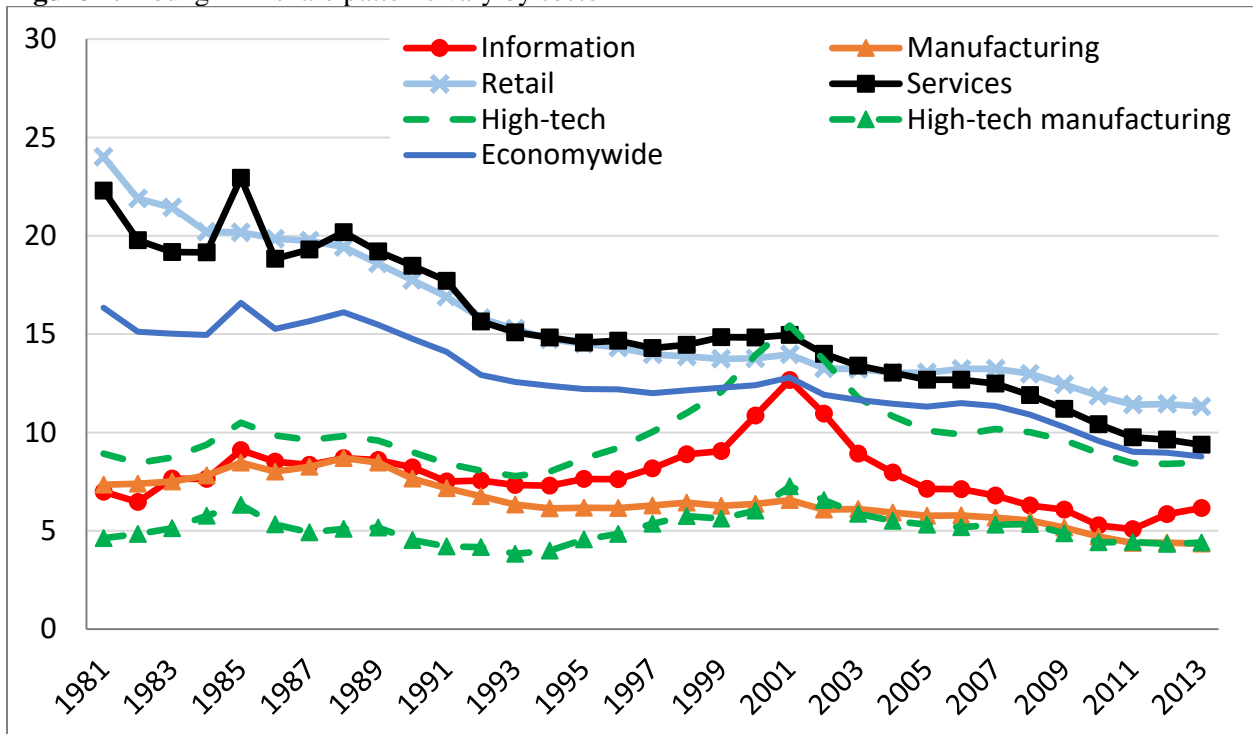
Note: Young firms have age less than 5. Industries are defined on a consistent NAICS basis; high-tech is defined as in Hecker (2005). Data include all firms (new entrants, exiters, and continuers). Author calculations from the LBD.

Figure 3: Most variation in job reallocation is not explained by changing startup rates
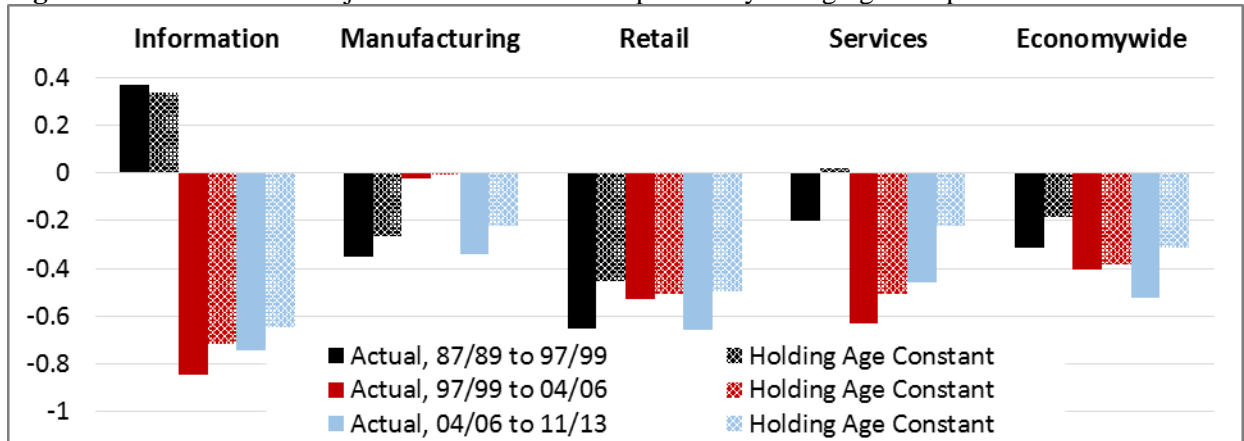
Note: Sectors are defined on a consistent NAICS basis. Author calculations from the LBD.

Figure 4a: Responses of key moments to changes in TFP dispersion in theoretical model
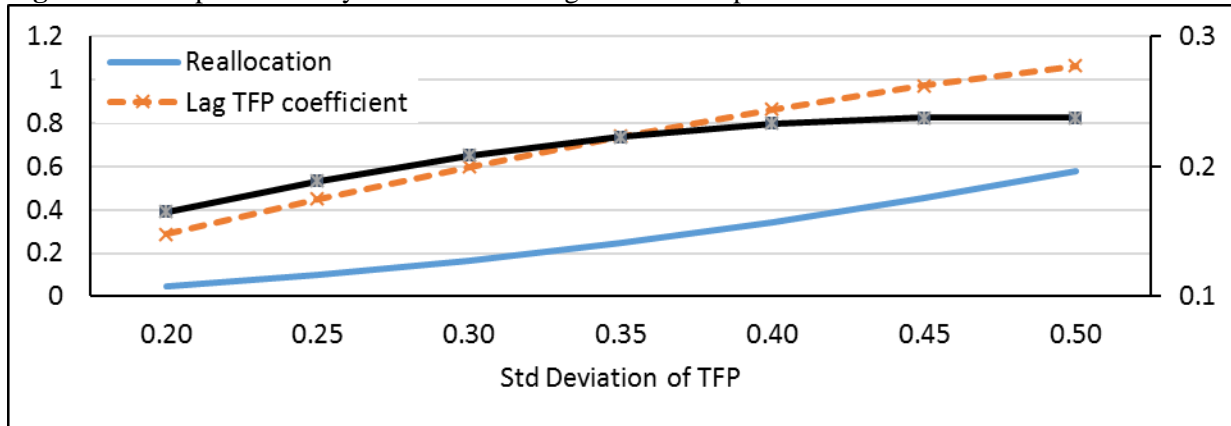
Note: Model with kinked adjustment costs ($F_+ = 0.85, F_- = 0$). General equilibrium model with flexible wage and inelastic labor supply. See Appendix B for details.

Figure 4b: Responses of key moments to changes in adjustment costs in theoretical model
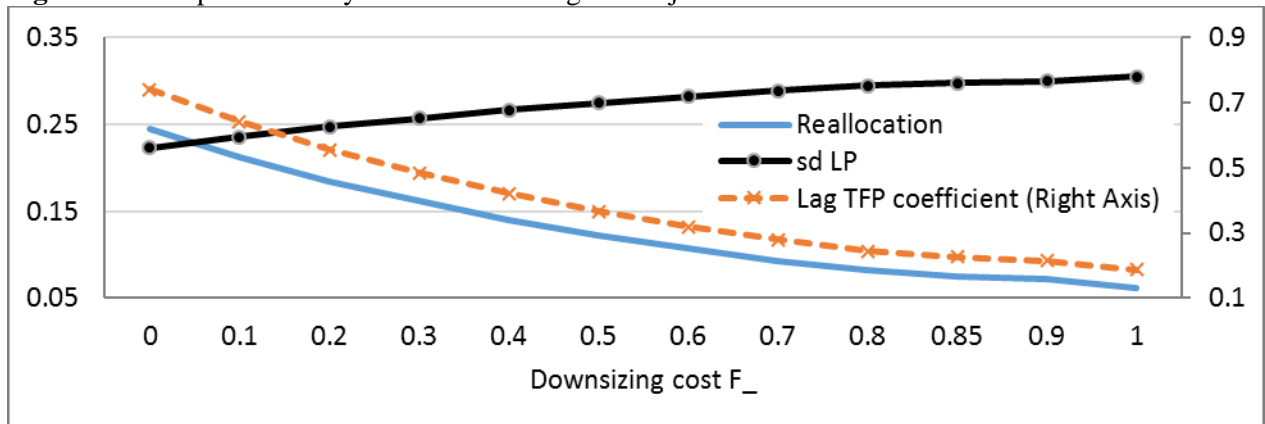
Note: The x axis reflects values of F_-, or the cost of reducing employment, holding the hiring cost F_+ fixed at $F_+ = 0.85$. General equilibrium model with flexible wage and inelastic labor supply. See Appendix B for details.

Figure 5: Within-industry TFP dispersion has risen (manufacturing)
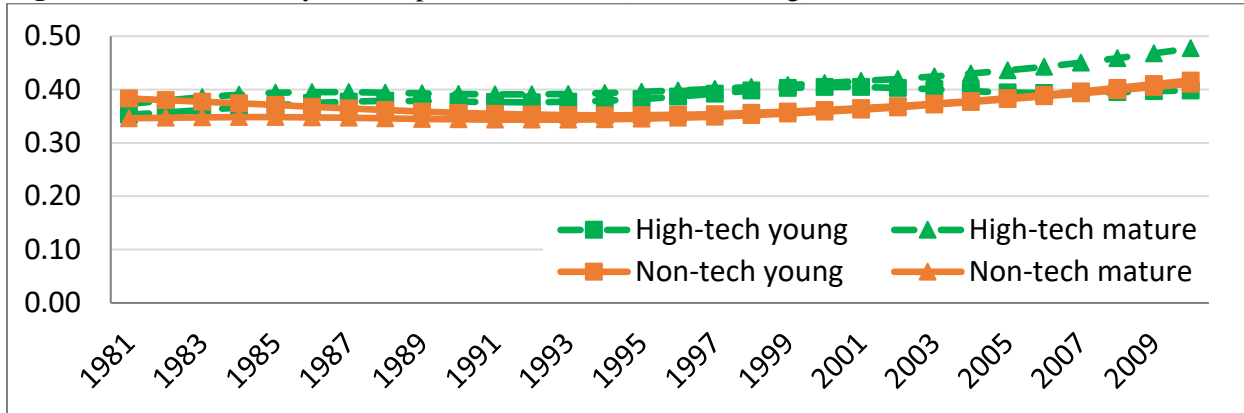
Note: Y axis does not start at zero. Young firms have age less than 5. Standard deviation of within-detailed industry log TFPR. High-tech defined as in Hecker (2005). Author calculations from the LBD, the Annual Survey of Manufacturers (ASM), and the Census of Manufacturers (CM). HP Trends.

Figure 6: Establishment job growth has become less responsive to TFP (manufacturing)
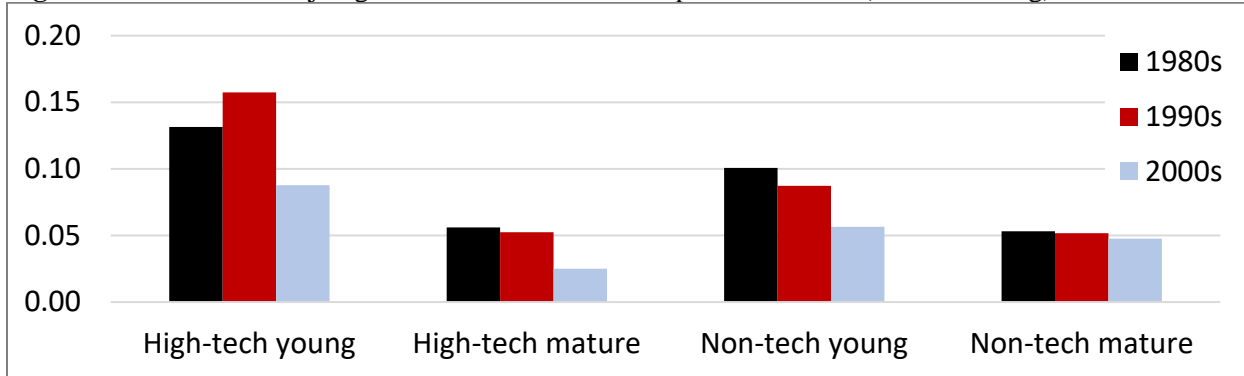
Note: Young firms have age less than 5. High-tech is defined as in Hecker (2005). Growth rate of plant with TFPR one std. dev. above industry mean vs. industry mean. Author calculations from the LBD, the ASM, and the CM.

Figure 7: Changing contribution of reallocation to aggregate TFP (manufacturing)
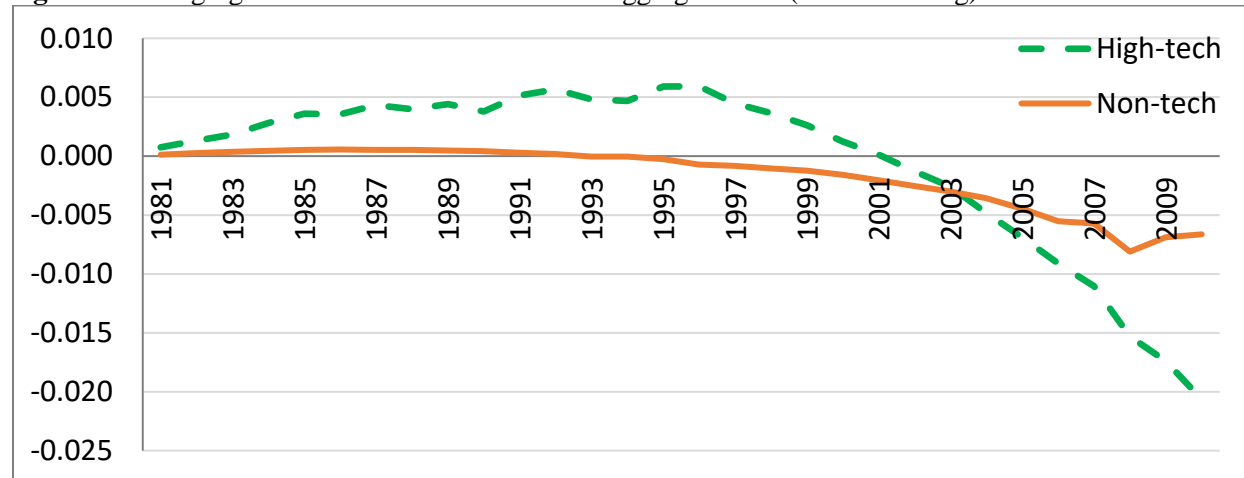
Note: Figure depicts diff-in-diff counterfactual as described in the text from TFPR concept. High-tech is defined as in Hecker (2005). Author calculations from the LBD, the ASM, and the CM.

Figure 8: Establishment investment rates have become less responsive to TFP (manufacturing)
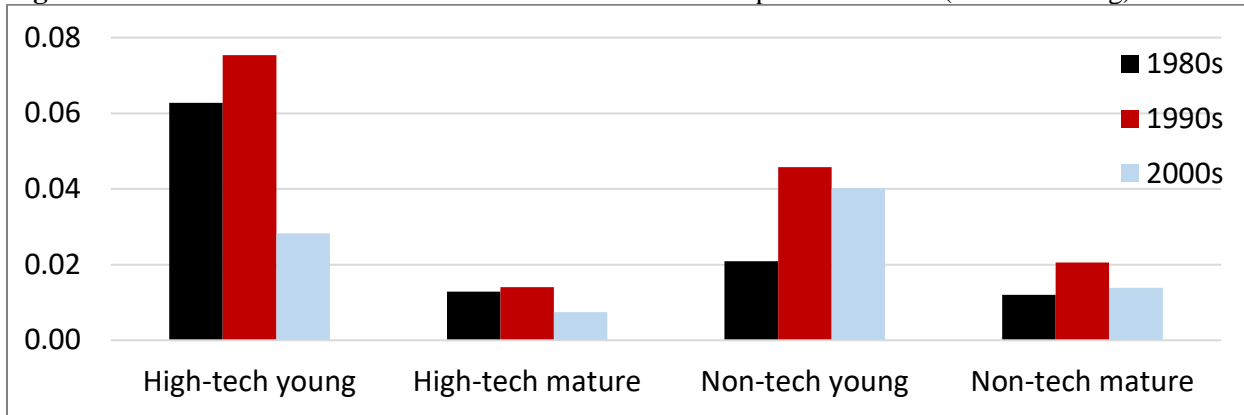
Note: Young firms have age less than 5. High-tech is defined as in Hecker (2005). Investment rate of plant with TFPR one std. dev. above industry mean vs. mean. Author calculations from the LBD, the ASM, and the CM.

Figure 9: Within-industry labor productivity dispersion has risen (economywide)
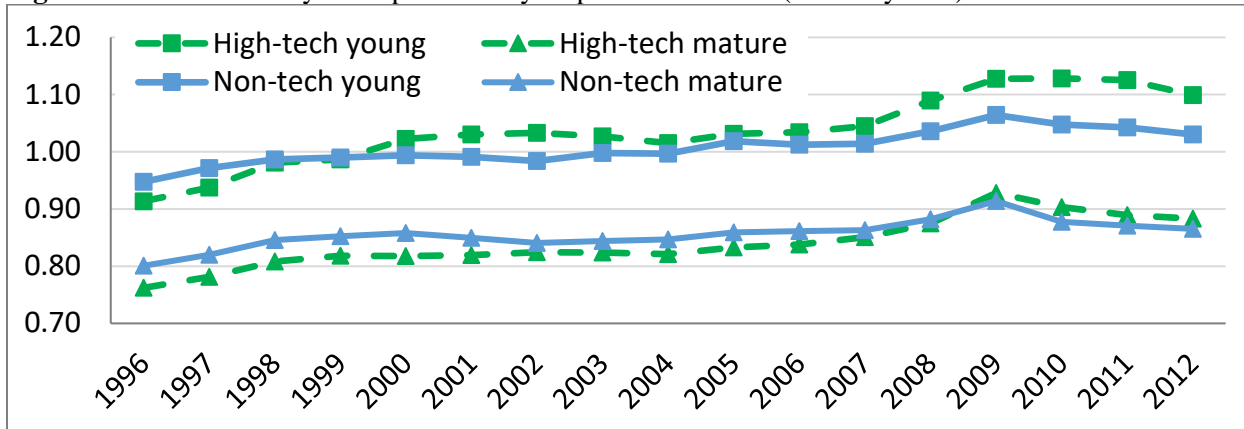
Note: Y axes do not begin at zero. Standard deviation of log labor productivity deviated from industry by year means. Young firms have age less than five. High-tech is defined as in Hecker (2005). Author calculations from the RE-LBD. Finance, Insurance and Real Estate (NAICS 52-53) omitted.

Figure 10: Firm growth has become less responsive to labor productivity (economywide)
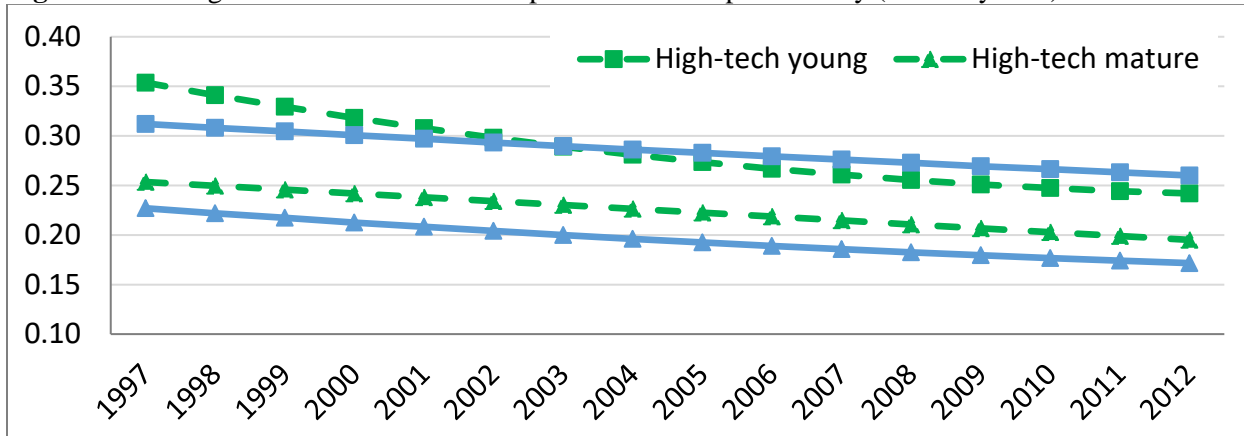
Note: Y axis does not start at zero. Growth rate of firm with labor productivity one std. dev. above industry mean vs. industry mean. Young firms have age less than five. High-tech defined as in Hecker (2005). Author calculations from the RE-LBD. Finance, Insurance and Real Estate (NAICS 52-53) omitted.

Figure 11: Changing contribution of reallocation to aggregate labor productivity (economywide)
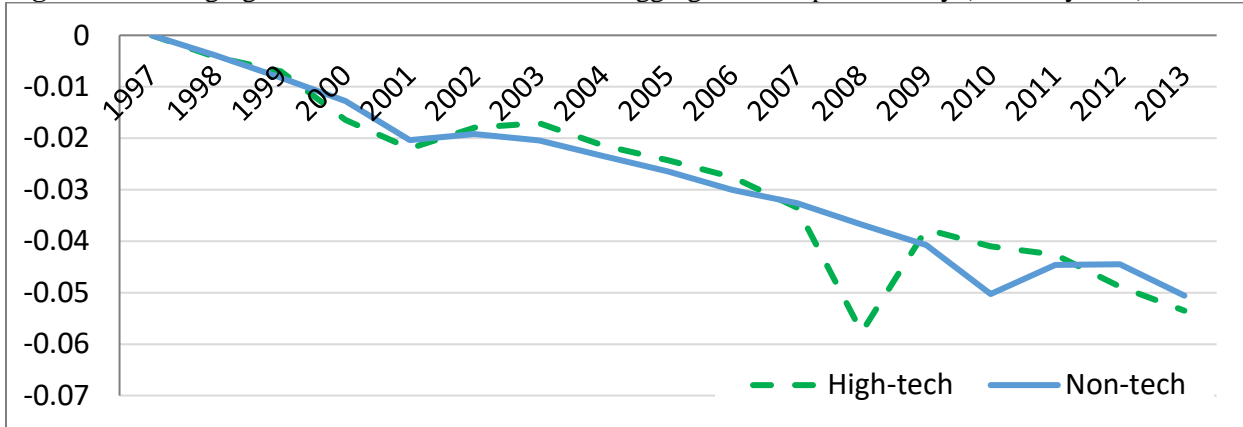
Note: Figure depicts diff-in-diff counterfactual as described in the text. High-tech is defined as in Hecker (2005). Author calculations from the RE-LBD. Finance, Insurance, and Real Estate (NAICS 52-53) omitted.

Figure 12: Within-firm productivity growth in the average industry (economywide)
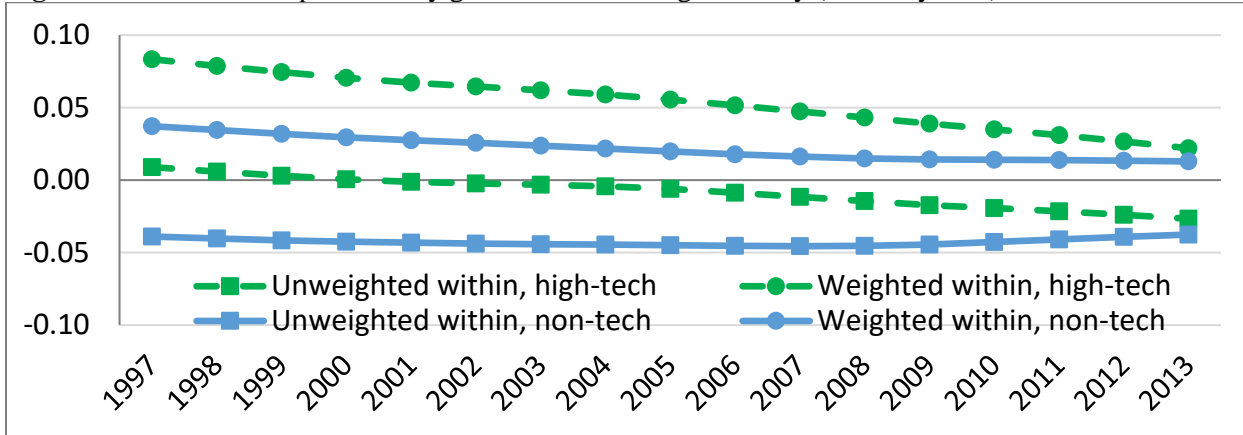
Note: Average within-firm productivity growth, with and without employment weights. Author calculations from the RE-LBD.

Table 1: Effect of Lagged Productivity on Plant-Level Employment Growth and Exit

	Growth including exit		Exit	
	High-tech	Non-tech	High-tech	Non-tech
TFP*Young	0.2025***	0.2732***	-0.0292*	-0.0905***
	(0.0390)	(0.0090)	(0.0162)	(0.0037)
TFP*Young*Trend	0.0317***	0.0016	-0.0160***	-0.0005
	(0.0061)	(0.0014)	(0.0025)	(0.0006)
TFP*Young*Trend^2	-0.0012***	-0.0002***	0.0005***	0.0001***
	(0.0002)	(0.00005)	(0.0001)	(0.00002)
TFP*Mature	0.1228***	0.1394***	-0.0403***	-0.0464***
	(0.0174)	(0.0043)	(0.0072)	(0.0018)
TFP*Mature*Trend	0.0054**	0.0016	-0.0016	-0.0012
	(0.0026)	(0.0007)	(0.0011)	(0.0003)
TFP*Mature*Trend^2	-0.0003***	-0.0001*	0.0001***	0.00005***
	(0.0001)	(0.00002)	(0.00003)	(0.00001)
N	125000	2055000	125000	2055000
R^2	0.059	0.032	0.098	0.054

Notes: Standard Errors in Parentheses. Dependent variable in Growth columns is DHS growth rate. Dependent variable in Exit columns is indicator=1 if exit, 0 Young firms have age less than 5. Unreported are estimates of controls including year effects, state effects, firm age dummies, firm size dummies, log plant level employment in period t, state cyclical indicators (change in state level unemployment rate), state cyclical indicators interacted with TFP. All variables that use TFP including all interactions are fully interacted with firm age dummies. Sample size rounded to nearest 5000 observations.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Estimated Effect of Productivity on Plant-Level Equipment Investment Rate

	High-tech	Non-tech
TFP*Young	0.0826***	-0.0125**
	(0.0236)	(0.0052)
TFP*Young*Trend	0.0189***	0.0156***
	(0.0037)	(0.0008)
TFP*Young*Trend^2	-0.0008***	-0.0004***
	(0.0001)	(0.0000)
TFP*Mature	0.0232**	0.0039
	(0.0105)	(0.0025)
TFP*Mature*Trend	0.0024	0.0067***
	(0.0016)	(0.0004)
TFP*Mature*Trend^2	-0.0001*	-0.0002***
	(0.00005)	(0.0000)
N	125000	20055000
R^2	0.068	0.047

Notes: Standard errors in parentheses. Young firms have age less than 5. Unreported are estimates of controls including year effects, state effects, firm age dummies, firm employment size dummies, log plant level employment in period t, dummies for initial capital, state cyclical indicators (change in state level unemployment rate), state cyclical indicators interacted with TFP. All variables that use TFP including all interactions are fully interacted with firm age dummies. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

**Table 3:** Lagged Labor Productivity and Firm-Level Employment Growth and Exit

	Growth including exit		Exit	
	High-tech	Non-tech	High-tech	Non-tech
LP*Young	0.3845***	0.3467***	-0.1258***	-0.1224***
	0.0020	0.0005	0.0009	0.0002
LP*Young*Trend	-0.0141***	-0.0043***	0.0026***	0.0014***
	0.0006	0.0001	0.0002	0.0001
LP*Young*Trend^2	0.0004***	0.0000***	-0.0001***	0.0000***
	0.0000	0.0000	0.0000	0.0000
LP*Mature	0.2755***	0.2522***	-0.0710***	-0.0758***
	0.0021	0.0004	0.0009	0.0002
LP*Mature*Trend	-0.0042***	-0.0056***	-0.0008***	0.0020***
	0.0006	0.0001	0.0002	0.0000
LP*Mature*Trend^2	0.0000	0.0001***	0.0001***	-0.0001***
	0.0000	0.0000	0.0000	0.0000
N	55385000	55385000	55385000	55385000
R^2	0.126	0.108	0.105	0.093

Dependent variable in all regressions is firm-level employment growth rate (DHS). All regressions include controls for state business cycle (change in state unemployment rate) and firm employment size in period t-1. Labor productivity is measured as the log difference from 6-digit NAICS industry mean. High-tech is defined as in Hecker (2005). Observations rounded to nearest five thousand.
*** $p<0.01$; ** $p<0.05$; * $p<0.10$

Appendix A.   Figures and tables to supplement the main text

Figure A1: Exit selection on TFP has weakened (manufacturing)
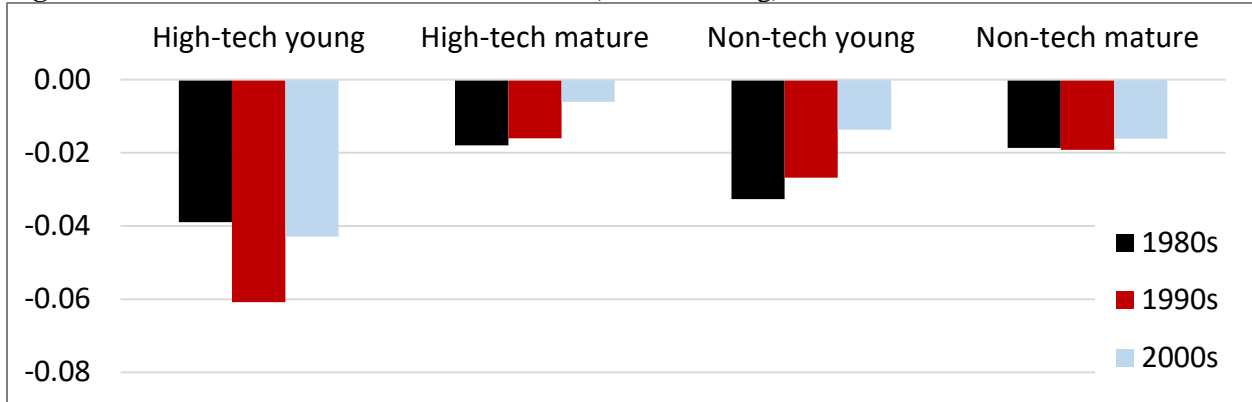
Note: Young firms have age less than 5. High-tech is defined as in Hecker (2005). Exit probability of plant with TFPR one std. dev. above industry mean vs. industry mean. Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure A2: Exit selection on labor productivity has weakened (economywide)
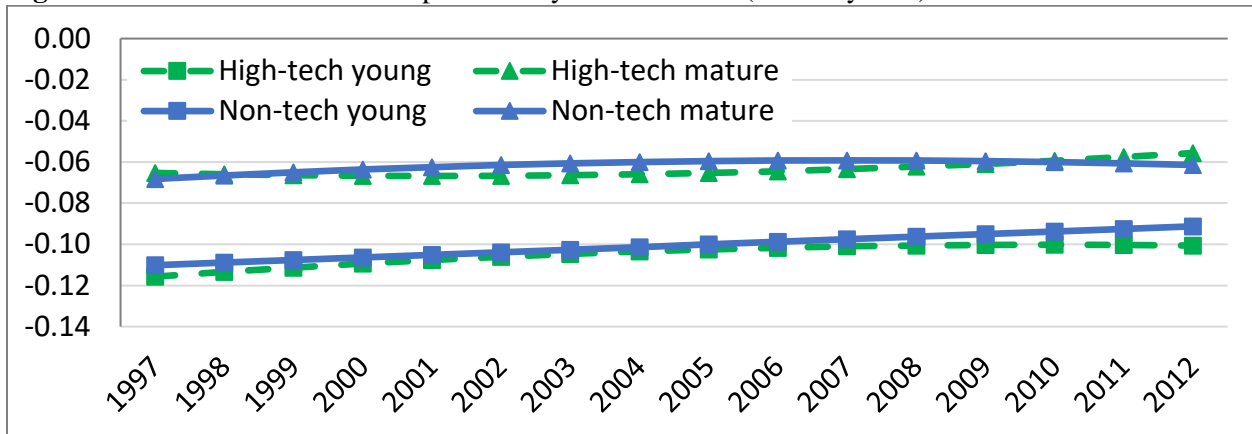
Note: Annual coefficients constructed from Table 3. Young firms have age less than five. High-tech defined as in Hecker (2005). Exit probability of plant with labor productivity one std. dev. above industry mean vs. industry mean. Author calculations from the RE-LBD. Finance, Insurance and Real Estate (NAICS 52-53) omitted.

Figure A3: Average industry-level productivity growth, BLS and aggregated microdata
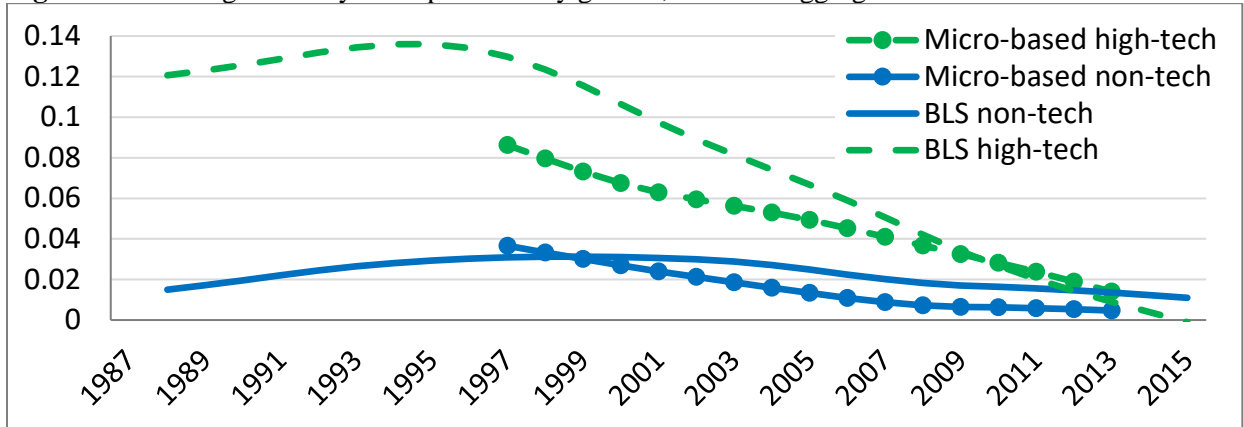
Source: BLS and author calculations from RE-LBD.

Figure A4: Little change in persistence of TFP (manufacturing)
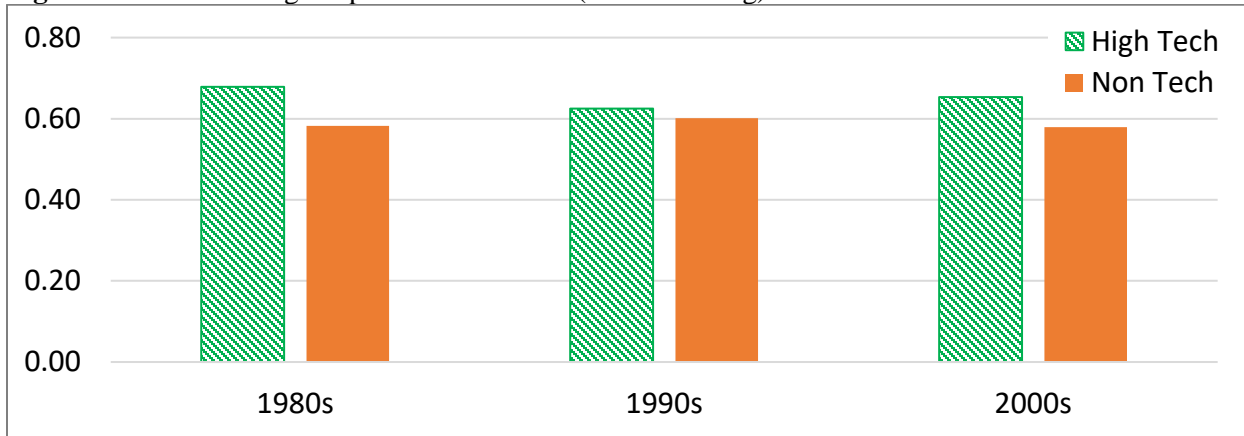
Note: High-tech is defined as in Hecker (2005). AR(1) coefficients for establishment TFPR, averaged by decade. The LBD-ASM-CM database is not ideally suited for estimating persistence since this requires relying on the longitudinal nature of the ASM/CM, which is less robust than the longitudinal properties of the LBD. That is, estimating productivity persistence parameters requires pairwise continuing plants in t and t+1 to be measured in the ASM/CM. The panel rotation of the ASM as well as Census years make this a challenge. That is, in the first years of a new ASM panel and in Census years we have a much smaller and less representative set of continuing plants than other years. For this exercise we exclude those years; even for other years, though, our propensity score weights are not ideally suited for making the sample of continuers representative. In principle, we can develop separate propensity score weights for this restricted sample of continuing plants. Doing so is more of a challenge, given the rotating nature of the ASM sample. See Figure B2 in appendix B for the same exercise on RPR productivity.

Figure A5: Standard deviation of innovations to plant-level TFPR
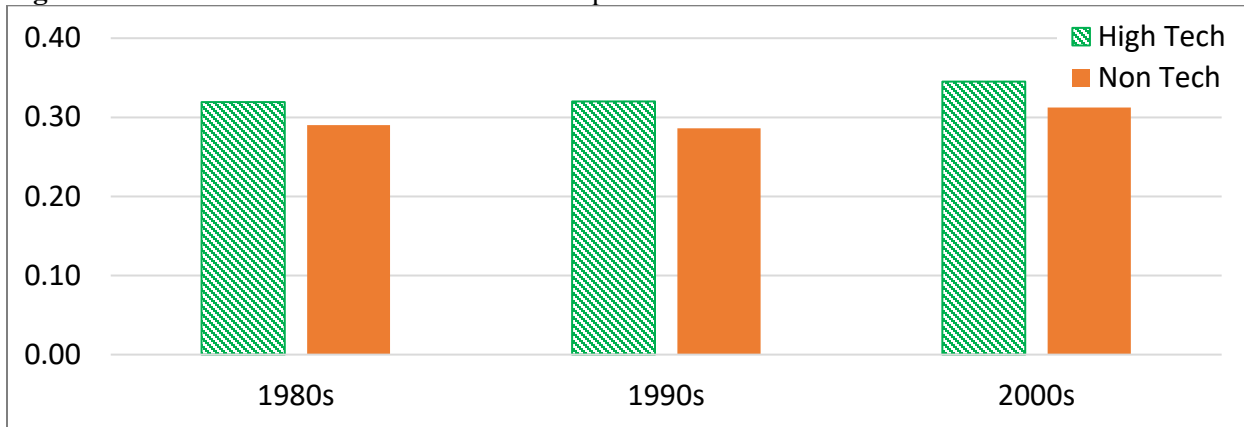
Note: High-tech is defined as in Hecker (2005). For the set of years where we can estimate the AR(1) process (see note for Figure A4), we can also recover the distribution of innovations to plant-level TFP for continuing plants. Since this is for selected years we report averages of standard deviation of innovations to TFP by decade as we did with persistence.

Figure A6: Rising labor productivity dispersion in survey and administrative data (manufacturing)
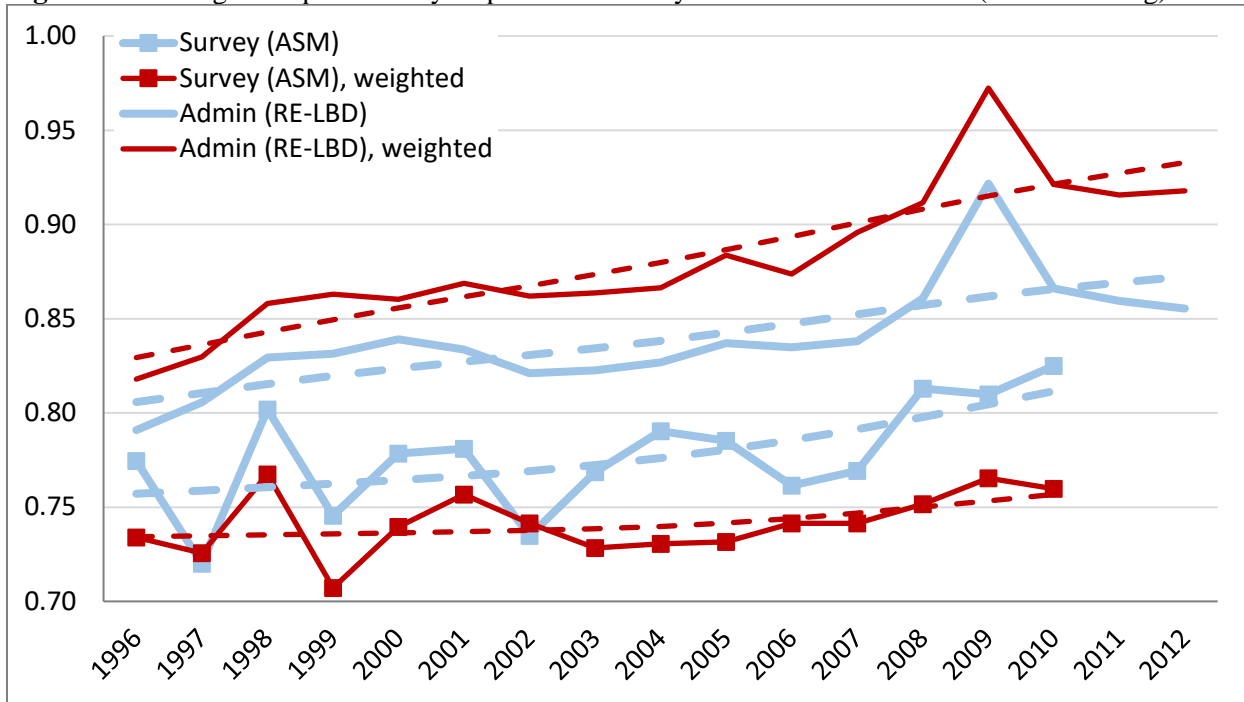
Note: Revenue labor productivity is measured as real revenue per employee. For the ASM, revenue and employment are from survey data. For the RE-LBD, revenue and employment are from administrative data. Dashed lines for each series depict HP-Filtered series. Two alternative types of weighting are used. For the employment-weighted series, the within industry dispersion at the 6-digit NAICS level are computed for each industry and year cell using the propensity score weights as described in the text. Then the employment shares for each industry are computed from the RE-LBD and the reported dispersion at the total manufacturing level uses such industry-level employment weights. Both the ASM and RE-LBD series use the RE-LBD industry-level employment weights. Results are similar using ASM-based employment weights. For the non-employment-weighted measures, the reported measure is based on first sweeping out 6-digit NAICS-by-year effects from the micro data and then computing the pooled dispersion measure using the propensity score weights. This is equivalent to a 2-step method weighting industries by the number of producers in the data. One source of discrepancy between the ASM and RE-LBD is the former is at the establishment level while the latter is at the firm level.

Appendix B.   Illustrative Model of Adjustment Costs

A. Model environment

Consider the following model of firm-level adjustment costs.[53] A firm maximizes the present discounted value of profits. The firm's value function and its components are specified as follows:

$$V(E_{et-1}; A_{et}) = A_{et}E_{et}^{\varphi} - w_t E_{et} - C(H_{et}) + \beta V(E_{et+1}; A_{et+1})$$

with:

$$C(H_{et}) = \begin{cases} \dfrac{\gamma}{2}\left(\dfrac{H_{et}}{E_{et}}\right)^2 + F_+\max(H_{et},0) + F_-\max(-H_{et},0) & if \ H_{et} \neq 0 \\ 0, & otherwise \end{cases}$$

$$a_{et} = \rho a_{et} + \eta_{et}$$
$$E_{et} = E_{et-1} + H_{et}$$

where $\varphi \leq 1$ due to product differentiation so that $A_{et}E_{et}^{\varphi}$ is the revenue function, E_{et} is employment for time t, H_{et} is net hires made at the beginning of time t, or $H_{et} = E_{et} - E_{e,t-1}$ (this can be positive or negative), w_t is the wage, and $a_{et} = \log(A_{et})$ is a revenue shock potentially reflecting TFPQ and demand shocks (for expositional convenience we focus on TFPQ). We interpret the revenue function curvature as reflecting product differentiation rather than decreasing returns to help draw out relations between revenue productivity and technical efficiency. That is, let firm-level prices be given by $P_{et} = Q_{et}^{\varphi-1}$ where $Q_{et} = \tilde{A}_{et}E_{et}$ is firm-level output subject to a CRTS technology. This implies that $A_{et} = \tilde{A}_{et}^{\varphi}$. In terms of the terminology of the literature and the main text, $\tilde{A}_{et}$ is TFPQ; and since labor is the only factor of production, both TFPR and revenue labor productivity (RLP) are given by $P_{et}\tilde{A}_{et}$. Note that this specification nests the price-taking version of the model with $\varphi = 1$. In that case, TFPQ, TFPR and RLP are equivalent. We focus on the $\varphi < 1$ case in our calibration but discuss some aspects of the price-taking case below.

---

[53] We use the term "firm" for expositional purposes; for modeling purposes we do not distinguish between firms and establishments. Our main empirical results focus on establishments.

This simple adjustment cost model is similar to Cooper, Haltiwanger, and Willis (2007), Elsby and Michaels (2013), and Bloom et al. (2016) and, in principle, accommodates both convex and non-convex costs. Under non-convex costs the solution has the following form:

$$V = \max(V^I, V^H)$$

where

$$V^I = A_{et}E_{et-1}^{\varphi} - w_t E_{et-1} + \beta V(E_{et}; A_{et+1}) \; if \; H_{et} = 0$$

$$V^H = A_{et}E_{et}^{\varphi} - w_t E_{et} - C(H_{et}) + \beta V(E_{et}; A_{et+1}) \; if \; H_{et} \neq 0$$

with the notation indicating that V^I is the value of inaction (i.e., zero net hiring), and V^H is the value of nonzero net hiring (in either positive or negative amounts).

We view this model as primarily designed to yield the reduced form hypotheses that we confront with the data, but we seek to use a reasonable baseline calibration that matches many of the features of the data and the parameters of the existing literature. Appropriate caution is needed since we do not model entry or exit, and we do not have any lifecycle learning dynamics or frictions that make young firms different from more mature firms.

Our main calibration exercise implements "general equilibrium" in the sense that the wage adjusts to clear the labor market; however, we fix the labor supply, so this calibration may be thought of as one relevant bound. In unreported exercises, we also consider the opposite bound in which labor supply is perfectly elastic and the wage is fixed (i.e., partial equilibrium). A limitation of the partial equilibrium exercise is that when the wage is fixed, adjustment frictions can have large effects on average firm size and therefore aggregate productivity via channels that are unrelated to reallocation. However, our key results on how adjustment costs affect reallocation rates, firm-level productivity responsiveness, and the OP covariance do not substantively depend on general vs. partial equilibrium. We report the more realistic general equilibrium (inelastic labor supply) results here.

B. Calibration

We set $\beta = 0.96$, consistent with annual data. We specify that $\varphi = 0.8$, consistent with a markup of 25 percent. For the shock process, we specify $\sigma_a = 0.35$, which is roughly the standard deviation of TFPR or RPR in U.S. manufacturing during the 1980s; and we set $\rho = 0.65$, broadly consistent with the AR(1) coefficient on TFPR and RPR that we find among

manufacturing establishments in the 1980s (see Figure A4 and Appendix C). These values of σ_a and ρ imply that innovations to TFP have a standard deviation of $\sigma_\eta = 0.26$. Strictly speaking, if plant-level prices are endogenous (which this model permits) the appropriate empirical moments are those from RPR since it reflects only the exogenous fundamentals. However, as discussed below, the benchmark calibration yields a very high correlation between RPR and TFPR. This in turn implies, as shown below, the critical inferences we use from the theory to take to the theory are robust to using either RPR or TFPR.

We calibrate the adjustment cost parameter(s) to target a job reallocation rate of 25 percent, roughly the rate for the U.S. manufacturing sector in the 1980s. Focusing only on kinked (non-convex) adjustment costs (i.e., setting $\gamma = 0$), we find that the target reallocation rate implies $F_+ = 0.85$ when $F_- = 0$. For quadratic adjustment costs (i.e., $F_+ = F_- = 0$) matching the job reallocation rate of 0.25 requires $\gamma = 1.3$. We make no attempt to jointly calibrate convex and non-convex adjustment costs for our illustrative purposes. For the key moments we study empirically, the model produces broadly similar predictions regardless of cost type. The literature suggests that non-convex costs are important for certain properties of microdata, so we focus on this cost type here and leave convex cost exercises unreported. However, the qualitative patterns we present are robust to consideration of variation of alternative adjustment cost parameters.

The model produces a (non-targeted) correlation between TFPQ and TFPR (which is the same as RLP in this one-factor setting) of 0.90, qualitatively similar to the 0.75 found by Foster, Haltiwanger, and Syverson. This strong correlation implies that the responsiveness of growth to realizations of productivity is essentially the same whether we use TFPQ or TFPR/RLP as the measure of productivity.

We consider two types of experiments in the simulation. The first is an increase in adjustment frictions starting from $F_+ = 0.85, F_- = 0$, our baseline non-convex cost calibration described above.[54] We increase F_- from zero to study rising adjustment costs from a starting point that matches the patterns of TFP and reallocation in the 1980s. In the second experiment, we fix adjustment costs at the baseline ($F_+ = 0.85, F_- = 0$) but vary TFP dispersion around the

---

[54] We also find similar patterns if we consider a case with F->0 and F+=0 and then begin increasing F+ Cooper, Haltiwanger, and Willis (2007) found that including a kinked adjustment cost to be important to match the shape of the growth rate distribution of employment but found similar results using either F- or F+ to generate the kink.

baseline to study the relation between TFP dispersion and responsiveness, RLP dispersion, and the OP covariance.

C. Changing adjustment costs

Here we describe results from the general equilibrium (i.e., flexible wage with inelastic labor supply) model. We consider several moments calculated on model-simulated data;[55] these moments are the job reallocation rate; the standard deviation of labor productivity (or, equivalently in this model, TFPR); the TFP coefficient from a regression of firm-level employment growth from time t to $t + 1$ on TFP in time t and (log) employment in time t (that is, we run a regression analogous to equation (4) in the paper; hereafter we call this moment the "productivity coefficient"); the Olley-Pakes (OP) covariance for TFP (hereafter "OP covariance"); and the OP covariance for labor productivity. As we discuss further below and in the main text, we do not use the OP covariance in our empirical analysis. However, we find it instructive to calculate the OP covariance in our various experiments in the theoretical model. This also is helpful to highlight the limitations of the OP covariance relative to our diff-in-diff counterfactual developed below.

Figures 4b of the main text and B1 of this appendix show the effect of increasing adjustment frictions (i.e., raising the downsizing cost F_- while holding fixed $F_+ = 0.85$). An increase in adjustment frictions yields: (i) a decline in the job reallocation rate; (ii) a decline in the productivity coefficient; (iii) an increase in the dispersion of RLP; and (iv) a decline in the OP covariance for both TFP and RLP (where employment serves as weights). Each of these relationships is monotonic in adjustment costs except the Olley-Pakes covariance for RLP (which we discuss below).

There are many possible moments relating growth to TFP that are similarly sensitive to adjustment costs. For example, in unreported regression results (in which we always control for period-t log employment), increasing adjustment costs yield (i) a decline in the estimated coefficient of a regression of firm-level growth between t and $t + 1$ on TFP in period $t + 1$ (rather than t); (ii) a decline in the estimated coefficient of a regression of firm-level growth between t and $t + 1$ on the change in TFP from period t to $t + 1$; and (iii) a decline in the estimated coefficient of a regression of firm-level growth between t and $t + 1$ on the innovation

---

[55] We simulate 2000 firms for 1000 periods then discard the first 100 periods.

of TFP (η_{et}). In principle, we could use any of these moments to detect a change in adjustment frictions. We focus on the specification in B1 to be consistent with our main econometric specification in equation (4), which is our preferred specification for a number of econometric and measurement reasons discussed briefly in the main text. For example, a limitation of using innovations of TFP with the ASM-LBD integrated data is that the panel rotation issues of the ASM imply that first panel years should be excluded in such specifications (see Foster et al. (2016) for details). Nevertheless, in unreported results with the ASM-LBD integrated data we have found that the main patterns of changing responsiveness hold when estimating equation (4) with innovations of TFP. We also note that the exact timing in the model vs. the data are different, so it is reassuring that the predictions on responsiveness hold equally well qualitatively in the numerical analysis using current or lagged productivity.

The model-based predictions in Figure 4b are the primary moments that we explore empirically in the main text. In the empirical analysis we also consider the estimated coefficient of firm-level growth between t and $t + 1$ on revenue labor productivity in t (with log period-t employment as a control as usual). Given the simple revenue functions in the model, this estimated coefficient is identical to the coefficient on TFP shown on Figure 4b (this is because the only production factor is labor, and initial employment is explicitly included in all regressions). This precise equivalence is model dependent, but the general inference is not. That is, in response to an increase in adjustment frictions, there should be a decline in the covariance between firm-level growth and realizations of labor productivity as firms find it more costly to equalize their marginal products.

D. Changing TFP dispersion

Figure 4a in the main text shows how key moments vary with TFP dispersion. Increased TFP dispersion yields: (i) an increase in the job reallocation rate; (ii) an increase in the productivity coefficient; and (iii) an increase in the dispersion of labor productivity. As before, the finding for the productivity coefficient also holds using real labor productivity as the regressor.

While the results in Figure 4a are generally intuitive, one finding merits further discussion—specifically, the finding that responsiveness increases in TFP dispersion. In the presence of non-convex adjustment costs, the net effect of TFP dispersion on responsiveness reflects two competing mechanisms (as discussed in the main text). The first is the "real

options" effect: Non-convex costs create "inaction bands" or regions of the TFP innovation range in which firms prefer inaction (i.e., zero hiring) to action. Inaction bands widen as shock dispersion or volatility rises (consistent with an "uncertainty" interpretation), which, *ceteris paribus*, reduces responsiveness. We observe this effect in our simulated data when we examine only the extensive margin: a given absolute change in TFP is more likely to induce action when TFP dispersion is smaller (holding initial employment constant). However, in the model this effect is dominated by the "volatility effect" in which adjustments—when they actually do occur—are larger when TFP is more widely dispersed. The dominance of the "volatility effect" is consistent with Barro (1972); Bloom et al. (2016) find, in a business cycle-focused model that is otherwise similar to ours, that the real options effect dominates for a few quarters at high frequency, but the volatility effect dominates after a year or more. Given our long-run focus, the steady-state result that the "volatility effect" dominates is the relevant intuition.

E. The frictionless case

The patterns in Figures 4a, 4b and B1 are, for the most part, robust to changes in the curvature of the revenue function, the shock space and the adjustment cost parameters. One exception highlights the importance of using multiple moments in our empirical exercises. Specifically, in a frictionless benchmark with zero adjustment costs (in contrast to our baseline above, in which hiring costs are set to $F_+ = 0.85$), there is zero labor productivity dispersion and, therefore, zero OP covariance for labor productivity (though still positive OP covariance for TFP). At first glance, this implies that the OP covariance for labor productivity may not be an informative moment, but we show here that the frictionless benchmark yields patterns that are very far from empirical plausibility.

Figure B2 shows the effects of increasing adjustment costs starting from zero, the frictionless case. As adjustment frictions rise above zero, labor productivity dispersion rises (Figure B2) and, consistent with the discussion above, reallocation and the productivity coefficient decline. But Figure B4 shows that the OP covariance for labor productivity initially rises as labor productivity begins to be dispersed, continuing to rise over the range of adjustment frictions that produce reallocation rates above 30 percent (compare Figure B4 to Figure B2). But since productivity responsiveness declines monotonically as adjustment costs rise, the OP covariance eventually declines as labor is increasingly "trapped" in unproductive firms while productive firms are starved of resources (i.e., employment weight). Thus, the OP covariance for

labor productivity is decreasing in adjustment costs (and increasing in misallocation) across the plausible range of costs. This pattern is related to that found in Bartelsman, Haltiwanger, and Scarpetta (2013), in whose model distortions reduce the OP covariance for labor productivity as long as the benchmark is characterized by sufficient frictions; we explore the OP covariance in more detail below.

As can be seen on Figure B2, the frictionless benchmark produces implausibly large rates of reallocation, above 100 percent, while empirical reallocation in the manufacturing sector was around 25 percent in the 1980s and has since declined. Here we demonstrate basic analytical intuition underlying the empirical implausibility of the frictionless case. In the frictionless case ($\gamma = F_+ = F_- = 0$), the first-order condition for labor is given by:

$$E_{et} = \left(\frac{\varphi A_{et}}{w_t}\right)^{\frac{1}{1-\varphi}}$$

Taking logs (indicated by lower case) and differences (indicated by Δ) and sweeping out year and industry effects yields:

$$\Delta e_{et} = \frac{1}{1-\varphi}\Delta a_{et}$$

which implies

$$std(\Delta e_{et}) = \frac{1}{1-\varphi}std(\Delta a_{et})$$

where $std()$ indicates standard deviation. That is, in the frictionless model, the dispersion of employment growth rates (i.e., log differences) is proportional to the dispersion of TFP with the factor of proportionality greater than one. The $std(\Delta a_{et})$ among continuing manufacturing plants is about 0.33 (this is from RPR—similar statistics emerge from TFPR). For $\varphi = 0.8$, corresponding to a markup of 25 percent, we should expect $std(\Delta e_{et}) = 1.65$; for a 33 percent markup we should expect $std(\Delta e_{et}) = 1.32$. Yet in U.S. manufacturing data $std(\Delta e_{et}) = 0.35$. This relatively low dispersion of plant-level employment growth rates compared to the dispersion in shocks illustrates that the frictionless model yields implausible empirical patterns.

Given the empirical implausibility of the frictionless model, then, we are comfortable drawing inference from key moments—including the OP covariance for labor productivity— along the plausible range of adjustment costs. We elaborate further on strengths and limitations of the OP decomposition below.

F. Markups

Even though the frictionless model exhibits implausible dispersion of employment growth rates, this analysis is useful for demonstrating that an increase in markups reduces reallocation and responsiveness but does not inherently affect dispersion in revenue labor productivity. The markup is given by $1/\varphi$. From the above, reallocation (dispersion in employment growth rates) and the responsiveness of employment growth are both decreasing in the markup, but the dispersion of revenue labor productivity is zero in the frictionless benchmark regardless of the markup.

This independence of labor productivity dispersion holds under adjustment costs as well. Figure B3 shows the effect of varying the markup around our benchmark calibration of a 25 percent markup and adjustment costs targeting the 1980s reallocation rate. As the markup increases, reallocation and responsiveness sharply decline. The standard deviation of labor productivity dispersion is flat, to numerical precision, over this wide range of markups.

G. The Olley-Pakes decomposition and covariance

While we show above that the OP covariance is declining in adjustment costs over plausible parameterizations, additional doubts about its usefulness may arise from limitations of the accounting-based decomposition on which it is based (Levinsohn and Petrin (2003); Petrin, White, and Reiter (2011); Hsieh and Klenow (2017)). We do not formally use the OP covariance in any of our empirical analysis but instead use diff-in-diff counterfactuals. However, it is instructive to compare and contrast our diff-in-diff counterfactuals with the OP covariance. The OP productivity decomposition is given by:[56]

$$A_t^{OP} = \sum_e \theta_{et} A_{et} = \bar{A}_t + cov(\theta_{et}, A_{et})$$

where A_t^{OP} is the Olley-Pakes concept of industry aggregate productivity, θ_{et} is the employment share of firm e, A_{et} is productivity of firm e, $\bar{A}_t$ is (unweighted) average productivity for the industry, and $cov(\theta_{et}, A_{et})$ is the OP covariance (which is proportional to true covariance).

---

[56] In empirical work, it is common to use a weighted average of log firm-level productivity to avoid index number problems. We use the weighted average of firm-level productivity here to highlight that under CRTS and price-taking behavior that this weighted average is equal to aggregate productivity. Empirical implementation still requires addressing the index number problems.

Define aggregate productivity as $A_t = Q_t/E_t$, that is, aggregate output divided by aggregate employment. A critical question is how well A_t^{OP} matches A_t; a related question is how to interpret the quantitative variation in the OP covariance term. Continuing with our model, characterized by constant returns to scale but potentially imperfect competition,[57] note that aggregate output is given by:

$$Q_t = \left(\sum_e Q_{et}^{\varphi} \right)^{\frac{1}{\varphi}}$$

where $\varphi \leq 1$. Aggregate productivity is as follows:

$$A_t = \frac{Q_t}{E_t} = \frac{\left(\sum_e A_{et} E_{et}^{\varphi} \right)^{\frac{1}{\varphi}}}{E_t}$$

It is straightforward to see that if $\varphi = 1$, that is, perfect competition, aggregate output of the final good is simply the sum of firm-level output, and the OP decomposition measures aggregate productivity exactly:

$$A_t = \frac{Q_t}{E_t} = \frac{\sum_e A_{et} E_{et}}{E_t} = \sum_e A_{et}\theta_{et} = A_t^{OP}$$

In this case the size distribution can only be determined in the presence of frictions; but the OP covariance term's effect on aggregate productivity is intuitive since the firm with the highest average productivity also has the highest marginal productivity. Moving resources to the most productive firm will always increase aggregate productivity (this inference is robust to a model with multiple inputs). The key is that with CRTS and perfect competition, the marginal revenue product of a firm does not change with scale but only varies with TFPQ.

For $\varphi < 1$, however, aggregate productivity is no longer equal to the OP-based weighted measure of firm productivity. Moreover, it is no longer the case that continually moving resources to the most productive firm (in terms of average productivity) will increase aggregate productivity. This is because with $\varphi < 1$, as resources are moved to the most productive firms marginal revenue products rise for the least productive firms and fall for the most productive firms. This implies that with $\varphi < 1$ the OP weighted average of firm productivity declines more

---

[57] A broader interpretation of our simple model is that the revenue function curvature reflects either (or both) imperfect competition or decreasing returns to scale. The effects of either on the accuracy of the OP decomposition are similar, so we proceed in this subsection with a focus on the imperfect competition.

rapidly with an increase in adjustment costs than does aggregate productivity. Intuitively, revenue function curvature (i.e., low or declining φ) dampens the marginal responsiveness of firms to their TFP realizations relative to linear revenue, or $\varphi = 1$ (which, in unreported results, we verify in our model); the OP covariance declines more quickly than industry productivity as adjustment costs rise because it does not capture the impact of curvature on marginal revenue products and optimal responsiveness.

To overcome this limitation of the OP covariance measure, in our empirical analysis we rely instead on a diff-in-diff counterfactual (equation (5) in the text) given by:

$$\Delta_t^{t+1} = \sum_f (\theta_{e,t+1}^T - \theta_{e,t+1}^{NT}) a_{et}$$

This diff-in-diff counterfactual isolates the impact of changing responsiveness on the weighted mean holding everything else constant. In particular, this diff-in-diff object implicitly relies on estimated policy functions that reflect optimal growth responses of firms to their productivity draws and initial size, taking the degree of revenue curvature as given and therefore implicitly acknowledging that equalization of marginal products is the driving force behind labor adjustment. As such, the diff-in-diff object more accurately tracks changes in aggregate productivity than does the raw OP covariance, which relies on linear revenue function assumptions and therefore sees all movement of labor toward high-productivity firms as productivity enhancing. Conveniently, this diff-in-diff object can be easily constructed empirically, with the main source of error being estimation error inherent in the process of using OLS to estimate nonlinear policy functions.

Interestingly, using the OP covariance for revenue labor productivity (RLP) captures the effect of curvature on industry-level productivity in a way that the OP covariance for TFPQ does not. RLP is given by $A_{et} E_{et}^{\varphi-1}$ so that the employment-weighted average of RLP is given by:

$$RLP_t^{OP} = \sum_e \theta_{et} A_{et} E_{et}^{\varphi-1} = \frac{\sum_e A_{et} E_{et}^{\varphi}}{E_t}$$

Comparing RLP_t^{OP} to A_t shows that both reflect the curvature of the revenue function in a manner that A_t^{OP} does not; this convenient feature of RLP arises because RLP (the average revenue product of labor) is proportional to the marginal product of labor.

Figure B4 illustrates how aggregate productivity, the OP covariance using TFPQ, the OP covariance using RLP, and the diff-in-diff counterfactuals using either TFPQ or RLP (equivalent

to TFPR in the one-factor model) evolve as adjustment costs rise. Consistent with the discussion above, as adjustment costs rise, true aggregate productivity declines less rapidly than the OP covariance using TFPQ. In contrast, the diff-in-diff counterfactuals track aggregate productivity very closely; moreover, the OP covariance using RLP also tracks aggregate productivity well over the reported range of adjustment costs of Figure B5. In unreported results, we have shown that the properties of Figure B5 are robust to alternative values of φ. In particular, the close quantitative correspondence between changes in aggregate productivity and the diff-in-diff counterfactuals is robust to alternative values of $\varphi < 1$.

The finding that the OP covariance for RLP tracks industry productivity well over the range of adjustment costs (Figure B6) is interesting since, as shown on Figure B4, the OP covariance for RLP exhibits a non-monotonic relationship with adjustment costs when starting from a frictionless benchmark. Figure B4 suggests caution in using the OP decomposition with RLP over a wide range of adjustment costs that includes zero costs. However, over the empirically relevant cost range the OP decomposition with RLP tracks industry productivity reasonably well. This finding is not directly related to the findings in this paper since we focus on the diff-in-diff counterfactuals that directly exploit estimated policy functions. However, it is relevant for the large literature that uses the OP decomposition or the dynamic OP decomposition with labor productivity (e.g., Decker et al. (2017) and Alon et. al. (2017)). The findings here suggest that using the OP approach with the RLP index yields a decomposition that quantitatively tracks industry-level productivity reasonably well under empirically relevant adjustment costs.

One might ask why we do not simply quantify changes in aggregate (industry) productivity in our empirical analysis. The reason is that simple comparisons cannot isolate the change in aggregate productivity that is due specifically to changing adjustment frictions and slowing reallocation. Our diff-in-diff counterfactual achieves this objective and, since it relies on actual (estimated) firm policy functions, is reasonably robust to the parameters of the economic environment.

Figure B1: Responses of OP covariances to changes in adjustment costs
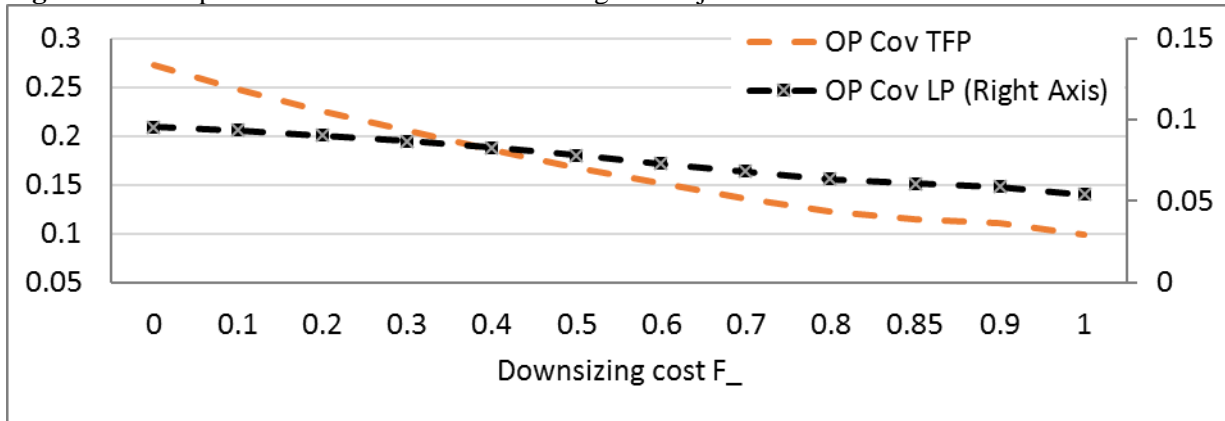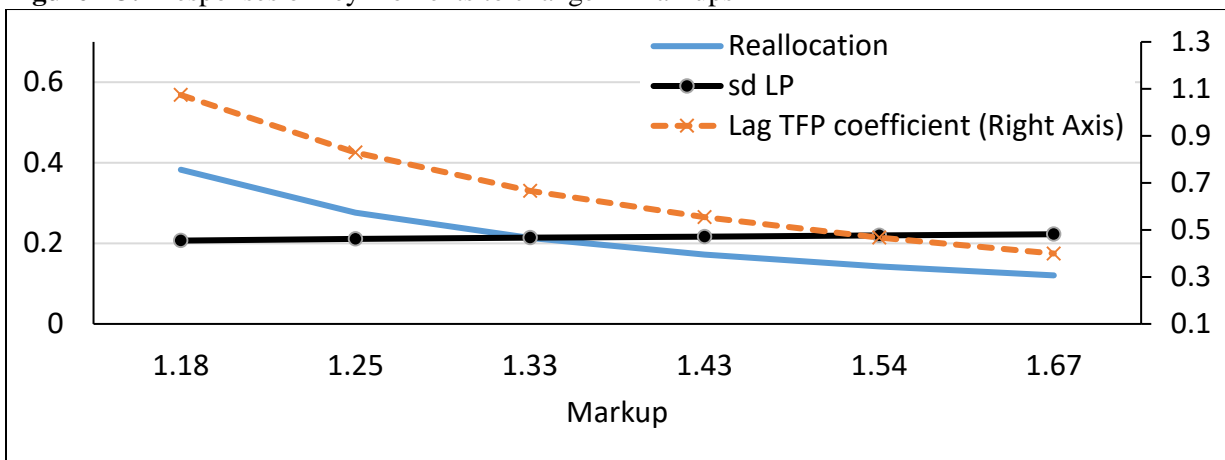

Note: The x axis reflects values of F_-, or the cost of reducing employment, holding the hiring cost F_+ fixed at $F_+ = 0.85$. General equilibrium model with flexible wage and inelastic labor supply.

Figure B2: Responses of key moments to changes in adjustment costs from frictionless benchmark
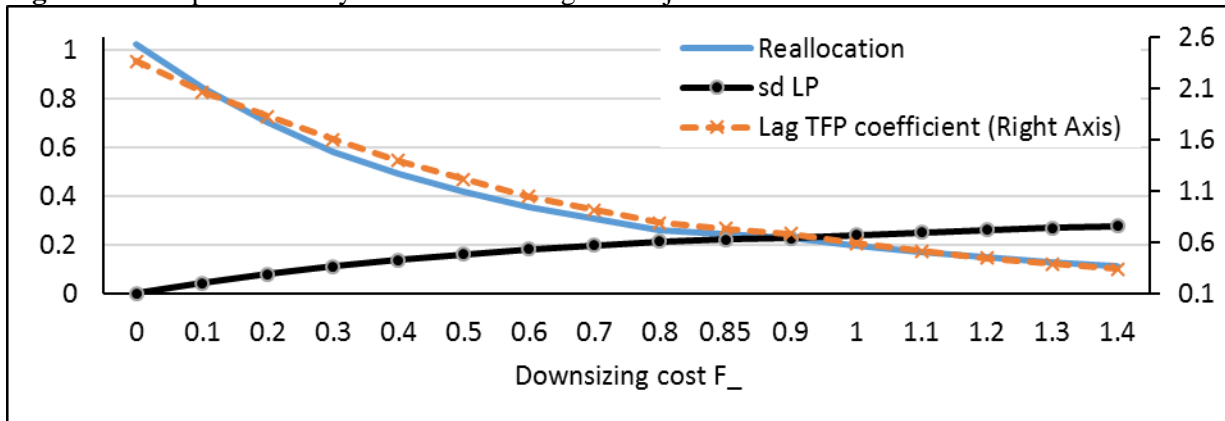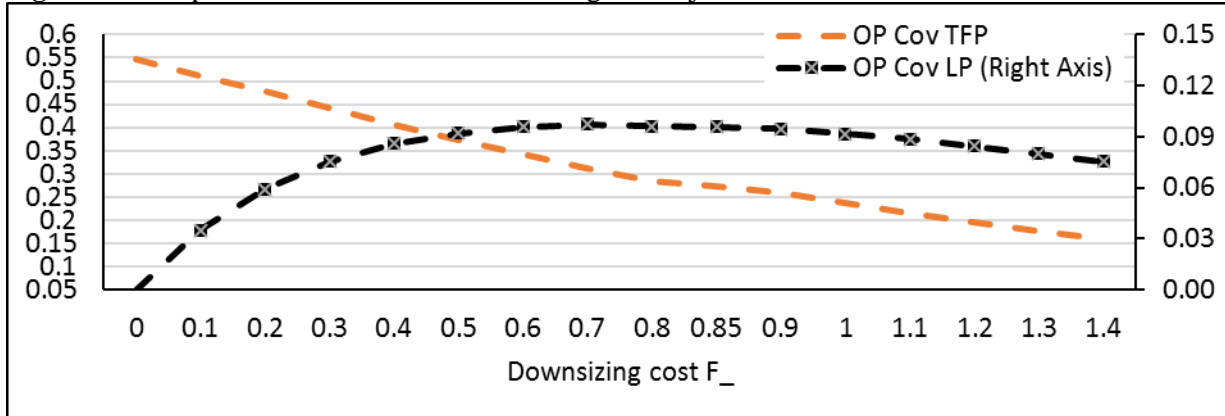

Note: Model with no upward or downward adjustment costs ($F_+ = 0$) with varying downward adjustment costs (F_-) indicated on the x axis. General equilibrium model with flexible wage and inelastic labor supply.

Figure B3: Responses of key moments to change in markups

Note: Model with upward adjustment costs ($F_+ = 0.85$) with varying markups (curvature in revenue function) indicated on the x axis. General equilibrium model with flexible wage and inelastic labor supply.

Figure B4: Responses of OP covariances to changes in adjustment costs from frictionless benchmark

Note: Model with no upward adjustment costs ($F_+ = 0$) with varying downward adjustment costs (F_-) indicated on the x axis. General equilibrium model with flexible wage and inelastic labor supply.

Figure B5: Response of Diff-in-Diff Counterfactuals to Increase in Adjustment Costs
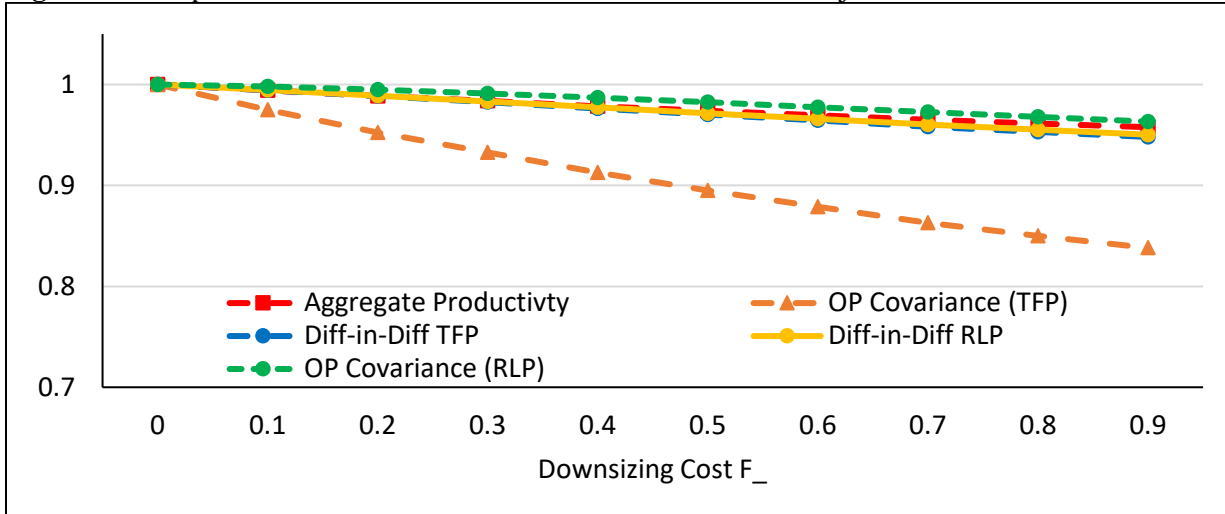
Note: The x axis reflects values of F_-, or the cost of reducing employment, holding the hiring cost F_+ fixed at $F_+ = 0.85$. General equilibrium model with flexible wage and inelastic labor supply.

Appendix C. Alternative TFP calculation

While our TFPR measure as a measure of TFP is common in related literature, as we discus in the main text we also consider an estimate of RPR using the proxy method of Wooldridge (2009). As we show in equation (3), RPR is only a function of exogenous TFPQ and demand shocks (even if plant-level prices are endogenous) because the elasticities recovered by revenue function estimation are revenue elasticities (not factor elasticities) capturing both production and demand parameters (Foster et al. (2017)). In this appendix, we discuss the estimation of RPR and the results using the RPR measure of TFP. Given the possible presence of demand shocks, RPR should be interpreted as reflecting both TFPQ and demand shocks.

Foster et al. (2017) find that the Woolridge residuals are sensitive to outliers; pooling across a large number of observations mitigates this sensitivity, so we estimate revenue elasticities that vary at the 3-digit NAICS level.[58] After estimating the elasticities, we compute the revenue productivity residuals and deviate the latter from 6-digit NAICS industry by year means. We find that RPR_{et} has a correlation of 0.76 with $TFPR_{et}$.

We replicate our main empirical exercises replacing TFPR with RPR. Figure C1 shows the evolution of within-industry dispersion in RPR_{et} for manufacturing plants. Consistent with Figure 5, we observe gradually rising RPR dispersion throughout the time period, with higher dispersion in high-tech than elsewhere. Figure C2 reports AR(1) coefficients for plant-level RPR (see note to Figure A4 in appendix A for a discussion of this measure and its limitations in our dataset). Again, RPR results are consistent with TFPR results, confirming that changes in the TFP distribution cannot explain aggregate job reallocation patterns.

We estimate equation (4) using RPR in place of TFP. Figure C3 reports growth differentials (between the plant with productivity one standard deviation above its industry mean and the mean) as discussed in the main text. The results are generally consistent with those reported for TFP on Figure 6, with young firm productivity responsiveness in high-tech that rises from the 1980s to the 1990s then falls in the 2000s. Among mature firms in high-tech,

---

[58] The cost share-based TFPR measure we use in the main text is constructed with cost shares that vary at the 4-digit SIC level prior to 1997 and the 6-digit NAICS level thereafter. The instability and outlier sensitivity of the RPR elasticity estimates precludes this level of detail. The cost share-based TFPR method therefore allows for more flexibility in elasticity values implying a better fit in detailed industries, while the RPR method avoids problems of price endogeneity and isolates exogenous TFPQ and demand shocks. This tradeoff is the reason we ensure robustness of our exercises to both productivity concepts.

responsiveness is somewhat flat from the 1980s to the 1990s before falling markedly in the 2000s.

Finally, Figure C4 reports the diff-in-diff counterfactual described by equation (5). Among high-tech plants, declining responsiveness produces a counterfactual that is broadly similar—both qualitatively and quantitatively—with the TFPR-based results from Figure 7, with a productivity "drag" that is only slightly smaller under RPR than under TFPR. Among non-tech plants, the counterfactual produces somewhat different results from those reported in Figure 7, with a gap opening up early in the sample then remaining stable (and negative) after the late 1990s. In general, the RPR results confirm the TFP-based findings suggesting a quantitatively significant change in the contribution of reallocation to aggregate productivity growth.

We conduct an additional robustness check addressing the concerns of Gandhi et al. (2016), who argue that if some factors are completely variable then the Wooldridge (2009) and related proxy methods may not be not identified. Ackerberg, Caves, and Frasier (2015) offer guidance about when their (and the equivalent Wooldridge (2009)) method is identified;[59] however, as a robustness check we consider a hybrid method that estimates revenue elasticities for materials and energy (the factors most likely to be completely variable) using non-parametric methods in a first step. In practice, this amounts to using the means of the revenue cost shares of materials and energy to estimate the revenue elasticities. Then the Wooldridge (2009) method is used to estimate the revenue elasticities for labor and capital in a second step. We do not report this robustness exercise for the sake of brevity, but the results are quite similar to those we report for RPR. This is not surprising since this hybrid measure yields an RPR measure with a correlation of 0.94 with the RPR measure estimated using the Wooldridge (2009) procedure that we describe above. The standard deviation of the two alternative RPR measures (sweeping out industry-by-year effects) are both about 0.41.

---

[59] If input prices of variable factors are serially correlated then this is one way that the ACF and Wooldridge estimators are still identified.

**Figure C1:** Within-industry dispersion in RPR (standard deviation), manufacturing
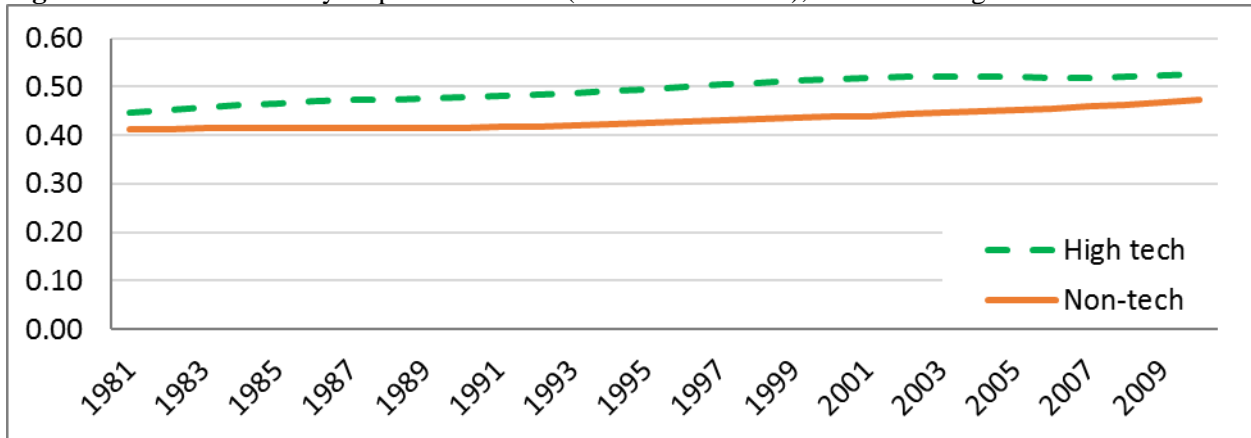
Note: The standard deviation is the based on within-detailed industry log revenue productivity residual. High-tech is defined as in Hecker (2005). Manufacturing is defined on a consistent NAICS basis. Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers. HP trends.

Figure C2: Persistence of plant-level RPR: High-tech vs. non-tech
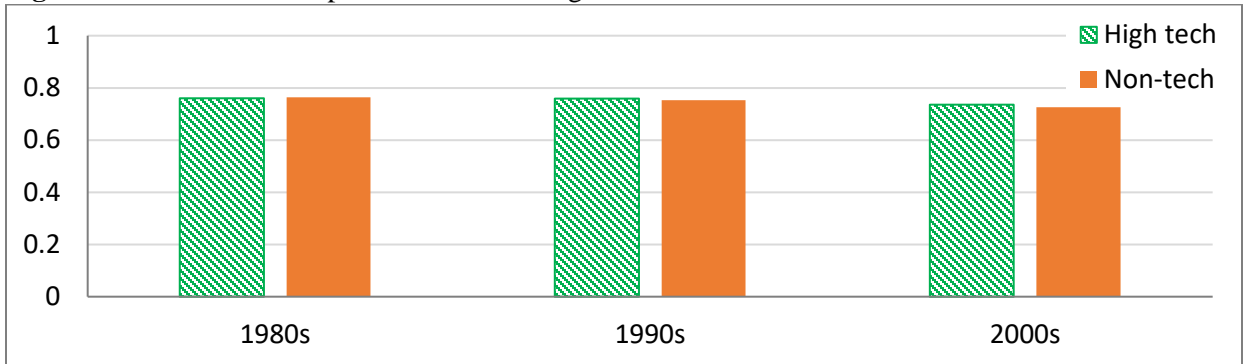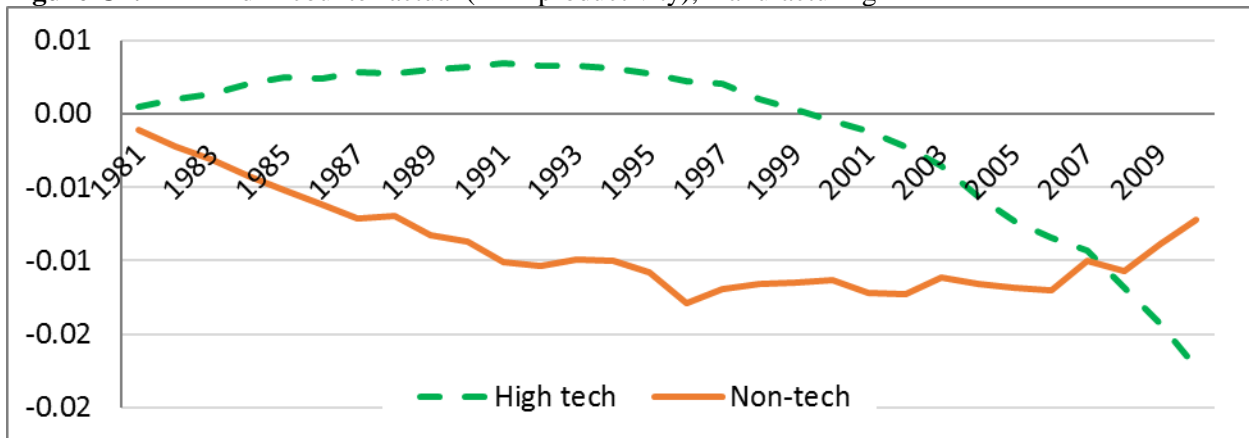

Note: High-tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure C3: Relative employment growth rates, high-productivity vs. average-productivity plant (RPR)
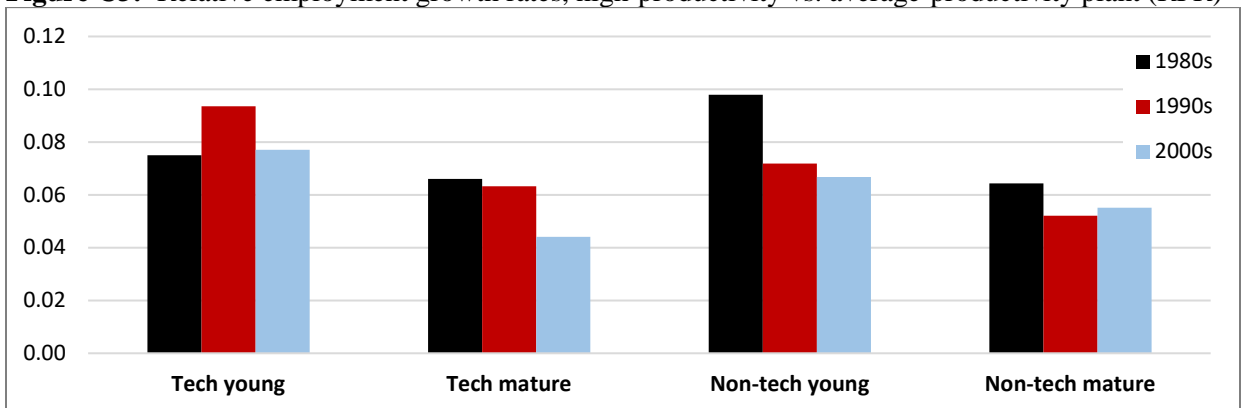
Note: Young firms have age less than 5. High-tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure C4: Diff-in-diff counterfactual (RPR productivity), manufacturing

Note: Figure depicts diff-in-diff counterfactual as described in the text. High-tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Appendix D.   Changing Business Models in Manufacturing

A. Investment

As noted in the text, Table 2 and Figure 6 report our standard responsiveness regression using establishment investment rates as the dependent variable in place of employment growth. We include the capital stock as an additional state variable in these regressions; we do not include the capital stock in our main employment growth regressions, but in unreported exercises we find our main results are robust to its inclusion. The timing is slightly different for equipment investment as opposed to the employment growth specifications: in the employment growth specifications, we measure growth from March of t to March of t+1 as a function of size in March t and productivity for year t. In the investment specification, we measure investment throughout year t as a function of size in March t, productivity for year t, and capital stock at the beginning of year t. Consistent with standard models, we have in mind a time-to-build assumption that investment during period t contributes to capital used during period t+1, which means the difference in timing from the employment growth regressions is not large. Moreover, our model exercises in Appendix B suggest that our theoretical framework is not heavily dependent on specific timing concerns.

B. Globalization

Globalization may be playing a role in declining responsiveness since increased exposure to foreign trade facilitates adjustment by scaling international operations. That is, it may be that rather than growing domestically, productive firms are more likely to expand and produce in other countries, a dynamic that could eliminate or even reverse the standard positive correlation between growth and productivity (since we do not observe employment outside the U.S.). There is substantial evidence already that the decline in US manufacturing employment is closely linked to rising import penetration of production activity from low wage countries (see, e.g., Bernard, Jensen, and Schott (2006), Schott (2008) and Pierce and Schott (2016)).

Bernard, Jensen and Schott (2006) and Schott (2008) develop measures of import penetration ratios from low wage countries. Their measures vary by 4-digit SIC industry from 1972-2005 and by 6-digit NAICS industry from 1989-2005; we extend the time series using the public domain information from Census on imports by country and industry.[60] We integrate

---

[60] To construct low-wage import penetration data by year and industry, Bernard, Jensen and Schott first construct domestic absorption for each industry. Next, they construct total imports of goods produced by each industry that are

these public domain data into our data infrastructure from 1981-2010. Our ability to integrate this is facilitated by our having 4-digit SIC codes in the micro level data from 1981-1996 and 6-digit NAICS codes from 1981-2010; hence, we need not rely on aggregate SIC/NAICS concordances.[61] Figure D1 shows aggregate import penetration ratios in and out of high-tech manufacturing.

Table D1 presents results of a modified version of our main regressions in equation (3). The additional regressors added are the 6-digit NAICS import penetration ratio for each year and the interaction of this ratio with lagged TFP. We permit the coefficients on this interaction effect to differ between plants belonging to young and mature firms. The main effect of the import penetration (not reported) is negative and significant: Consistent with Bernard, Jensen, and Schott (2006), plants in industries with especially large increases in import penetration have lower net employment growth.

The last two rows of Table D1 show that the interaction effect for young plants of lagged TFP and the import penetration ratio is estimated to be negative and significant. This implies that young-firm plants in industries with especially large increases in import penetration ratios have larger decreases in responsiveness. In Figure D2, we quantify the effect of changing import penetration ratios using the estimated effects from Table D1. The overall effects show, consistent with Table 1, that the marginal effect of productivity on employment growth among young high-tech firms increased from the 1980s to 1990s then declined in the post-2000 period. We compute the fraction of these patterns accounted for by the changing import penetration ratios by using the coefficients from Table D1 along with the aggregate pattern of import penetration ratios for high-tech manufacturing. The role of rising penetration is very modest in the 1980s to 1990s. However, the rapid rise in import penetration during the 2000s accounts for a substantial share (about 16 percent) of the overall decline in responsiveness over that period. We also perform these analyses using the Wooldridge (2009) RPR productivity measure (unreported), finding no significant role for import penetration, so we consider this evidence mixed. More research on globalization and dynamism is needed; promising avenues include

---

sourced in a low-wage country, which are defined as countries whose GDP per capita is less than 5 percent of the U.S. Import penetration is the ratio of low-wage imports to total domestic absorption, by industry and year. We thank Peter Schott for providing the import data and guidance necessary for extending the dataset.

[61] We integrate the SIC-based import penetration ratios from 1981-88 and the NAICS-based ratios from 1989-2010 into the micro data. We use the internally consistent NAICS codes in the micro data from 1981-2010 to conduct our analysis. (see Fort and Klimek (2016)).

specific policy variation, distinction between intermediate and final goods competition, and differences between TFP concepts.

C. Composition Effects

In the high-tech manufacturing sector, another possible cause of declining productivity responsiveness after 2000 is the transition from "general-purpose" to "special-purpose" equipment manufacturing in the U.S documented by Byrne (2015).[62] Perhaps manufacturers of special-purpose products are less responsive to productivity shocks due to demand constraints or uncompetitive environments that reduce adjustment imperatives. Figure D3 shows that during the 1990s the share of employment in among general purpose technology producers grew rapidly but, consistent with Byrne (2015) (who examined revenue shares), the general purpose share has fallen substantially since the late 1990s. Given these compositional changes, it is possible that changing average responsiveness reflects differential responsiveness across industries.

We estimate equation (4) separately for each 6-digit industry in high-tech manufacturing but, importantly, we omit time trend interactions from the specification. With the estimated responsiveness coefficients for each 6-digit industry, we compute the employment-weighted aggregate responsiveness in each year using the actual annual 6-digit employment weights.[63] If any industry composition effect—including the shift between general and special purpose electronics—is driving our results, we should see these aggregated responsiveness patterns mimicking the result from Figure 6.

Figure D4 shows the implied changing responsiveness over time due to composition effects within high-tech manufacturing. There is no implied increase in responsiveness due to composition effects from the 1980s to the 1990s (which would have been expected if general-purpose producers were more responsive on average), and there is actually a modest increase in responsiveness from the 1990s to the 2000s rather than a decline. Declining responsiveness must therefore be a within-category phenomenon with respect to the general-purpose/special-purpose taxonomy and other industry characteristics.

[62] We thank Christopher Foote for this insight.
[63] We use employment weights given our interest in the implications of changing responsiveness for job reallocation.

**Figure D1:** Import penetration ratios from low-wage countries
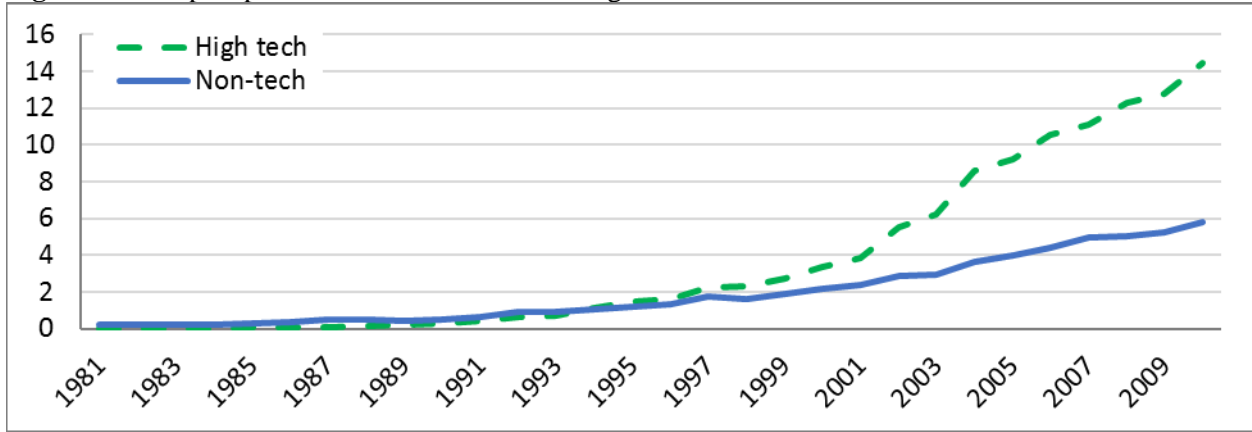
Source: Extended versions of Import Penetration Ratios from Bernard, Jensen and Schott (2006) and Schott (2008). Reported statistics are averages across 6-digit NAICS industries for high-tech and non-tech industries.

Figure D2: The role of globalization in changing responsiveness (high-tech Manufacturing)
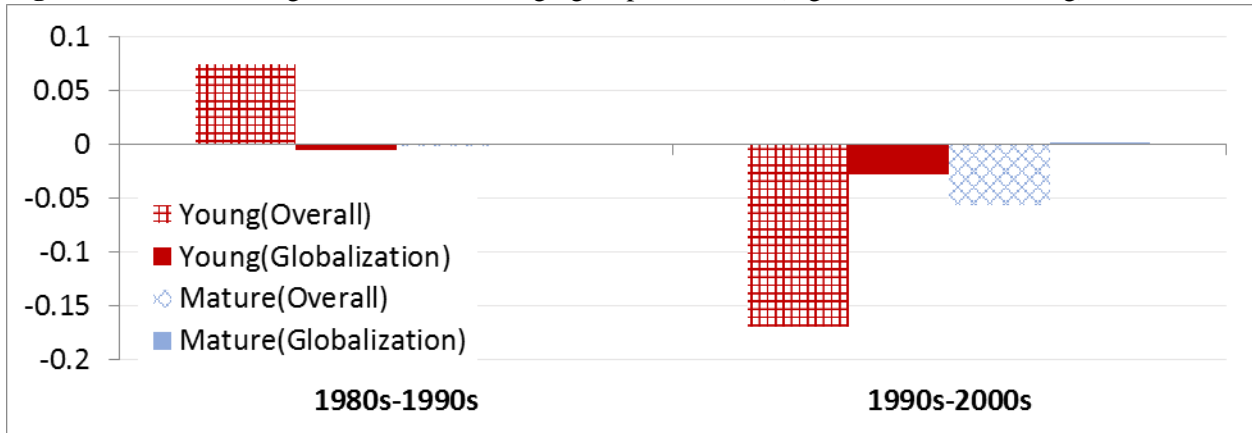

Note: "Overall" bars for young and mature are the change in marginal responsiveness of employment growth to productivity across decades. Globalization reflects implied change in marginal responsiveness accounted for by changes in import penetration ratios from low wage countries.

Figure D3: General purpose technology share of high-tech Manufacturing
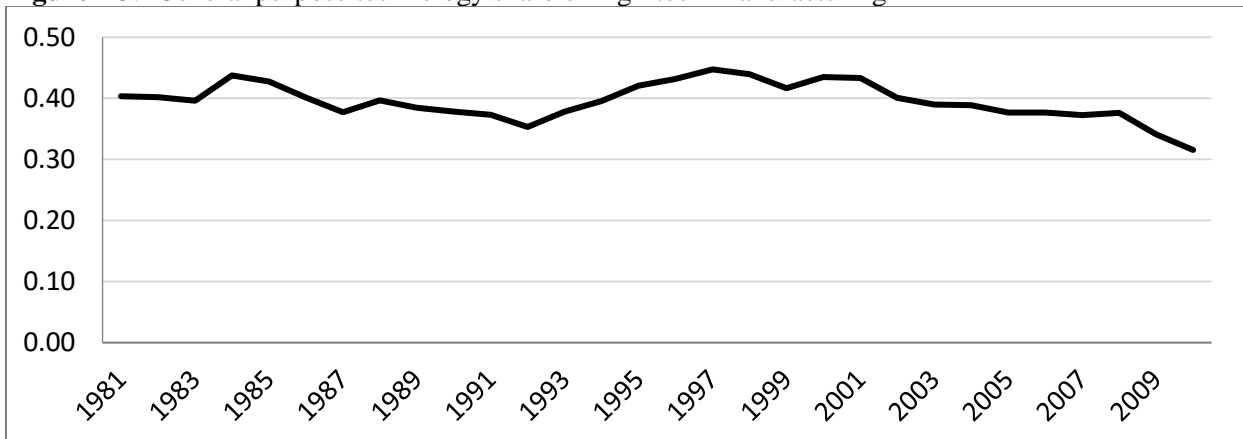
Note: Tabulations from the LBD by authors. General purpose high-tech 4-digit industries are NAICS 3341 (Computers), NAICS 3342 (Communication Equipment) and NAICS 3344 (Semiconductors).

Figure D4: Change in responsiveness due to industry composition changes (high-tech)
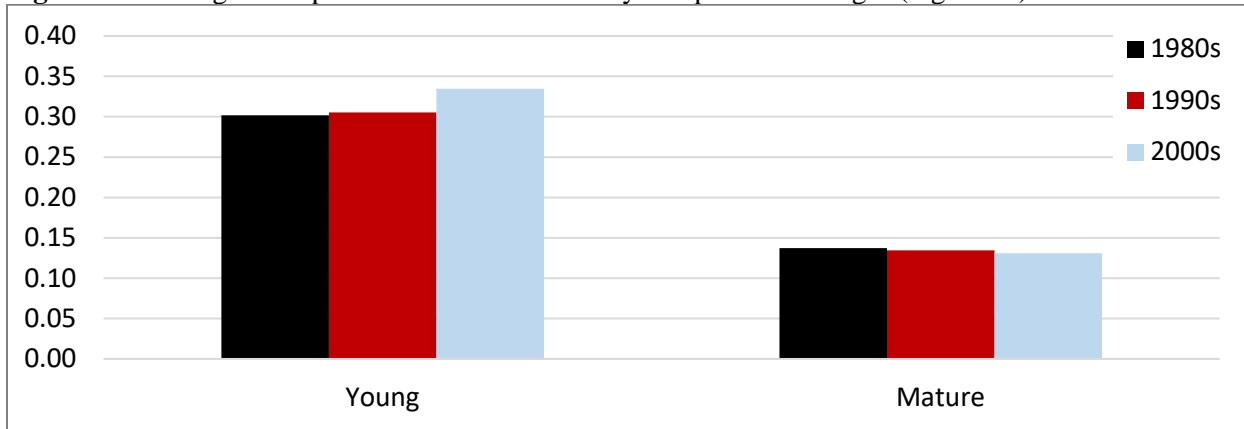
Note: Specification (2) as in Table 1 estimated for every 6-digit NAICS industry but without any trend effects. Reported coefficients are employment-weighted averages of the 6-digit NAICS industry estimated coefficients. Employment-weights vary by year.

Table D1: Employment growth, lagged productivity, and import penetration

TFP*Young	0.2085***
	(0.0390)
TFP*Young*Trend	0.0298***
	(0.0061)
TFP*Young*Trend^2	-0.0011***
	(0.0002)
TFP*Mature	0.1246***
	(0.0174)
TFP*Mature*Trend	0.0052**
	(0.0026)
TFP*Mature*Trend^2	-0.0003***
	(0.0001)
TFP*Young*Import Penetration	-0.0037***
	(0.0011)
TFP*Mature*Import Penetration	0.0002
	(0.0004)

Notes: Standard Errors in Parentheses. High-tech sample used. See notes to Table 1. Young firms have age less than 5. Unreported are estimates of controls including year effects, state effects, firm age dummies, firm size dummies, log plant level employment in period t, state cyclical indicators (change in state level unemployment rate), state cyclical indicators interacted with TFP, and a main effect for the 6-digit import penetration ratio. All variables that use TFP including all interactions are fully interacted with firm age dummies. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.