Changing Business Dynamism and Productivity: Shocks vs. Responsiveness

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

June 2017

Abstract

The pace of business dynamism as measured by indicators such as job reallocation has declined in recent decades in the U.S., and theory suggests that this may have implications for productivity. At first glance the timing of the patterns of changes in the pace of reallocation do not appear to match the changes in aggregate productivity. Productivity surged in the 1990s, led by the ICT sector, and has declined in the post-2000 period in ICT and more broadly, while overall reallocation and the entry rate of new firms declined throughout the 1980s, 1990s, and 2000s. However, the High Tech sectors of the economy have a hump-shaped pattern of job reallocation and entry that broadly mimics the patterns of productivity. Moreover, the economywide startup and reallocation trends of the 1980s and 1990s largely reflected the productivity-enhancing consolidation of the Retail Trade sector. Even taking into account different sectoral patterns, however, the changing patterns of reallocation pose difficult identification challenges. Much of the literature on declining dynamism broadly has proposed hypotheses focused primarily on accounting for variation in the startup rate and the age composition of firms. While the startup decline is undoubtedly important, we show that most of the variation in the patterns of reallocation both economywide and within sector is *not* accounted for by the changing age structure of firms. This prompts us to focus on changes in dynamism within industries and firm age classes, and we focus especially on the High Tech sectors of the economy given the outsized role these sectors have played in aggregate productivity dynamics. We find changing patterns of establishment-level and firm-level responsiveness to productivity realizations within narrow industries and firm age classes. These changes imply a drag on productivity from a reduced pace of reallocation. During the productivity slowdown of the post-2000 period, we find not only declining responsiveness but rising within-industry productivity dispersion in the High Tech sectors. Taken together these findings are consistent with an increase in frictions or distortions in the U.S. economy.

*Board of Governors of the Federal Reserve System, University of Maryland and NBER, U.S. Census Bureau, and U.S. Census Bureau, respectively. We gratefully acknowledge financial support from the Kauffman Foundation. Cody Tuttle provided excellent research assistance. We thank John Abowd, Rudi Bachmann, Martin Baily, Dave Byrne, Chris Foote, Lucia Foster, Clément Gorin, Bronwyn Hall, Kristin McCue, and conference or seminar participants at the 2015 Atlanta Fed Conference on Secular Changes in Labor Markets, the ASSA 2016 meetings, the 2016 Brookings "productivity puzzle" conference, the 3rd International ZEW conference, the 2016 ICC conference, Drexel University, the Federal Reserve Board, George Mason University, Georgetown University, the New York Fed, Princeton University, the University of Chicago, the spring 2017 Midwest Macro meetings, and the 2017 UNC IDEA conference for helpful comments. We are grateful for the use of the manufacturing productivity database developed in Foster, Grim and Haltiwanger (2016) as well as the revenue productivity database developed in Haltiwanger et al. (2016), and we thank Peter Schott for providing guidance and data for constructing import penetration data. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau or of the Board of Governors or its staff. All results have been reviewed to ensure that no confidential information is disclosed.

I. Introduction and motivation

A hallmark of market economies is the continual reallocation of resources from less-valued or less-productive activities to more-valued or more-productive ones. Business dynamics—the process of business birth, growth, decline and exit—is a critical driver of the reallocative process. An optimal pace of business dynamics balances the benefits of productivity and economic growth against the costs associated with reallocation—which can be high for certain groups of firms and individuals. While it is difficult to prescribe what the optimal pace should be, there is accumulating evidence from multiple datasets and a variety of methodologies that the pace of business dynamism in the U.S. has fallen over recent decades and that this downward trend accelerated after 2000.[1]

Canonical models of firm dynamics and empirical evidence imply that there is a tight link between business dynamism and productivity growth. As highlighted by Hopenhayn and Rogerson (1993), increases in the dynamic frictions of adjustment on the extensive or intensive margins will reduce the pace of reallocation and lower productivity. Thus, a *prima facie* concern arising from these trends in business dynamism is that they may have had adverse effects on aggregate productivity growth. The question is particularly important in light of the growing body of evidence showing that aggregate productivity growth in the U.S. has been declining since the early 2000s (Fernald (2014)).[2]

At first glance, medium-run fluctuations in economywide productivity growth do not match up with patterns of declining business formation and business dynamism. Productivity growth accelerated in the 1990s through the early 2000s before slowing down after 2003, while aggregate startup activity and job reallocation fell throughout the 1980-2014 period. However, a more careful review of theory and evidence resolves the inconsistency: during the 1980s and 1990s, the decline in entrepreneurship and reallocation was dominated by the Retail Trade

---

[1] See Davis et al. (2007), Haltiwanger, Jarmin and Miranda (2011), Reedy and Litan (2011), Hyatt and Spletzer (2013), Davis and Haltiwanger (2014) and Decker et al. (2014).

[2] The Hopenhayn and Rogerson (1993) finding is about the impact of an increase in adjustment frictions on the level of productivity. This can translate into a decline in productivity growth if there is an ongoing increase in adjustment frictions over time. Our findings below are consistent with a trend increase in adjustment frictions in the post 2000 period.

sector, where evidence suggests that falling dynamism was actually consistent with rising productivity growth.[3]

Fernald (2014) highlights that the surge in productivity from the late 1980s to early 2000s and the subsequent decline were both led by the ICT-producing and intensive ICT-using sectors. Interestingly, the High Tech sector exhibits a rise in business formation and job reallocation over the first period and a sharp decline in the post-2000 period, with the period since 2000 also being characterized by a decline in high-growth firm activity throughout the US economy more generally (Haltiwanger, Hathaway and Miranda (2014)).[4]

In this paper, we find that changes in how businesses respond to their idiosyncratic productivity conditions are an important driver of the evolution of aggregate job reallocation and productivity in recent decades, especially in the High-Tech sector. We argue that the observed decline in responsiveness is consistent with models of firm dynamics in which increases in adjustment frictions can reduce the pace of reallocation and, consequently, productivity growth. As noted above, the canonical model is Hopenhayn and Rogerson (1993), but this theme is consistent with a wide class of firm-level adjustment cost models (e.g., Cooper and Haltiwanger (2006), Cooper, Haltiwanger and Willis (2007, 2016), and Elsby and Michaels (2013)). The core hypothesis is intuitive. An increase in adjustment frictions makes firms more cautious in responding to idiosyncratic productivity shocks. This yields a decline in the pace of job reallocation (as firms' hiring and downsizing decisions become more sluggish), an increase in the dispersion of marginal revenue products and a decline in aggregate productivity.

There are other possible sources of changes in the responsiveness of firms that relate to the observed decline in the pace of startups and the accompanying changes in the age structure of firms. In particular, if young firms are more responsive to productivity and profitability shocks, changes in the age composition of the firm distribution would reduce overall responsiveness. Young firms are more volatile arguably because they are engaged in resolving uncertainty about their type through the learning dynamics hypothesized by Jovanovic (1982). Young firms facing uncertainty about their prospects will enter small, but as they learn about their type they will

[3] There is an extensive literature documenting the shift away from single unit establishment firms ("Mom and Pop" firms) to large national chains (see, for example, Foster et al. (2006) and Jarmin, Klimek and Miranda (2009)). For evidence that establishments of large national chains are more productive and more stable see Foster et al. (2006) and Foster et al. (2015). We discuss this evidence further below.

[4] For this purpose, we follow Hecker (2005) and construct a High-Tech sector based on the 14 four-digit NAICS industry groups with the largest shares of STEM workers. The 14 industry groups are listed in the appendix.

respond substantially to actual realizations of productivity. Greater responsiveness of young firms compared to mature firms combined with changes in the age structure of firms yields changes in overall responsiveness. The explanation of such changes in responsiveness depend on the origins of the changes in the pace of startups and young firms.

We mitigate the identification challenges of the sources of changes in responsiveness by focusing on changes in patterns of reallocation dynamics within firm age groups. Evidence suggests that the changing pace of startups and the accompanying change in the age composition of firms is an important contributing factor to the changing overall patterns of reallocation, but it is far from the dominant factor. Consistent with earlier findings (e.g., Davis et al. (2006), Decker et al. (2014)), we find that the changing firm age structure accounts for about one quarter of the overall decline in the job reallocation rate. This fraction varies substantially across sectors, but in the Information sector the changing age structure accounts for about one fifth of the increase in reallocation in the 1990s and about one fifth of the decrease in the post-2000 period. The dominant role of within-age group variation in changes in reallocation motivates our focus on that variation in this paper. This approach is not meant to suggest that variation in the pace of startups and the age structure of businesses is not relevant, but as we discuss below there are a number of confounding factors that may underlie such changes. Focusing on within-firm age variation provides cleaner identification.

We investigate these issues with a particular focus on the High Tech component of the Manufacturing sector where we can construct measures of establishment-level total factor productivity (TFP) to empirically describe idiosyncratic shocks. Given measured productivity shocks faced by businesses, we can evaluate both changes in the distributional characteristics of idiosyncratic productivity ("changing shocks") and changes in business-level responses (in terms of growth and survival) to their own productivity ("changing responsiveness"). We compare and contrast our results for High Tech Manufacturing with the remainder of the Manufacturing sector. Then, using new data on firm-level labor productivity for the entire U.S. private sector, we perform similar analyses of the relationship between firm-level productivity and growth for businesses outside of Manufacturing. While the evolution of the dispersion in labor productivity is inherently endogenous, our labor productivity exercises strongly suggest that the patterns we observe in Manufacturing generalize more broadly to other areas of the economy.

With reference to the model-based framework of changing shocks vs. changing responsiveness to shocks, we find that the dispersion of idiosyncratic productivity within High Tech Manufacturing is rising over time while the persistence of productivity exhibits little or no trend. Thus, there is little evidence that changes in patterns of idiosyncratic productivity shocks can account for the changing pattern of reallocation in the High Tech sector. Instead, we find that business-level responsiveness to productivity shocks has changed significantly over time within firm age groups. In High Tech Manufacturing, during the 1980s and 1990s plant-level survival and growth became more responsive to idiosyncratic TFP differences across plants, while in the post-2000 period responsiveness declined substantially. These results are particularly notable when we study exit alone, where we find that the intensity of productivity-based exit selection has weakened markedly. We conclude that the declining pace of aggregate reallocation in the U.S. is not being driven by changes in the distribution of business-level productivity shocks but is instead related to a dampening of the marginal responses of individual businesses to those shocks.

In addition to shedding light on the puzzle of declining business dynamism generally, our results have important implications for aggregate productivity. Counterfactual exercises show that the increased responsiveness of the 1980s and 1990s yielded as much as half a log point annual boost in industry-level TFP in the High Tech sector by the second half of the 1990s. 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 rise and fall of aggregate productivity growth in the U.S., which was concentrated in ICT-related industries (Fernald (2014)).

We find that the post-2000 decline in the responsiveness of businesses to their idiosyncratic productivity realizations is remarkably robust and widespread. Within Manufacturing, the post-2000 decline in responsiveness occurs among both young and mature businesses and both inside and outside of the High Tech sector. When productivity is measured as output per worker, allowing us to expand our analysis outside Manufacturing, we see that the change in the productivity/growth relationship is widespread across the major sectors of the U.S.

economy. A particularly striking finding for the wider economy is that we observe an increase in within-industry dispersion of labor productivity (within firm age groups) for both the High Tech and Non Tech sectors of the economy. Such increases in within-industry productivity dispersion are consistent with the decline in responsiveness that we observe being generated by an increase in adjustment frictions.

We do not identify the increase in adjustment frictions driving our results. However, we conduct some brief auxiliary exercises that provide guidance about the nature of the changes we have detected. For example, we show that the responsiveness decline among High Tech Manufacturing businesses can also be observed when we measure productivity responsiveness in terms of equipment investment instead of employment growth. The decline in employment-growth responsiveness is not the result of a key composition shift in High Tech Manufacturing—the shift of production from general-purpose to special-purpose technology production documented by Byrne (2015)—but is occurring within these categories of producers. We also find that the decline in responsiveness in the High Tech Manufacturing sector is especially large in the industries that have seen the greatest increase in import penetration from low-wage countries in the post-2000 period.

The paper proceeds as follows. Section II describes key facts about the declining pace of business dynamism that are relevant for our analysis. A key result here is that while the changing firm age distribution is important, most of the variation in the patterns of reallocation is not accounted for by the age distribution. Section III describes the various datasets we employ. In section IV, we use establishment-level data for Manufacturing, with a particular focus on High Tech Manufacturing, to study whether the evidence implies a change in the distribution of shocks or a change in the response to those shocks, 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. Section VI explores possible mechanisms underlying the changing responsiveness patterns to productivity shocks in the Manufacturing sector. Concluding remarks are in section VII.

II. Business Dynamics: Basic facts

A. Sectoral Patterns of Reallocation and Young Firm Activity

Various studies have documented a decline in the annual pace of establishment-level job reallocation and other indicators of business dynamism in recent decades (see, e.g., Davis et al. (2007), Decker et al. (2014, 2016), Davis and Haltiwanger (2014), Hyatt and Spletzer (2013) and Malloy et al. (2016)).

Underlying the aggregate decline in reallocation is substantial variation across sectors. Figure 1 shows the trends in job reallocation (using HP trends) for selected NAICS sectors along with the economywide trend. Retail Trade exhibits the sharpest decline in job reallocation rates during the 1980s and 1990s. In contrast, Information and FIRE exhibit increases in the pace of reallocation until about 2000 and then sharply decline thereafter. In a related fashion, Figure 2 shows the share of employment accounted for by young firms for the same sectors and economywide. Neither FIRE nor Information exhibits the declines in young firm activity through 2000 exhibited by sectors such as Services and Retail Trade. The share of employment accounted for by young firms in the Information sector rises in the second half of the 1990s then declines after 2000. Figures 1 and 2 together highlight that not all sectors have exhibited a monotonic decline in indicators of business dynamism and entrepreneurial activity.

The changing patterns of the share of young activity in Figure 2 account for an important fraction of the changing patterns of reallocation in Figure 1. Figure 3a shows the annualized change in reallocation rates for the same broad NAICS sectors from the business cycle peak in the late 1980s to the business cycle peak in the late 1990s, and Figure 3b shows the decline from the business cycle peak in the late 1990s to the mid 2000s (we use three-year averages in 1987-89, 1997-99, and 2004-06 for this purpose). Also depicted are the annualized changes holding the age composition of businesses constant (at its initial state) within each of these sectors. For this purpose, we use seven firm age groups (firm age 0,1,2,3,4,5,and 6+ where firm age 0 are startups). During the 1990s, the sharp decline in reallocation in the Retail Trade sector and the increase in reallocation in the Information sector are evident. The changing age composition helps account for both of these patterns. In the 1990s, the declining share of young business activity accounts for 32 percent of the decline in reallocation in Retail Trade, and the rising share of young business activity in Information accounts for 23 percent of the rise in reallocation in that sector. The Services sector exhibited a relatively smaller decline in reallocation rates in the 1990s, but interestingly 100 percent of the 1990s decline is accounted for by the declining share of young business activity in that sector. Turning to the post-2000 period, it is evident that the pace of decline in job reallocation accelerated. During the post-2000 period all broad sectors exhibited a decline (unlike the 1990s). The Information sector exhibits the sharpest decline in the post-2000 period, with 18 percent of the decline accounted for by a decline in young business activity. The main inference we draw from Figure 3 is that most of the variation in the reallocation patterns is within firm age groups. This finding encourages us to proceed by focusing on changing patterns of responsiveness of firms within firm age classes.

B. Possible Sources of Decline in Startups

While it is not our focus, the changing patterns of startup activity and, in turn, the changing age structure of businesses is also of great interest; there are a number of competing hypotheses to account for such variation. For example, variation in startup rates may endogenously reflect changes in the pace of innovation in an industry for the reasons hypothesized by Gort and Klepper (1982): a period of rapid innovation leads to a surge in entry, reallocation and subsequent productivity growth from the innovation. [5] Moreover, Gordon (2016) has hypothesized that most of the pathbreaking developments in the ICT sector and ICT-using sectors were implemented during the 1990s, and the productivity slowdown is due to a declining pace of innovation and implementation. The implication of these combined hypotheses is that the changing pace of startups and the accompanying change in reallocation in the High Tech sector may be related to productivity dynamics, but with the causality running from innovation variation to the associated changes in startups, reallocation and productivity growth.

The changing age structure of firms is also connected to changes in the business model in some sectors. In Retail Trade, there has been a well-documented transition from single unit establishment firms to large national chains (see, e.g., Foster et al. (2006), Jarmin et al. (2009) and Foster et al. (2016)). These studies show that the share of sales and employment from single unit establishment firms fell from about 50 to 35 percent from 1977 to 2007. Almost all of this is accounted for by a rapid rise in the share of sales and employment accounted for by large, national chains. These studies also find establishments in large, national chains are more productive (by about 30 log points) and more stable than single unit establishment firms (exit rates for single unit establishments are more than 15 times higher than that for establishments of

[5] Foster et al. (2017) provide supportive empirical evidence for these dynamics for the 1990s in the U.S. High Tech sector.

large, national chains). As argued in these studies, retail consolidations was likely facilitated by advances in information technology and globalization that have permitted large, national chains to build large and efficient supply chains and distribution networks. For our purposes, two key points are of interest. First, Retail Trade provides an example where the decline in reallocation is associated with a change in the business model that has been productivity enhancing. Second, this change is reflected in the age structure of businesses in Retail Trade from which we abstract in our main analysis.

Yet another contributing factor to the decline in startups may be the reduction in the population growth rate in the U.S. Karahan, Pugsley, and Sahin (2015) argue that as population growth slows so will the growth in the number of firms, an adjustment accommodated by fewer new firms.

Finally, the hypothesis that is the main focus of this paper—rising frictions inducing lower responsiveness—may also be a contributing factor for declining startups. An increase in frictions raises the cost of business activity and reduces the expected discounted value of profits for entrants. In this respect, our focus on within-firm age variation may be understating the contribution of the decline in responsiveness to the decline in productivity in the post-2000 period. However, more work is needed to sort out the reasons for the decline in startups which we leave for future work.

C. Reallocation Patterns for High Tech Industries

Before proceeding to the main analysis, it is instructive to show the patterns of reallocation for the High Tech industries of the economy. The Information sector includes some (but not all) of the sectors that are typically considered the High-Tech sectors of the economy. Included in Information are industries such as Software Publishing (NAICS 5112) and Internet Service Providers and Web Search Portals (NAICS 5161), but other High Tech industries are classified as Manufacturing, such as Computer Hardware and Peripherals (NAICS 3341), or Services. Moreover, Information includes sectors that are not considered part of the High-Tech sector such as Newspaper, Periodical and Book Publishing (NAICS 5111) and Radio and Television Broadcasting (excluding cable) (NAICS 5151). For this purpose, we follow Hecker (2005) and construct a High Tech sector based on the 14 four-digit NAICS sectors with the largest shares of STEM workers. The 14 sectors are listed in Table A1 in the appendix.

For a core part of our empirical investigation we focus on the Manufacturing sector, with a particular emphasis on High Tech Manufacturing. Figure 4 shows the Hodrick-Prescott trends for the Information sector, the High Tech sector as defined above, the manufacturing component of the High Tech sector, and the overall Manufacturing sector. Information, High Tech, and High Tech Manufacturing all exhibit very similar patterns highlighting that there was a rising pace of business dynamism in the High Tech part of the economy through 2000, but this has declined sharply in the post-2000 period. Focusing on the High Tech sector is of interest since it is a critical sector for innovation and productivity growth as highlighted by Fernald (2014). Interestingly, Figure 4 shows a sharp decline in the pace of job reallocation in the post-2000 period in the High-Tech sector coinciding with the trend slowdown in productivity driven by a slowdown in IT-producing and using industries.

In what follows, we focus on the High Tech sector of the economy. This focus is motivated by the outsized role that the High Tech sector has had in productivity dynamics. We also report results for other sectors for purposes of comparison. As emphasized above, we focus on dynamics within firm age groups in these sectors. This helps us abstract from the many interesting but confounding factors that may underlie the changing pattern of startups and the associated age dynamics.

III. Data and Measurement

The backbone datasets for our analysis are the Longitudinal Business Database (LBD), the Annual Survey of Manufactures (ASM) and revenue data from the Census Business Register that have recently been integrated into a firm-level version of the LBD. The LBD includes annual observations beginning in 1976, and we use the LBD through 2013 (this version of the data has consistent NAICS codes for the entire period as constructed by Fort and Klimek (2016)). It provides information on detailed industry, location and employment for every establishment in the private, non-farm sector. Employment observations in the LBD are for the payroll period covering the 12th day of March in each calendar year. For a full description of the LBD, see Jarmin and Miranda (2002).

A unique advantage of the LBD is its comprehensive coverage of both firms and establishments. In the LBD, firm activity is captured up to the level of operational control instead of being based on an arbitrary taxpayer ID. The ability to link establishment and firm information allows firm characteristics such as firm size and firm age to be tracked for each establishment and firm. Firm size measures are constructed by aggregating the establishment information to the firm level using the appropriate firm identifiers. The construction of firm age follows the approach adopted for the BDS and based on our prior work (see, *e.g.*, Becker et al. (2006), Davis et al. (2007) and Haltiwanger, Jarmin and Miranda (2013)). Specifically, when a new firm ID arises for whatever reason, we assign the firm an age based on the age of the oldest establishment that the firm owns in the first year in which the new firm ID is observed. The firm is then allowed to age naturally (by one year for each additional year it is observed in the data) regardless of any acquisitions and divestitures, as long as the firm as a legal entity continues operations. We utilize the LBD to construct annual establishment-level and firm-level employment growth rates.

A. The ASM: Plant-Level TFP Measures

For the main analyses in the paper, we focus on the Manufacturing sector where we can construct measures of establishment-level TFP. To do so, we supplement the LBD with a consistent and representative plant-level TFP database for all plants in the Annual Survey of Manufactures (ASM) and the Census of Manufactures (CM) from 1981 to 2010.[6] The strength of these data is that we are able to measure plant-level TFP for over 2 million plant-year observations. A limitation of the ASM in non-Census years is that, while it is representative in any given year, it is a rotating sample so its longitudinal properties are inferior to those of the LBD. Following FGH we integrate the ASM/CM TFP data into the LBD. For the LBD we have the outcomes in terms of establishment-level growth for all manufacturing establishments. For the integrated ASM/CM/LBD we have the subset of establishments from the LBD for which we can measure TFP. We use propensity score weights to adjust the ASM/CM/LBD sample so that it matches the complete LBD for manufacturing in terms of the detailed industry, size and age distributions (see FGH for details).[7] A second key advantage of integrating the ASM/CM data

[6] We are building on the data infrastructure developed by Foster, Grim and Haltiwanger (2016)—hereafter FGH. Our empirical specification also is closely related to FGH. The latter examined the changing responsiveness of reallocation to productivity over the cycle. We use the same terms FGH used for this purpose as controls in our analysis.

[7] The propensity score approach is based on a logit model that estimates the probability that a plant in the LBD (the universe) is in the ASM/CM as a function of detailed industry, firm size and firm age. It allows us to make the cross sectional distribution of plants in any given year be representative of the LBD on these dimensions. Note that these

with the LBD is that it allows us to avoid using the ASM to construct measures of growth and survival. The rotating panel nature of the ASM makes this difficult on many different dimensions.

The plant-level TFP measure we use is an index similar to that used in Baily, Hulten and Campbell (1992) and a series of papers that built on that work.[8] The index is given by:

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

where Q^R is real output, K is real capital, L is labor input, M is materials, E is energy, α denotes factor elasticities, the subscript e denotes individual establishments and the subscript t denotes time. Details on measurement of output and inputs are in FGH, but we provide a brief overview here. Nominal output is measured as total value of shipments plus the total change in the value of inventories. Output is deflated using an industry-level deflator from the NBER-CES Manufacturing Industry Database. Capital is measured separately for structures and equipment using a perpetual inventory method. Labor is measured as total hours of production and non-production workers. Materials are measured separately for physical materials and energy where each is deflated by an industry-level deflator. Outputs and inputs are measured in constant 1997 dollars. Factor elasticities are estimated using industry-level cost shares (of total factor costs).[9] A Divisia index approach is used for the latter so that industry-level cost shares are permitted to vary over time.

Given the large differences in output and input measures as well as the production technology across industries, we focus on a relative measure of TFP within industries. We do this by creating measures of (log) TFP that are deviations from the detailed industry-by-year average. We use detailed (e.g., 6-digit NAICS) industry effects for this purpose. We refer to this as "TFP" in the remainder of the paper, but it should be interpreted as the log deviation of establishment-level TFP from the industry-by-year average. Given our focus on within-industry-

weights are appropriate for making the cross sectional distribution in any given year representative but are not the ideal weights for using samples of ASM/CM that are present in both t and t+1. We discuss this further below.

[8] Syverson (2011) provides an excellent summary.

[9] Cost shares yield factor elasticities under the assumptions of cost minimization and full adjustment of factors. We are not assuming the latter 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. We consider an alternative approach as described below using proxy methods to estimate revenue function elasticities.

by-year idiosyncratic shocks, this implies we are abstracting from the direct influence of aggregate and industry-specific shocks on firm growth dynamics. Importantly, this measure of establishment-level productivity allows us to avoid problems associated with the mismeasurement of industry-level prices, such as those documented by Byrne and Corrado (2015, 2016). We refer to this idiosyncratic TFP measure as the productivity shock.

The framework we have in mind is that the idiosyncratic component of TFP is a persistent process, and we model this below as an AR(1) process. The current-period realization of the idiosyncratic component of TFP is the shock, and we also consider innovations to these shocks by estimating the AR(1) process below.

Our measure of TFP is based on revenue. In this respect, we are using a TFPR measure of productivity.[10] This means differences in establishment-level prices are embedded in our measure of productivity. Unfortunately, the Census Bureau does not collect establishment-level prices on a wide scale in the ASM and CM. However, as Foster, Haltiwanger and Syverson (2008) (henceforth FHS) have shown, it is possible to measure establishment-level prices and physical quantities for a limited set of homogeneous commodity-like products in Economic Census years (years ending in "2" and "7"). FHS create a physical quantity measure of TFP (which they denote as TFPQ) for a set of 11 homogeneous goods (for example, white pan bread). The within-industry correlation between TFPR and TFPQ is high (about 0.75). However, FHS also find an inverse relationship between physical productivity and prices consistent with establishments facing a differentiated product environment. In addition, FHS find establishment-level prices are positively related to establishment-level demand shocks and that such demand shocks are positively correlated with TFPR. As such, our measure of establishment-level productivity should be interpreted as reflecting both technical efficiency and demand factors (including product quality factors that may be embedded in prices). For our purposes, a key finding from FHS is that the relationship of growth and survival with TFPQ and demand shocks is quite similar to the relationship of growth and survival with TFPR. It is time variation in the relationship between growth and survival and our revenue-based measure of TFP that we are exploring.

[10] As defined by Foster, Haltiwanger and Syverson (2008), TFPR = P*TFPQ where the latter is physical productivity or technical efficiency. There are alternative revenue productivity concepts that we consider below.

While our specific TFP measure is commonly used in related literature, we also consider an alternative productivity measure based on the revenue function estimation proxy method described by Wooldridge (2009). As emphasized by Foster et al. (2017), this approach has the advantage that even in the face of endogenous plant-level prices, this alternative revenue productivity measure is only a function of exogenous TFPQ and demand shocks. The reason is that in estimating the revenue function, the estimates are of revenue elasticities and not factor elasticities so that they capture both production and demand parameters. In Appendix C, we describe exercises based on this revenue productivity residual (RPR) in place of TFPR. As shown there, our main results are robust to this alternative productivity specification. The results reported in the main text, however, rely on the TFPR measure described above. In our theoretical analysis below (both in section IV.A and in appendix D), we also show that in the presence of adjustment costs, TFPR and revenue labor productivity are inherently highly correlated with fundamentals. For example, in our simulated analysis in appendix D, the correlation of TFPR and revenue labor productivity with TFPQ is about 0.90 for the baseline calibration that matches the pace of reallocation in the 1980s in the US manufacturing sector.

B. Revenue Data from the Business Register: Measuring Revenue Labor Productivity.

While our main empirical investigation focuses on establishment-level data for the Manufacturing sector where TFP is measurable, in Section IV we extend our analysis to the entire economy by constructing measures of firm-level labor productivity. For this purpose we make use of a new revenue-enhanced version of the Longitudinal Business Database (RE-LBD). The RE-LBD integrates employment measures from the LBD with revenue-based output measures from administrative files. We describe basic features of its construction and underlying data here (see Haltiwanger, Jarmin, Kulick and Miranda (2016) for additional details).

The RE-LBD is a firm level file with revenue and employment measures. The RE-LBD incorporates firm-level revenue measures by collapsing the EIN revenue measures contained in the Business Register that belong to the same firm. Firm-level employment measures result from collapsing employment in the LBD across all establishments that belong to the same firm. Revenue coverage begins in 1996 and runs through 2013. A critical aspect of the RE-LBD is that the procedure for calculating the current-year and previous-year employment variables is adjusted so that all employment growth represents organic changes in establishment-level employment rather than artificial growth created by mergers and acquisitions (M&A).

Several additional aspects of the RE-LBD are worth mentioning at this point. The Census Bureau is not always able to tie revenue and employment measures belonging to the same firm in the BR.[11] This results in employment observations with missing revenue data. The RE-LBD sample excludes observations with missing revenue.[12] To account for the missing data and the potential selection effects that might arise we use inverse propensity scores to weight the data. Propensity scores are developed using the full LBD compared with the RE-LBD using models that include firm size, firm age, the employment growth rate, broad industry and a multi-unit status indicator. We use inverse propensity scores that are generated independently based on models for continuers, births and deaths (details about the filters used and the construction of the propensity scores can be found in the data appendix to Haltiwanger et al. (2016)).

We construct a relative measure of firm-level labor productivity within detailed industries that mimics our TFP concept described above. Specifically, the measure we use is (the log of) revenue per worker deviated from detailed (6-digit NAICS) industry-by-year means. Again, by deviating from detailed industry-year means we control for several factors. First, we control for relative price differences across detailed industries so our measure of labor productivity is consistent with a relative gross output per worker measure within detailed industries. Second, prior research (see, e.g., Foster, Haltiwanger and Krizan (2001, 2006)) has shown that relative gross output per worker within industries has a high correlation with relative value added per worker within industries and a strong correlation with relative TFP measures within industries. This reflects the fact that there are similar patterns of, for example, materials shares across firms in the same industry as well as other factor shares such as capital shares.

For this firm-level analysis, we assign a firm to a detailed industry based on its modal employment. This is a potential source of measurement error since large multi-unit establishment firms may be operating in multiple 6-digit NAICS industries. We note, however, that over 95 percent of all firms and more than 99 percent of all young (less than five year old) firms are single unit establishment firms. This implies that the industry-by-year mean we deviate

[11] This is because firms can report income and payroll activities under different EINs. When this happens the income EIN may fall outside of the set of EINs that the Census considers part of that firm when accounting for employment.

[12] Approximately 20 percent of the firm-year observations do not have output measures.

from is very similar in practice to the industry-by-year mean for single unit establishment firms. As discussed in appendix E, we have conducted sensitivity analysis that uses a more sophisticated approach to controlling for industry effects for multi-unit establishment firms that operate in multiple industries. We find our results are robust to this alternative approach.

Finally, given the difficulties associated with measuring the output and productivity of firms in the Finance, Insurance, and Real Estate sectors, we omit firms in those industries (NAICS 52-53) from all analysis below.

IV. Change in shocks vs. change in responsiveness

A. Theoretical motivation

Having described basic facts about business dynamism and our data, we now proceed with a framework guided by canonical models of firm dynamics. Models of firm[13] dynamics suggest that a within-sector decline in the pace of reallocation is either due to a change in the volatility of shocks faced by firms or a change in the response to those shocks. A classic reference for our purposes is Hopenhayn and Rogerson (1993). In that paper, firms face idiosyncratic productivity shocks and adjustment frictions for labor; an increase in adjustment frictions reduces the dispersion of firm-level growth rates and reduces aggregate productivity because productivity-enhancing reallocation is reduced.

To be more specific, canonical models of firm dynamics with adjustment frictions yield decision rules for net hiring rate (similar remarks apply to the growth of other inputs, such as investment in capital) that are a function of the two key state variables each period: the realization of productivity and the initial employment in the period.[14] A reduced form representation of this decision rule is given by: $g_{ft+1} = f_{t+1}(A_{ft+1}, n_{ft})$ where g_{ft+1} is the net hiring rate (or equivalently the net employment growth rate) of firm f between t and $t+1$, n_{ft} is employment in t, and A_{ft+1} is the realization of firm (idiosyncratic) productivity (which is typically observed prior to the growth decision). A_{ft+1} is considered to be a persistent process, and consistent with the literature we specify this as an AR1 process in our empirical work below.

[13] We use the term "firms" loosely in this subsection. Much of the literature focuses on establishment-level dynamics but we use the term "firm" in this section for expositional ease. Our empirical work focuses on both establishment-level and firm-level dynamics.

[14] For net hiring rate dynamics, see, e.g., Cooper, Haltiwanger and Willis (2007, 2015) and Elsby and Michaels (2013). For investment dynamics, see, e.g. Cooper and Haltiwanger (2006).

Decision rules for the net hiring rate that take this form imply that changes in the distribution of firm employment growth rates stem from either changes in the distribution of A_{ft+1} or the responsiveness of the firm to a given realization of A_{ft}.

Our empirical approach in the next section is motivated by the insights of this class of firm dynamic models with adjustment costs. Building on this literature, in Appendix D we consider a kinked adjustment cost model of labor adjustment costs to illustrate the features and predictions of this class of models that are especially relevant for our empirical analysis. The model we specify is consistent with the discussion above. Firms are subject to idiosyncratic productivity shocks where the realization of productivity in the current period is drawn from an AR1 process. Because our empirical analysis uses alternative measures of productivity including revenue productivity-based measures, we specify that firms face downward sloping demand curves with an isoelastic demand structure. This specification yields that the revenue function exhibits curvature and is subject to idiosyncratic revenue shocks from productivity and potentially demand shocks. We don't formally model entry and exit but discuss below the implications of considering these extensive margins based on the insights of the existing literature.

We conduct numerical analysis of a calibrated version of this model to motivate the empirical specifications and moments we consider below. We calibrate the dispersion and persistence of the productivity processes to be consistent with the empirical estimates for Manufacturing in the 1980s detailed below. We calibrate the adjustment cost parameters to match the pace of job reallocation in the 1980s. One of the findings in the baseline calibration that is relevant for the empirical approach we take below is that revenue productivity measures—such as TFPR or revenue labor productivity—are highly correlated with TFPQ (pairwise correlations of about 0.90) in the presence of a baseline level of adjustment costs.

Using this as a starting point, we consider two types of exercises. First, we consider how key moments change with an increase in adjustment frictions. This yields the following predictions that we take to the data in our analysis. We find that as adjustment frictions increase (holding constant the shock process and all other parameters): (i) job reallocation declines; (ii) the response of firm-level employment growth from t to $t+1$ to the realization of productivity in t (conditional on employment in t) declines;[15] (iii) the standard deviation of labor productivity increases; and (iv) the Olley-Pakes (OP) covariance for both TFP and labor productivity declines. The latter implies a decline in aggregate productivity consistent with a higher extent of misallocation in the economy due to the increase in adjustment frictions.[16] Most of these predictions are intuitive and are consistent with predictions in the literature.[17] These are the primary hypotheses that we investigate in our empirical analysis. From the basic facts above, we already know that there are changing patterns of job reallocation, and our objective is to explore whether these changes are consistent with a change in adjustment frictions.

In considering the "increasing frictions" hypothesis, we have in mind a potentially broader interpretation than the simple adjustment frictions in our illustrative model. Given the recent literature on idiosyncratic distortions affecting the allocation of factors (e.g., Restuccia and Rogerson (2008), Hsieh and Klenow (2009) and Bartelsman et al. (2013)), an increase in the dispersion of distortions will both have a similar impact on the relationship between firm growth (and survival) and fundamentals and have adverse implications for productivity. Moreover, while our simple model has only employment dynamics, we have in mind any type of increased friction that may impede adjusting the scale of operations at a firm.

[15] As we discuss in Appendix D, there are a variety of candidate moments relating growth to realizations of productivity that yield similar qualitative predictions with respect to changes in adjustment frictions. We focus on this specific one for our empirical analysis for measurement and econometric reasons as discussed below. We also show in Appendix D that this predicted decline in responsiveness with an increase in adjustment costs holds whether we use TFPQ, TFPR or RLP (revenue labor productivity). This is because, as we have noted above, in the presence of adjustment costs all of these measures are highly correlated.

[16] In standard models, including our illustrative model, the relationship between adjustment frictions and the covariance between size and labor productivity (the "Olley-Pakes covariance") is nonlinear, as we show in Appendix D. This is because, under strong but common assumptions, an absence of adjustment frictions would imply that firms equate their marginal revenue products of labor, producing zero covariance between size and labor productivity. The covariance rises as adjustment costs initially rise (as labor productivity dispersion moves above zero), but the covariance eventually falls over the range of plausible adjustment cost parameterizations as misallocation rises. The frictionless case is not a realistic or useful benchmark as it produces extremely high rates of job reallocation. The intuition on which we rely—with increasing frictions causing a decline in covariance and therefore, *ceteris paribus*, a decline in aggregate productivity—is reasonable under empirically plausible parameterizations of adjustment costs. We note that the exercise we use in the empirical analysis conducts a diff-in-diff between the OP covariance with the estimated responsiveness in the current period vs. the counterfactual OP covariance if the estimated responsiveness was the same as at the beginning of the sample. We show in Appendix D that the diff-in-diff exercise in the simulated data shows that when adjustment frictions rise the diff-in-diff OP covariances for both TFP and labor decline.

[17] The decline in reallocation, the rise in dispersion in labor productivity and the decline in productivity from a reduction in allocative efficiency are found in Hopenhayn and Rogerson (1993) and Cooper, Haltiwanger and Willis (2007).

In contrast to the "increasing frictions" hypothesis for the empirical fact of declining job reallocation and related moments, we also investigate an alternative hypothesis suggested by the existing firm dynamics literature as well as our illustrative model. Our model has the property that a decline in the dispersion of TFP (holding the persistence of TFP and all other model parameters constant) yields: (i) a decline in job reallocation; (ii) a decline in the response of firm-level growth from t to $t + 1$ to the realization of productivity in t (conditional on employment in t); and (iii) a decline in the standard deviation of labor productivity.[18]

These predictions highlight that a change in the dispersion of firm-level growth rates (e.g., job reallocation) can be accounted for either by changes in the distribution of productivity/profitability of shocks or by changes in the marginal response of firm-level growth to productivity/profitability shocks (or by some of both).[19] Moreover, a change in the shock distribution can change firm-level responsiveness; to disentangle these forces we need to capture the empirical evolution of both the shock distribution and changing responsiveness.

One important distinguishing feature of the shocks vs. responsiveness framing of our study of the change in dynamism is that the two hypotheses have opposite predictions for the dispersion of labor productivity. That is, if a decline in the pace of job reallocation is due to increasing labor adjustment frictions then we should also observe an accompanying increase in the dispersion of labor productivity. This is because increased frictions reduce the speed with which firms move their marginal revenue products toward equalization. In contrast, if a decline in the pace of job reallocation is due to a decline in the dispersion of shocks then we should observe an accompanying decline in the dispersion of labor productivity. We exploit these starkly different predictions about within-industry productivity dispersion in the empirical analysis that follows.

[18] The first and the third predictions are intuitive, but a few remarks are in order about the second prediction. It stems in part from our consideration of a kinked adjustment cost model so that there is a range of inaction in adjustment. As dispersion in TFP decreases 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). This prediction, then, is consistent with the findings of Vavra (2014) who argues that the standard finding in the literature is that the volatility effect dominants the real options effect in the steady state, a general result extending back to Barro (1972).

[19] 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.

This latter point on implications for changes in the patterns of labor productivity dispersion highlights a more general feature of our empirical analysis—specifically, examining multiple moments rather than focusing on the predictions for a single moment. As we note above, the decline in the pace of reallocation could be due to a decline in the dispersion of shocks, but this would also imply a decline in the dispersion in labor productivity. As we also discuss further in Appendix D, the Olley-Pakes covariance moment for labor productivity exhibits a non-monotonic pattern with respect to increases in adjustment costs over the full range of the parameter space. However, over the range where increases in adjustment frictions yield a decline in the OP labor covariance, such changes yield declines in the responsiveness of firm growth to shocks, a decline in the OP covariance for TFP and an increase in the dispersion of labor productivity. Exploring these multiple margins simultaneously is therefore important for identification.

There are additional forces that may be at work beyond changes in frictions and shocks that are not apparent from our illustrative model and abbreviated discussion of the firm dynamics literature. First, our illustrative model in Appendix D neglects firm entry and exit, and we already know from the basic facts that there have been striking changes in entry dynamics in recent decades. Hopenhayn and Rogerson (1993) incorporate entry and exit dynamics and find that a rise in adjustment frictions will reduce entry and exit. In their model, the lower bound of productivity necessary for survival will decline with an increase in frictions. The empirical prediction, then, is that not only will firm growth for continuers become less responsive to productivity but so will exit. We explore this prediction in the empirical analysis below.

As discussed in the introduction, another related factor that is likely quite important is that firm dynamics of young firms differ from mature firms. In the empirical analysis that follows, we consider changing dynamics of responsiveness within firm age groups to abstract from changes in the prevalence of young firms in the U.S.

As will become apparent, we confront the theory largely by characterizing the evolution of key moments as well as the reduced form relationships discussed above. Unlike some of the literature cited here we do not seek to identify a structural model of adjustment frictions. Given our findings below, we think this is a rich area for future research. For example, relative to the discussion above, we do not take a stand on the exact form of adjustment costs such as convex vs. non-convex adjustment costs (this topic has been under active investigation in the literature).

One potential use of our empirical findings would be as moments to discipline such analysis.[20] A benefit of our reduced form approach is that it readily permits controlling for many different factors in a panel regression environment and allowing estimates to vary systematically by key firm characteristics such as detailed industry and firm age. In addition, we use this reduced form approach to explore potential explanations for changes in the responsiveness to shocks that we detect.

B. Empirical Analysis of U.S. Manufacturing

In this section, we investigate these issues for the U.S. manufacturing sector with a focus on the High-Tech component of manufacturing. For our purposes, the High-Tech sector includes the 4-digit sectors in Table A.1 that are in manufacturing.[21] To help provide perspective on our findings for High-Tech Manufacturing we also consider the rest of the Manufacturing sector which for ease of exposition we call Non Tech Manufacturing.

The first exercise we consider is to explore the evolution of the within-industry dispersion in (log) TFP, where dispersion is quantified as the standard deviation of the within-industry plant-level (log) TFP distribution (that is, we measure plant productivity deviated from its industry-year mean). This is our measure for the dispersion of idiosyncratic productivity shocks faced by establishments. Figure 5 shows the evolution of within-industry productivity dispersion for plants in High-Tech Manufacturing (top panel) and Non Tech Manufacturing (bottom panel). We report separate results for young and mature firms since the evidence presented above suggests that plants of young firms exhibit different paces of reallocation. We use the LBD and its firm age measures for each plant to classify plants in the ASM/CM/LBD integrated data for this purpose. Given our interest in low-frequency variation, we report HP trends of this measure of dispersion.

Consistent with the literature, there is large dispersion in TFP across plants in the same industry (see, *e.g.*, Syverson (2004, 2011)). We find that the levels of within-industry TFP

---

[20] Indeed, Cooper and Haltiwanger (2000) used reduced form regressions similar to those we estimate in an indirect inference estimation of structural parameters of adjustment costs (in their case the application was capital adjustment). They also show in the numerical analysis of their structural model that the marginal responsiveness of investment to profit shocks declines with increases in adjustment costs, either convex or non convex.

[21] These include manufacturers in 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).

dispersion are about the same for plants of young and mature firms in both the High-Tech and Non Tech areas of Manufacturing. In High-Tech Manufacturing, plants of young and mature firms exhibit a positive trend in dispersion that roughly mimics the overall. The same holds for Non Tech Manufacturing.[22]

To help understand the implications of this rising within-industry dispersion of productivity, it is also useful to study the patterns of persistence in plant-level TFP. Much of the literature on plant-level productivity has found that plant-level productivity shocks exhibit considerable persistence but are far from a unit root process. In terms of implications of productivity shocks for plant-level dynamics, the adjustment cost literature (*e.g.*, Cooper and Haltiwanger (2006) and Cooper, Haltiwanger and Willis (2007)) shows that the implied patterns of plant-level growth dynamics depend on the persistence of the idiosyncratic shocks. This is intuitive since in the face of adjustment costs plants are more likely to respond to persistent shocks. Our data infrastructure is not ideally suited for estimating persistence and the innovations to the process, but in Appendix B we find that persistence is reasonably stable with an AR(1) estimated coefficient of about 0.6 to 0.7, and the pattern of dispersion in innovations matches those of overall dispersion.

The findings presented thus far suggest that the changing patterns of reallocation are not driven by changing patterns in the dispersion of TFP or the persistence in TFP. Consider plants in High-Tech Manufacturing. Figure 4 shows a pattern of rising reallocation during the 1990s and then falling reallocation in the post-2000 period. For dispersion and persistence of TFP to account for these patterns we would expect dispersion and/or persistence to mimic these patterns. The patterns we present suggest that, if anything, we should see a rising pace of reallocation in High Tech and Non Tech industries in the post-2000 period, which is exactly the time during which we observe a decline in the pace of reallocation.

We now turn to investigating whether there is a change in the responsiveness of growth

[22] Figure A1 in the appendix reports overall dispersion (for all ages) for High Tech Manufacturing, Non Tech Manufacturing, and the entire Manufacturing sector. The standard deviation of log TFP averages about 36 log points for the all manufacturing and Non Tech manufacturing samples (the lines overlap in the chart). It averages 40 log points for the plants in the High-Tech part of manufacturing. For High-Tech, trend dispersion in TFPR rose mildly through the 1990s and then more substantially in the post-2000 period. For the remainder of manufacturing, trend dispersion of TFP was relatively constant through the 1990s but rose in the post-2000 period. In Appendix C, we also report a version of Figure A1 that is based on the Wooldridge (2009) revenue productivity residual (RPR) estimation instead of our TFPR measure; Figure C1 reports these results, confirming that the general upward trend of TFPR is robust to the RPR productivity concept.

and survival to idiosyncratic differences in TFP across plants. We estimate establishment-level regressions relating employment growth from t to $t+1$ to measured TFP in period t for all incumbents in period t with appropriate controls. We define employment growth for establishments between periods t and $t+1$ using the Davis, Haltiwanger and Schuh (1996) (hereafter DHS) concept that can accommodate entry and exit.

Equation (2) is our basic specification:

$$\begin{aligned} g_{e,t+1} = \lambda_{t+1} + \beta_y * TFP_{et} * Young_{et} + \delta_{1y} * TFP_{et} * Young_{et} * Trend_t \\ + \delta_{2y} * TFP_{et} * Young_{et} * Trend_t^2 \\ + \beta_m * TFP_{et} * Mature_t + \delta_{1m} * TFP_{et} * Mature_{et} * Trend_t \\ + \delta_{2m} * LP_{et} * Mature_{et} * Trend_t^2 + X'_{et}\Theta + \varepsilon_{e,t+1} \end{aligned} \tag{2}$$

where $g_{e,t+1}$ is the DHS employment growth rate for establishment e between time t and time $t+1$, TFP_{et} is total factor productivity for establishment e at time t deviated from (six-digit NAICS) industry-by-year means, and $Trend_t$ is a simple linear time trend (which we also include as a quadratic term in many specifications). $Young_{et}$ is a dummy equal to 1 if the plant's firm is young (age less than five) in year t, $Mature_{et}$ is a dummy equal to 1 if the plant's firm is mature in year t, and X_{et} is a set of controls discussed further below. Note that trend terms are not entered as main effects since there is a full set of year effects that capture general trends as well as national cyclical effects. We estimate equation (2) using our propensity score weights. All of the terms involving TFP are fully saturated with young and mature dummies, as the evidence in the prior sections suggest that the dispersion of plant-level growth dynamics differs systematically across plants owned by young and more mature firms.

While this is a reduced form specification, it is broadly consistent with the specifications of selection and growth dynamics from the literature we discussed above. First, it is consistent with the adjustment cost model calibration exercises in Appendix D that show, estimating the equivalent of equation (2) in the simulated data, that an increase in adjustment frictions will yield a decline in the responsiveness of plant-level growth to lagged realizations of TFP. Second, by using DHS growth rates, we can incorporate both the extensive margin (exit) and the intensive margin of plant-level growth. Standard empirical specifications of exit in the literature (see, e.g., Syverson (2011)) relate the decision to exit between t and $t+1$ to the realization of TFP in period t along with other controls (*e.g.*, endogenous state variables such as size, which is part of our X_{et} as described below). As we have already noted, adjustment cost models of employment growth yield predictions that relate the growth in employment from period t to $t+1$ to the realization of TFP in period t along with period-t size. In this sense, equation (2) produces a reduced-form yet direct estimate of policy functions generated by canonical models.[23] In interpreting the timing of equation (2), it is useful to note that 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 D than might first appear.[24]

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. In the exercises described following our discussion of the regression analysis, we exploit this time-varying productivity responsiveness estimation to understand the changing contribution of job reallocation to aggregate productivity growth.

We estimate specification (2) for 1981-2010 with the following controls as captured by X_{et}: the young firm dummy, establishment size, firm size, state effects and a state-level business cycle indicator (the change in state-level unemployment rate).[25] We also interact the state-level cyclical indicator with plant-level TFP following FGH. The cyclical variables are all interacted with the young and mature dummies. Since we are interested in the changing response and the Great Recession is at the end of our sample, we do not want our estimates of the changing trend responses (the main coefficients of interest) to be driven by the changes in the response to TFP over the cycle.

---

[23] There is also a measurement/econometric justification for (2) given the timing of the data. TFP in period t is measured for calendar year t while establishment growth is measured from March of t to March of $t+1$. Similar remarks apply to our labor productivity analysis below.

[24] In Appendix D, we show declining responsiveness of firm-growth to current or lagged realizations in productivity from an increase in adjustment frictions.

[25] For firm size effects, we use firm size classes in period *t*. For establishment size effects, we have considered both establishment size classes and log employment at the establishment level in period *t*. We obtain very similar results for both cases, and in the paper we use log employment at the establishment level.

The first column of Table 1 shows the estimates for the plants in High-Tech Manufacturing while the second column shows the estimates for the plants in Non Tech Manufacturing. We only report the main effects for TFP by firm age group and the interactions with the trend terms. All of the effects of interest for the High Tech sector in column 1 of Table 1 are statistically significant at the five percent level. For plants in Non Tech Manufacturing, four of the six coefficients in column 2 of Table 1 are statistically significant at the 10 percent level or better. Columns 3 and 4 of Table 1 report results for the changing responsiveness of exit that underlies part of the patterns of the first two columns.

The estimates for parameters β_y and β_m are given by the "TFP*Young" and "TFP*Mature" rows of Table 1. These positive coefficients show that, consistent with previous literature, growth and productivity are positively related. On average, establishments that are more productive than others in their industry are more likely to grow, while those that are less productive are more likely to contract or exit. This relationship is stronger among establishments of young than mature firms, consistent with stronger selection dynamics at work among recently entered businesses. The positive relationship between productivity and growth at the establishment level is consistent with a positive contribution of reallocation to aggregate productivity growth.

The estimates for parameters δ_{1y} and δ_{1m} are given by the "TFP*Young*Trend" and "TFP*Mature*Trend" rows of Table 1, respectively. These are our key coefficients of interest as they show how the marginal responsiveness of establishments to their own idiosyncratic productivity has changed with time. Notably, in the High Tech Manufacturing sector these coefficients are positive and significant for plants of both young and mature firms, suggesting that productivity responsiveness strengthened in the early years of the sample (which begins in 1980), while the coefficients are close to zero in the Non Tech Manufacturing sector. Both inside and outside of High Tech, however, the coefficients on the quadratic term are negative, suggesting interesting dynamics in productivity responsiveness that we now describe graphically.

Figure 6 shows the pattern of the marginal effect of TFP on plant-level growth for young and mature plants by decade. To compute these statistics, we set the cyclical indicator (the state level change in unemployment) to zero so the effects reflect controlling for the cycle but are evaluated at a neutral cyclical state. The top and bottom panels of Figure 6 show the patterns for High Tech Manufacturing and Non Tech Manufacturing plants, respectively. As mentioned above, the figure makes apparent that plants of young firms are more responsive to productivity than are plants of mature firms, particularly in High-Tech. Taken together with earlier findings, the high pace of reallocation of young plants is not driven by a high variance of TFP but rather by a high responsiveness to TFP differences. This is consistent with, for example, a learning model in which young plants are especially responsive to TFP as they learn where to find themselves in the productivity distribution.

Our main focus is the variation in the responsiveness over time. First, consider High Tech Manufacturing, shown on Figure 6a. The difference in responsiveness between plants in young and mature firms implies that overall responsiveness will change given changes in the age composition. For example, the increase in the share of activity accounted for by young businesses in High-Tech during the 1990s implies an increase in overall responsiveness, while the decrease during the post-2000 period implies a decrease in overall responsiveness. We also find interesting patterns within age groups. For plants in young firms, responsiveness increases from the 1980s to the 1990s and then declines in the post-2000 period. For plants in mature firms, responsiveness decreases throughout the time sample but accelerates during the post-2000 period.

Continuing our study of Figure 6a, note that the magnitudes of the coefficients readily lend themselves to quantitatively meaningful interpretations since productivity is deviated from its industry-year mean (and hence has mean zero). Recall that Figure 5a shows the standard deviation of within-industry productivity for High-Tech Manufacturing establishments; multiplying the relevant standard deviation by the relevant coefficient on Figure 5a suggests that a young-firm plant with productivity that was one standard deviation above its industry-year mean had an employment growth rate that was about 12 percentage points higher than the industry average in the 1980s, and it was about 17 percentage points higher in the 1990s but only about 10 percentage points higher in the 2000s. Among mature firms, the growth differential was about 5 percentage points in the 1980s, 6 percentage points in the 1990s, and 3 percentage points in the 2000s. For plants of both young and mature firms, if productivity dispersion had remained constant throughout the period then the decline in relative growth from the 1990s to the 2000s would have been even larger.[26]

---

[26] More precisely, to calculate the growth advantage of the plant that is a standard deviation above the mean versus the average plant in an industry, we simply multiply the standard deviation by the responsiveness coefficient. Rising

Figure 6b shows the analogous patterns for Non Tech Manufacturing. Here again we find that plants in young firms are more responsive to TFP than plants in mature firms. For Non Tech Manufacturing there has been a decline in the share of young business activity throughout the period implying a decline in overall responsiveness due to composition effects. Within age groups, we find a decline in responsiveness throughout the period with an acceleration of the decline in the post-2000 period. As above, we can construct rough estimates of the changing growth gap between high- and average-productivity plants (within industries) using standard deviation estimates from Figure 5b. Among young firms, the growth rate gap between plants with productivity one standard deviation above the mean and plants at the mean was roughly 10 percentage points in the 1980s, 9 percentage points in the 1990s, and only 6 percentage points in the 2000s. Among mature firms, the growth advantage was about 5 percentage points throughout the time period, though the advantage would have fallen had productivity dispersion remained flat.

In Figure C3 of Appendix C, we report analogous results based on the Wooldridge (2009) revenue productivity residual (RPR) estimation method. Results are quite similar both qualitatively and quantiatively.[27] Thus, using a measure of revenue productivity that more plausibly only reflects idiosyncratic productivity and demand shocks, we find evidence of changing responsiveness rather than changing shocks driving the decline in dynamism in the post 2000 period.

Putting the pieces together, the patterns imply an overall increase and then decline in responsiveness of growth to TFP for plants in High Tech Manufacturing. This is driven by a number of factors: (i) the higher responsiveness of plants in young firms and the shifting age

productivity dispersion means that the gap in productivity separating the average plant and the plant that is one standard deviation above average grows over time; for example, Figure 6a shows that a young-firm plant with productivity one standard deviation above the industry-year mean was about 35 log points more productive than average at the beginning of the time series and almost 50 log points more productive than average by 2010. The net growth differential figures calculated in the text reflect offsetting effects from rising dispersion (and hence rising productivity advantage) and declining marginal responsiveness, with the declining responsiveness factor being so large that it typically outweighs the rising dispersion factor.

[27] Figure C3 reports decade averages of coefficients from the regression in (3) but using RPR in place of TFP. We again observe a pattern of responsiveness for young-firm plants in High-Tech Manufacturing that rises from the 1980s to the 1990s, then falls in the 2000s. 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 in the TFPR-based regressions, but it is still significant as we show below in aggregate productivity counterfactuals. High-Tech Manufacturing plants of mature firms see broadly steady responsiveness from the 1980s to the 1990s followed by a sharp decline in the 2000s. The bottom panel of Figure C3 likewise shows that the RPR responsiveness patterns among Non Tech Manufacturing plants are quite similar to those based on TFPR and reported on in Figure 5.

composition; (ii) the increase and then decrease in the responsiveness of plants in young firms; and (iii) the decline in responsiveness for plants of both young and mature firms in the post-2000 period.

There may be interactions between the effects we have detected. The rising dispersion of TFP (and its innovations) in the post-2000 period could be contributing to declining responsiveness through expansion of inaction bands, as noted in our theoretical framework discussion above. But this cannot account for all of the declining responsiveness since rising dispersion of TFP typically should yield an increase in the pace of reallocation and we find the opposite in the post-2000 period. In addition, the 1990s exhibited a mild increase in productivity dispersion accompanied by an increase in responsiveness for plants in young firms (and rising reallocation).

Another potential source of interaction is the role of selection in influencing the observed dispersion in TFP. Columns 3 and 4 of Table 1 report results from regressions in which a binary exit indicator serves as the dependent variable (making this a linear probability model). We find that part of our declining responsiveness in the post-2000 period is due to a declining responsiveness of exit to productivity shocks. Figure A2 provides the graphical representation of the changing responsiveness of exit. For High Tech Manufacturing, we find exit for young plants became more responsive during the 1990s and then responsiveness declined. For Non Tech Manufacturing, exit for young plants became less responsive throughout the period. For mature plants (both tech and non tech), exit responsiveness does not change much from the 1980s to the 1990s but then declines in the post-2000 period.

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, which will be a drag on aggregate productivity. The reduced covariance between survival and productivity can also contribute to rising dispersion since low-productivity plants are more likely to survive. But again the patterns over time suggest this cannot be a dominant part of the story for changing dispersion. During the 1990s, we find increased responsiveness of exit but mild increases in dispersion in productivity. This also holds for the alternative RPR measure of productivity (see appendix Figure C.1).

Taken together, these results have important implications for the evolution of firm dynamics in recent decades. We find significant evidence that the way in which individual businesses respond to their idiosyncratic realizations of productivity has changed over time. The positive relationship between realized productivity and subsequent employment growth remains robust, but it has weakened over time (particularly since 2000). From the standpoint of the firm dynamics models described above, our results can be viewed as evidence that policy functions have changed over time, particularly for young businesses but also for older ones.[28] 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.

We emphasize that it is especially in the post-2000 period that there is evidence of increasing frictions. During the 1990s in High Tech, responsiveness of plants of young firms increased. This is not consistent with rising frictions, but it does not necessarily follow that this pattern is driven by a decrease in frictions over this period since this is a period with rising TFP dispersion especially for plants of young firms (see Figure 5).

C. Implications for aggregate (industry-level) productivity

How important are the changes in responsiveness for aggregate fluctuations in productivity? For this purpose, we exploit counterfactual differences in the weighted average of plant-level productivity based on the estimated models of changing responsiveness. Recall that in section IV.A (and more formally in Appendix D) we showed that the Olley-Pakes covariance indices computed in our numerical analysis of a calibrated model decline with a rise in adjustment costs. Motivated by these findings, we begin with the index of industry-level productivity underlying the Olley-Pakes decomposition:

$$P_t = \sum_e \theta_{et} P_{et} \tag{3}$$

[28] These results contrast with Karahan, Pugsley and Sahin (2016), who argue that the dynamics of incumbent firms (in a Hopenhayn framework) have not changed over this same time period. These authors point to aggregated average growth rates for various incumbent age classes as evidence for stable incumbent firm dynamics. We differ from their approach by directly estimating incumbent firm 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, though the labor force demographic evidence is an important component of the final explanation.

where θ_{et} is the weight for establishment e in year t and P_{et} is establishment-level (log) productivity.[29] The Olley-Pakes decomposition of this index is given by $P_t = \bar{P}_t + cov(\theta_{et}, P_{et})$ where $\bar{P}_t$ is the unweighted average productivity and $cov(\theta_{et}, P_{et})$ is the Olley-Pakes covariance. Our objective in this section is to isolate empirically the impact of the change in responsiveness on the patterns of the second term of this index. We do this by using the estimated model of responsiveness using a diff-in-diff counterfactual. Specifically, we use the estimated changing responsiveness model to compute:

$$\Delta_t^{t+1} = \sum_e (\theta_{e,t+1}^{T} - \theta_{e,t+1}^{NT}) P_{et} \tag{4}$$

where θ_{et+1}^{T} is the predicted employment share for establishment e in period $t+1$ based upon the estimated model with the parameters reflecting the changing pattern of responsiveness (where the T superscript refers to "trend"), and θ_{et+1}^{NT} is the predicted employment share for establishment e in period $t+1$ based upon the estimated model with the parameters reflecting the pattern of responsiveness at the beginning of the sample period (that is, we set the trend terms equal to zero, or NT means "no trend").[30] In computing the predicted employment share in each period, we use the actual realizations of productivity and initial employment for each establishment each period and then use the estimated model parameters to predict employment growth into the following period. This counterfactual diff-in-diff isolates the difference in the change in the productivity index accounted for by changes in responsiveness; it is the implied difference in the productivity index in the current period implied by the difference in responsiveness. Note that in making the calculation in (4), we hold constant plant-level productivity and hence the unweighted productivity distribution is unchanged. Since this implies that the unweighted mean productivity is the same in the two counterfactuals of the diff-in-diff, this diff-in-diff isolates the impact of the changing responsiveness on the OP covariance term[31]

[29] FGH show that the index of industry-level productivity in (3) yields fluctuations in industry-level productivity that mimic the patterns of productivity from aggregated statistics. For example, FGH construct measures of industry-level productivity growth using traditional growth accounting techniques with aggregated plant-level data yielding industry output and input measures. The correlation between such traditional measures and the industry-level measures that emerge from (3) from pooled industry by year data is 0.70 when employment weights are used in (3) and is 0.80 when output weights are used in (3).

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

[31] This approach is related to, but distinct from, the accounting productivity decompositions in the literature (see, e.g., Foster, Haltiwanger and Krizan (2001) for a review). Our present approach differs since it focuses only on model-driven reallocation that is identified to be the reallocation arising specifically from variation in productivity

A key advantage of this diff-in-diff counterfactual analysis is that it will only capture the effect of time-varying responsiveness within firm age groups. Our empirical results show that young firms are more responsive than mature firms on average, consistent with Jovanovic (1982) lifecycle dynamics of firms. Those differences will be present in both the counterfactual with and the counterfactual without the trend, as will the changing age structure of firms overall. Thus, our diff-in-diff counterfactual will yield a non-zero productivity contribution if and only if there are changes in responsiveness within firm age groups.

The results of this counterfactual exercise are depicted on Figure 7; the figure can be interpreted as follows. Each annual observation reports Δ_t^{t+1} from (4). For example, the observation for $t+1=1981$ has $\Delta_{1980}^{1981}=0$ because the trend variable begins then, and for High-Tech Manufacturing the year 2001 again gives $\Delta_{2000}^{2001}=0$. But the 2004 observation for High-Tech shows that 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 7 is driven by fitting a quadratic trend. But we regard our findings as implying that the drag on this index of industry level productivity due to the decline in responsiveness may be quite substantial.[32]

It is interesting that the changing responsiveness starts to be a drag on productivity in

across businesses. In addition, this approach focuses only on the reallocation components since the exercise holds constant the productivity distribution at the micro level between t and t+1. Still, in any given period this approach is close to quantifying the change in the Olley-Pakes covariance for a given distribution of productivity across plants. Decker et al. (2017) use the Dynamic Olley-Pakes decomposition developed by Melitz and Polanec (2015) to show that these accounting decompositions also show a decline in the contribution of the change in the covariance terms (often interpreted as indicators of allocative efficiency) in the post-2000 period. Alon et al. (2017) likewise use the Dynamic Olley-Pakes decomposition but focus on the cumulative contribution of changes in entry rates. In appendix D, we show that in the simulated data this diff-in-diff Counterfactual declines with a rise in adjustment costs.

[32] Another source of caution is that since we are using a TFPR measure the rising dispersion of TFPR may reflect endogenous factors in a manner similar to the labor productivity measures we use in the next section. Under this view, the counterfactual may partially reflect transitory or one-time gains. We note however that we obtain similar implications using the RPR measure in appendix C.

2003, about the time that Fernald (2014) finds a trend break in productivity growth in the IT sector. For Non Tech Manufacturing plants, the changing responsiveness has relatively little impact until the post-2000 period, consistent with the estimates in Table 1 that are near zero on the linear trend terms but negative on the quadratic terms. By 2005 the acceleration of the decline in responsiveness in this part of manufacturing yields as much as a half a log point drag on productivity per year.

In Appendix C, we report results of this counterfactual exercise that rely on the Wooldridge (2009) revenue productivity residual (RPR) method. The results for High-Tech Manufacturing are broadly consistent with the TFPR-based results, as shown on Figure C4, with rising responsiveness implying an increasing in the index of productivity in the 1990s followed by a decline in the 2000s. Interestingly, among Non Tech Manufacturing plants we find that the RPR method results in an earlier decline in the index from declining responsiveness than does the TFPR approach.

Given that we use the actual distribution in TFP in each year for these counterfactuals, the changing patterns of dispersion we have shown are also potentially contributing factors. However, since we are examining a "diff-in-diff" comparison, the changing pattern of dispersion influences both the counterfactual with and without the changing trend response. We also note that 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. If young firms play a critical role in the innovative process, then a decline in the share of young firm activity can also contribute to a declining pace of within-firm productivity growth. We include some further discussion and analysis of within-firm productivity growth patterns below.

V. Beyond Manufacturing

Thus far, our analysis has focused on the Manufacturing sector. As mentioned above, this is a sector in which we have high-quality data on TFP, which allows us to map the theoretical framework discussed above to the data on employment growth. While we find the evidence from Manufacturing—and High-Tech Manufacturing in particular—to be compelling, an important question is whether the patterns of productivity dispersion and responsiveness we have documented are present in other areas of the economy. Much of the innovation that characterizes the U.S. economy, particularly after 2000, has been in the Information sector (e.g., software publishing, internet portals, and so on). Moreover, the patterns of changes in business formation and in the dispersion and skewness of firm growth rates are even more dramatic in non-manufacturing components of the High-Tech sector (see Decker et al. (2016)). Relatively little is known about the connection between reallocation dynamics and productivity growth outside of Manufacturing.[33]

While the question of productivity responsiveness and aggregate productivity growth is difficult to answer in the absence of precise concepts and data for TFP in the non-manufacturing sectors, we can provide evidence using output per worker as our productivity concept. For this purpose, we employ the RE-LBD dataset (described in Section III above), which permits the measurement of revenue per worker at the firm level. A key strength of the RE-LBD is its comprehensive coverage of the U.S. private, non-farm sector from the mid-1990s to 2013. We emphasize that our analysis in this section differs from the previous section not only in its labor-based definition of productivity but also in the use of firms as the unit of analysis (rather than establishments).

When studying TFP and Manufacturing, we adopted a "shock responsiveness" interpretation of the relationship between business-level growth and productivity; here we adopt a looser interpretation with a ready acknowledgement that gross output per worker (the measure of labor productivity we use in this section) is endogenous reflecting not only TFP but also changes in adjustment frictions. Recall from section IV that we noted that an increase in adjustment frictions implies an increase in the within-industry dispersion of labor productivity across firms. The relevant mechanism is that an increase in adjustment frictions slows down the

[33] Significant exceptions include the numerous studies of Retail Trade cited above that highlight the shift in the business model toward large, national chains that has been productivity enhancing. See, for example, Foster, Haltiwanger and Krizan (2006), Jarmin, Klimek and Miranda (2009) and Foster et al. (2015).

tendency for marginal revenue products to be equalized which, in turn, will yield an increase in the dispersion of measured labor productivity.[34] Thus, observed increases in the within-industry dispersion of labor productivity may reflect increases in the dispersion of shocks, increases in adjustment frictions or both.[35]

Even using labor productivity measures it is still the case that an increase in adjustment frictions will reduce the covariance between firm employment growth and labor productivity for the empirically plausible range of adjustment costs. As we show in Appendix D, an increase in adjustment frictions yields a decrease in the responsiveness of firm-level employment growth from t to $t + 1$ to the realization of revenue labor productivity in t (holding constant the level of employment in t). We note for this exercise that the timing of the measures is closely related to that for the prior analysis: specifically, we are measuring employment growth from March t to March $t - 1$, and the productivity measure is for the calendar year t while we control for the (log) level of employment in March t.

A. Productivity and growth at the firm level

Figure 8a (the top panel of Figure 8) reports the interdecile range of within-industry labor productivity differences across firms in the High-Tech sector. Note that our definition of High-Tech is broader in this section than it was in our discussion of Manufacturing results; in our labor productivity exercises, High-Tech includes not only High-Tech Manufacturing but also certain industry groups from Services and Information. The dispersion of labor productivity in High Tech increases substantially over our sample period within firm age groups. Notably, however,

[34] It is this type of insight that has led researchers such as Hsieh and Klenow (2009) to interpret increases in the dispersion of the average product of labor as representing an increase in the distortions or wedges that are impeding the equalization of marginal revenue products. The strict interpretation by Hsieh and Klenow (2009) depends on strong functional form assumptions (see Haltiwanger (2016) and Haltiwanger, Kulick and Syverson (2016)), but this perspective is potentially relevant in our setting. We take a more agnostic position here acknowledging that changes in dispersion of labor productivity may reflect changes in shocks, changes in frictions or both. If we used the Hsieh and Klenow (2009) assumptions of Cobb-Douglas production with constant returns to scale and iso-elastic demand with common markups, then we would interpret the rising dispersion in labor productivity within sectors as reflecting increasing wedges or distortions. We cannot rule out this interpretation, and it is consistent with our findings of smaller responsiveness of growth to productivity differences within sectors. The simple model we consider in Appendix D elaborates on these issues.

[35] There may be other factors at work as well. For example, changes in the dispersion of capital intensity across firms in the same industry will yield changes in the dispersion of measured labor productivity. This could arise if there is some change in the production structure of firms and may be relevant, for example, in the presence of biased technological change where not all firms adopt new technologies. Alternatively, changes in capital intensities may reflect changes in the adjustment frictions for capital. Thus, changes in responsiveness of employment growth detected by our analysis may in fact reflect changes in the responsiveness of capital accumulation.

the dispersion of labor productivity is much higher among young firms than mature firms, while TFP dispersion is roughly equal for young and mature (see Figure 5). This pattern is consistent with the discussion in section IV that young firms face greater uncertainty, learning or other frictions. This pattern also highlights the importance of controlling for firm age in empirical exercises. The declining share of young firms in High Tech over this period acts as a dampening factor on overall rising within-industry dispersion in High Tech. Figure 8b shows that labor productivity dispersion is also rising among firms outside the High Tech sector, again importantly controlling for firm age, and young firms have much higher dispersion than more mature firms.

Our finding of rising within-industry productivity dispersion is broadly 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, concluding that the pace of technological diffusion has slowed. While the authors do not provide direct evidence for the hypothesis that slowed technological diffusion is the cause of increasing productivity dispersion, the diffusion hypothesis could play a role; however, our estimates of TFP persistence (in Appendix B) suggest that the group of "frontier firms" is sufficiently fluid to somewhat limit the diffusion story's explanatory power. Weakening responsiveness of growth and survival to productivity that we document in the next section is an alternative, but not mutually exclusive, explanation. Both explanations allow for a decoupling of technological progress and aggregate productivity growth.[36] However, as we have discussed above, declining responsiveness is potentially consistent with increased adjustment frictions or other forces that slow down the tendency for marginal revenue products to be equalized.

The evidence on labor productivity dispersion suggests that changes in productivity shock patterns are not the cause of declining aggregate job reallocation. If the only factor at

[36] 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). While ORBIS coverage of the U.S. is weaker than its coverage of other countries, we view our evidence as strongly supportive of the notion that gaps between the most productive firms and other firms have increased since the late 1990s. In that sense our work is complementary to the work of ACG, though we caution that their measures of productivity dispersion are sufficiently conceptually different from ours as to make direct quantitative comparisons difficult. 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)). Our measure of dispersion is defined within detailed (6-digit NAICS) industries and is based upon percentiles rather than the selection of an absolute number of businesses.

work were changing shock patterns, we should have observed a decline in labor productivity dispersion. Instead we observe the opposite. To explore these issues further, we now study the relationship between labor productivity and growth using an approach similar to our analysis in Section IV. We measure DHS employment growth, but in this case we use firm-level growth rather than establishment-level growth. We then estimate the following equation:

$$\begin{aligned} g_{f,t+1}^{E} = \lambda_{t+1} &+ \beta_y * LP_{ft} * Young_{ft} + \delta_{1y} * LP_{ft} * Young_{ft} * Trend_t \\ &+\delta_{2y} * LP_{ft} * Young_{ft} * Trend_t^2 \\ &+\beta_m * LP_{ft} * Mature_t + \delta_{1m} * LP_{ft} * Mature_{ft} * Trend_t \\ &+\delta_{2m} * LP_{ft} * Mature_{ft} * Trend_t^2 + X_{ft}'\Theta + \varepsilon_{\mathrm{f,t+1}} \end{aligned} \tag{5}$$

Equation (5) is almost identical to (2) except that growth, productivity, and size are all measured at the firm level (versus the establishment level as in (2)). The variable LP_{ft} is our measure of firm-level labor productivity, which is (log) deviated from detailed industry-year mean. We estimate this equation over a shorter time period (1997-2013) reflecting the limitations of our revenue dataset, but we include the same controls in X_{ft} (see Section IV). As we discuss above, this specification stretches the “shocks vs. responsiveness” interpretation we adopted in our TFP-based regressions, but it is still useful for studying the relationship between growth and productivity and can be similarly used in counterfactual exercises. We estimate (5) using propensity score weights given that we have only about 80 percent of the firms in the LBD have revenue productivity measures.

Table 2 reports results of the regression from (5) for all firms and for the High Tech and Non Tech sectors separately. The first three columns report regressions using the DHS growth rate denominator inclusive of exit; the last three columns report results using only a binary exit outcome as the dependent variable. Figure 9 graphically shows the time series pattern of the coefficients for overall growth.

As shown on Figure 9a (the top panel of Figure 9), we find a significant positive relationship between productivity in year t and employment growth from t to $t+1$ among both young and mature firms in both High-Tech and Non Tech industries, consistent with our TFP results and existing literature (e.g., Syverson (2011)). The productivity/growth relationship is stronger for young than for mature firms. All categories of firms see a reduction in the strength of the productivity/growth relationship. The reduction in strength is greater for young High Tech firms than young Non Tech firms.

The probability of exit is decreasing strongly in productivity, with young firms being particularly sensitive to productivity selection. Figure 9b shows that the negative relationship between productivity and exit is moderating among all young firms, with a particularly strong effect among young High Tech firms. Broadly speaking, the evidence suggests that the survival and growth differential between high- and low-productivity firms in the High Tech sector is declining over time.

In summary, we find that the relationship between productivity and growth at the firm level, while robust, is weakening over time in most sectors of the economy and among both young and mature firms. The same is true for exit with a general weakening of the productivity selection mechanism across sectors and age classes. The decline in this relationship is particularly pronounced in the High-Tech sector. The results on both labor productivity dispersion and the relationship between labor productivity and growth broadly confirm our TFP-based results from Manufacturing and suggest that the patterns documented in the latter sector are likely to hold in other areas of the economy. As we argued above, the framework of firm dynamics models, applied to our evidence, suggests that slowing reallocation indicates the presence of increased frictions rather than changes in the distribution of idiosyncratic productivity shocks. Our evidence shows rising dispersion of idiosyncratic shocks in manufacturing and rising dispersion of labor productivity more broadly. The former, holding other factors constant, should have yielded an increase in the pace of reallocation. The latter is inherently endogenous but is consistent with rising frictions and distortions.

B. Reallocation and aggregate labor productivity

Following the approach from Section IV, we can quantify the labor productivity regression results by relating them to aggregate productivity growth using shift-share analysis. With the only differences being the use of labor productivity (versus TFP) and firm-level data (versus establishment), we again construct a base-year index using the analogue to equation (3) and a model-based diff-in-diff counterfactual using equation (4).[37] We report results for High Tech and Non Tech firms on Figure 10.

While this exercise is mechanically identical to that performed with TFP in Section IV, we caution that the economic interpretation is more complicated. As mentioned above, labor productivity dispersion is endogenous to productivity responsiveness even beyond the selection effects that influence the distribution of TFP. While selection effects can make the dispersion of both TFP and labor productivity endogenous to responsiveness, in the case of labor productivity dispersion there are additional mechanisms of slowed responsiveness yielding a slowed pace of marginal revenue product equalization and related issues associated with capital deepening and other factors. Labor productivity dispersion rose during the 1996-2013 period, and this increase is likely at least partially endogenous to the decline in responsiveness we document. This complicates the shocks/responses intuition that guided our discussion of related exercises in Section IV.

In particular, these considerations imply that the productivity counterfactuals in Figure 10 partially reflect cumulative factors in addition to simple annual responsiveness mechanisms. The lower productivity we estimate for each year partly reflect both current and cumulative past shocks to which firms still have yet to fully adjust. The diff-in-diff calculation each year is thus an estimate of the gains in productivity there would be if adjustment dynamics suddenly reverted to their more responsive 1997 rates using the current dispersion of productivity, which is in part driven by previous adjustment dynamics. We note again that our diff-in-diff approach yields a contribution only through changing responsiveness within firm age groups.

With the above details in mind, the effects we find are quite large. By 2012, the weakening productivity/growth relationship accounts for about 6 log points in the productivity index. 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, Figure A4 in

[37] The analogue to the productivity index in (3) yields industry-level indices that are very similar to a traditional index of gross output per worker at the industry level. The correlation between the industry level indices from the analogue to (3) and traditional indices is 0.76. At the economy-wide level, the correlation between the growth rate in traditional gross output per worker labor productivity measures and the employment-weighted micro based estimate is 0.81. We also show in Appendix Figure A.4 that the industry-level indices using (3) yield patterns very similar to the traditional index of gross output per worker using BLS industry level indices.

Appendix A shows diff-in-diff results for only the Manufacturing sector, with results separated by High Tech and Non Tech Manufacturing as in our TFP-based exercises. Within Manufacturing, we do indeed find stronger results for High-Tech than for other firms, which is qualitatively consistent with our TFP-based investigation.

C. Changing Patterns of Within-Firm Productivity Growth

One missing piece from our analysis thus far is that we have provided little evidence on patterns of within-firm productivity growth. Our focus has been on the changing patterns of reallocation and the potential implications of these changes for productivity growth. To help put those results into perspective, it is instructive to examine the patterns of within-firm productivity at the micro level alongside the changing patterns of the contribution of reallocation. For this purpose, we use the firm labor productivity database and construct two related but distinct measures of within-industry productivity growth for each 6-digit NAICS industry. The first measure is the simple unweighted mean of within-firm productivity growth for each industry. The second is the employment-weighted mean of within-firm productivity growth using the employment weights at time t for the productivity growth between t and $t + 1$. We compute these measures at the 6-digit industry level and then show averages across industries using time-invariant employment weights for each industry.

Figure 11 shows within-firm productivity growth for the average industry for industries in both High Tech and Non Tech using both the weighted and the unweighted approaches. We show HP filtered series given our focus on low-frequency variation. Several patterns are worth noting. First, for High Tech, within-firm productivity growth declines using both the weighted and the unweighted measures. Second, for Non Tech, within-firm productivity growth declines for the weighted measure but exhibits less systematic variation for the unweighted measure. Third, for both High Tech and Non Tech, the weighted measure is much larger than the unweighted measure. Moreover, the unweighted measure is always negative for Non Tech and turns negative for High Tech early in the sample. This latter finding might at first glance be surprising since it implies negative productivity growth for the average firm in Non Tech throughout and even for High Tech in the second half of the sample. However, as discussed by Decker et al. (2017), when interpreting the unweighted mean within-firm productivity growth it is important to emphasize that it overwhelmingly reflects the contribution of very small firms. More than 90 percent of firms have fewer than 20 employees.

Decker et al. (2017) shed further light on this issue by noting that the difference between the weighted and the unweighted measure is given by:

$$\sum_f (\theta_{ft} - 1/N)\Delta p_{ft}$$

where θ_{ft} is the employment share of firm f in period t, Δp_{ft} is the growth (log difference) in firm-level productivity from t to $t+1$ and N is the number of firms in the industry. The implication is that a positive difference between the weighted and the unweighted mean reflects a positive relationship between within-firm productivity growth and initial shares. Returning to Figure 11, weighted within-firm productivity growth is always positive for both High Tech and Non Tech—but it is declining.

For the purposes of the current analysis, the main inference we draw is that there is no evidence of increasing within-firm productivity growth over the period during which we have observed a declining contribution of reallocation.[38] Instead, during the period of a declining contribution of reallocation there is also declining within-firm productivity growth. This is important since one possibility is that some structural change had occurred so that there was a substitution away from productivity-enhancing reallocation to greater within-firm productivity growth. For example, it might have been that as the High Tech sector matured in the post-2000 period, the productivity gains in High Tech came increasingly from within-firm innovations rather than through reallocation dynamics. Figure 11 suggests this is not the case.

VI. Investigating mechanisms in Manufacturing

In this section we return to our analysis of the Manufacturing sector using TFP data from the ASM/CM/LBD linked data employed in Section IV. While a thorough examination of the causes of the empirical patterns we describe above is beyond the scope of the present study, we can provide evidence on several common hypotheses with a particular focus on High Tech

[38] 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.

Manufacturing businesses: capital substitution, globalization, and composition changes within High-Tech.

A. The response of equipment investment

The changing responsiveness of employment growth to productivity shocks, especially by plants of young firms, may be due to changing margins of adjustment. One possible change in the margin of adjustment is that establishments that have high productivity may be expanding through capital accumulation rather than employment growth. To investigate the possible role of capital-labor substitution, we examine the changing responsiveness patterns of equipment investment to productivity shocks. The ASM/CM data we are using has equipment investment flows in each year, so it is straightforward to construct equipment investment rates as the real investment in equipment divided by beginning-of-period equipment capital.

We estimate the analogue to equation (2), this time using as the dependent variable the equipment investment rate instead of the employment growth rate.[39] Both investment and employment growth are endogenous margins of adjustment, so we can think about our specifications as reduced form specifications relating both of these margins of adjustments to productivity shocks, controlling for initial size in the period. When estimating equation (2) for employment growth rates, we controlled for initial size by using lagged plant- and firm-level employment controls. In this analysis of equipment investment, we additionally control for the beginning-of-period (log) capital stock. Thus, our specification includes the key state variables: the realization of productivity and initial capital and employment.[40]

[39] The timing is slightly different for the equipment investment as opposed to the employment growth specifications. In the employment growth specification, the dependent variable is employment growth between March of t to March of t+1 as a function of initial size in t and the realization of productivity in period t. In the investment specification, the dependent variable is the investment rate throughout period t as a function of initial size (measured by both capital and labor) and the realization of productivity in period t. There is a time to build assumption in capital accumulation with investment in period t contributing to capital to be used in period t+1 (available for use at the beginning of period t+1). Given this time to build assumption, the difference in timing is not large – employment is from March t to March t+1 while the investment rate instead captures the capital accumulation from January of t to January t+1. Moreover, our model exercises in Appendix D suggest that our responsiveness framework is not sensitive to specific timing concerns. Note that Cooper and Haltiwanger (2006) used as a key moment to estimate and calibrate their adjustment cost model the correlation between the investment equipment rate in period t and the realization of the profitability in period t. Our specification can be interpreted as identifying how this covariance between investment and profitability is changing over time.

[40] As a robustness check, we have re-estimated specification (2) with the employment growth rate as the dependent variable but adding lagged capital as a control. This alternative specification makes the employment growth rate specification consistent with the investment equipment rate specification. In unreported results, we find the coefficients reported in Table 1 are robust to including lagged capital as a control.

Table 3 reports the estimated coefficients plants in High Tech Manufacturing. Consistent with theory, literature, and the employment-based results we describe above, equipment investment is positively related to plant productivity. Moreover, plants of young firms are more responsive than plants of mature firms, as in our employment growth results. However, it is apparent the responsiveness of equipment investment is changing over time, particularly for young-firm plants.

Figure 12 shows the changing marginal responsiveness of equipment investment to productivity for plants of young and mature firms in recent decades. Strikingly, the patterns in Figure 12 mimic the patterns of employment growth responsiveness shown in Figure 6. Just as with employment growth, equipment investment responsiveness increases from the 1980s to the 1990s but then declines sharply in the post-2000 period. We can again quantify the magnitude of these changes in responsiveness by referencing productivity dispersion numbers from Figure 5a. In the 1990s, a young-firm plant with TFP one standard deviation above the industry-year mean had an equipment investment rate 8 percentage points higher than a plant with the industry-by-year mean productivity; this difference is less than 4 percentage points in the post-2000 period. The growth differential among mature firms falls as well.

The strong similarity between Figures 12 and 6 implies that, among High Tech Manufacturing businesses, the post-2000 declining responsiveness of employment growth in Figure 6 is not accounted for by rising responsiveness of equipment investment over this same period of time. In unreported results, however, we find that among Non Tech Manufacturing businesses there is rising investment responsiveness from the 1980s to the 1990s, with responsiveness remaining elevated in the 2000s. Given the employment-responsiveness evidence from Table 1 and Figure 6, it does appear that capital-labor substitution may play a role in changing employment growth dynamics outside of the High Tech sector.

B. Globalization

Globalization may be playing a role 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. 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)). We build on that research to explore the impact of rising import penetration for changing responsiveness of U.S. establishments to productivity differentials.

Bernard, Jensen and Schott (2006) and Schott (2006) 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 latter using the public domain information from Census on imports by country and industry.[41] We integrate 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.[42]

Using these extended import data, we find that import penetration ratios from low-wage countries are very small in the 1980s, rise slightly in the 1990s and then rise dramatically in the post-2000 period (as shown on Figure A5 in Appendix A). Of particular interest for current purposes is that the rise is especially pronounced for the 6-digit NAICS industries in High-Tech.

We exploit the 6-digit NAICS variation in import penetration ratios to explore the possible role of this dimension of globalization in accounting for the patterns of changing responsiveness we have detected. Table 4 presents results of a modified version of estimating specification (2). 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. We report the same coefficients as in Table 1 with the addition of these two interaction effects. We note that the main effect of the import penetration (not reported) is negative and significant. Consistent with Bernard, Jensen and Schott (2006), we find that plants in industries with especially large increases in import penetration have lower net employment growth.

[41] 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 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.

[42] 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. The latter provides micro based concordances between NAICS and SIC for the 1981-88 period.

Our interest is in the role of globalization for changing responsiveness. The last two rows of Table 4 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 13, we quantify the impact of changing import penetration ratios using the estimated effects from Table 4. The overall effects show, consistent with Figure 6, that the marginal effect of productivity on employment growth increased from the 1980s to 1990s for plants of young High-Tech firms and 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 4 along with the aggregate pattern of import penetration ratios for High-Tech Manufacturing. The impact of rising penetration is very modest in terms of accounting for the change in the 1980s to 1990s. It goes the "wrong way" but it is small. However, in the change in responsiveness from the 1990s to 2000s, the rapid rise in the import penetration ratios accounts for a substantial share (about 16 percent) of the overall decline in responsiveness.

While this investigation does not uncover specific micro mechanisms that yield the connection between import penetration to changing dynamism and productivity responsiveness, it does suggest this is a promising area for future research.

C. Changing Specialization of High-Tech

We also investigate the hypothesis that the declining responsiveness of growth to productivity during the post 2000s is the transition from "general-purpose" to "special-purpose" equipment manufacturing in the U.S as documented by Byrne (2015). Businesses manufacturing these special-purpose products might be less responsive to productivity due to demand constraints or uncompetitive environments that reduce adjustment imperatives. We investigate this hypothesis in Appendix F. We verify evidence of the changing structure of High-Tech manufacturing but find no evidence that this accounts for the declining responsiveness observed in High-Tech Manufacturing.

VII. Conclusion

The post-2000 period in the U.S. has been characterized by declines in both gross job flows and aggregate productivity growth. The decline in both has been most dramatic in the High Tech areas of the economy.

A key insight from previous literature on declining business dynamism, that patterns of young firm activity and gross job flows differ markedly across sectors, is of particular relevance to the productivity question. During the 1980s and 1990s, the decline in dynamism was dominated by the Retail Trade sector related to the well-documented shift away from "Mom and Pop" retail businesses to large, national chains. Since the evidence also shows that the productivity of large, national chains is substantially greater than that of Mom and Pop businesses, this decline in dynamism in Retail Trade arguably reflects benign changes in retail business models so that the typical Retail Trade establishment has become both more productive and more stable over time. In the post-2000 period, however, there has been an acceleration of the overall decline in indicators of business dynamism that has been led by declines in key innovative sectors such as High Tech where reallocation and entrepreneurship had previously been rising. Since 2000, this is the sector with the largest declines in these indicators.

We focus our attention in this paper to changing patterns of responsiveness within firm age groups. We first show that while the changing age structure of firms is an important contributing factor to the observed variation in reallocation dynamics, it is far from the dominant factor. Focusing on dynamics within firm age groups, we study the role of changes in the idiosyncratic productivity or profitability "shocks" faced by individual businesses vs. changes in how businesses respond to those shocks. Focusing on the Manufacturing sector and using a large dataset on establishment-level TFP, we find that the distribution of productivity has not become less volatile or otherwise evolved in a way that would be expected to reduce job reallocation. However, by directly estimating the employment growth policy functions governing business-level responses to productivity, we learn that marginal productivity responses at the business level have tracked patterns of aggregate reallocation. In the High Tech sector, responsiveness rose in the 1990s then fell in the 2000s, consistent with reallocation data for that sector. For the manufacturing sector broadly, responsiveness declined steadily throughout 1980-2010. We conclude that aggregate patterns of reallocation are closely related to changes in the marginal responses of individual businesses to their own productivity; canonical models of firm dynamics suggest that this declining marginal responsiveness is likely related to increased costs or frictions on adjustment.[43] These TFP-based results for the Manufacturing sector are broadly echoed by an alternative approach based on a revenue productivity residual (RPR) definition of productivity in place of TFP.

We find further support for the notion that responsiveness has declined by studying firm-level labor productivity for the entire U.S. economy and the High-Tech sector (including High Tech businesses outside of Manufacturing). Consistent with our TFP-based results from Manufacturing, we find evidence of a sharp weakening of the relationship between firm-level growth and productivity in the post-2000 period with particularly strong declines in the High-Tech sector. We also find rising dispersion of within industry firm-level labor productivity dispersion consistent with declining responsiveness that emerges from some type of increase in frictions or distortions.

The changing pattern of responsiveness of plant-level growth and survival to TFP has implications for aggregate (industry-level) productivity growth. We find that increased responsiveness of growth and survival to idiosyncratic differences in TFP in High Tech Manufacturing during the 1990s yielded an increase in the an industry-level Olley-Pakes covariance productivity index of as much as half a log point per year. In turn, we find that the decline in responsiveness of plant-level growth to idiosyncratic TFP differences in High Tech yields a decline in the Olley-Pakes covariance index of as much as two log points per year. Our economywide evidence based on labor productivity finds similarly substantial declines in the covariance indices from declining responsiveness.

Declines in responsiveness could potentially be balanced by increased within-firm productivity growth. A plausible hypothesis is that there is a tradeoff between productivity growth from reallocation and within productivity growth (e.g., technological and organizational changes might induce frictions in adjustments that are leveled by within-firm growth). We find no evidence of increasing within-firm productivity growth over the period during which we have observed a declining contribution of reallocation. Instead, we find declining within-firm productivity growth during this period as well.

[43] Using industry variation, Goldschlag and Tabarrok (2014) find no evidence that federal regulation counts relate with changes in the pace of gross flows. Davis and Haltiwanger (2014) find evidence relating employment protection policies to lower rates of reallocation. Molloy et al. (2016) find no evidence of a role for land use regulations or improved worker-firm matches in declining worker flows.

Though the primary focus of this paper is to document declining productivity responsiveness and analyze its implications using multiple data sources and methodologies, we also briefly investigate three possible mechanisms underlying these changing patterns of firm dynamics. First, we investigate the hypothesis that the changing responsiveness reflects plants changing their margin of adjustment from employment to capital; we find that during the post-2000 period the High Tech sector saw a decline in investment responsiveness while responsiveness was flat outside of High Tech. Second, we explore the possibility that globalization may be playing a role since increased exposure to foreign trade facilitates adjustment by scaling international operations. We find evidence in support of this hypothesis in that it is especially in detailed industries with large increases in import penetration from low-wage countries that young High Tech plants have exhibited large declines in responsiveness in the post-2000 period. Such changes in import penetration account for about 16 percent of the decline in responsiveness of young plants in High-Tech manufacturing in the post-2000 period. Finally, we investigate compositional changes in Manufacturing reflecting movement from general-purpose to special-purpose technology production, finding no evidence of a role of these shifts in declining responsiveness.

In addition to shedding light on the drivers of declining business dynamism, these findings comprise a novel and important contribution to the literature on the productivity slowdown in the U.S. in the post-2000 period. We interpret the micro evidence as implying that the productivity slowdown in the post-2000 period in High Tech industries cannot be explained solely by a decline in the pace of innovation and technological change. We draw this inference from multiple pieces of evidence. First, our findings focus on changes in responsiveness within firm age groups. Second, in the post-2000 period we find rising within-industry productivity dispersion in High Tech. A slowdown in the pace of innovation should, other things equal, yield a decline in within-industry productivity dispersion. We find the opposite in the post-2000 period in High Tech.

References

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

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.*

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.

Baily, Martin Neil, Eric Bartelsman, and John Haltiwanger. 2001. "Labor Productivity: Structural Change and Cyclical Dynamics." *Review of Economics and Statistics*, 83(3), 420-433.

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

Becker, Randy A., John Haltiwanger, Ron S. Jarmin, Shawn D. Klimek, and Daniel J. Wilson. 2006. "Micro and Macro Data Integration: The Case of Capital." Chap. 12 in *Studies in Income and Wealth* edited by Dale W. Jorgenson, J. Steven Landefeld, and William D. Nordhaus, vol. 66, *A New Architecture for the U.S. National Accounts*, pp. 541-609. Chicago, IL: National Bureau of Economic Research, University of Chicago Press.

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.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica*, 623-85.

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.

Cairó, Isabel. 2013. "The Slowdown in Business Employment Dynamics: The Role of Changing Skill Demands." *Mimeo.* At http://www.econ.upf.edu/eng/graduates /gpem/jm/cairo.html.

Cooper, Russell and John Haltiwanger. 2000. "On the Nature of Capital Adjustment Costs." NBER Working Paper No. 7925.

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.

Cooper, Russell, John Haltiwanger, and Jonathan Willis. 2015. "Dynamics of labor demand: Evidence from plant-level observations and aggregate implications," *Research in Economics*, 69: 39-50.

Davis, Steven J. and John Haltiwanger. 1991. "Wage Dispersion Between and Within U.S. Manufacturing Plants, 1963-1986." *Brookings Papers on Economic Activity: Microeconomics*, 1991, pp. 115-200.

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.

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) (forthcoming).

Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth R. Troske. 2004. “Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment.” *Journal of Labor Economics* 22 no. 2:397-430.

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

Ericson, Richard, and Ariel Pakes. 1995. “Markov-Perfect Industry Dynamics: A Framework for Empirical Work.” *Review of Economic Studies* 62(1): 53–82.

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. "Innovation, Productivity Growth and Productivity Dispersion," mimeo.

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.

Foster, Lucia., John Haltiwanger, and Chad Syverson, (2016), "The Slow Growth of New Plants: Learning about Demand?" Economica, 83: 91–129.

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

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.

Haltiwanger, John. 2016. "Firm Dynamics and Productivity: TFPQ, TFPR and Demand Side Factors." *Economia*, Forthcoming.

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. 2016. In *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges* edited by John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar. Chicago: University of Chicago Press.

Haltiwanger, John, Jarmin, Ron, and Javier Miranda. 2009. “Business Dynamics Statistics: An Overview.” Kauffman Foundation: Other Research. At http://www.kauffman.org/uploadedFiles/BDS_handout_011209.pdf.

Haltiwanger, John, Ron Jarmin, and Javier Miranda. 2011. “Historically Large Decline in Job Creation from Startups and Existing Firms in the 2008-09 Recession.” Business Dynamics Statistics Briefing No. 5, The Kauffman Foundation.

Haltiwanger, John, Ron S. Jarmin, and Javier Miranda. 2013. “Who Creates Jobs? Small vs. Large vs. Young.” *Review of Economics and Statistics* 95(2): 347-61.

Haltiwanger, John, Robert Kulick, and Chad Syverson, 2016. “Misallocation Measures: Glowing Like the Metal on the Edge of a Knife.” Mimeo.

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.

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

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. 2016. "Understanding the 30-Year Decline in the Startup Rate: A General Equilibrium Approach." Mimeo.

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

Lucas, Robert E. 1978. "On the Size Distribution of Business Firms." *The Bell Journal of Economics* 9 no. 2:508-23.

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.

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.

Reedy, E.J. and Robert E. Litan. 2011. "Starting Smaller; Staying Smaller: America's Slow Leak in Job Creation." Kauffman Foundation Research Series: Firm Foundation and Economic Growth. At http://www.kauffman.org/research-and-policy/starting-smaller-staying-smaller-americas-slow-leak-in-job-creation.aspx.

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

Song, Jae, David J. Price, Fatih Guvenen, Nicholas Bloom and Till von Wachter. 2016. "Firming Up Inequality." Working paper.

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: Sectoral trends in job reallocation

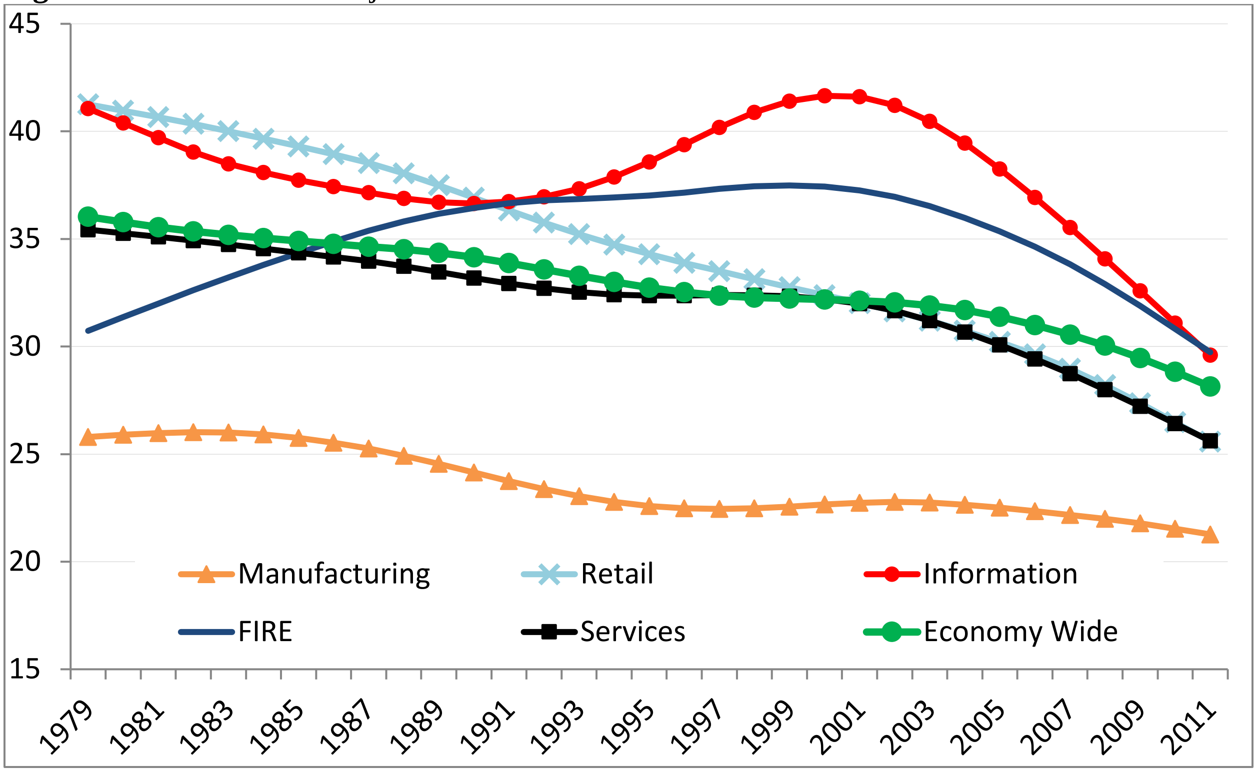

Note: Y axis does not start at zero. Data are HP trends using parameter set to 100. Industries are defined on a consistent NAICS basis. Data include all firms (new entrants, continuers, and exiters). Author calculations from the Longitudinal Business Database.

Figure 2: Employment shares for young (<5) firms by broad sector

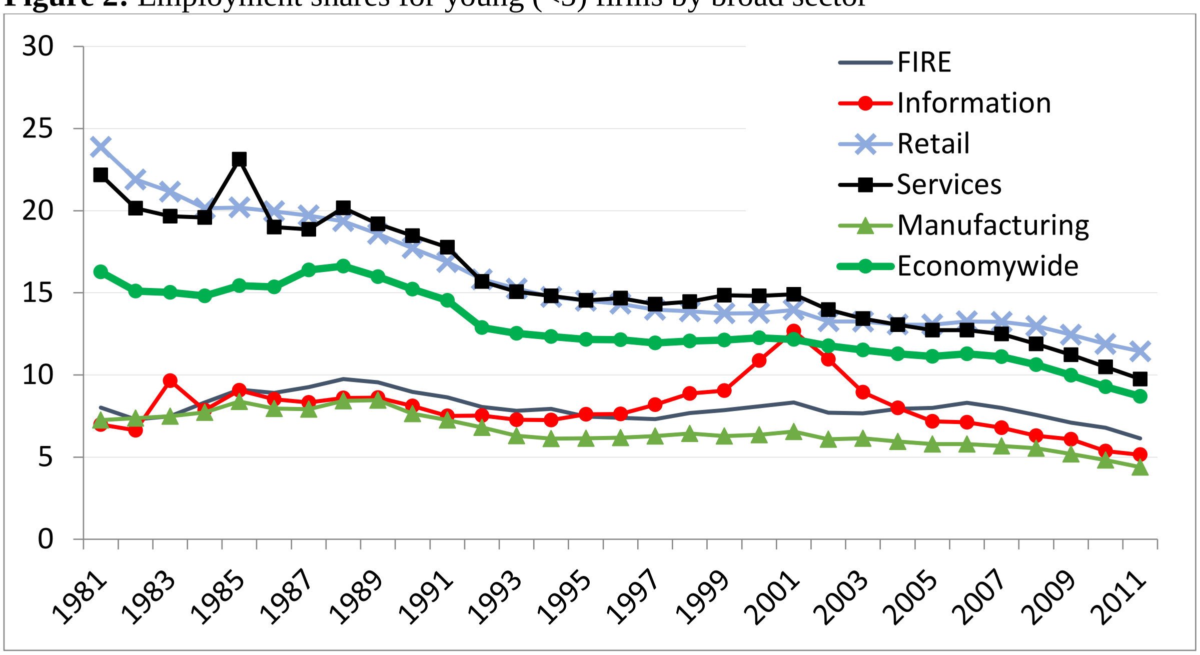

Note: Young firms have age less than 5. Industries are defined on a consistent NAICS basis. Data include all firms (new entrants, exiters, and continuers). Author calculations from the Longitudinal Business Database.

Figure 3: Annual change in reallocation rate: Actual and holding age composition constant

(a) 1987-1989 to 1997-1999

(b) 1997-1999 to 2004-2006

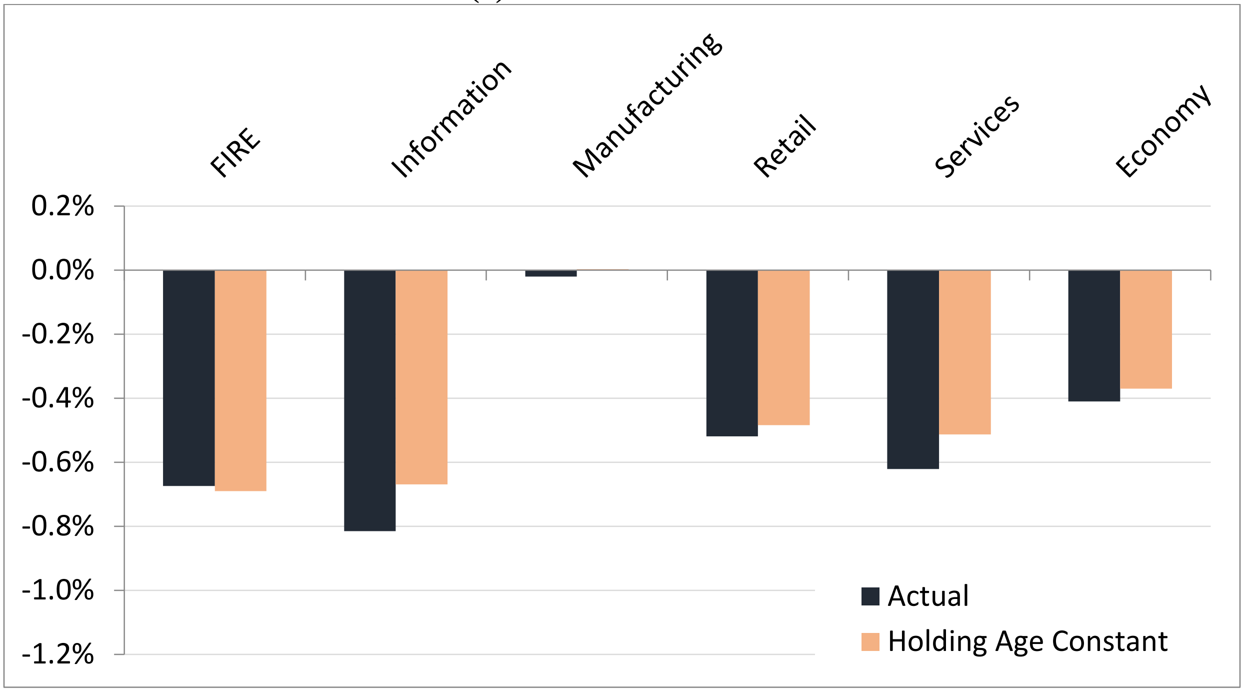

Note: Sectors are defined on a consistent NAICS basis. Author calculations from the Longitudinal Business Database.

Figure 4: Job reallocation in High-Tech, Information, Manufacturing and Tech Manufacturing

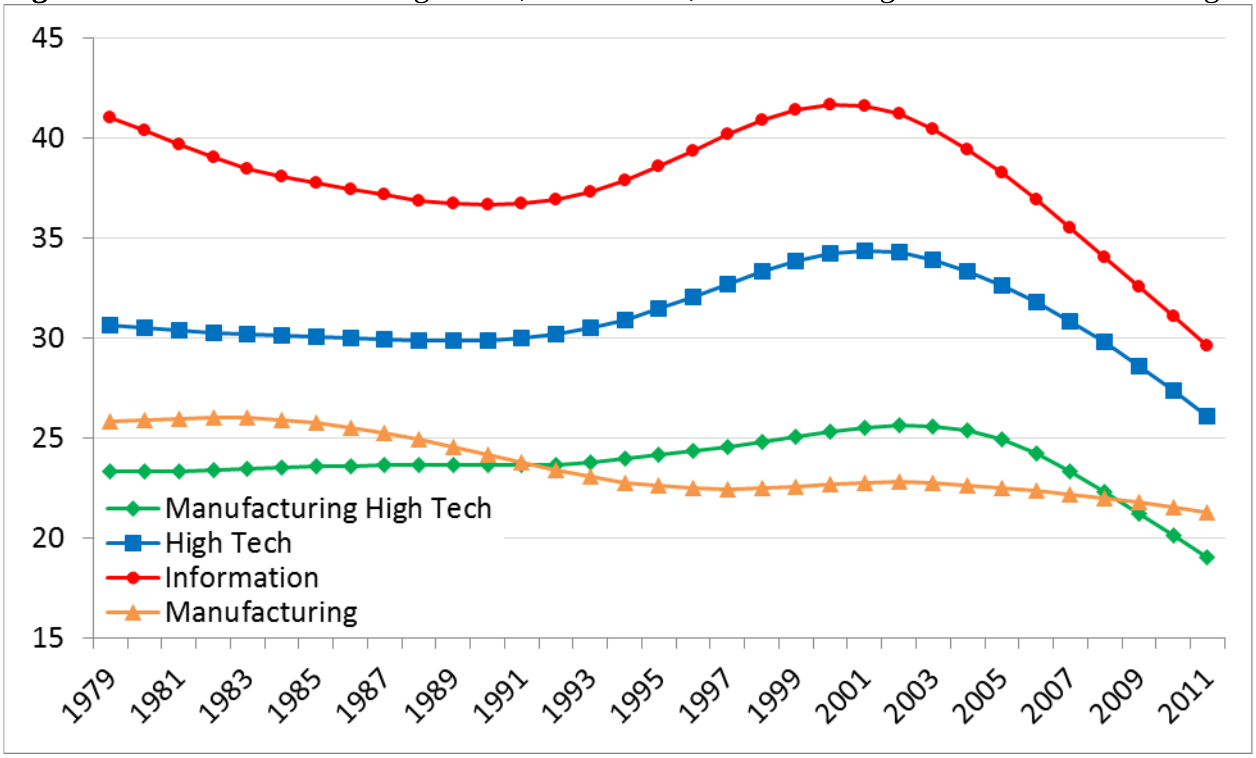

Note: Y axis does not start at zero. High-Tech is defined as in Hecker (2005). Information and Manufacturing sectors are defined on a consistent NAICS basis. Author calculations from the Longitudinal Business Database.

Figure 5: Within-industry TFP dispersion (standard deviation), young vs. mature

(a) High-Tech Manufacturing

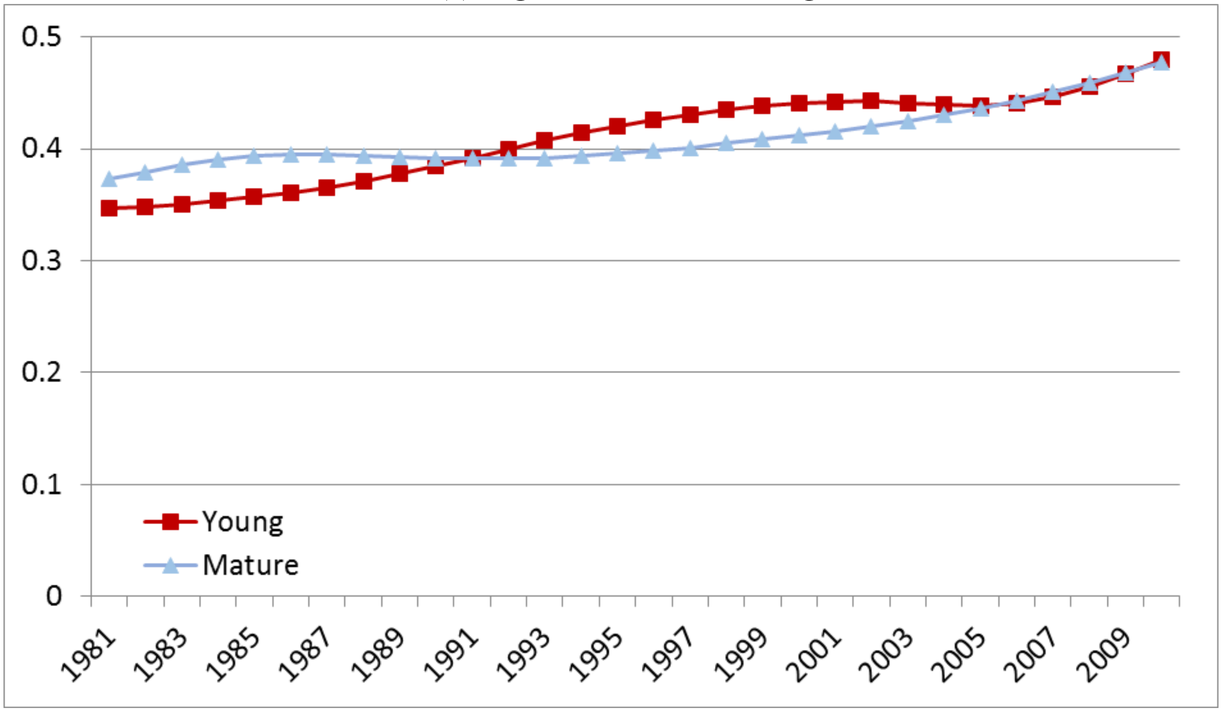

(b) Non Tech Manufacturing

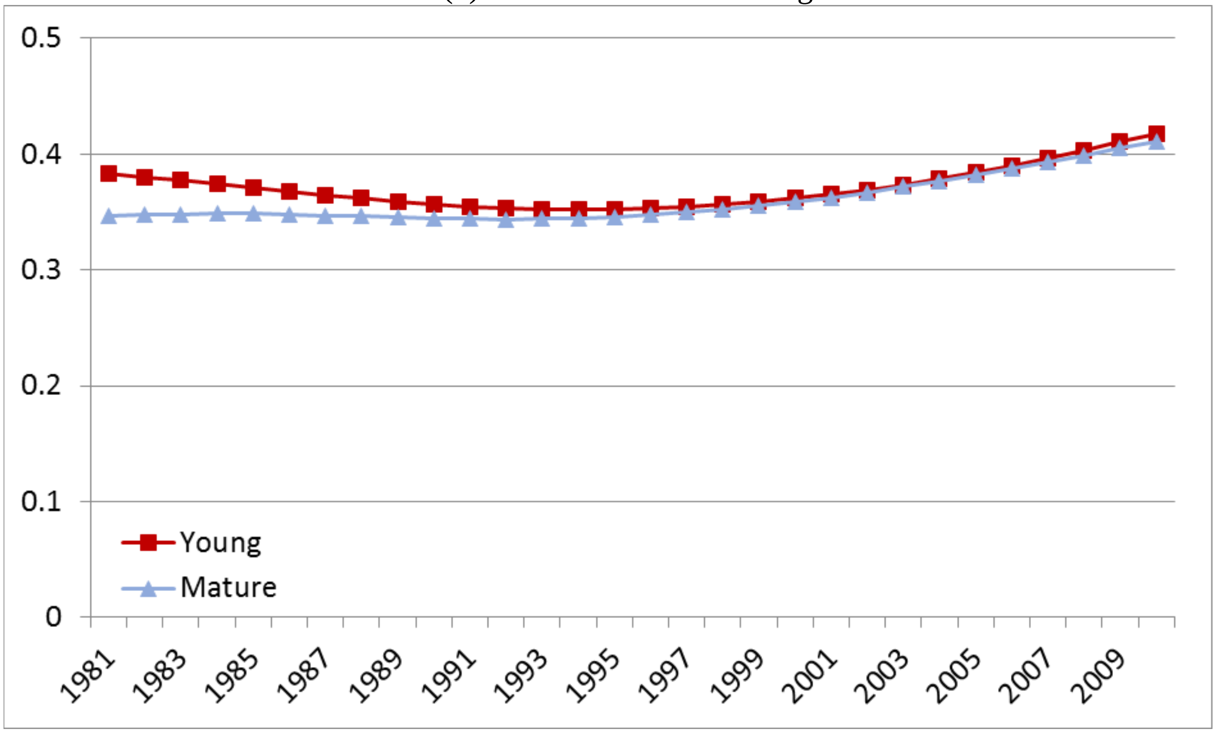

Note: Young firms have age less than 5. Manufacturing plants. The standard deviation is the based on within-detailed industry log TFP. High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers. HP Trends Depicted.

Figure 6: Marginal effect of TFP on plant-level net employment growth: young vs. mature

(a) High-Tech plants

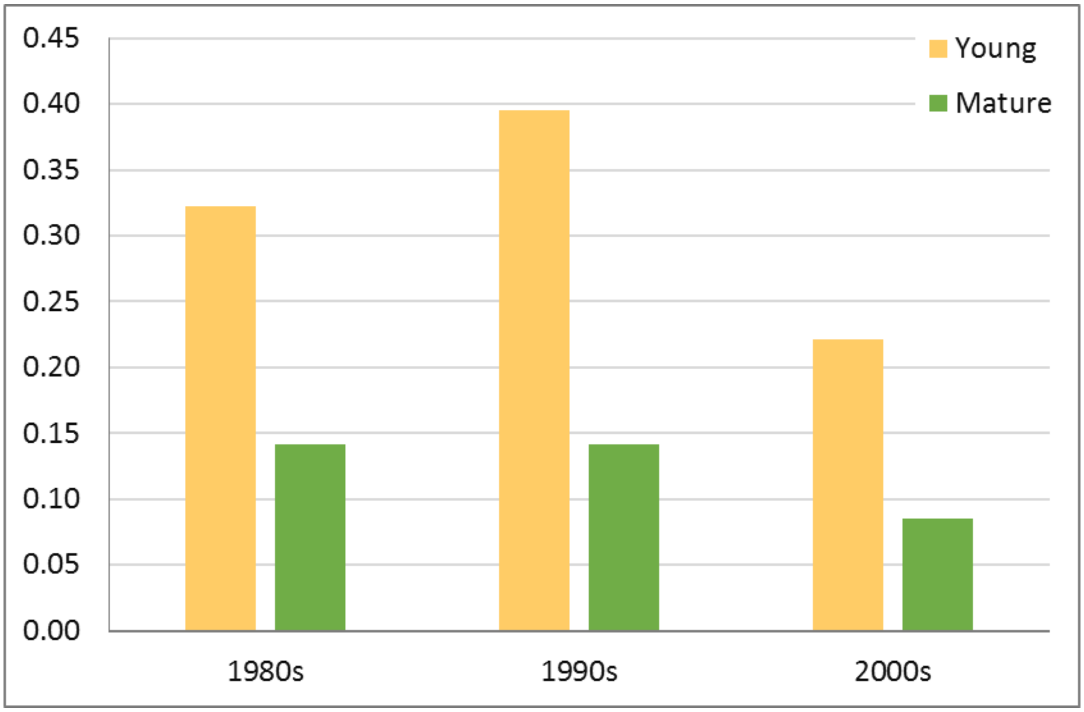

(b) Non Tech plants

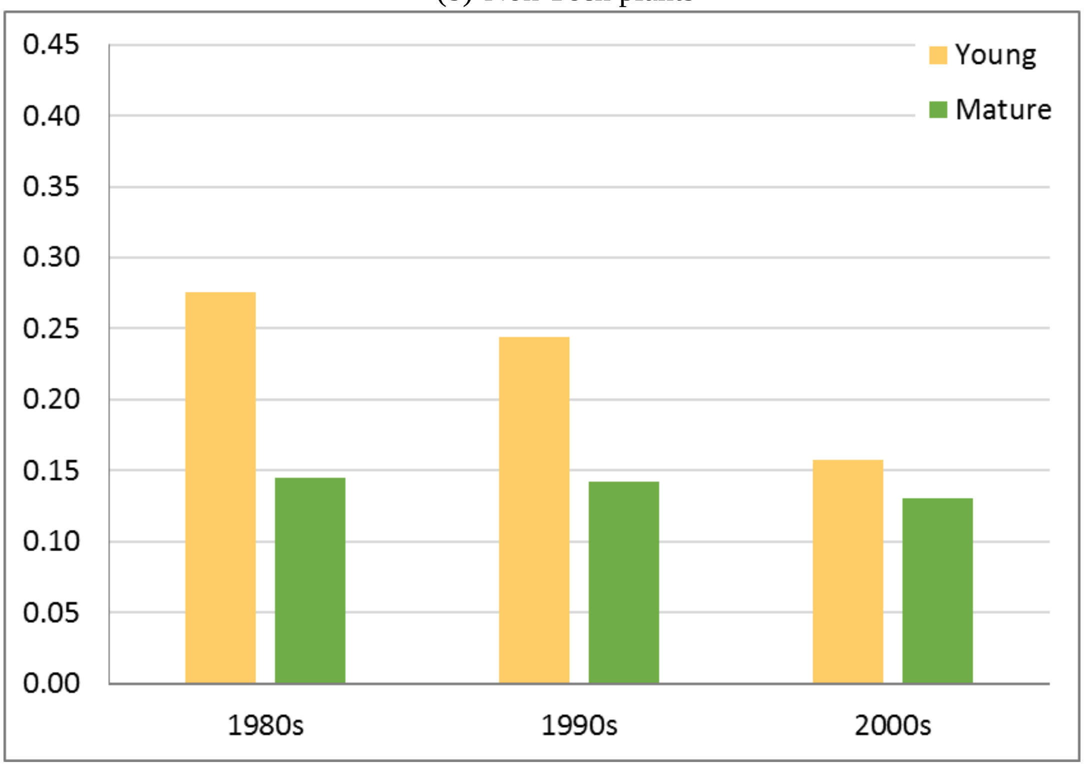

Note: Young firms have age less than 5. High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure 7: Diff-in-diff counterfactual (TFP), Manufacturing

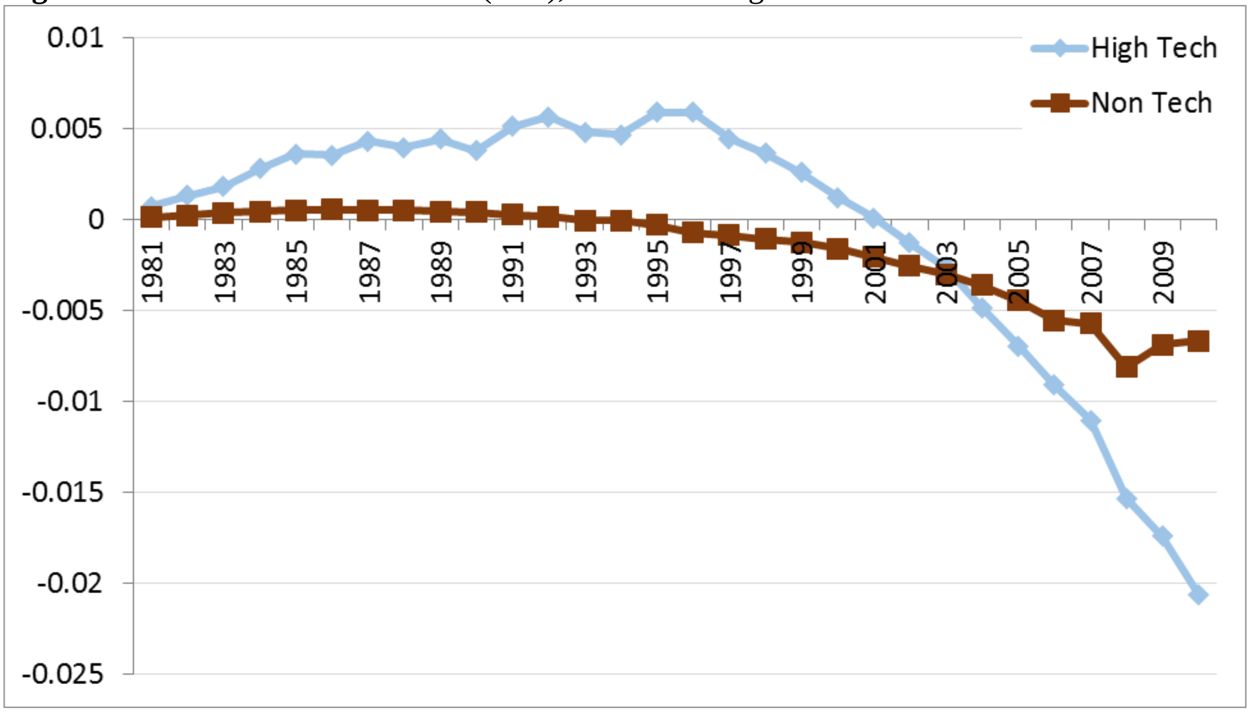

Note: Figure depicts counterfactual change in reallocation contribution to aggregate TFP growth. High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure 8: Within-industry dispersion in labor productivity, young vs. mature

(a) High-Tech firms

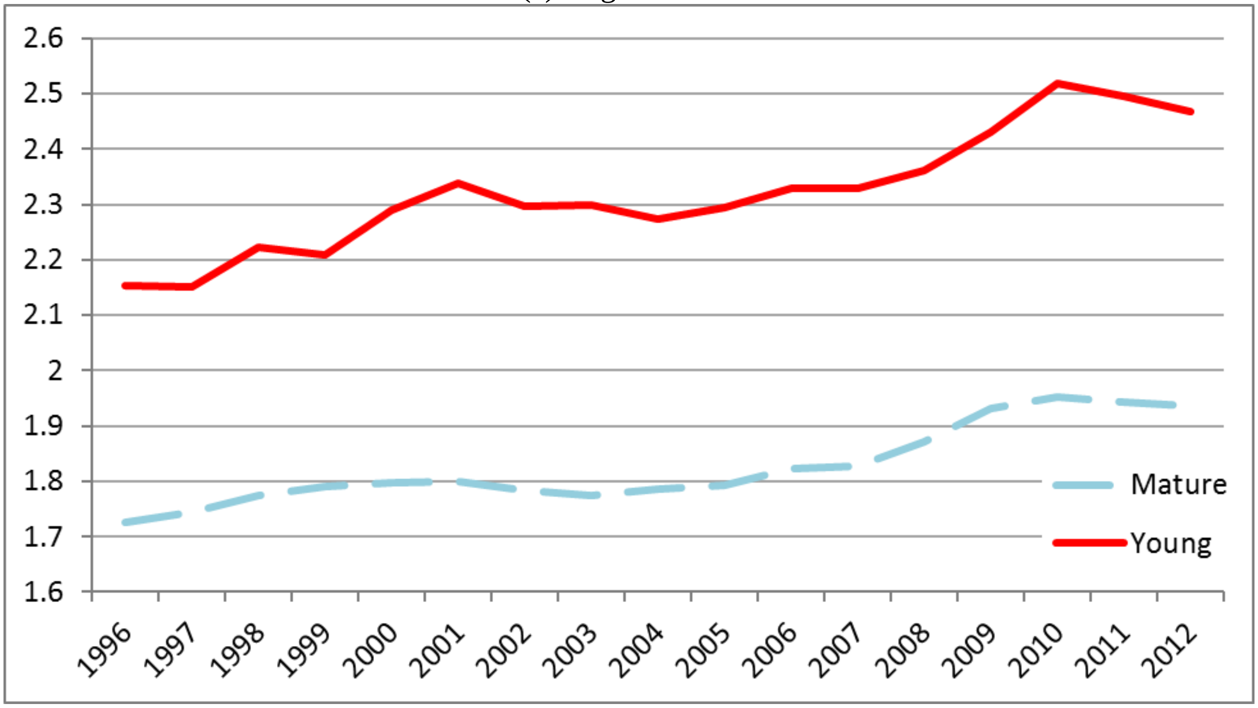

(b) Non Tech firms

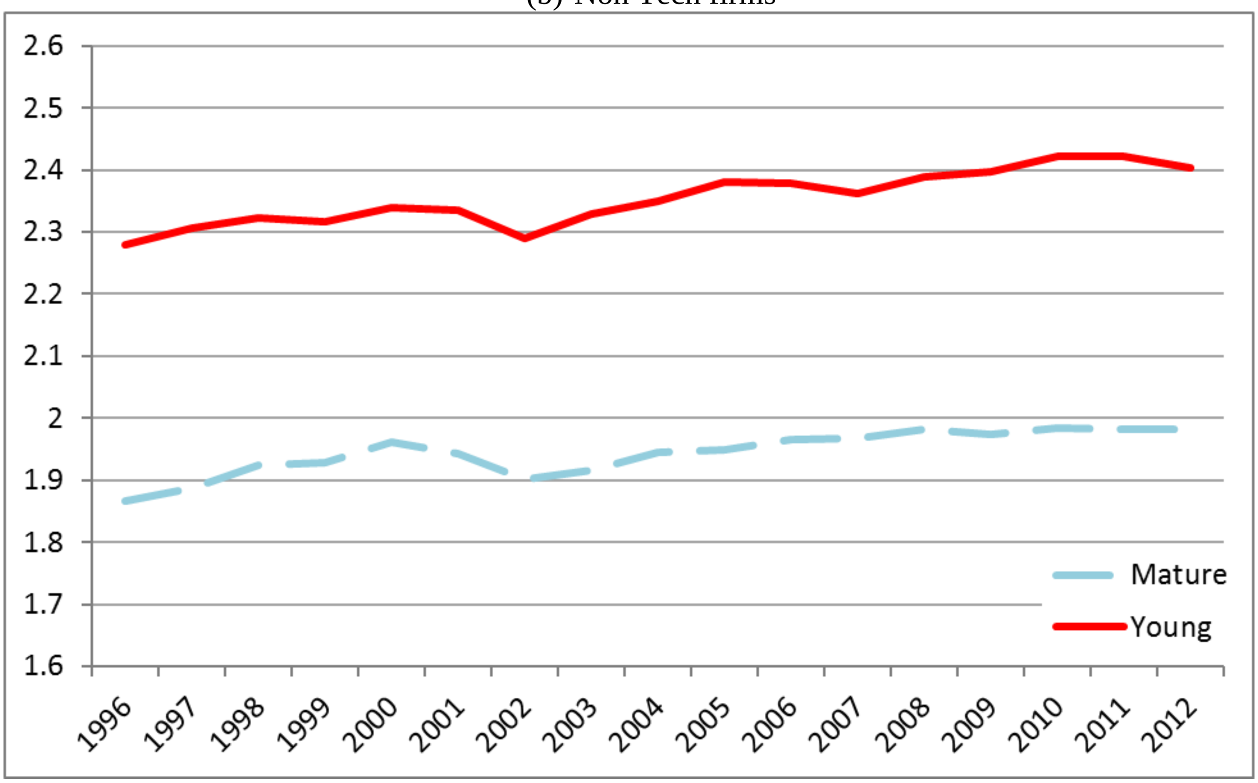

Note: Y axes do not begin at zero. Data reflect interdecile range of log labor productivity deviated from industry by year means. Young firms have age less than five. High-Tech is defined as in Hecker (2005). Author calculations from the RE-LBD.

Figure 9: Labor productivity and growth at the firm level (economywide)

(a) Overall DHS employment growth (including exit)

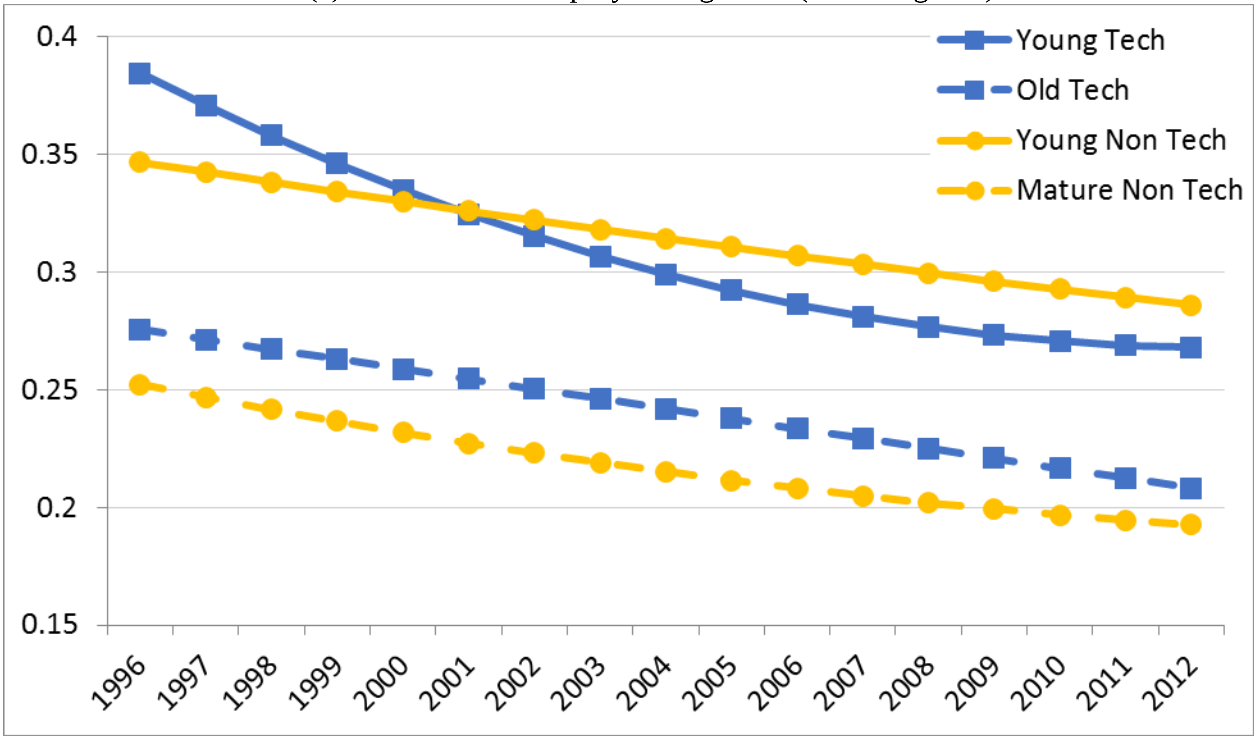

(b) Firm exit

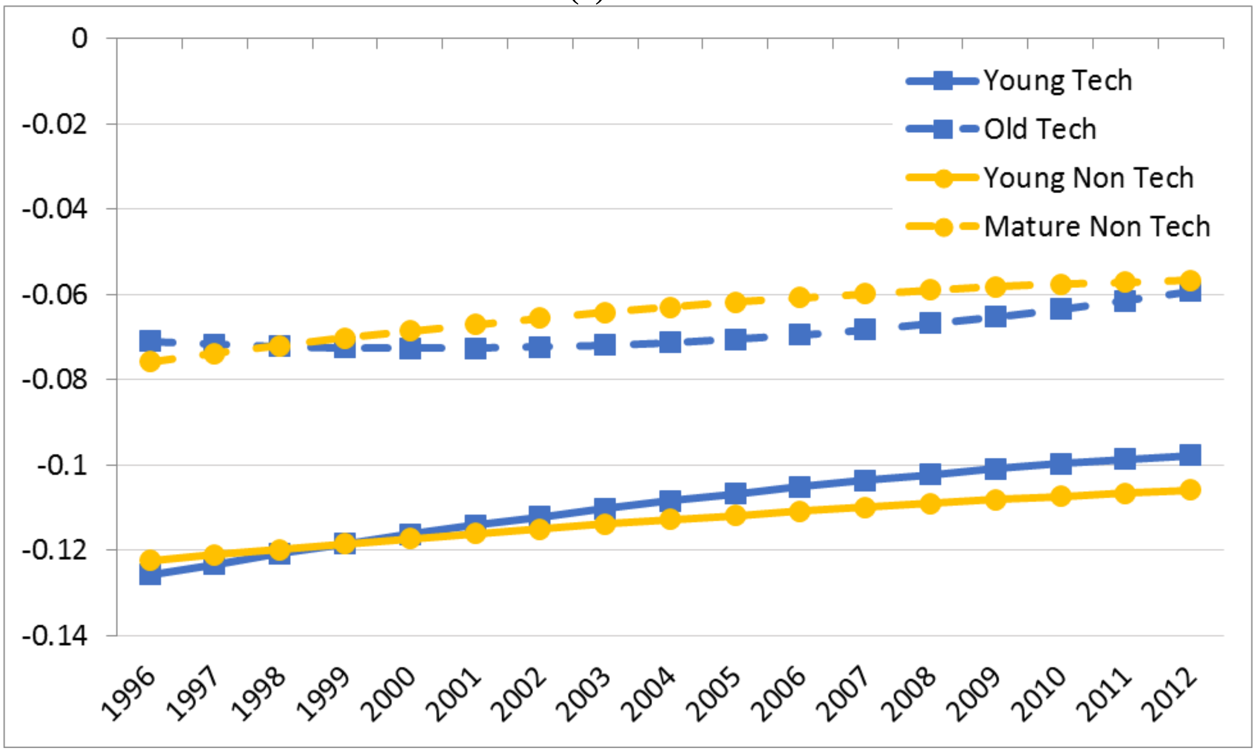

Note: Annual coefficients constructed from Table 2. Young firms have age less than five. High-Tech defined as in Hecker (2005). Author calculations from the RE-LBD. Finance, Insurance and Real Estate (NAICS 52-53) omitted.

Figure 10: Diff-in-diff counterfactual (labor productivity)

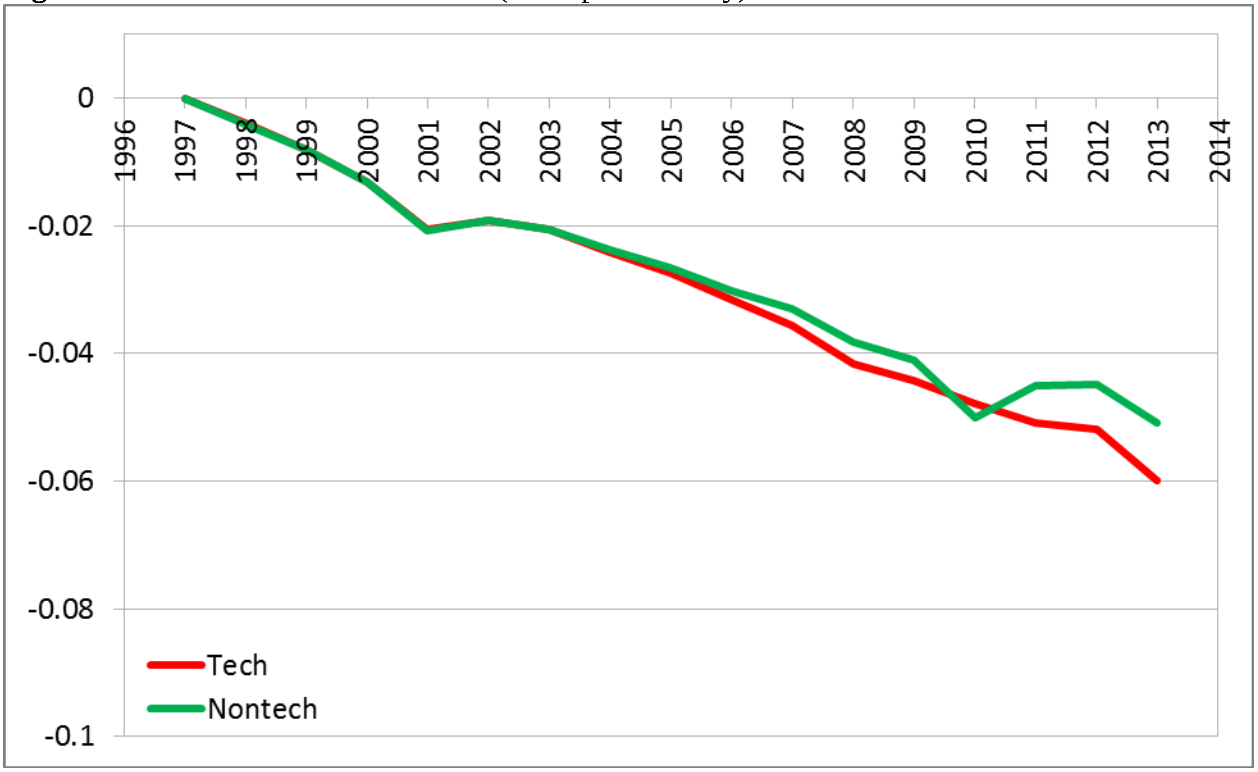

Note: Figure depicts counterfactual change in reallocation contribution to aggregate labor productivity growth. High-Tech is defined as in Hecker (2005). Author calculations from the RE-LBD.

Figure 11: Within-firm productivity growth in the average industry

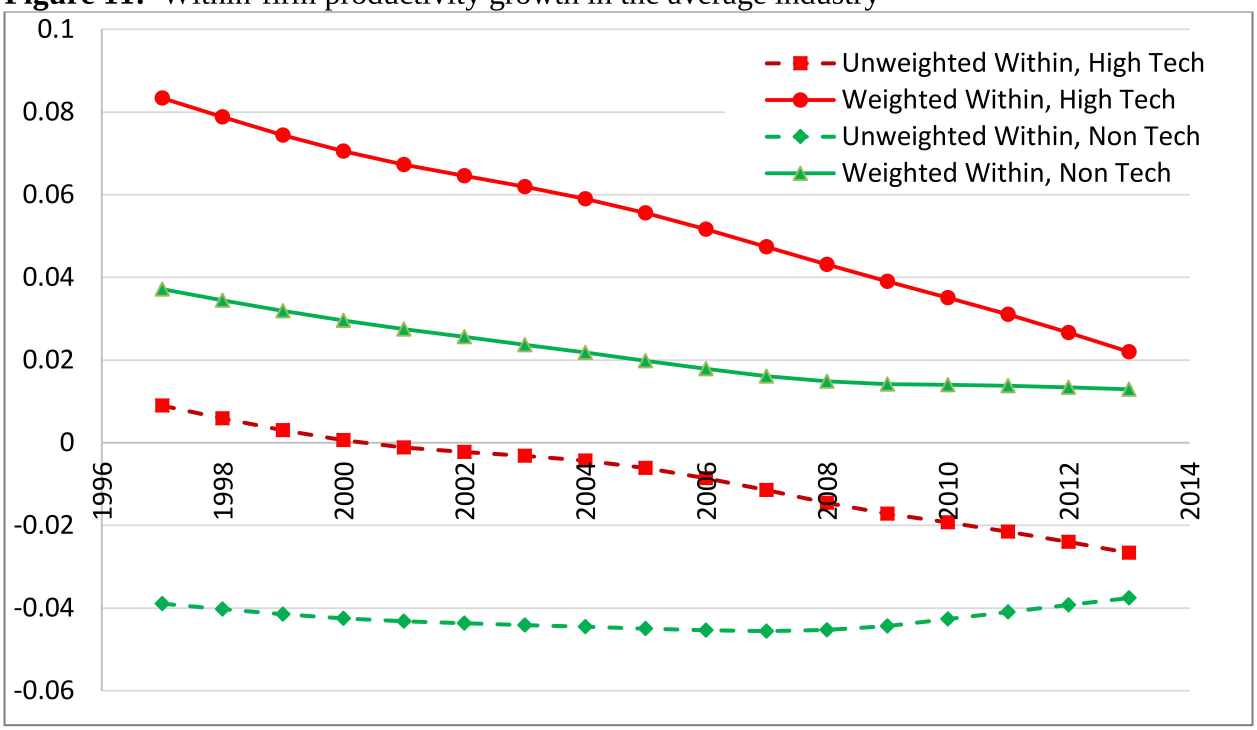

Note: Average within-firm productivity growth, with and without employment weights.

Figure 12: Marginal effect of TFP on plant-level equipment investment in High-Tech

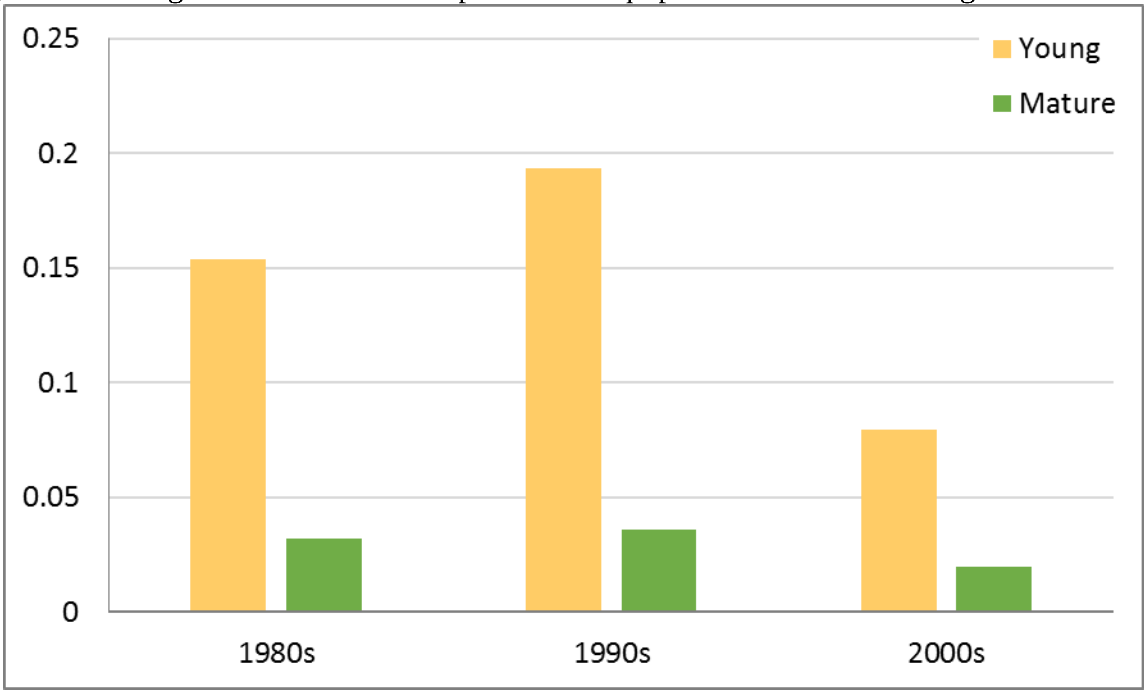

Note: Young firms have age less than 5. High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure 13: The role of globalization in changing responsiveness (High-Tech Manufacturing)

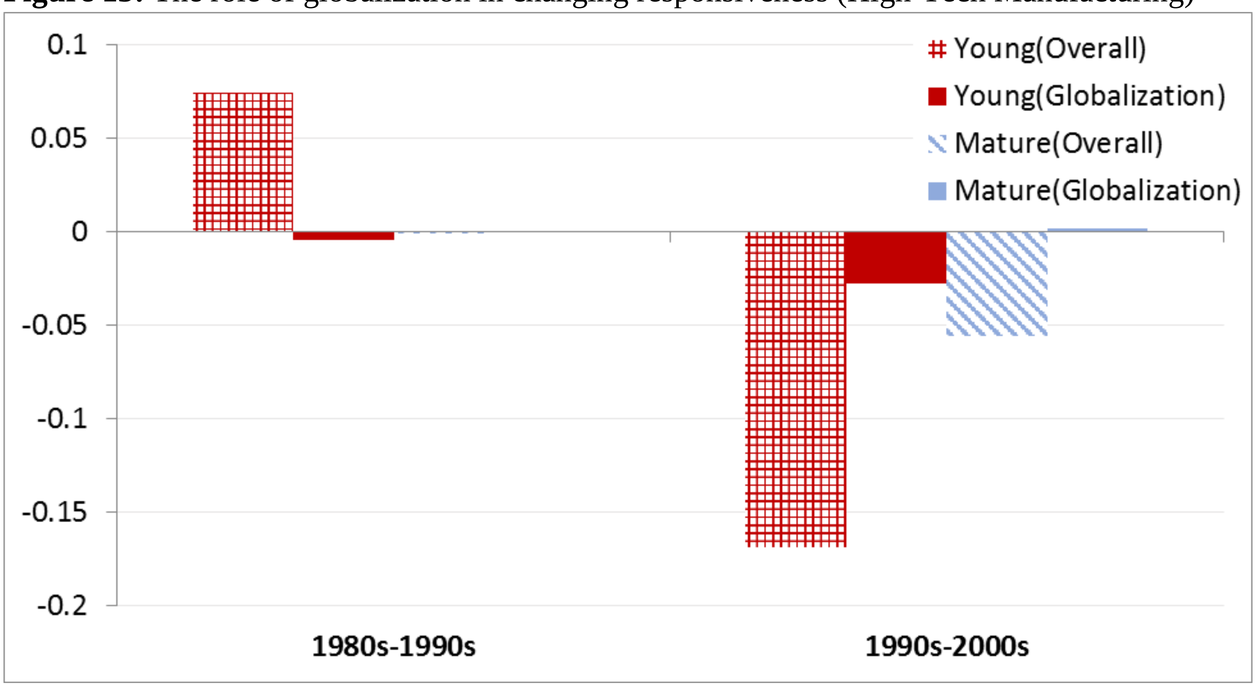

Note: "Overall" bars for young and mature are the change in marginal responsiveness of employment growth to productivity across decades. Globalization reflects implied change in marginal responsiveness accounted for by changes in import penetration ratios from low wage countries.

Table 1: Estimated Impact 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***	0.2767***	-0.0292*	-0.0905***
	(0.0390)	(0.0090)	(0.0162)	(0.0037)
TFP*Young*Trend	0.0317***	0.0014	-0.0160***	-0.0005
	(0.0061)	(0.0014)	(0.0025)	(0.0006)
TFP*Young*Trend^2	-0.0012***	-0.00024***	0.0005***	0.0001***
	(0.0002)	(0.00005)	(0.0001)	(0.00002)
TFP*Mature	0.1228***	0.1439***	-0.0403***	-0.0464***
	(0.0174)	(0.0043)	(0.0072)	(0.0018)
TFP*Mature*Trend	0.0054**	0.0005	-0.0016	-0.0012
	(0.0026)	(0.0007)	(0.0011)	(0.0003)
TFP*Mature*Trend^2	-0.0003***	-0.00004*	0.0001***	0.00005***
	(0.0001)	(0.00002)	(0.00003)	(0.00001)

Notes: Standard Errors in Parentheses. Dependent variable in Overall Growth columns is DHS growth rate. Dependent variable in Exit columns is indicator=1 if exit, 0 otherwise (linear probability Tech Sample is more than 120000 plant-year observations from 1981-2010. Non Tech Sample has more than 2 million observations. 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.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Estimated Relationship Between Firm-level Employment Growth and Exit and Labor Productivity

	Growth including exit			Exit		
	All firms	High-Tech	Nontech	All Firms	High-Tech	Nontech
LP*Young	0.3484***	0.3845***	0.3467***	-0.1225***	-0.1258***	-0.1224***
	0.0004	0.0020	0.0005	0.0002	0.0009	0.0002
LP*Young*Trend	-0.0047***	-0.0141***	-0.0043***	0.0014***	0.0026***	0.0014***
	0.0001	0.0006	0.0001	0.0001	0.0002	0.0001
LP*Young*Trend^2	0.0000***	0.0004***	0.0000***	0.0000***	-0.0001***	0.0000***
	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
LP*Mature	0.2530***	0.2755***	0.2522***	-0.0756***	-0.0710***	-0.0758***
	0.0004	0.0021	0.0004	0.0002	0.0009	0.0002
LP*Mature*Trend	-0.0055***	-0.0042***	-0.0056***	0.0019***	-0.0008***	0.0020***
	0.0001	0.0006	0.0001	0.0000	0.0002	0.0000
LP*Mature*Trend^2	0.0001***	0.0000	0.0001***	0.0000***	0.0001***	-0.0001***
	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
N	55383000	55383000	55383000	55383000	55383000	55383000
R^2	0.1090	0.1263	0.1083	0.0937	0.1053	0.0931

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 log difference from 6-digit NAICS industry mean. High-Tech is defined as in Hecker (2005). Observations rounded to nearest thousand.
*** p<0.01; ** p<0.05; * p<0.10

Table 3: Estimated Impact of Productivity on Plant-Level Equipment Investment Rate

TFP*Young	0.0826***
	(0.0236)
TFP*Young*Trend	0.0189***
	(0.0037)
TFP*Young*Trend^2	-0.0008***
	(0.0001)
TFP*Mature	0.0232**
	(0.0105)
TFP*Mature*Trend	0.0024
	(0.0016)
TFP*Mature*Trend^2	-0.0001*
	(0.00005)

Notes: Standard Errors in Parentheses. Tech Sample is more than 120000 plant-year observations from 1981-2010. 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 4: Estimated Impact of Lagged Productivity on Plant-Level Employment Growth with Import Penetration Ratio Effects (High-Tech)

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. Tech Sample is more than 120000 plant-year observations from 1981-2010. 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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Appendix A. Figures and tables to supplement the main text

Figure A1: Within-industry TFP dispersion (std deviation) in Manufacturing, High-Tech Manufacturing and Non Tech Manufacturing (HP Trends)

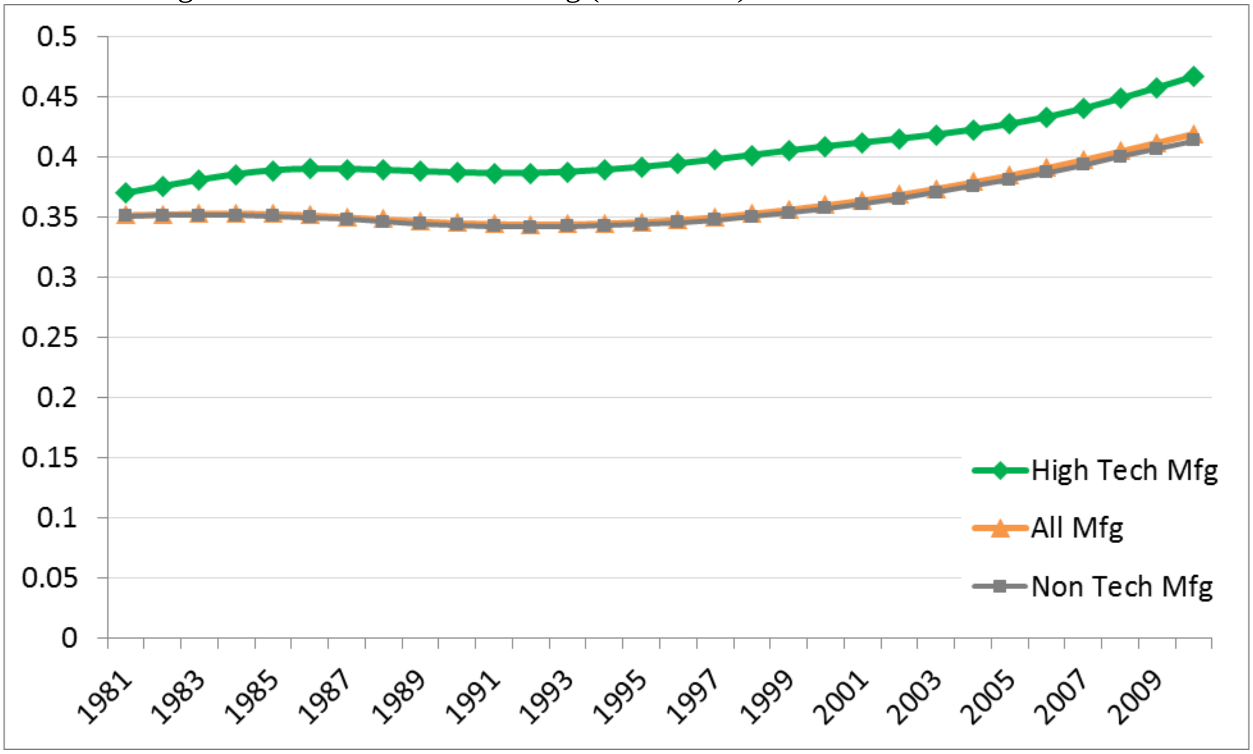

Note: The standard deviation is the based on within-detailed industry log TFP. High-Tech is defined as in Hecker (2005). Manufacturing is defined on a consistent NAICS basis. Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers. Hodrick Prescott Trends depicted.

Figure A2: Marginal effect of TFP on plant exit: young vs. mature

(a) High-Tech plants

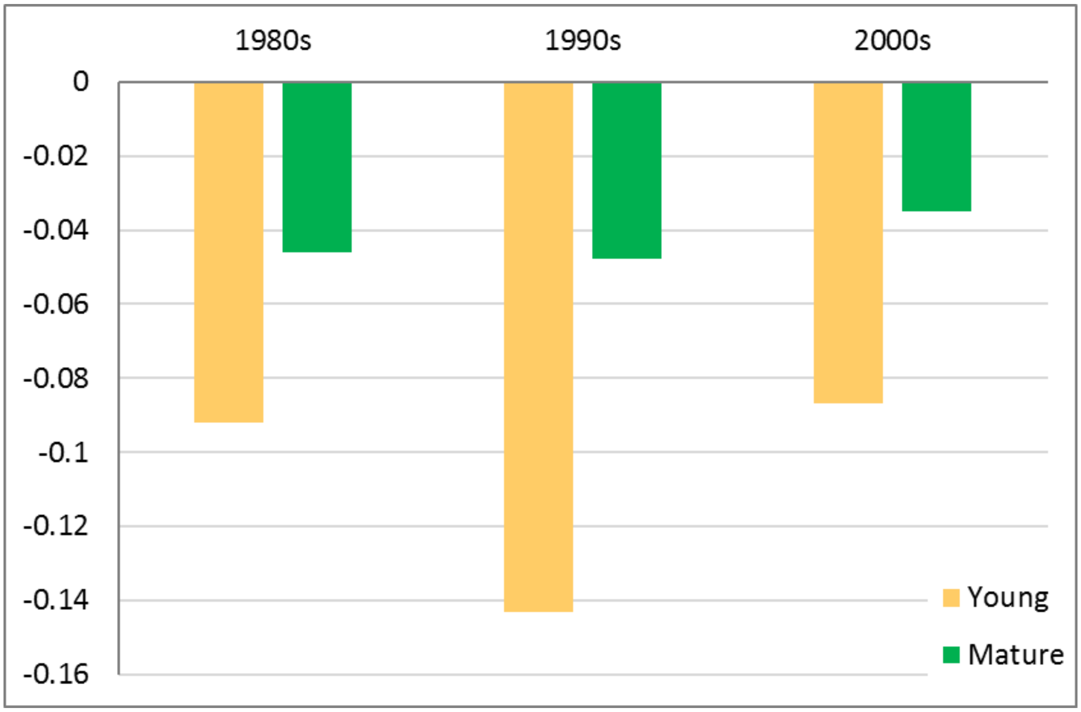

(b) Non Tech plants

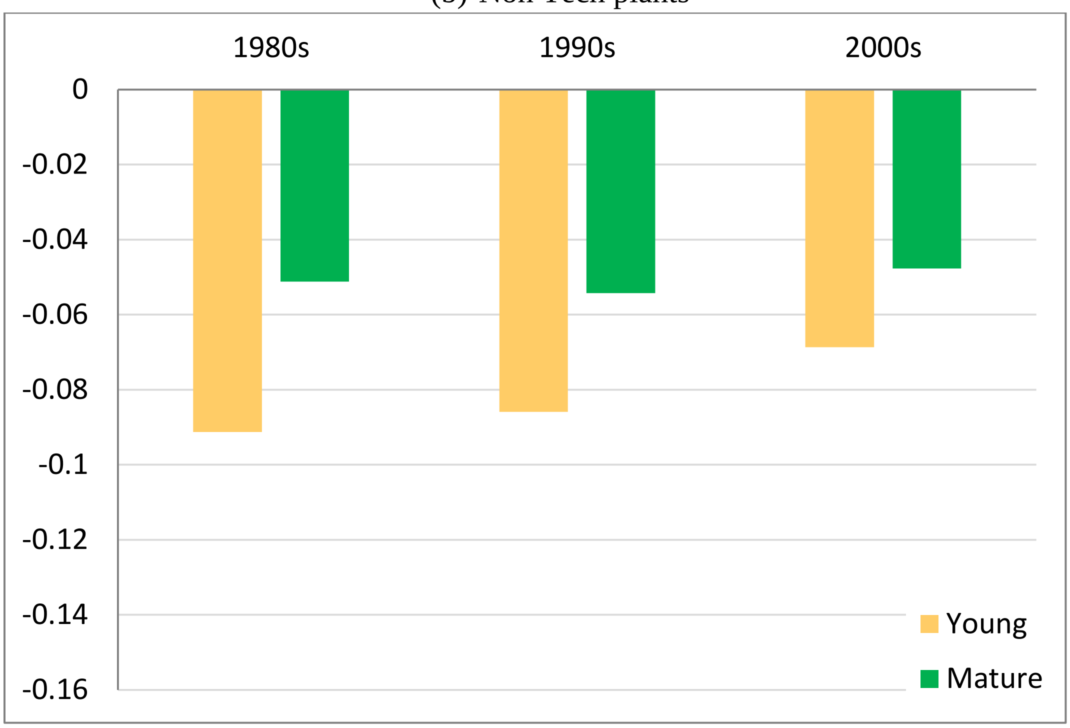

Note: Young firms have age less than 5. High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure A3: Diff-in-diff counterfactual (Revenue Labor Productivity), Manufacturing

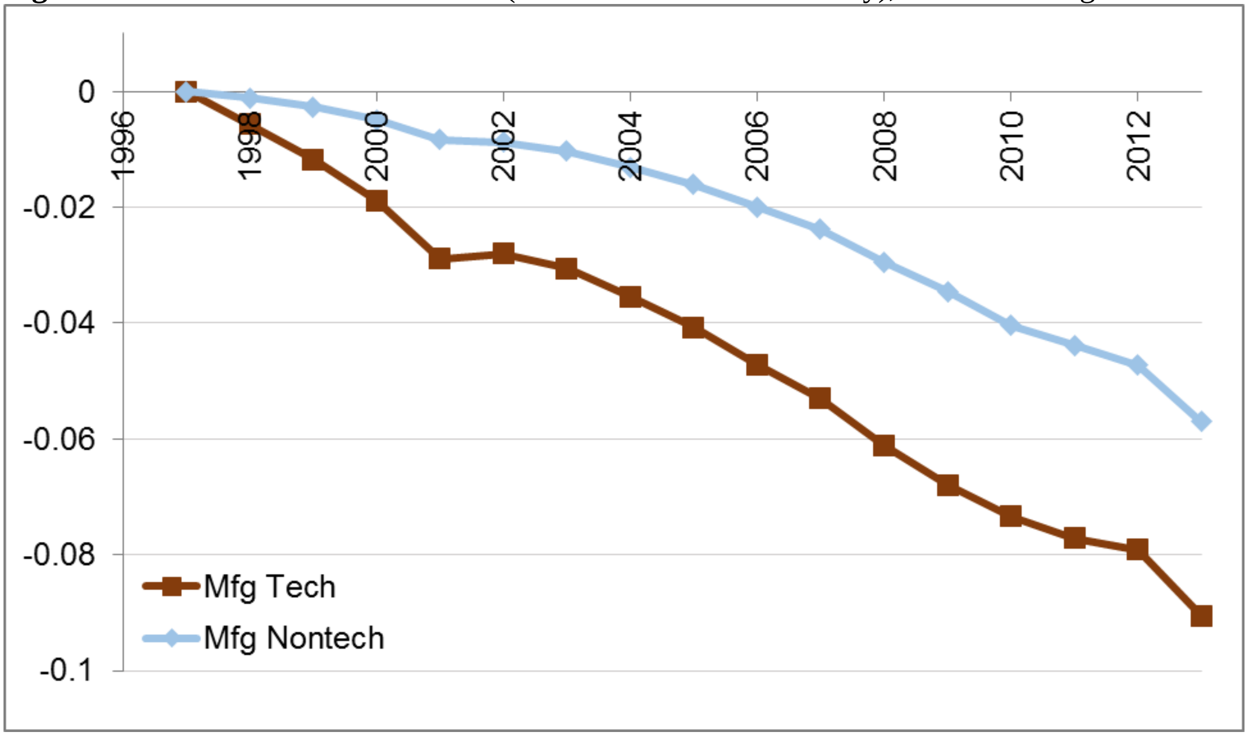

Note: Figure depicts counterfactual change in reallocation contribution to aggregate labor productivity growth. High-Tech is defined as in Hecker (2005). Author calculations from the RE-LBD.

Figure A4: Average industry-level productivity growth, BLS and aggregated microdata, Tech and Non Tech (HP filtered)

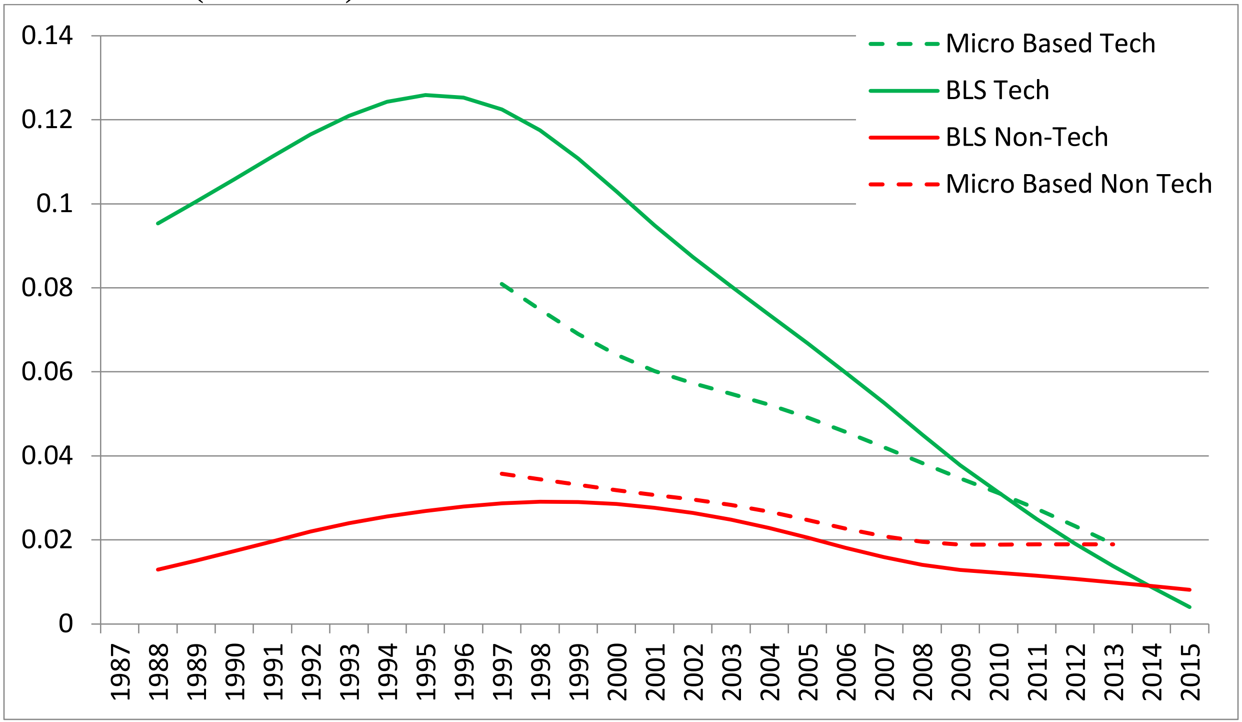

Source: BLS and author calculations from RE-LBD.

Figure A5: Import penetration ratios from low-wage countries

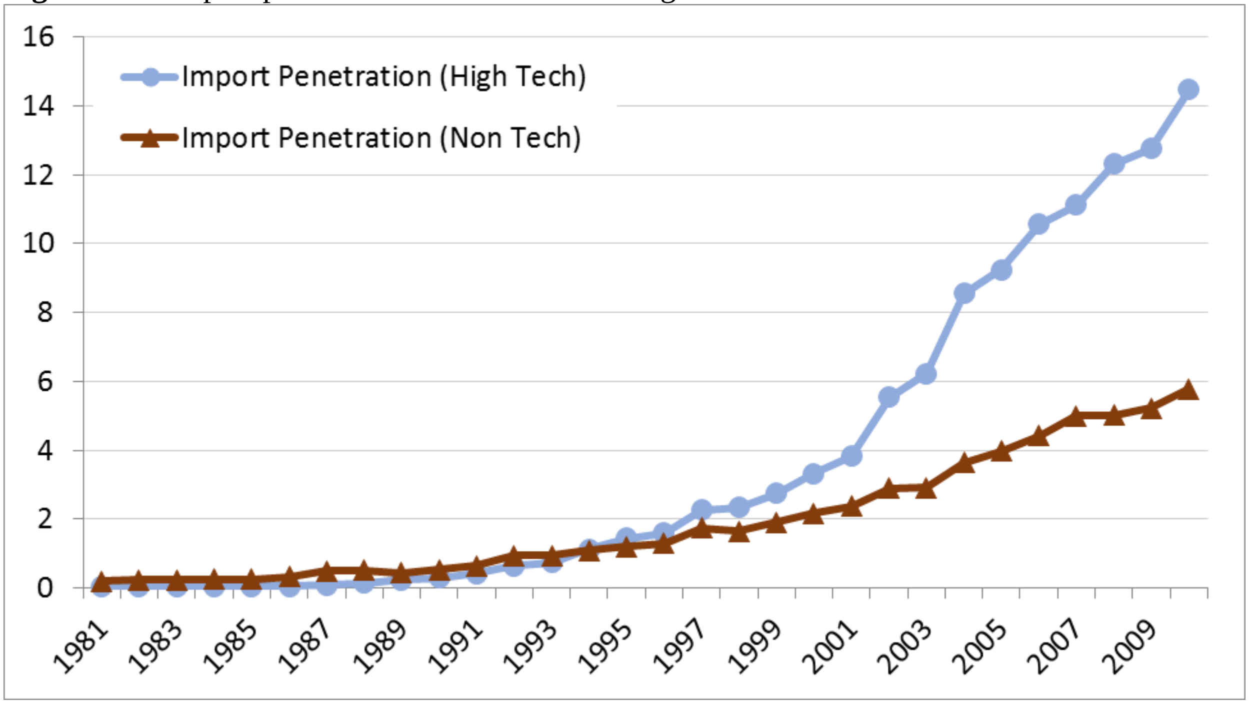

Source: Extended versions of Import Penetration Ratios from Bernard, Jensen and Schott (2006) and Schott (2008). Reported statistics are averages across 6-digit NAICS industries for High-Tech and Non Tech industries.

Table A.1: High-Technology Industries

NAICS Code	Industry
Information and Communications Technology (ICT) High-Tech	
3341	Computer and peripheral equipment manufacturing
3342	Communications equipment manufacturing
3344	Semiconductor and other electronic component manufacturing
3345	Navigational, measuring, electromedical, and control instruments manufacturing
5112	Software publishers
5161	Internet publishing and broadcasting
5179	Other telecommunications
5181	Internet service providers and Web search portals
5182	Data processing, hosting, and related services
5415	Computer systems design and related services
Miscellaneous High-Tech	
3254	Pharmaceutical and medicine manufacturing
3364	Aerospace product and parts manufacturing
5413	Architectural, engineering, and related services
5417	Scientific research-and-development services

Source: Bureau of Labor Statistics, Hecker (2005)

Appendix B. Persistence and Innovation in TFP in Manufacturing

Our data infrastructure 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.[44] With these caveats in mind, we estimate an AR(1) model of TFP applied to continuing plants. Figure B1 reports these AR(1) coefficients, averaged by decade, separately for plants in High-Tech Manufacturing and plants in Non Tech Manufacturing. The estimates for both categories are in the 0.6 to 0.7 range and are reasonably stable over time. For High-Tech Manufacturing, there is a slight decrease in the persistence of plant-level TFP from the 1980s to the 1990s, but it rebounds in the 2000s.[45]

For the set of years where we can estimate the AR(1) process, 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; these are reported on Figure B2. We find patterns that mimic the pattern of dispersion in TFP. That is, the dispersion of innovations for High-Tech rises mildly in the 1990s and then rises more substantially in the post-2000 period, but the magnitude is reasonably stable over time.

[44] Even for other years, 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.
[45] We conduct this persistence exercise using the Wooldridge (2009) revenue productivity residual (RPR) estimation method. The RPR persistence pattern is generally consistent with the TFPR result.

Figure B1: Persistence of plant-level TFP: High-Tech vs. Non Tech Manufacturing

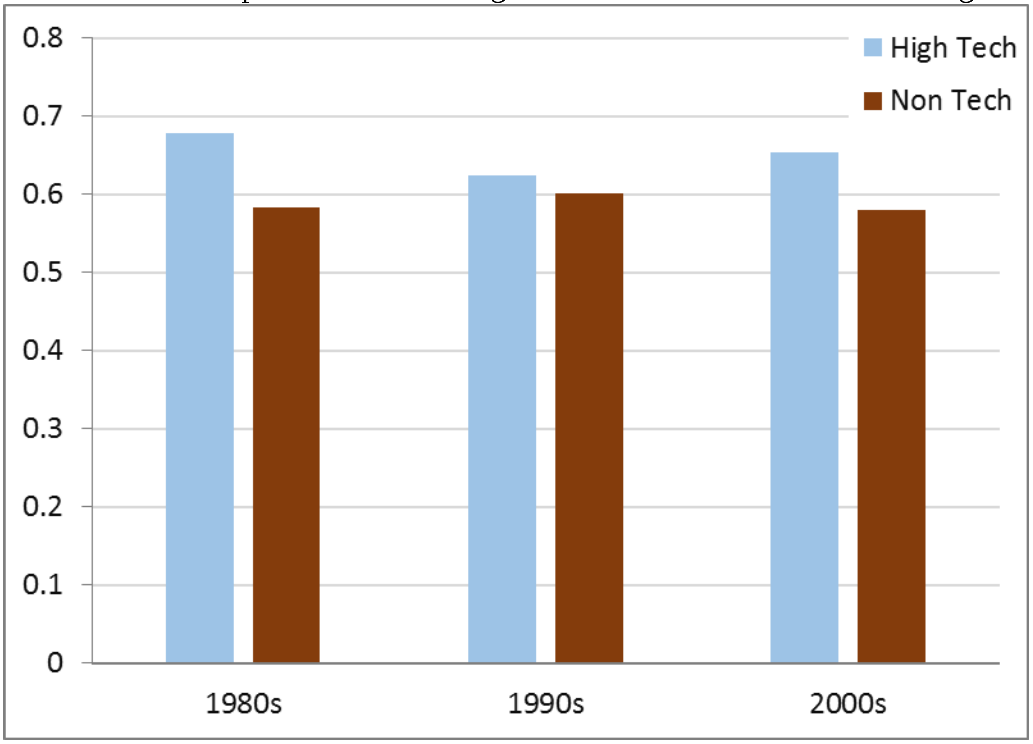

Note: High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure B2: Standard deviation of innovations to plant-level TFP: High-Tech vs Non Tech

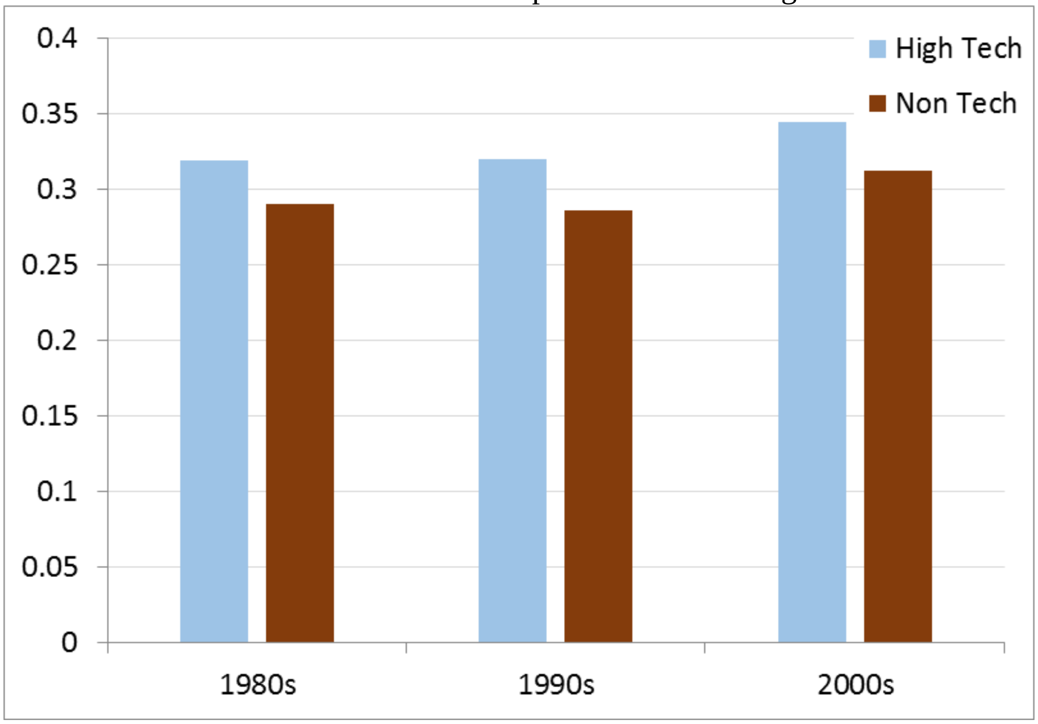

Note: High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Appendix C. Alternative TFP calculation

Here we consider an alternative revenue productivity residual measure that explicitly incorporates potentially endogenous plant-level prices. Let $P_{et} = D_{et}Q_{et}^{\varphi-1}$ where D_{et} is an idiosyncratic demand shock and $\varphi - 1$ is the inverse demand elasticity. Then plant-level revenue is given by (lower case variables are in logs):

$$p_{et} + q_{et} = \beta_k k_{et} + \beta_l l_{et} + \beta_m m_{et} + \beta_e e_{et} + \varphi a_{et} + d_{et} \tag{C1}$$

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. The revenue elasticities can be estimated consistently using the Wooldridge (2009) one-step GMM method. The latter builds on the proxy methodology of Olley and Pakes (1996) and Levinsohn and Petrin (2001). With estimates of the revenue function elasticities, the revenue productivity residual can be recovered which is given by:

$$RPR_{et} = \varphi a_{et} + d_{et} \tag{C2}$$

The revenue productivity residual is a function of idiosyncratic demand and TFPQ shocks. We note that the RPR_{et} is distinct conceptually from $TFPR_{et}$. The former is based on revenue per unit input using revenue elasticities that incorporate demand parameters while the latter is revenue per unit input using factor elasticities from cost shares. This implies that even with endogenous prices in the face of plants facing downward sloping demand functions that RPR will exhibit dispersion regardless of any frictions or distortions.

We implement this methodology following the guidance of Foster et al. (2016) in estimating the revenue elasticities to vary at the 3-digit NAICS level. They find that the proxy method estimates using the Woolridge method are sensitive to outliers and pooling across a large number of observations mitigates this sensitivity. 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 the exercises we have conducted with $TFPR_{et}$ with the alternative measure of RPR_{et}. Figure C1 shows the evolution of within industry dispersion in RPR_{et} for all manufacturing industries, manufacturing High-Tech and non-manufacturing High-Tech. As with Figure A3 of Appendix A (which depicts analogous patterns for TFP), we observe gradually rising RPR dispersion throughout the time period, with higher dispersion in High-Tech Manufacturing than in the rest of the sector. Figure C2 reports AR(1) coefficients for plant-level RPR (see Appendix B for a discussion of this measure and its limitations in our dataset). Again, consistent with the TFP persistence results (depicted on Figure B1 of Appendix B), we find little change in productivity persistence over the time period. Taken together, the dispersion and persistence results shown here confirm our conclusions from TFP data: changes in the distributional characteristics of plant-level productivity cannot explain aggregate patterns of job reallocation.

We estimate equation (2) using RPR in place of TFP. Figure C3 reports annual regression coefficients relating productivity and growth, averaged by decade. The results are generally consistent with those reported for TFP on Figure 5 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, 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 (4). Among High-Tech plants, declining responsiveness produces a counterfactual that is broadly similar—both qualitatively and quantitatively—with the TFP-based results from Figure 6, with a productivity "drag" that is only slightly smaller under RPR than under TFP. Among Non Tech plants, the counterfactual produces somewhat different results from those reported in Figure 6, 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.

Figure C1: Within-Industry Dispersion in the Revenue Productivity Residual (Std Deviation) in Total Manufacturing, High-Tech Manufacturing and Non Tech Manufacturing (HP Trends)

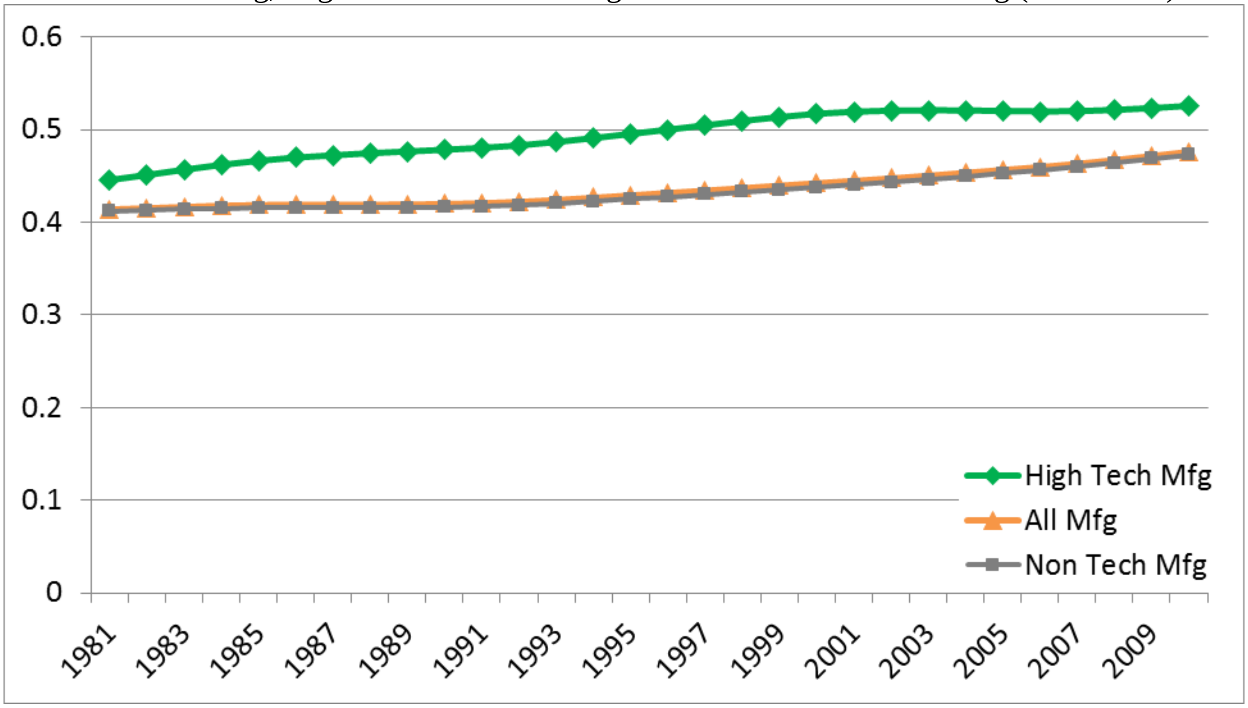

Note: The standard deviation is the based on within-detailed industry log revenue productivity residual. High-Tech is defined as in Hecker (2005). Manufacturing is defined on a consistent NAICS basis. Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers. Hodrick Prescott Trends depicted.

Figure C2: Persistence of Revenue Productivity Residual for Plants: High-Tech vs. Non Tech

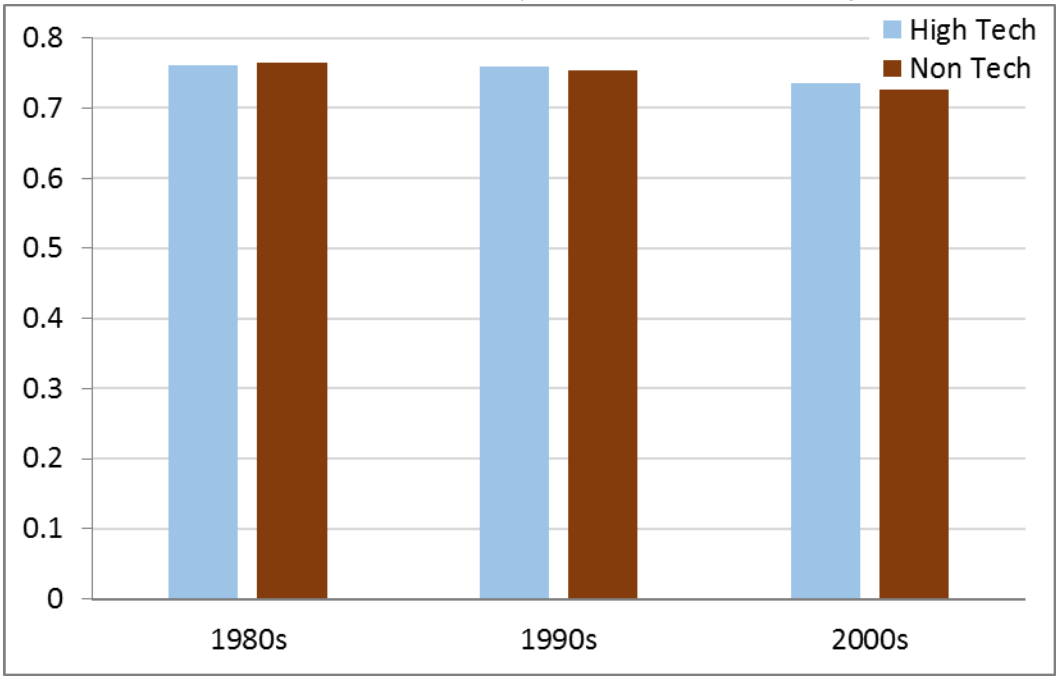

Note: High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Figure C3: Marginal effect of revenue productivity residual on plant-level net employment growth

(a) High-Tech plants

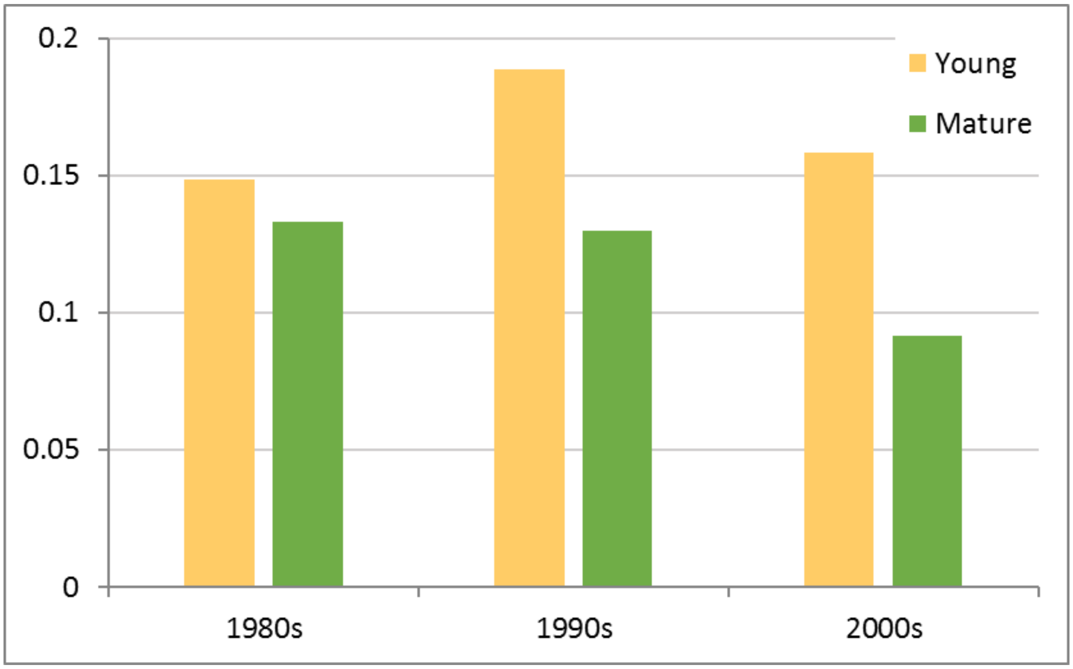

(b) Non Tech plants

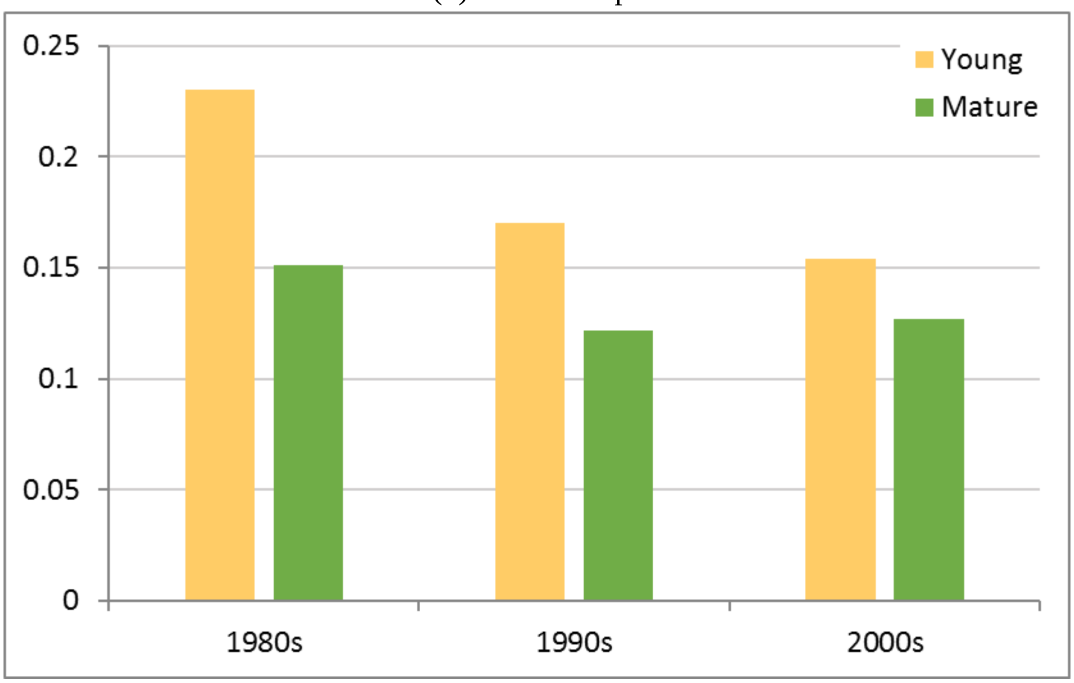

Note: Young firms have age less than 5. High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

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

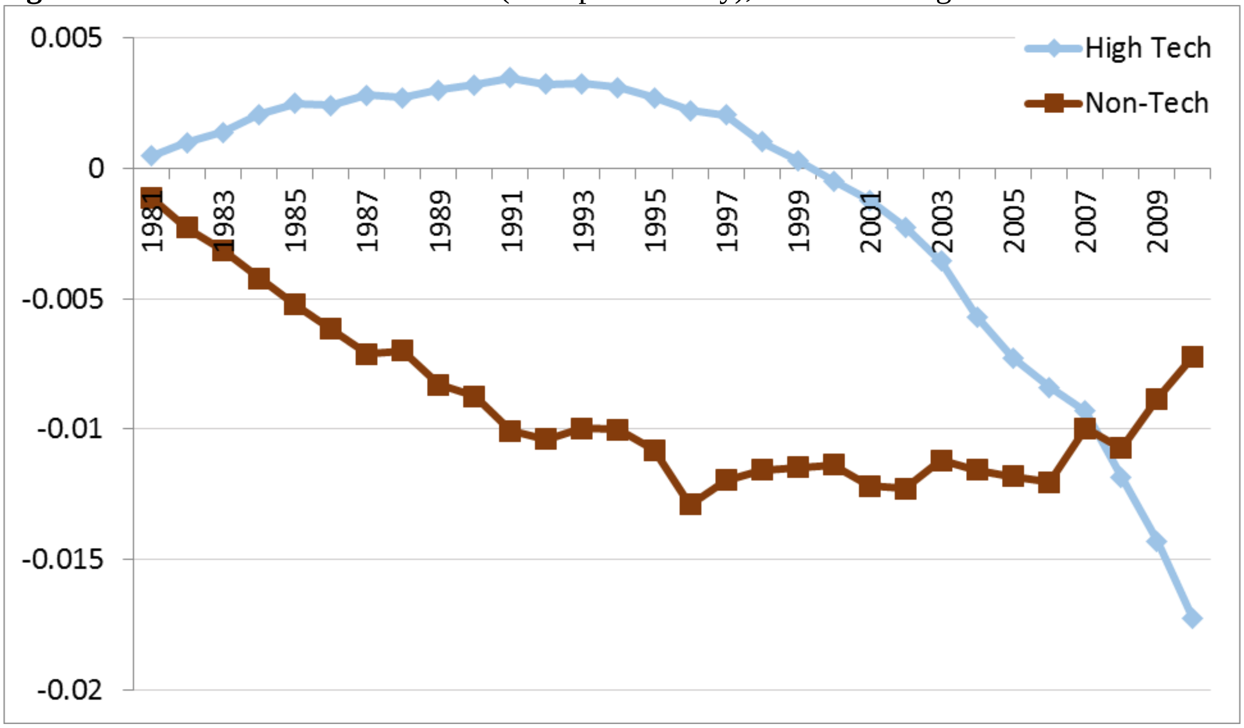

Note: Figure depicts counterfactual change in reallocation contribution to aggregate TFP. High-Tech is defined as in Hecker (2005). Author calculations from the Longitudinal Business Database, the Annual Survey of Manufacturers, and the Census of Manufacturers.

Appendix D. Illustrative Model of Adjustment Costs

Consider the following model of firm-level adjustment costs. We use the term "firm" here for expositional convenience although much (but not all) of our empirical analysis is at the plant-level. A firm maximizes the present discounted value of profits. The firm's value function and its components are specified as follows:

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

with:

$$C(H_{et}) = \begin{cases} \frac{\gamma}{2}\left(\frac{H_{et}}{E_{et-1}}\right)^2 & + F_{+}\max(H_{et}, 0) + F_{-}\max(-H_{et}, 0)\ if\ H_{et} \neq 0 \\ 0,\ otherwise \end{cases}$$

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

where $\varphi < 1$ due to product differentiation so that $A_{et}E_{et}^{\varphi}$ is the revenue function, E_{et} is employment used in production during 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 and is chosen prior to production), w_t is the wage rate, and $a_{et} = \log(A_{et})$ is the revenue shock reflecting potentially TFPQ and demand shocks (although as we discuss below we neglect the latter in the current discussion for expositional convenience). We focus on the interpretation of curvature in the revenue function arising from product differentiation rather than decreasing returns to help draw out the issues relating revenue productivity to 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 RLP (revenue labor productivity) are given by $P_{et}\tilde{A}_{et}$.

This simple adjustment cost model is similar to Cooper, Haltiwanger and Willis (2007, 2016) and Elsby and Michaels (2013) and, in principle, accommodates both convex and non-convex costs. The latter make it so that the solution has the following form:

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

where

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

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

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 calibrated this model with shock processes and parameters that are consistent with relevant data and the existing literature. We do this to make the calibration as realistic as possible. However, the model is missing some key features of the data. First, we do not model entry or exit. Second, as discussed in the main text, we do not have any life cycle learning dynamics or frictions that make young firms different from more mature firms. Given these limitations, we regard the calibration as mostly providing guidance about the qualitative predictions about key moments that we explore in the data as described in the text.

In our initial calibration we conduct partial equilibrium exercises taking the real wage as constant. We also consider a simple way of accounting for general equilibrium considerations by specifying a perfectly inelastic aggregate labor supply and allowing the wage to adjust to clear the labor market. Since our partial equilibrium analysis is consistent with a perfectly elastic labor supply, these two calibrations provide perspective on the potential impact of general equilibrium considerations. As will become clear, for our moments of interest the results are very similar in the perfectly elastic and inelastic cases. This is not surprising given our focus on second moments such as variances and covariances in the cross section across firms in any given period.

Our calibration approach is as follows. We set $\beta = 0.96$, consistent with annual data. We specify that $\varphi = 0.8$; this is within the range of estimates of markups estimated in the literature (for the latter this is consistent with a markup of 25 percent). For the shock process, we specify $\sigma_a = 0.35$, consistent with the standard deviation of TFP in Manufacturing during the 1980s that we report in Figure A.1; and we set $\rho = 0.65$, consistent with the AR(1) coefficient on TFP that we find among Manufacturing establishments in the 1980s (see Figure B.1). These values are also broadly consistent with the RPR based measures of productivity residuals from Appendix C.

These assumptions imply that TFP innovations have an implied standard deviation of $\sigma_\eta = 0.26$. We then calibrate the adjustment cost parameter(s) to target the value of 0.25 for the job reallocation rate, which is about the value of job reallocation in the 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 $F_- = 0$. Here we have arbitrarily set $F_- = 0$ in our baseline, but we consider alternatives in the calibrations discussed below. For quadratic adjustment costs (i.e., $F_+ = F_- = 0$) we find that matching the job reallocation rate of 0.25 requires $\gamma = 1.3$. We are not in a position to identify convex vs. non-convex costs since, unlike Cooper, Haltiwanger and Willis (2007, 2014), we are not calibrating or estimating moments that would enable this identification. For the moments we focus on in our empirical analysis, we find broadly similar patterns of increases in adjustment costs either from convex or non-convex costs. Since the empirical evidence in the literature points towards including non-convexities to account for some properties of the micro evidence, we focus the remainder of this analysis on the model with kinked adjustment costs.

An important property of our baseline calibration that matches the job reallocation rate in the 1980s is that TFPQ, TFPR and RLP (the latter two are the same in this setting given one factor of production) are highly correlated. The correlation is about 0.90 (this also holds in our general equilibrium analysis below). It is this property that underlies our finding that the responsiveness estimates of growth to realizations of productivity are 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. The second is a change in the dispersion of TFP. For this analysis, we focus on the kinked adjustment costs starting with $F_+ = 0.85, F_- = 0$, that is, our baseline non-convex cost calibration described above. We then increase F_- from zero to study an increase an adjustment costs. This permits us to obtain perspective on increasing adjustment costs from a starting point where we match the patterns of shocks and reallocation in the 1980s. In our second experiment, we start from the same baseline but consider a decline in the dispersion of TFP. Thus, in a similar fashion we obtain perspective on the implications of a decline in TFP dispersion from a starting point that matches the patterns in the 1980s.

A. *Changing adjustment costs*

In what follows, we describe results from the partial equilibrium version of the model (or, equivalently, the model setup in which labor supply is perfectly elastic). Figures D.1 and D.2 show the impact of increasing adjustment frictions (i.e., the downsizing cost F_{-}, holding fixed $F_{+} = 0.85$) for the key moments of interest. We find that an increase in adjustment frictions yields: (i) a decline in the job reallocation rate; (ii) a decline in the estimated coefficient of a regression of firm-level growth between t and $t+1$ on TFP in period t (where we include log employment in period t as a control; (iii) an increase in the dispersion of labor productivity (defined as revenue per worker); and (iv) a decline in the Olley-Pakes covariance (between size and productivity) for both TFP and revenue labor productivity (where employment serves as the weights). Each of these relationships are generally monotonic, with the exception of the Olley-Pakes covariance for revenue labor productivity (which we discuss further below). For the Olley-Pakes covariances, we use the standard Olley-Pakes decomposition given by:

$$P_{it} = \sum_{e \in i} \theta_{et} p_{et} = \bar{p}_{it} + cov(\theta_{et}, p_{et}) \tag{D1}$$

where P_{it} is industry aggregate productivity defined as the weighted average of firm-level productivity, $\bar{p}_i$ is the unweighted average of (log) firm-level productivity for the firms in the industry, θ_e is the share of industry employment accounted for by firm e, and p_e is the (log) labor productivity of firm e. We use the simulated data to implement this decomposition and report the covariance terms for both TFP and revenue labor productivity in Figure D.2.

There are many possible moments relating growth to realizations of TFP that are similarly sensitive to adjustment costs. For example, in unreported regression results (in which we always include period-t log employment as a control) we find that increasing adjustment costs yields (i) a decline in the estimated coefficient of a regression of firm-level growth between t and $t+1$ on TFP in period $t+1$; (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 of TFP (η_{it}). In principle, we could use any of these moments to detect a change in adjustment frictions. We use the specification reported in Figure D.1 for measurement and related econometric reasons as we discuss in the main text. But as that discussion notes, the exact timing in the model vs. the data are different so that 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 D.1 and D.2 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 under consideration, the model implies that this estimated coefficient is identical to the estimated coefficient on TFP that we report in Figure D.1. This precise equivalence is model dependent, but the general inference should not be. 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 (holding initial employment constant).

B. Changing TFP dispersion

Figure D.3 shows how key moments change in response to changes in the dispersion of TFP. Figure D.3 shows that in response to an increase in dispersion of TFP: (i) job reallocation increases, (ii) the dispersion in real labor productivity increases, and (iii) the estimated coefficient of a regression of firm level growth between t and t+1 on TFP in period t increases (where we include as a control log employment in period t). As before this last finding also holds using real labor productivity as the regressor given this simple model.

While the results in Figure D.3 are generally intuitive, one of the findings merits further discussion—specifically, the finding that responsiveness increases with the standard deviation of shocks to TFP. 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" documented by previous literature on non-convex adjustment costs. Non-convex costs create "inaction bands" or regions of the range of productivity innovations in which firms prefer inaction (i.e., zero hiring) to action. Inaction bands tend to 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 more than offset by the “volatility effect” or the notion that adjustments—when they actually do occur—are larger when TFP is more widely dispersed. In the model the volatility effect dominates, so overall responsiveness as measured by our regressions increases with dispersion.

C. General equilibrium

We have also examined a version of the same calibration exercise that permits the market wage to clear the labor market under the assumption of perfectly inelastic labor supply. A limitation of the partial equilibrium exercises is that with a fixed wage and perfectly elastic labor supply, the experiments create large changes in the average size of firms. This is particularly relevant when studying aggregate productivity: one of the terms in the Olley-Pakes productivity decomposition is the unweighted average of firm-level productivity, which can fall when some adjustment costs increase as firms hoard labor. In general equilibrium, the wage should adjust to dampen this hoarding incentive. For this and other reasons, we examine a case in which general equilibrium effects are extremely salient to provide insights into the degree to which general equilibrium concerns may affect our intuition. The results of this exercise are reported in Figures D.4-D.6. Comparing the latter to the partial equilibrium results suggests that the patterns in Figures D.1 through D.3 reported above are actually quite robust to this alternative.

D. The Olley-Pakes Covariance and the Diff-in-Diff Counterfactual

The patterns we have emphasized in Figures D.1-D.6 are, for the most part, robust to changes in key parameters of the curvature of the revenue function, the shock space and the adjustment cost parameters. There is one important exception that highlights the importance of using multiple moments in our evaluation of the empirical patterns in the data. Specifically, we note that in a frictionless benchmark with zero adjustment costs (in contrast to our benchmark above, in which hiring costs are set to $F_{+} = 0.85$), there will be zero labor productivity dispersion and in turn a zero Olley-Pakes covariance term for labor productivity. This frictionless benchmark yields predictions that are far from reality (e.g., the pace of reallocation is over 100 percent of employment), which can be seen on Figure D7. As adjustment frictions rise from this frictionless benchmark, labor productivity dispersion rises (Figure D7) as does the OP labor productivity covariance (Figure D8). However, since the responsiveness of firm growth between t and $t + 1$ on real labor productivity in t (with log period-t employment as a control) declines monotonically with an increase in adjustment frictions (Figure D7), as adjustment frictions rise sufficiently, the Olley-Pakes labor productivity covariance declines with further increases in adjustment frictions (Figure D8). This pattern is related to that found in Bartelsman, Haltiwanger and Scarpetta (2013) who found that rising distortions reduce the Olley-Pakes labor productivity covariance as long as the benchmark is characterized by sufficient frictions.

In our empirical analysis, we use a diff-in-diff approach for our analysis of Olley-Pakes covariance terms that makes this discussion largely moot. That is, as described in the main text we compute the counterfactual diff-in-diff that isolates the impact of the changing responsiveness on the OP covariance terms. Since this diff-in-diff has as a starting point the responsiveness at the beginning of our sample, our starting point is far from the frictionless environment but rather incorporates the adjustment costs consistent with the pace of job reallocation at the beginning of our sample.

For the sake of completeness, we compute the same diff-in-diff counterfactuals in the simulated data. Following our empirical approach, we use our model to generate simulated panels of firm-level data then run the same responsiveness regressions on those simulated data. Using the responsiveness coefficients we obtain from this exercise, we conduct a diff-in-diff exercise analogous to our empirical analysis. Specifically, we construct regression-predicted OP covariance terms for a variety of adjustment cost specifications; then, we apply the regression coefficients from the data generated with baseline adjustment costs to the data generated with higher costs and again construct predicted OP covariance terms. For any given adjustment cost specification, we then calculate the diff-in-diff generated by the two OP covariance terms just described: that is, the predicted OP covariance term using the appropriate regression coefficient for a given cost specification and the predicted OP covariance term generated by applying the baseline responsiveness coefficient to the higher-cost data scenario. Figure D.9a shows the patterns of the diff-in-diff counterfactuals for the partial equilibrium case, and Figure D.9b shows the counterfactuals for the case where we incorporate general equilibrium considerations.

Figure D1: Responses of key moments to changes in adjustment costs

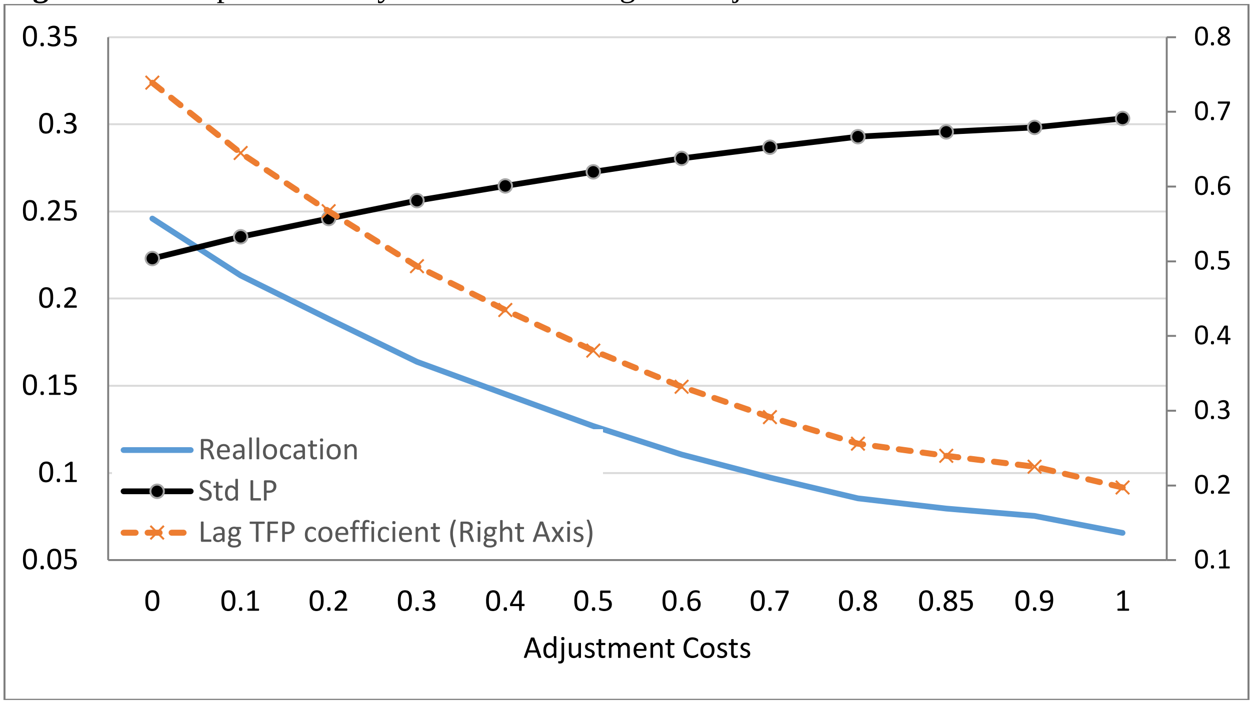

Note: Kinked adjustment costs. The x axis reflects values of F_{-}, or the cost of reducing employment, holding the hiring cost F_{+} fixed at $F_{+} = 0.85$. Partial equilibrium model with fixed wage and perfectly elastic labor supply.

Figure D2: Responses of Olley-Pakes covariances to changes in adjustment costs

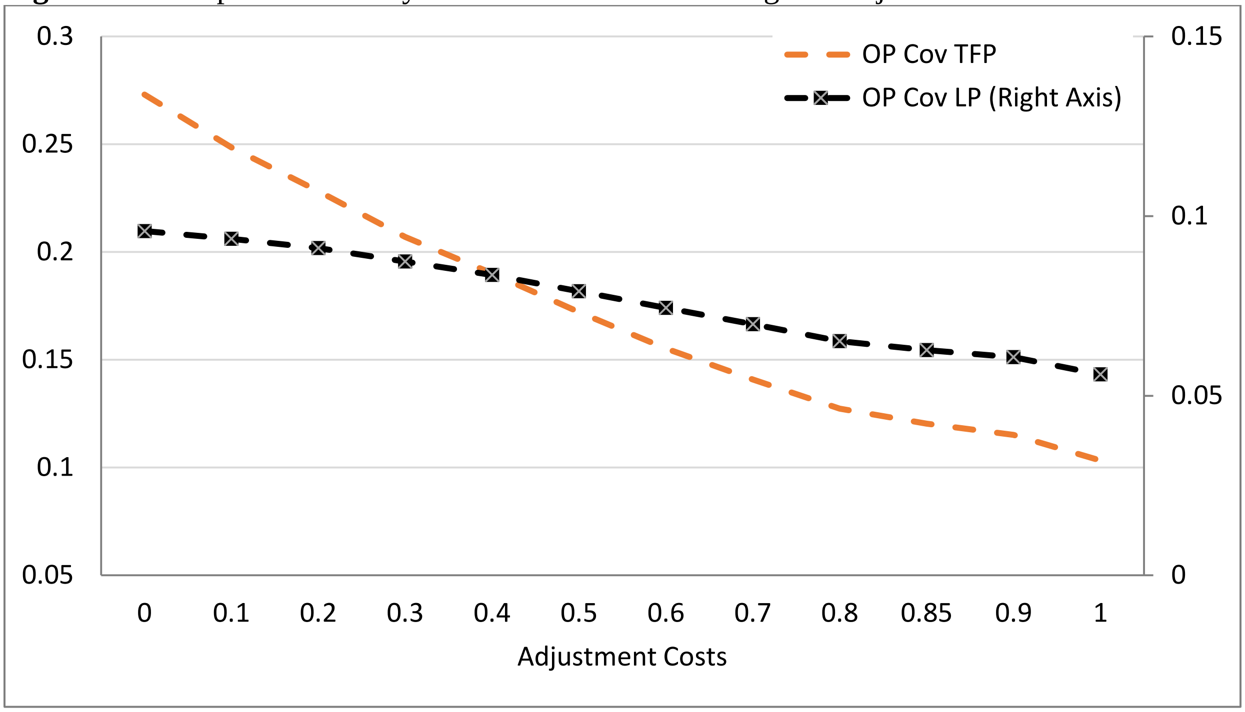

Note: Kinked adjustment costs. The x axis reflects values of F_{-}, or the cost of reducing employment, holding the hiring cost F_{+} fixed at $F_{+} = 0.85$. Partial equilibrium model with fixed wage and perfectly elastic labor supply.

Figure D3: Responses of key moments to changes in TFP dispersion

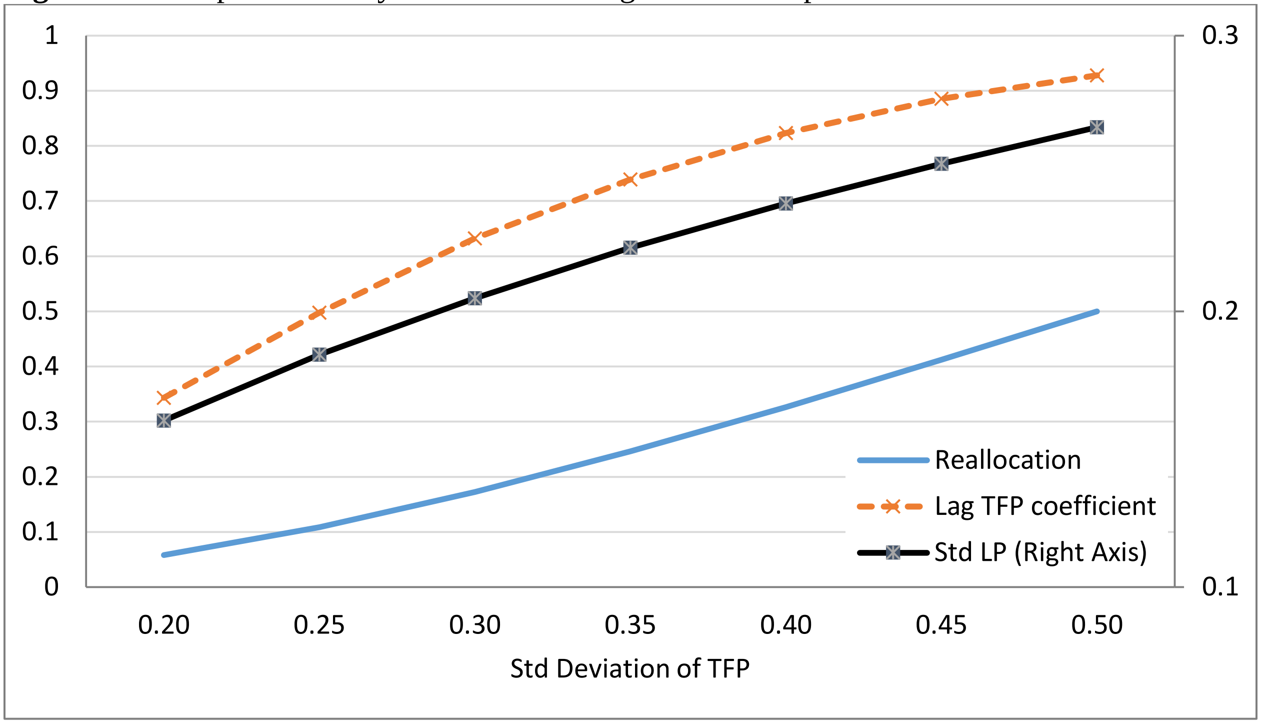

Note: Model with kinked adjustment costs ($F_{+} = 0.85$). Partial equilibrium model with fixed wage and perfectly elastic labor supply.

Figure D4: Responses of key moments to changes in adjustment costs (inelastic labor supply)

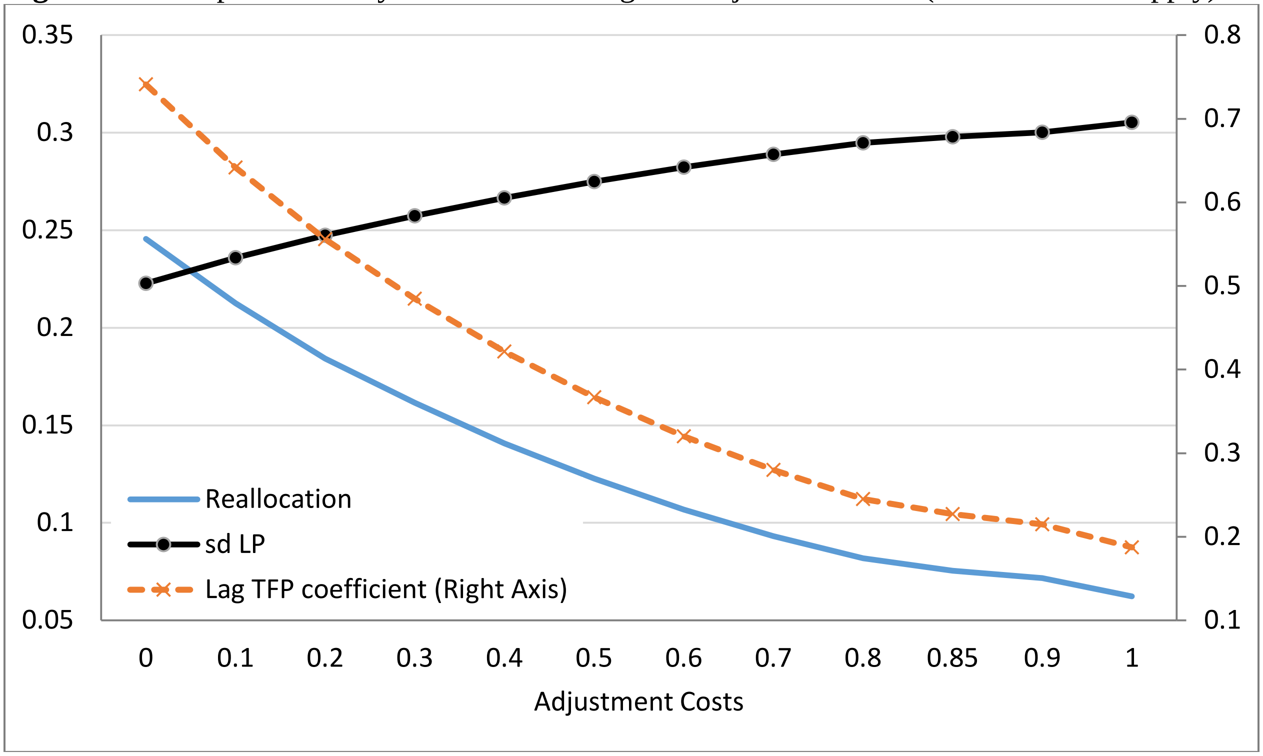

Note: The x axis reflects values of F_{-}, or the cost of reducing employment, holding the hiring cost F_{+} fixed at $F_{+} = 0.85$. General equilibrium model with flexible wage and inelastic labor supply.

Figure D5: Responses of Olley-Pakes covariances to changes in adjustment costs (inelastic labor supply)

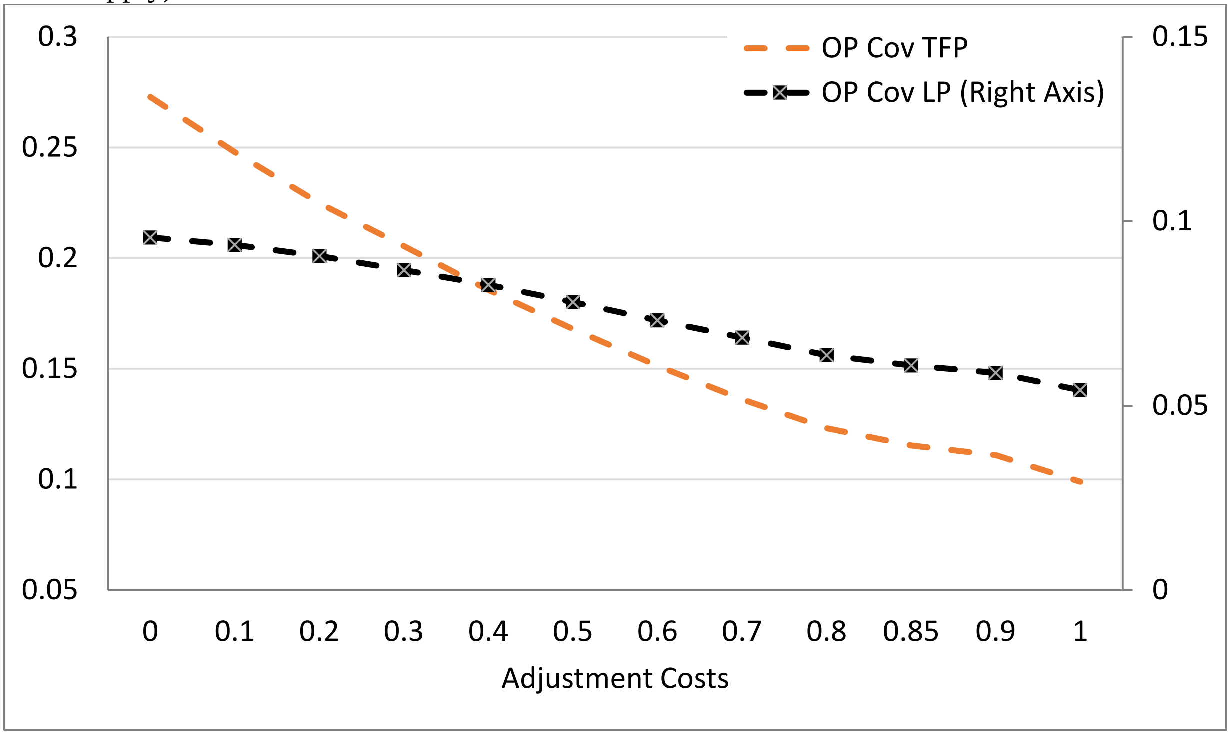

Note: Model with kinked adjustment costs. The x axis reflects values of F_{-}, or the cost of reducing employment, holding the hiring cost F_{+} fixed at $F_{+} = 0.85$. General equilibrium model with flexible wage and inelastic labor supply.

Figure D6: Responses of key moments to changes in TFP dispersion (inelastic labor supply)

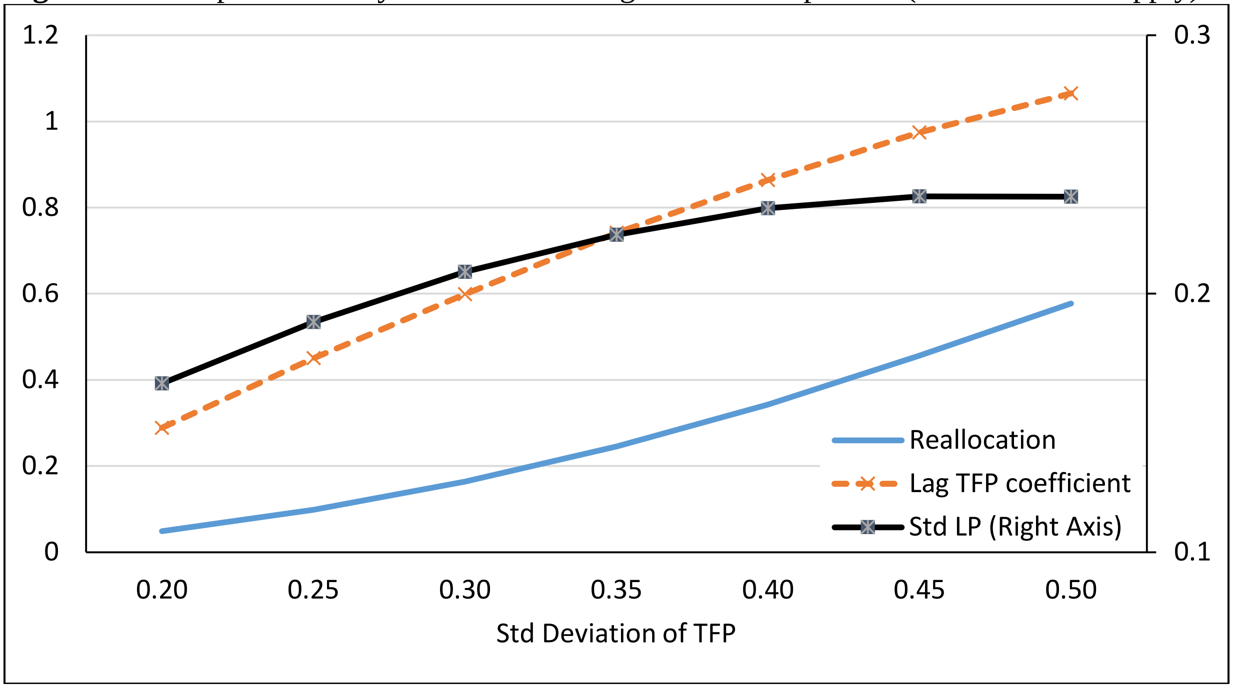

Note: Model with kinked adjustment costs ($F_{+} = 0.85$). General equilibrium model with flexible wage and inelastic labor supply.

Figure D.7: Responses of key moments to changes in adjustment costs from frictionless benchmark (inelastic labor supply)

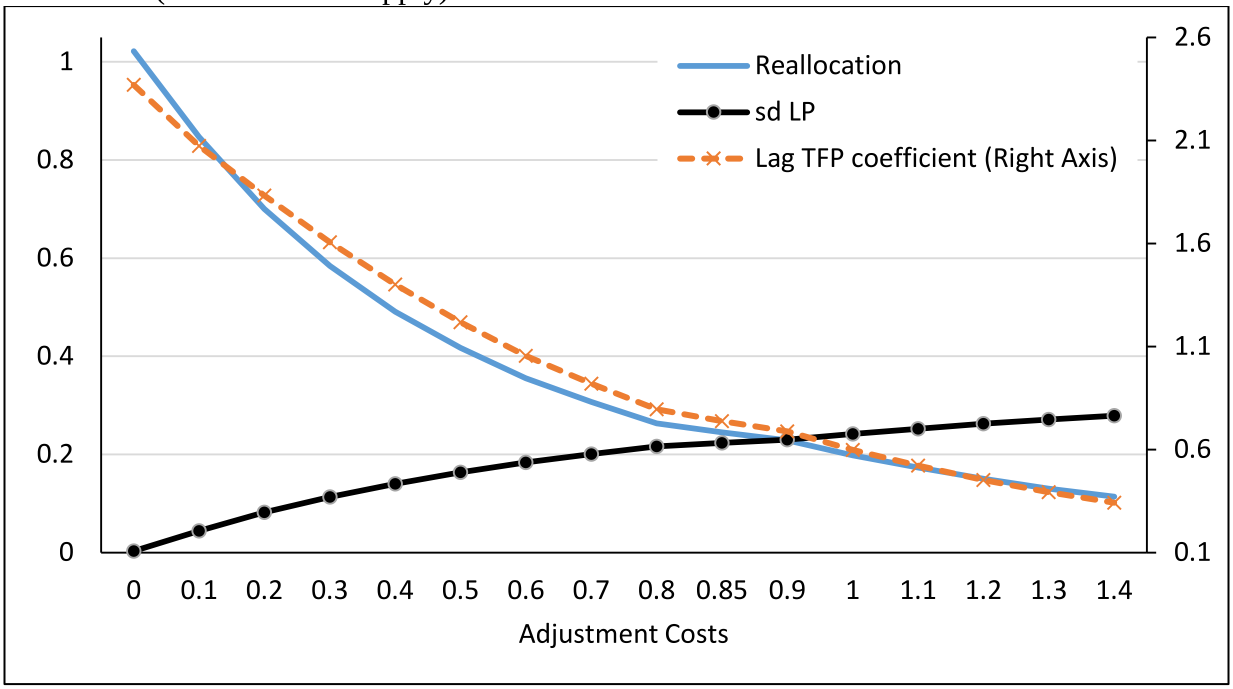

Note: Model with no upward adjustment costs ($F_+ = 0$) with varying downward adjustment costs (F_-) indicated on the x axis. General equilibrium model with flexible wage and inelastic labor supply.

Figure D8: Responses of Olley-Pakes covariances to changes in adjustment costs from frictionless benchmark (inelastic labor supply)

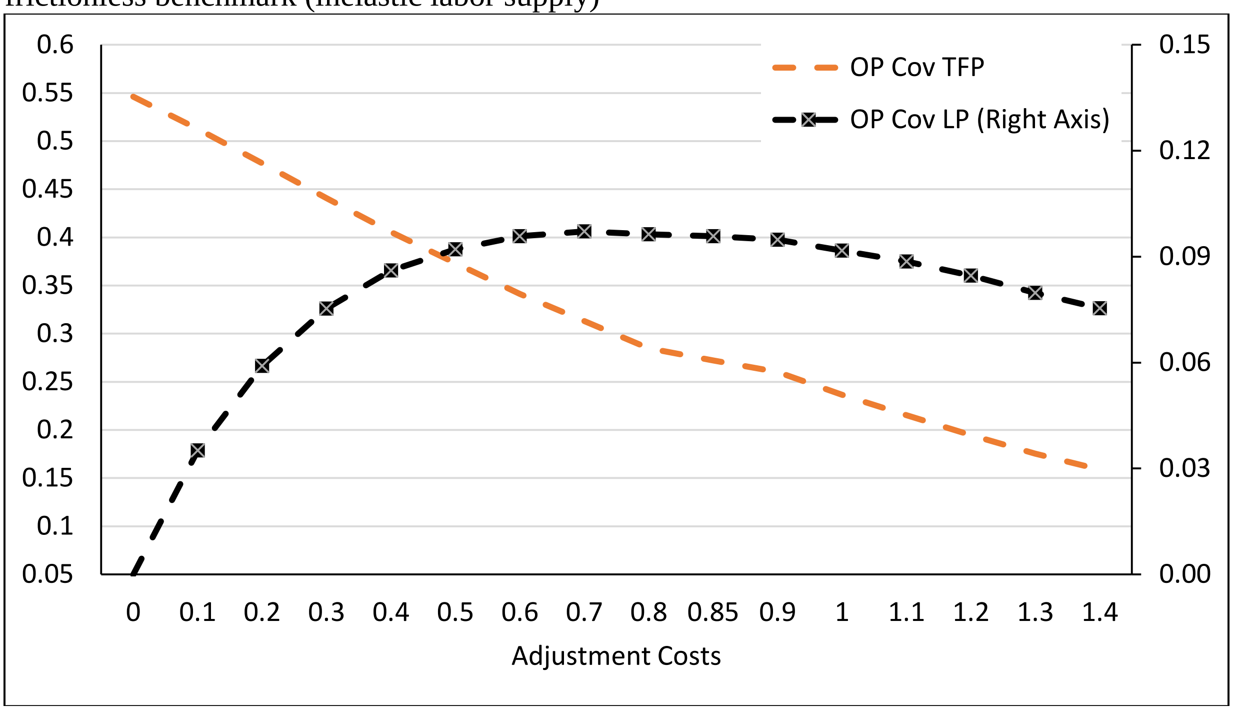

Note: Model with no upward adjustment costs ($F_+ = 0$) with varying downward adjustment costs (F_-) indicated on the x axis. General equilibrium model with flexible wage and inelastic labor supply.

Figure D.9a: Response of Diff-in-Diff Counterfactuals to Increase in Adjustment Costs (Partial Equilibrium)

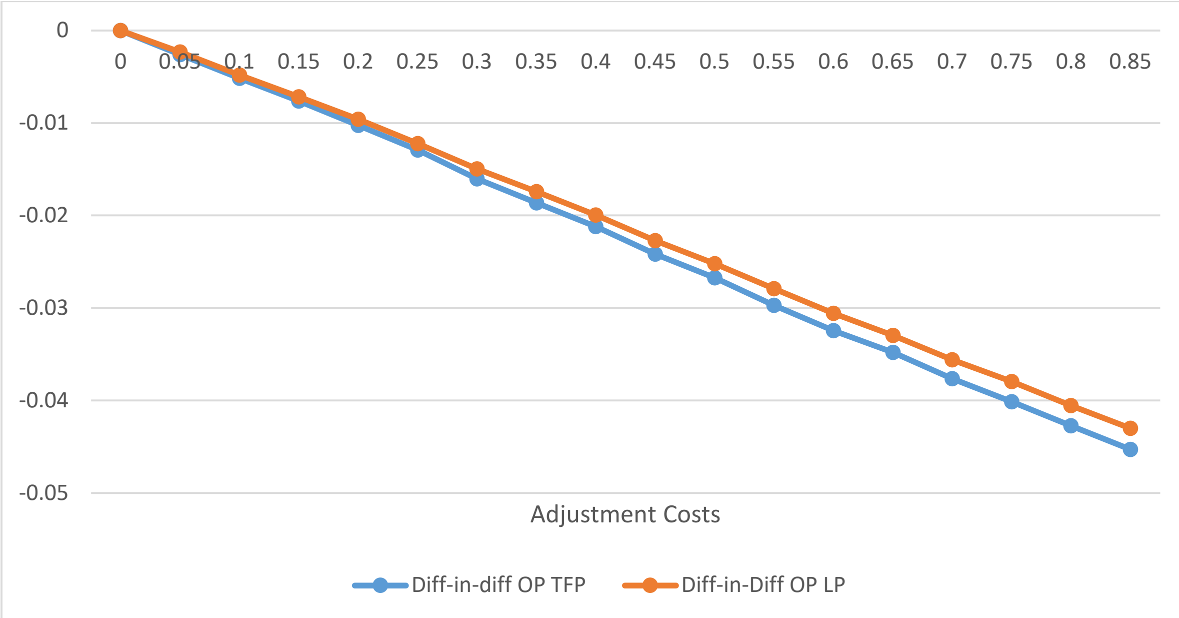

Note: Kinked adjustment costs. The x axis reflects values of F_{-}, or the cost of reducing employment, holding the hiring cost F_{+} fixed at $F_{+} = 0.85$. Partial equilibrium model with fixed wage and perfectly elastic labor supply.

Figure D.9b: Response of Diff-in-Diff Counterfactuals to Increase in Adjustment Costs (General Equilibrium)

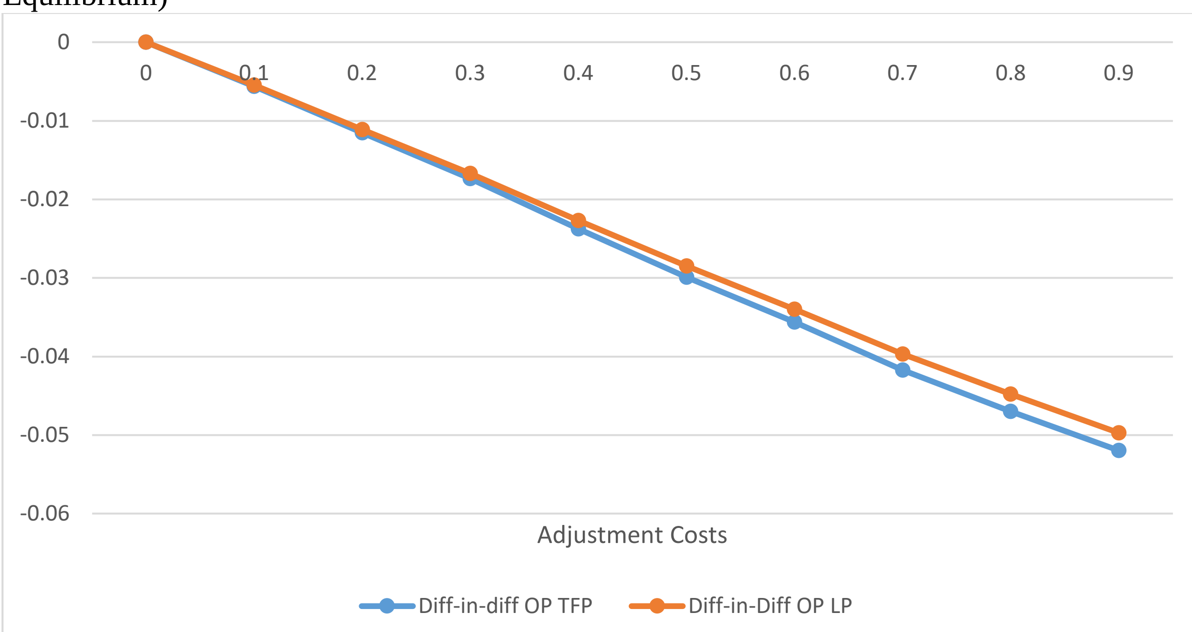

Note: Kinked adjustment costs. The x axis reflects values of F_{-}, or the cost of reducing employment, holding the hiring cost F_{+} fixed at $F_{+} = 0.85$. General equilibrium model with flexible wage and inelastic labor supply.

Appendix E. Multi-Unit Firms Operating in Multiple Industries

In the firm-level analysis, we deviate firm-level log revenue per worker from the firm's industry-by-year mean. This requires assigning a firm to a single industry. For this purpose, we use the modal industry of the firm based upon employment. Since 95 percent of firms are single-unit establishment firms this is unlikely to have large implications for empirical exercises that pool across single-unit and multi-unit firms. Indeed, in unreported results we have repeated the exercises using labor productivity for single-unit establishment firms only and obtain very similar results. In interpreting this, it is useful to note that the dispersion of within-industry productivity and the regressions on changing responsiveness use firm-level data without activity weighting. The one exercise that might be more sensitive to this issue is the diff-in-diff counterfactual since this is an employment-weighted exercise.

In this appendix, we consider an alternative way of treating multi-unit firms operating in multiple industries that will also potentially impact the single unit establishment firms as well. Specifically, the method used in the main text for generating the relative within-industry measure of productivity is equivalent to estimating a simple fixed effects regression of log revenue per worker on interacted industry-by-year effects and then using the residual as the relative productivity measure. As an alternative, we consider a regression of log revenue per worker on industry-by-year effects where each firm has up to five (6-digit NAICS) industry effects. Moreover, each of the (up to) five industry-by-year effects are interacted with a share variable that is equal to the share of employment within the firm for that industry. For a single unit establishment firm, only one industry effect is included, with the share variable equal to one.

Figure E1 is the analogue to Figure 7 showing within-industry dispersion for young and mature firms for tech and Non Tech firms. For computing statistical measures by industry group we still use the predominant modal industry but note that the industry grouping in Figure E1 is much broader than a 6-digit NAICS industry. Figure E2 is the analogue to Figure 8 in the main text, reporting the time-varying coefficients relating firm-level employment growth to firm-level labor productivity, based on firm-level productivity generated with this alternative demeaning method. Finally, Figure E3 is the analogue to Figure 9 in the main text, reporting diff-in-diff counterfactual results for the changing contribution of reallocation to aggregate productivity growth, again based on the alternative demeaning method. Observe that each figure based on the alternative demeaning method broadly confirms the results in the main text (that are based on a simple one-industry-per-firm approach). While the productivity counterfactual is modestly smaller in magnitude under the alternative approach, it still reveals large effects of the weakening relationship between productivity and growth.

Figure E1: Within-industry dispersion in labor productivity, young vs. mature (alternative method for demeaning)

(a) High-Tech firms

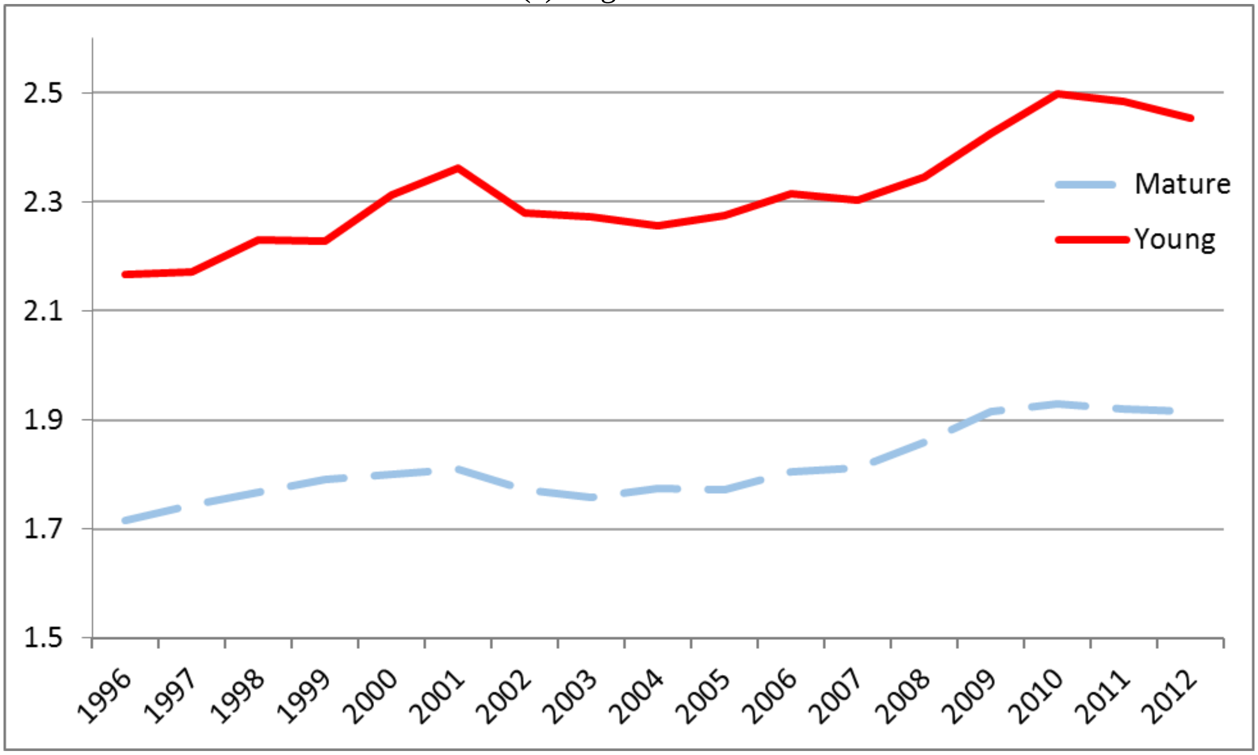

(b) Non Tech firms

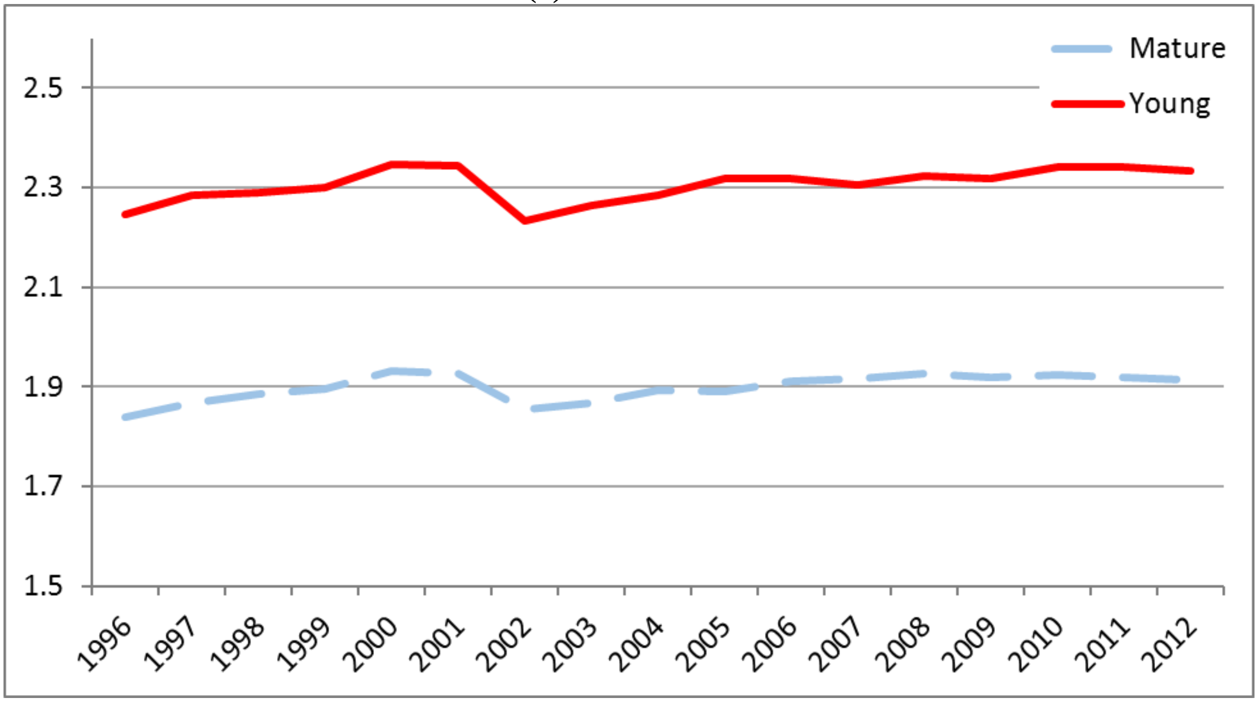

Note: Y axes do not begin at zero. Data reflect interdecile range of log labor productivity deviated from industry by year means. Young firms have age less than five. High-Tech is defined as in Hecker (2005). Author calculations from the RE-LBD.

Figure E2: Labor productivity and growth at the firm level (economywide, including exit), alternative demeaning method

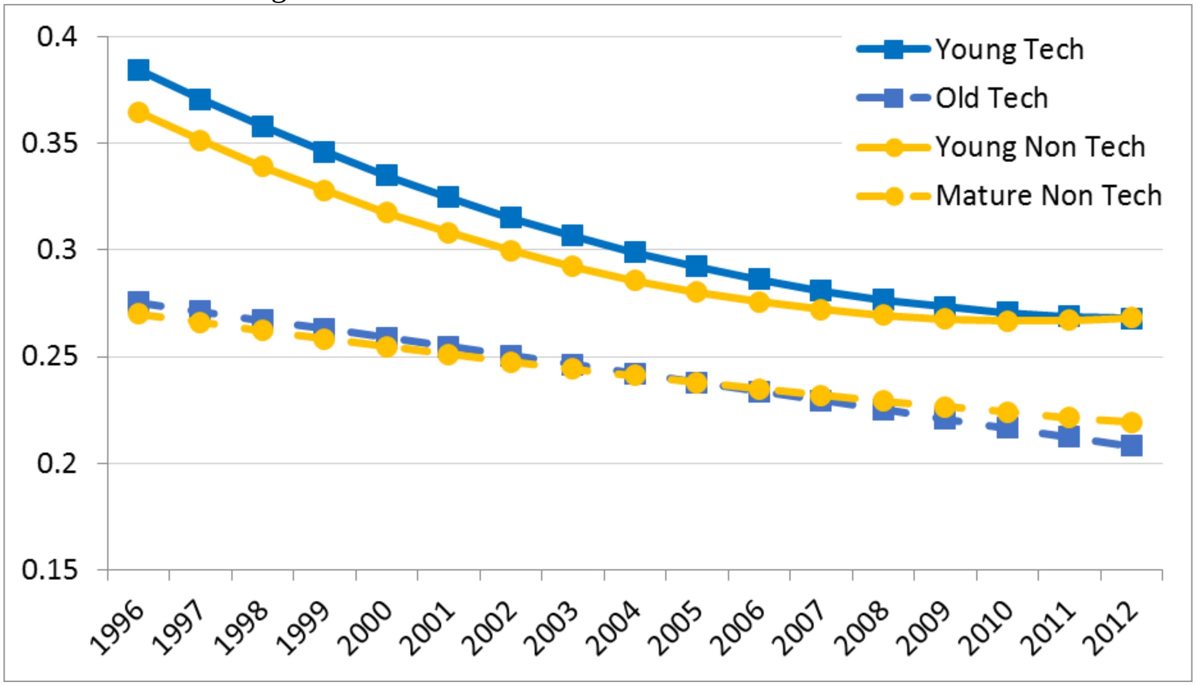

Note: Annual coefficients from analogue of Table 2 using the alternative measure. Young firms have age less than five. High-Tech defined as in Hecker (2005). Author calculations from the RE-LBD. Finance, Insurance and Real Estate (NAICS 52-53) omitted.

Figure E3: Diff-in-diff counterfactual (labor productivity using alternative demeaning)

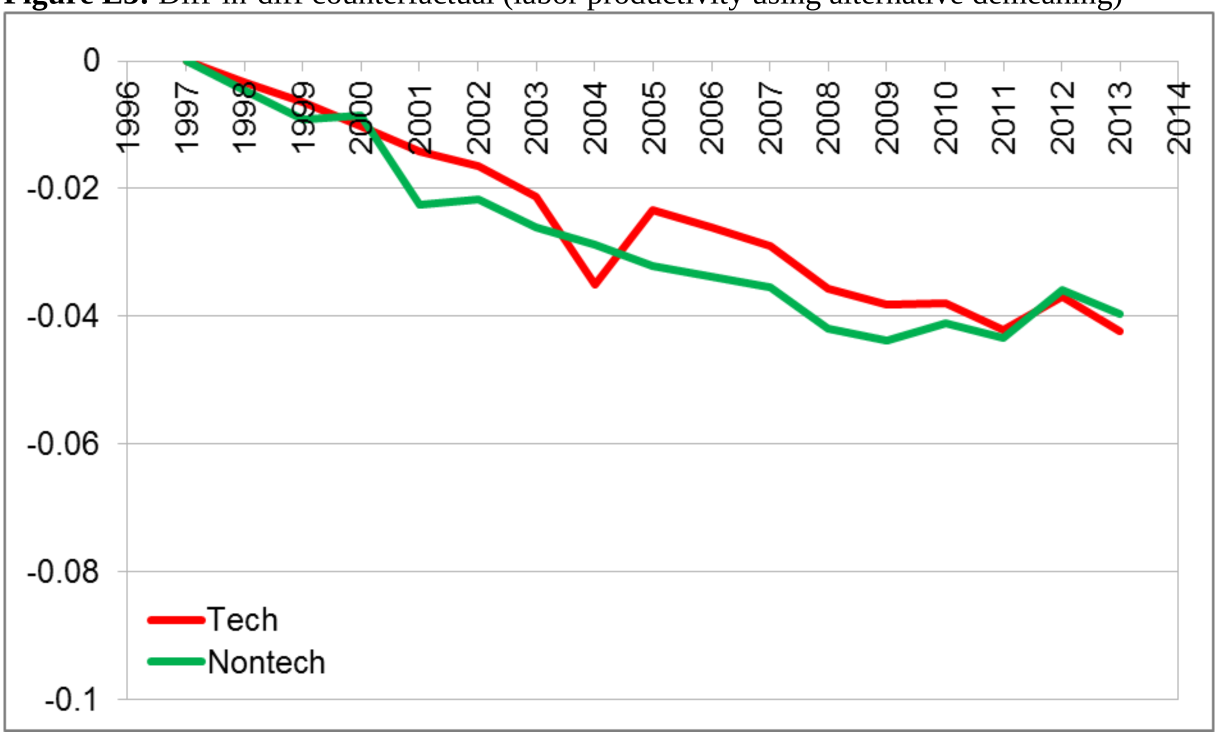

Note: Figure depicts counterfactual change in reallocation contribution to aggregate labor productivity growth. High-Tech is defined as in Hecker (2005). Author calculations from the RE-LBD

Appendix F. The changing composition of High-Tech Manufacturing

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).[46] One hypothesis is that manufacturers of special-purpose products may be less responsive to productivity shocks due to demand constraints or uncompetitive environments that reduce adjustment imperatives.
To investigate this hypothesis, we begin by examining the share of employment accounted for by the general purpose industries identified by Byrne (2015) in the High-Tech Manufacturing data. Figure F1 shows that during the 1990s the share of employment in among general purpose technology producers grew rapidly but, consistent with Byrne (2015) (which examined revenue shares), the general purpose share has fallen substantially since the late 1990s. Given these compositional changes, it is possible that the changing responsiveness reflects differential responsiveness across industries.

To explore the role of this composition effect, we estimate specification (3) 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 year 6-digit employment weights.[47] If the shift from general-purpose to special-purpose tech products is driving a decline in average productivity responsiveness, we would expect general-purpose plants to be more responsive than special-purpose plants, and this would be manifest in shift-share analyses that hold responsiveness constant but allow employment shares to shift over time as in the data.

Figure F2 shows the implied changing responsiveness over time due to composition effects. It is apparent 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 from composition effects rather than a decline. Interestingly, the average responsiveness of producers of special-purpose technologies is slightly higher than that of

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

general-purpose technology producers, which is why the aggregate responsiveness in our shift-share analysis rises slightly from the 1990s to the 2000s. Declining responsiveness must therefore be a within-category phenomenon with respect to the general-purpose/special-purpose taxonomy. These findings suggest that the rising and then declining pace of job reallocation in High-Tech Manufacturing cannot be accounted for by the changing composition of High-Tech.

Figure F1: General purpose technology share of High-Tech Manufacturing

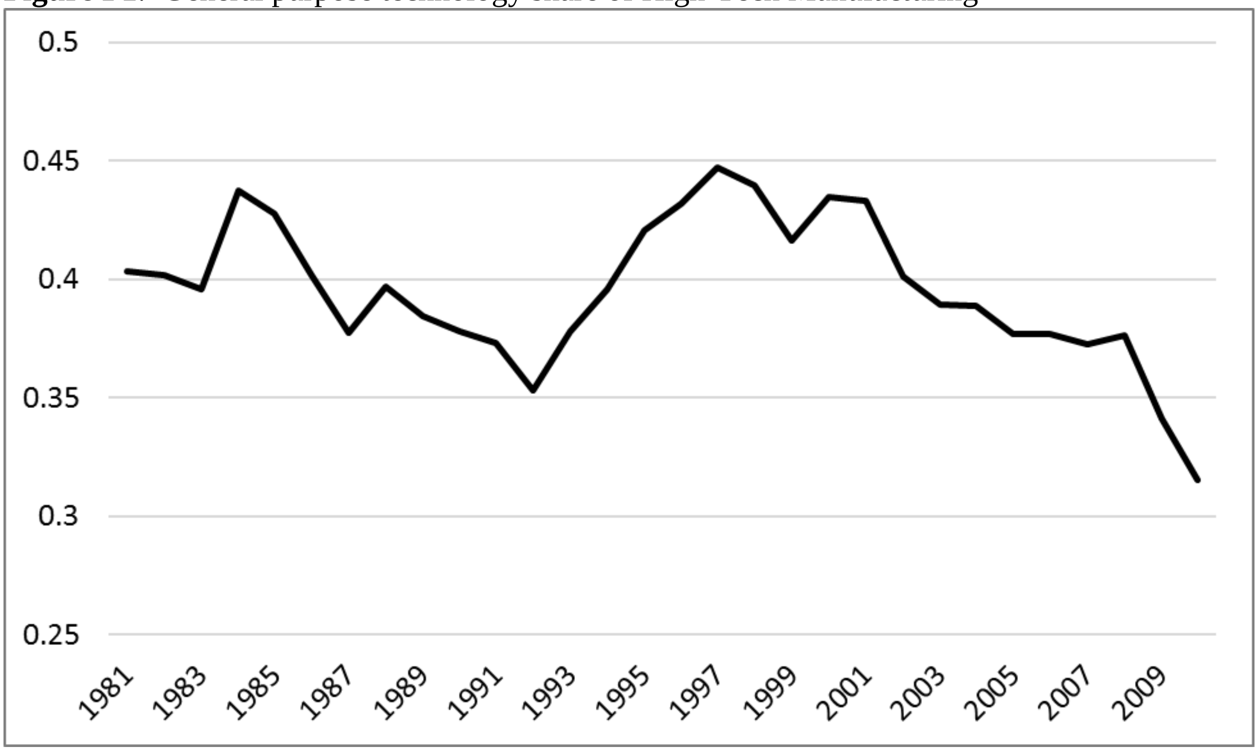

Note: Tabulations from the LBD by authors. General purpose High-Tech 4-digit industries are NAICS 3341 (Computers), NAICS 3342 (Communication Equipment) and NAICS 3344 (Semi-conductors).

Figure F2: Change in responsiveness due to industry composition changes (High-Tech)

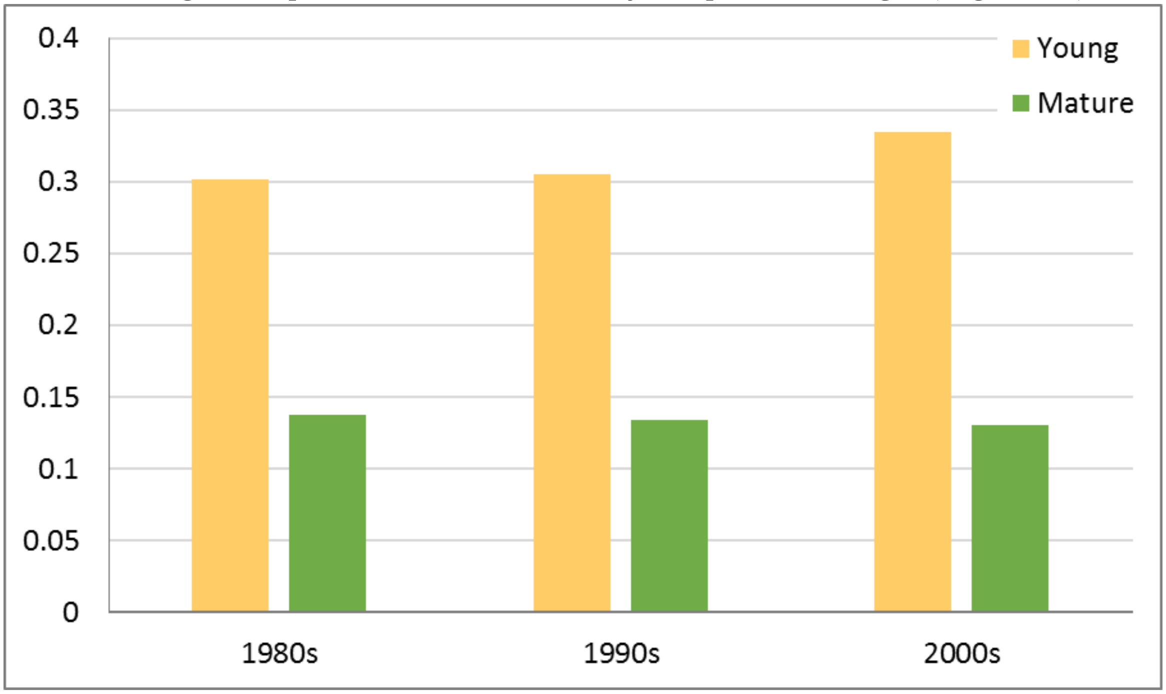

Note: Specification (2) as in Table 1 estimated for every 6-digit NAICS industry but without any trend effects. Reported coefficients are employment-weighted averages of the 6-digit NAICS industry estimated coefficients. Employment-weights vary by year.