Changing Business Dynamism and Productivity: Shocks vs. Responsiveness

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

The pace of job reallocation has declined in the U.S. in recent decades. We draw insight from canonical models of business dynamics in which reallocation can decline due to (a) lower dispersion of idiosyncratic shocks faced by businesses, or (b) weaker marginal responsiveness of businesses to shocks. We show that shock dispersion has actually risen, while the responsiveness of businesslevel employment to productivity has weakened. Moreover, declining responsiveness can account for a significant fraction of the decline in the pace of job reallocation, and we find suggestive evidence this has been a drag on aggregate productivity.

* Ryan Decker: Federal Reserve Board, ryan.a.decker@frb.gov; John Haltiwanger: University of Maryland and NBER, halt@umd.edu; Ron S Jarmin: U.S. Census Bureau, ron.s.jarmin@census.gov; Javier Miranda, U.S. Census Bureau, javier.miranda@census.gov. John Haltiwanger is also a Schedule A part-time employee of the U.S. Census Bureau at the time of the writing of this paper. We gratefully acknowledge financial support from the Kauffman Foundation. Cody Tuttle provided excellent research assistance. We thank John Abowd, Rudi Bachmann, Martin Baily, Jonathan Baker, Dave Byrne, Chris Foote, Lucia Foster, Clément Gorin, Bronwyn Hall, Matthias Kehrig, Pete Klenow, Kristin McCue, three anonymous referees, the journal editor, 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, BYU, University of Chicago, Drexel University, the Federal Reserve Board, the Richmond Fed, the spring 2017 Midwest Macro meetings, the 2017 UNC IDEA conference, and the 2017 NBER Summer Institute meetings 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. (2017). 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. The statistics reported in this paper have been reviewed and approved by the Disclosure Review Board of the Bureau of the Census (DRB-B0029-CED-20190314, DRB-B0056-CED-20190529, CBDRB-FY20-138).

Changing patterns of business dynamics—the entry, growth, decline, and exit of businesses have attracted increasing attention in recent literature. In particular, since the early 1980s the U.S. has seen a decline in the pace of "business dynamism" across many measures including the rate of business entry, the prevalence of high-growth firm outcomes, the rate of internal migration, and the rate of job and worker reallocation.¹ Declining business dynamism has attracted attention in part because business dynamics are closely related to aggregate productivity growth in healthy market economies, reflecting the movement of resources from less-productive to more-productive uses (Hopenhayn and Rogerson (1993), Foster, Haltiwanger, and Krizan (2001)). We draw insights from canonical models of firm dynamics in which declining reallocation reflects either a decline in the responsiveness of individual businesses to their underlying productivity or a decline in the dispersion of firm-level productivity shocks. Empirically, we find robust evidence of declining firm-level responsiveness accounts for a significant share of the aggregate decline in job reallocation and has been a drag on aggregate productivity.

"Job reallocation" measures the pace of job flows across businesses and is defined as total job creation by entering and expanding establishments plus total job destruction by downsizing and exiting establishments. Figure 1 shows the pace of aggregate job reallocation for the U.S. overall, the manufacturing sector, and the high-tech sector. The U.S. experienced an overall decline in the pace of job reallocation since the early 1980s. Even the high-tech sector saw a decline starting in the early 2000s. Understanding the causes of declining job reallocation has proven difficult. Decker et al. (2016b) show that it reflects in part a decline in firm-level growth rate skewness, or high-growth firm activity generally, but they do not investigate underlying causes. Young firms tend to exhibit a higher pace of job reallocation, and the share of activity accounted for by young firms has declined (Decker et al. (2014)), so some decline in reallocation is to be expected given composition effects. However, most of the variation in reallocation rates in recent decades has occurred within narrow age classes.²

¹ For job reallocation and employment volatility see Davis et al. (2007), Decker et al. (2014), and Decker et al. (2016b). For business entry see Decker et al. (2014) and Karahan, Pugsley, and Sahin (2018). For worker reallocation see Hyatt and Spletzer (2013) and Davis and Haltiwanger (2014). For migration see Molloy et al. (2016). For high-growth firms see Decker et al. (2016b), Haltiwanger, Hathaway, and Miranda (2014), and Guzman and Stern (2016).

 $^{^{2}}$ In Appendix III we describe a shift-share exercise to study the role of composition effects across firm age for explaining the overall decline in job reallocation. Figure A5 in the appendix reports the results. We are sympathetic to the view that studying the sources of the decline in startups and young firms is important for understanding the decline in job reallocation (e.g., Pugsley, Sedlacek and Sterk (2017)). Our focus is on the decline in job reallocation within firm age groups, which Figure A5 shows is quite important.

We study changing job reallocation patterns motivated by the framework of standard models of firm dynamics following Hopenhayn (1992) and a rich subsequent literature. In such models, reallocation arises from businesses' responses to their constantly shifting individual productivity and profitability environment. Businesses facing strong idiosyncratic productivity and profitability conditions expand (job creation), while those facing weak conditions downsize or exit (job destruction). The reallocation rate reflects the aggregation of these individual decisions. As such, a decline in the pace of reallocation can arise from one of two forces. First, the dispersion or volatility of idiosyncratic (business-level) conditions (which we call "shocks") could decline; in other words, a more tranquil business environment could reduce the need and incentives for businesses frequently and significantly to change their size or operating status. Second, or alternatively, the business-level responsiveness to those shocks could weaken; that is, businesses may hire or downsize less in response to a given shock (conditional on their initial level of employment), perhaps due (for example) to rising costs of factor adjustment.

These model-based considerations give rise to two competing hypotheses for declining job reallocation rates: the "shocks" hypothesis, in which the dispersion of idiosyncratic productivity or profitability realizations has declined; and the "responsiveness" hypothesis, in which businesses have become more sluggish in responding to realized shocks.³ We explore these hypotheses in high-quality business microdata for the U.S. We show that the dispersion of "shocks" faced by individual businesses has not in fact declined but has risen. However, business-level responsiveness to those shocks has declined markedly in the manufacturing sector and in the broader U.S. economy.⁴

These changes in responsiveness largely account for the observed decline in aggregate job reallocation. In the manufacturing sector, where we have high-quality measures of establishment-level productivity, we find that declining responsiveness accounts for virtually all of the decline in the pace of job reallocation from the 1980s to the post-2000 period (holding constant the age distribution of businesses). Even outside of manufacturing, where we have a more limited measure of firm-level productivity, declining responsiveness can account for about half of the late-1990s to post-2010s decline in job reallocation.

³ The general "shocks vs. responsiveness" framework has proven useful elsewhere; see Berger and Vavra (2019).

⁴ Rising labor productivity dispersion outside manufacturing was first documented in a related working paper, Decker et al. (2016a), and in Barth et al. (2016). Andrews et al. (2015) documented rising gaps in the growth of firm-level labor productivity in several OECD countries. Kehrig and Vincent (2017, 2020) present related evidence of a decline in responsiveness to a shock concept we denote as TFPR in our analysis below.

Business-level responses to productivity also facilitate productivity selection, and weaker responsiveness is indicative of weaker selection. We isolate the effect of changing responsiveness on an index of aggregate productivity using a simple counterfactual exercise. Aggregate total factor productivity (TFP) increased by about 30 percent in the U.S. manufacturing sector from the 1980s to 2000s, but our counterfactual exercise suggests TFP would have increased by about 33 percent if responsiveness in the 2000s were the same as in the early 1980s. We find similar effects on aggregate labor productivity for the U.S. private, nonfarm sector.

Taken together, our results suggest that declining reallocation is not simply a benign result of a less turbulent economy. Rather, declining reallocation appears to reflect weaker responses of businesses to their own economic environment, and the consequences of weaker responsiveness for aggregate living standards are nontrivial due to the important role of productivity selection. Determining the causes of weakening responsiveness is beyond the scope of this paper; however, we describe several possible avenues of investigation that are suggested by the model framework we employ.

Section I describes our conceptual framework and its empirical predictions for the "shocks" and "responsiveness" hypotheses. Section II describes our data, including our measures of productivity. Section III describes our empirical approach and results on "shocks" and "responsiveness." Section IV quantifies the implications for aggregate reallocation and productivity. Section V describes robustness exercises, and section VI concludes.

I. Conceptual Framework

A. General formulation

We begin by specifying, in quite general terms, the relationship of firm-level employment growth to firm-level productivity realizations (shocks) and initial employment in a one-factor (labor) model of business dynamics.⁵ Consider the employment growth policy function given by:

(1)
$$g_{jt} = f_t(A_{jt}, E_{jt-1})$$

⁵ We use the term "firm" loosely in this section. Our empirics feature both firm- and establishment-level data.

where g_{jt} is employment growth for firm j from t - 1 to t, A_{jt} is the productivity (or, more generally, profitability) realization in time t, and E_{jt-1} is initial employment. More concretely, we can motivate the canonical formulation in (1) with a model in which firms have a revenue function given by $A_{jt}E_{jt}^{\phi}$, where $\phi < 1$ due to either decreasing returns or imperfect competition; and the productivity process is $lnA_{jt} = \rho_a lnA_{jt-1} + \eta_{jt}$ (so η_{jt} is the innovation to productivity in period t). In typical models of this nature, $\partial f / \partial A > 0$; that is, among any two firms, the one with higher A_{jt} —holding initial employment constant—will have higher growth. We also include a time subscript t in $f_t(\cdot)$ to allow the relationship between growth and the underlying state variables to change over time (due, e.g., to changing employment adjustment costs).

While some expositions of this class of models specify g_{jt} as a function of the *change* in productivity (or of the innovation η_{jt}), we deliberately feature the *level* of A_{jt} in (1). Empirically, it is easier to relate the growth rate of firms (for which we have universe data) to productivity levels (for which we have cross-sectionally representative samples) than to changes or innovations (which require productivity data that are longitudinally representative). Moreover, the formulation in which A_{jt} is specified in levels, as in (1), is quite general since, under minimal assumptions, the inclusion of E_{jt-1} along with A_{jt} in the policy function fully incorporates information contained in A_{jt-1} and, therefore, the difference between A_{jt} and A_{jt-1} .⁶ That is, we can specify (1) using the level of A_{jt} and initial employment without significant loss of generality while improving the model's empirical comparability.

For empirical purposes, we focus on a log-linear approximation of (1) given by.

(2)
$$g_{jt} = \beta_0 + \beta_{1t}a_{jt} + \beta_{2t}e_{jt-1} + \varepsilon_{jt}$$

where the lowercase variables *a* and *e* refer to the logs of productivity and employment, respectively. The parameter β_{1t} is our measure of productivity responsiveness—it measures the marginal response of firm employment growth to firm productivity. In the typical model setting $\beta_{1t} > 0$, but the magnitude of this relationship depends on model parameters, distortions, adjustment frictions, and potentially firm characteristics (as we describe below). We refer to a

⁶ We formally show this in the appendix. The employment growth policy function can be specified in terms of levels of A_{jt} even in a frictionless model. Nevertheless, our empirical exercises are all robust to specifying the growth policy function in terms of changes in, or innovations to, A_{jt} .

change in β_{1t} as a change in responsiveness. Since the policy function specified in (1) and approximated by (2) determines firm-level employment changes, it also determines the aggregate job reallocation rate (which is simply the employment-weighted average of the absolute value of firm-level growth). Therefore, a decline in reallocation can be caused by either a decline in marginal responsiveness (β_{1t}) or a change in the distribution of A_{jt} shocks, a quite general result.

We next explore two concrete model specifications to illustrate numerically the implications of the shocks vs. responsiveness hypotheses.

B. Labor adjustment costs

Consider a canonical model of firm dynamics with labor adjustment costs in the tradition of Hopenhayn and Rogerson (1993). For simplicity we abstract from firm entry and exit. Faced with costs on labor adjustment, firms no longer adjust their labor demand to reach the firm size that would be implied by A_{jt} in a frictionless environment. Moreover, an increase in adjustment costs reduces responsiveness to A_{jt} (conditional on initial employment).

In Appendix I, we describe this model in detail under both non-convex cost and convex cost specifications; here, we initially summarize the results of numerical simulations using the model with non-convex adjustment costs (with the kinked adjustment costs explored in Hopenhayn and Rogerson (1993), Cooper, Haltiwanger and Willis (2007), and Elsby and Michaels (2013)). We then provide an overview of the analogous results with convex adjustment costs. We solve the model then simulate a panel of firms, allowing us to study job reallocation and productivity responsiveness (β_{1t} from equation 2) as well as another key moment, the dispersion of revenue per worker. The results using non-convex adjustment costs are in Figure 2.

Figures 2a and 2b illustrate our central hypotheses: declining reallocation can result from rising adjustment costs (i.e., lower responsiveness), as shown in 2a, or from declining shock (TFP or a_{jt}) dispersion, as shown in 2b. We focus first on adjustment costs. As these costs rise, job reallocation falls (Figure 2a) because the responsiveness coefficient weakens (the red long-dashed line on Figure 2c). As a result, dispersion of revenue per worker rises (the short-dashed green line in Figure 2c). In the absence of adjustment costs, equalization of marginal products would imply zero dispersion of revenue per worker; with adjustment costs, revenue per worker is positively

correlated with a_{jt} and exhibits positive dispersion. Additionally, as we show in the appendix (Figure A4), aggregate productivity declines as adjustment costs rise.

Alternatively, declining reallocation can result from declining dispersion of a_{jt} , as shown on Figure 2b. Figure 2d shows that, in this scenario, the dispersion of revenue productivity falls, as does the responsiveness coefficient. Therefore, this model can generate a decline in job reallocation if a_{jt} dispersion falls; other symptoms of declining a_{jt} dispersion are weaker responsiveness and lower revenue productivity dispersion.

All results in Figure 2 also hold under convex adjustment costs except for the dependence of the responsiveness coefficient on the dispersion of a_{jt} (see appendix discussion and Figure A1).⁷ Under convex adjustment costs, decision rules are approximately linear (see, e.g., Caballero, Engel and Haltiwanger (1997)) such that changes in second moments of shocks do not affect marginal responsiveness; responsiveness does not decline with a_{jt} dispersion under convex adjustment costs.⁸ The non-convex case therefore introduces some interaction between the shocks and responsiveness hypotheses, which we discuss further below.

C. Correlated "wedges"

The properties of the formulation in (1) and (2) are more general than the specific adjustment cost specifications just described. For example, consider a model with revenue function given by $S_{jt}A_{jt}E_{jt}^{\phi}$, where S_{jt} is a firm-specific distortion or "wedge" that can be thought of as a tax (when $S_{jt} < 1$) or a subsidy (when $S_{jt} > 1$). Let wedges be related to fundamentals such that log wedges (lower case) are determined by $s_{jt} = -\kappa a_{jt} + v_{jt}$ where, consistent with much of the recent literature, we assume $\kappa \in (0,1)$, and v_{jt} is independent of a_{jt} with $\mathbb{E}(v_{jt}) = 0.9$ In Appendix I

 $^{^{7}}$ For both the non-convex (Figure 2) and convex (Figure A1) setups we arrange the baseline scenario to target the pace of job reallocation in the 1980s among continuing U.S. manufacturing establishments; we also use the 1980s moments for the dispersion and persistence of shocks (see Table A1 in the appendix). The qualitative pattern of the impact of changing adjustment costs on responsiveness is similar across cost specifications, but the convex cost case generates a baseline responsiveness coefficient that is more quantitatively similar to our empirical results. While we do not generate structural estimates of adjustment costs in this paper, Cooper, Haltiwanger and Willis (2007) find that both convex and non-convex costs are needed to match the patterns in the data.

⁸ Non-convex adjustment costs give rise to inaction ranges. As productivity dispersion falls there is a decrease in the fraction of firms that make zero adjustment (i.e., the "real options" effect). But declining productivity dispersion also implies smaller adjustments among those firms that do adjust (i.e., the "volatility" effect). Vavra (2014) argues that the volatility effect dominates the real options effect in the steady state, a general result extending back to Barro (1972). Bloom (2009), Bloom et al. (2018), and others use a similar model to study the effects of uncertainty on business cycles; even in their model, the volatility effect dominates *at annual frequency* (see also Bachmann and Bayer (2013)).

⁹ A common finding in the literature is that indirect measures of wedges (i.e., revenue productivity measures like TFPR) are positively correlated with measures of fundamentals (technical efficiency and demand shocks) and have lower variance than fundamentals; see Foster, Haltiwanger, and Syverson (2008), Hsieh and Klenow (2009), and Blackwood et al. (forthcoming).

we show that an increase in κ acts in the same fashion as an increase in adjustment costs: reallocation declines, responsiveness declines, revenue labor productivity dispersion rises, and aggregate productivity declines (throughout the paper, we refer to increasing κ as "increasingly correlated wedges"). A decline in the dispersion of productivity shocks also yields a decline in reallocation and revenue productivity dispersion but, as in the convex adjustment cost model, responsiveness is not sensitive to dispersion in a_{jt} . The properties of the correlated wedge model are illustrated in Figure A2 of the appendix.

This wedge specification could be viewed as a reduced form encompassing the adjustment cost specification discussed in subsection B above (albeit with some important subtle differences given the explicitly dynamic components of an adjustment cost model). But this interpretation also may capture other possible changes in the distribution of wedges. For example, rising dispersion in variable markups that are correlated with fundamentals can play a similar role (see, e.g., De Loecker, Eeckhout, and Unger (2020); Edmond, Midrigan, and Xu (2018); Autor et al. (2019, forthcoming)).

D. Additional considerations

Our discussion thus far has focused on the intensive margin of responsiveness. However, related predictions apply for the extensive margin: Hopenhayn and Rogerson (1993) find that a rise in adjustment frictions reduces entry and exit. The empirical prediction of increased adjustment costs, then, is that not only will the growth of continuing firms become less responsive to firm productivity, but so will exit, a prediction we explore empirically.

Our motivating discussion also neglects post-entry dynamics from learning that can influence the responsiveness of both the extensive and intensive margins by firm age (see, e.g., Jovanovic (1982)). We consider this possibility in our empirical analysis. This variation by firm age is interesting in its own right but also permits us to abstract from changes in average responsiveness due to the changing age structure of firms. Given the decline in the U.S. firm entry rate in recent decades, if young firms have different average responsiveness from mature firms, aggregate responsiveness could have changed due to composition effects. We control for potentially exogenous changes in entry rates in the U.S. by studying responsiveness within firm age groups.

We consider additional nuances in our empirical work. For changes in the shock process, we consider not only the evolution of the dispersion in a_{it} but also, for restricted samples, the

evolution of the dispersion of innovations to the shock process and the persistence of this process.¹⁰ We also estimate responsiveness to innovations or changes in productivity, and we explore how changes in responsiveness vary across industries that have undergone different trends in productivity dispersion and persistence.

E. Summing up

In general, reallocation declines when either responsiveness or shock (a_{jt}) dispersion decline. A decline in responsiveness can be generated by, for example, an increase in adjustment costs (convex or non-convex) or, more generally, an increase in the correlation between reduced form wedges and the a_{jt} fundamental; in these cases, the dispersion of revenue labor productivity rises. A decline in the dispersion of the a_{jt} shock, while capable of generating a decline in reallocation, also reduces the dispersion of revenue labor productivity (and, in the case of non-convex costs, reduces responsiveness as well). These model predictions provide sufficient empirical moments for distinguishing between the shocks and responsiveness hypotheses. A critical point here is that we empirically examine both responsiveness and shock dispersion.

The gold standard empirical test of the responsiveness hypothesis is to estimate the changing relationship between the growth rate of employment and a_{jt} , controlling for initial employment. For the manufacturing sector, we can construct measures of a_{jt} (and η_{jt} , the innovation to a_{jt} , for restricted samples). For other sectors, we can only measure revenue per worker. However, in the adjustment cost models, an increase in adjustment frictions also implies a declining covariance between growth and the realization of revenue per worker.¹¹ Given this auxiliary prediction, we also explore changing "responsiveness" for non-manufacturing businesses using the changing relationship between employment growth and revenue per worker.

II. Data and measurement

The main database for our analysis is the U.S. Census Bureau's Longitudinal Business Database (LBD), to which we attach other data as detailed below. The LBD includes annual location, employment, industry, and longitudinal linkages for the universe of private non-farm

 $^{^{10}}$ In the adjustment cost framework, a decline in shock persistence can reduce responsiveness.

¹¹ See Appendix Figure A3a. These remarks also hold for revenue productivity measures such as TFPR as we discuss in our empirical analysis.

establishments, with firm identifiers based on operational control (not an arbitrary tax identifier). Employment measures in the LBD come from payroll tax and survey data. We use the LBD for 1981-2013 (during which consistent establishment NAICS codes are available from Fort and Klimek (2016)). For some exercises we focus on the high-tech sector; we define high-tech on a NAICS basis following Hecker (2005).¹² As in previous literature, we construct firm age as the age of the firm's oldest establishment when the firm identifier first appears in the data, after which the firm ages naturally.

For both our manufacturing and private sector economy analysis, we use the LBD to measure employment growth, initial employment, and exit (characterized as an establishment or firm that has positive employment activity in March of calendar year t and zero activity in March of calendar year t + 1). We use these LBD measures of growth and exit even when we merge in productivity measures from elsewhere (described next); that is, we have measures of growth and exit for the universe of businesses.

A. Manufacturing: Measuring establishment-level productivity

We construct establishment-level productivity for over 2 million plant-year observations (1981-2013) using updated data following the measurement methodology of Foster, Grim, and Haltiwanger (2016) (hereafter FGH) combining the Annual Survey of Manufacturers (ASM) with the quinquennial Census of Manufacturers (CM); see Appendix II for detail. The resulting ASM-CM is representative of the manufacturing sector in any given year, but it is based on a rotating sample and thus lacks the complete longitudinal coverage of the LBD; this is why we use LBD measures of employment and employment growth.¹³ Thus, a critical feature of our empirical approach (for manufacturing) is integrating the high-quality longitudinal growth measures from the LBD in any given year with the cross-sectional measures of productivity from the ASM-CM.

The productivity shocks we measure are intended to capture variation in both technical efficiency and demand or product appeal. To make our measurement approach transparent, it is helpful to

¹² Hecker (2005) defines industries as high-tech based on the 14 four-digit NAICS industries with the largest share of STEM workers. This definition includes industries in manufacturing (NAICS 3254, 3341, 3342, 3344, 3345, 3364), information (5112, 5161, 5179, 5181, 5182), and services (5413, 5415, 5417).

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

be explicit about the assumed production and demand structure. Consider establishment-level demand function $P_{jt} = D_{jt}Q_{jt}^{\phi-1}$ (where D_{jt} is an idiosyncratic demand shock, $\phi - 1$ is the inverse demand elasticity, and *j* indexes establishments) with Cobb-Douglas production, that is, $Q_{jt} = \tilde{A}_{jt} \prod X_{jt}^{\alpha_x}$ for inputs X_{jt} (where \tilde{A}_{jt} is technical efficiency, or TFPQ). A composite measure of productivity "TFP" reflecting idiosyncratic technical efficiency and demand shocks can be defined as $A_{jt} = D_{jt}\tilde{A}_{jt}^{\phi}$. The ASM-CM data provide survey-based measures of revenue, capital (*K*), employee hours (*L*), materials (*M*), and energy (*N*).¹⁴ Then establishment revenue is given by (lower case variables are in logs):

(3)
$$p_{jt} + q_{jt} = \beta_k k_{jt} + \beta_l l_{jt} + \beta_m m_{jt} + \beta_n n_{jt} + \phi \tilde{a}_{jt} + d_{et}$$

where $\beta_x = \phi \alpha_x$ for factor X, and t denotes time (in years).¹⁵ The β_x coefficients are factor revenue elasticities that reflect both demand parameters and production function factor elasticities. The implied *revenue function residual*, which we denote as TFP, is given by:

(4)
$$TFP_{jt} = p_{jt} + q_{jt} - \left(\beta_k k_{jt} + \beta_l l_{jt} + \beta_m m_{jt} + \beta_n n_{jt}\right) = \phi \tilde{a}_{jt} + d_{jt},$$

that is, this measure of TFP is a composite of idiosyncratic technical efficiency and demand *shocks*. In terms of the conceptual framework described previously (and in Appendix I), this is the relevant measure of fundamental shocks consistent with demand and technology assumptions made in this section. With estimates of the revenue elasticities, this measure of TFP can be computed from observable establishment-level revenue and input data. We refer to this measure as "TFP" or "productivity" in what follows, but it should be viewed as the composite shock reflecting both technical efficiency and product demand or appeal. Our use of the revenue function residual to capture fundamentals is not novel to this paper. Cooper and Haltiwanger (2006) estimate the revenue function residual in their analysis of capital adjustment costs. Hsieh and Klenow (2009)

¹⁴ Labor input is total hours measured from the survey responses in the ASM/CM. We estimate factor elasticities for equipment and structures separately but refer only to generic "capital" for expositional simplicity here. See Appendix II for more discussion of production factor measurement in the data.

 $^{^{15}}$ Output (q) is total value of shipments plus total change in the value of inventories, deflated by industry deflators from the NBER-CES Manufacturing Industry Database. Capital is measured separately for structures and equipment using a perpetual inventory method. Labor is total hours of production and non-production workers. Materials are measured separately for physical materials and energy (each deflated by an industry-level deflator). Outputs and inputs are measured in constant 1997 dollars. More details are in Appendix II.

use a closely related measure as their empirical measure of "TFPQ."¹⁶ Blackwood et al. (forthcoming) use a similar measure in their analysis of allocative efficiency as a proxy for "TFPQ."¹⁷

Below we discuss two alternative approaches to estimating the revenue function residual concept for TFP in (4), but first we describe another productivity concept that is widely used in the literature, "TFPR," which is given by:

(5)
$$TFPR_{jt} = p_{jt} + q_{jt} - (\alpha_k k_{jt} + \alpha_l l_{jt} + \alpha_m m_{jt} + \alpha_n n_{jt}) = p_{jt} + \tilde{\alpha}_{jt}.$$

The key conceptual and measurement distinction between TFP in (4) and TFPR in (5) is using revenue versus output elasticities; under the assumptions made in this section, TFPR confounds technical efficiency and endogenous price factors. As emphasized by Foster, Haltiwanger, and Syverson (2008), Hsieh and Klenow (2009), and Blackwood et al. (forthcoming), this implies TFPR is an endogenous measure in this context (i.e., when prices are idiosyncratic and endogenous). Without frictions or wedges, TFPR will exhibit no within-industry dispersion and is therefore not an appropriate measure of fundamentals. With adjustment costs or correlated distortions, however, TFPR will be positively correlated with fundamentals. Empirically, TFPR and fundamentals are strongly positively correlated (Foster, Haltiwanger, and Syverson (2008) and Blackwood et al. (forthcoming)). The high correlation in practice helps rationalize the widespread use of TFPR as a measure of TFP in the empirical literature.¹⁸ For our purposes, TFPR is a useful measure since, in our model, an increase in adjustment costs, or increasingly correlated wedges, yield a decline in the responsiveness of growth to TFPR and a rise in dispersion of TFPR. In this respect, TFPR has properties similar to revenue per worker. We emphasize that, given the potential endogeneity limitation of TFPR, we do not consider it to be as clean a measure of "shocks" as is

¹⁶ The empirical measure of TFPQ used by Hsieh and Klenow (2009) is proportional to the revenue function residual measure of TFP given by (4). The measure they use for TFPQ (in logs) is $(p_{jt} + q_{jt})\frac{1}{\phi} - (\alpha_k k_{jt} + \alpha_l l_{jt} + \alpha_m m_{jt} + \alpha_n n_{jt}) = \tilde{a}_{jt} + \frac{d_{jt}}{\phi}$ (see their equation 19); that is, their TFPQ measure is equal to our measure of TFP from (4) divided by ϕ . While proportional, it is more challenging to construct their measure of TFPQ since it also requires an estimate of ϕ , which requires decomposing the revenue elasticities into their demand and output elasticities components (see Blackwood et al. (forthcoming)). Both the Hsieh and Klenow empirical measure and our measure are inclusive of any idiosyncratic demand shocks. Foster, Haltiwanger, and Syverson (2008) define TFPQ to be technical efficiency.

¹⁷ The gold standard is to use establishment- or firm-level prices permitting separation of technical efficiency and demand (and also alternative estimate approaches for output and demand elasticities). However, such prices are available for only limited products in the Economic Censuses (see Foster, Haltiwanger, and Syverson (2008)).

¹⁸ It is also a measure of fundamentals if plants are price takers. Moreover, De Loecker et al. (2016) suggest that TFPR might be a preferred measure in the presence of unmeasured differences in materials prices and other inputs that reflect quality; output prices are likely correlated with such measures, so TFPR helps capture such variation.

the TFP concept from (4). Rather, TFPR is a measure of revenue productivity (reflecting the product of prices and technical efficiency).

We now describe how we estimate our various manufacturing productivity measures. The construction of TFP from (4) requires estimates of the β_x revenue elasticities. We obtain estimates in two different ways, resulting in two alternative TFP-based measures. Our first and preferred TFP measure relies on the first-order condition (for factor X) from static profit maximization:

(6)
$$\alpha_x \phi = \beta_x = \frac{W_{xt} X_{jt}}{P_{jt} Q_{jt}},$$

where W_{xt} is the price of factor X such that β_x is the share of that factor's costs in total revenue. The condition in (6) will not hold for all establishments at all times if there are adjustment frictions or wedges, but we only need (6) to hold on average when pooled through time and over establishments within industries, an assumption commonly used in the literature (e.g., Syverson (2011)). We obtain factor shares of revenue from the NBER-CES database (at the 4-digit SIC level prior to 1997 and the 6-digit NAICS level thereafter) then extract revenue function residuals using equations (3) and (4). We call this measure TFPS (for "TFP-Shares").

Our second TFP measure is based on estimation of the revenue function in (4) using the proxy method GMM approach of Wooldridge (2009), allowing elasticities to vary at the 3-digit NAICS level (see Appendix II for details; see other applications in, e.g., Gopinath et al. (2017) and Blackwood et al. (forthcoming)). We refer to this measure as TFPP (for "TFP-Proxy"). The TFPP method allows us to avoid reliance on first-order conditions, but the estimation process involves high-order polynomials and so requires large samples. Following the literature, then, we use higher levels of aggregation for estimating industry elasticities—we use 3-digit NAICS compared to the 6-digit NAICS used for TFPS.¹⁹ This limitation of the proxy methods makes TFPS our preferred measure, but our results are robust to using TFPP.

For the TFPR measure from (5), we construct output elasticities as cost shares of inputs out of total costs (under the assumption of constant returns to scale).²⁰ We use the NBER-CES

¹⁹ Blackwood et al. (forthcoming) find that the Wooldridge (2009) method residuals are highly sensitive to outliers, but pooling across more observations mitigates this problem.

²⁰ See, e.g., Baily, Hulten, and Campbell (1992); Foster, Haltiwanger, and Krizan (2001); Syverson (2011); Ilut, Kehrig, and Schneider (2018); and Bloom et al. (2018). We construct time-invariant elasticities; in unreported exercises, we allow elasticities to vary over time with a Divisia index and find similar results.

productivity database to recover factor cost shares. Cost shares equal factor elasticities under the assumptions of cost minimization and full adjustment of factors; again, however, one need not assume full adjustment for each establishment in each time period but rather that this holds approximately when pooling across all plants in the same industry over time. Like TFPS, our TFPR measure avoids the noisiness of estimation and allows us to use output elasticities that vary at the detailed industry level.

Our data are not ideally suited for tracking the persistence of, and innovations to, the TFP measures given the panel rotation of the ASM and our use of CM data. However, for results requiring us to estimate persistence of innovations, we exclude years for which we do not have a representative sample of continuing plants in t and t - 1 in our ASM-CM data (first panel years and Census years). Additionally, we acknowledge that our establishment-level measures of TFP are vulnerable to errors arising from omitted factors. For example, use of intangible capital in production is a potential source of measurement error (discussed further below).

Our preferred "shock" measures are TFPP and TFPS, which are measures of fundamentals under the demand structure and production function assumptions made in this section. TFPR is a closely related measure but, under the same assumptions, reflects both fundamentals and endogenous prices. Exploring richer demand and production structures is an open area for future research.²¹

B. Total economy revenue labor productivity

While TFP is the preferred concept in a shocks vs. responsiveness framework, we can only estimate TFP in the manufacturing sector. For the economy generally, we rely on revenue per worker ("revenue labor productivity" or RLP), which is necessarily a firm-level (rather than establishment-level) concept in our data. As discussed above, rising adjustment frictions or increasingly correlated distortions also imply rising RLP dispersion and declining "responsiveness" of growth with respect to RLP.

²¹ Much attention has recently been given to the possibility of variable markups across producers in the same industry (e.g., De Loecker, Eeckhout, and Unger (2020)). The De Loecker, Eeckhout, and Unger (2020) approach uses the dispersion of the cost shares of fully flexible production factors in total revenue to indirectly identify markups as a residual. An alternative approach to identification of variable markups is to maintain the CES demand structure assumption but consider oligopolistic competition (e.g., Hottman, Redding, and Weinstein (2016)). In this approach, while firm-level markups are increasing in market share, our revenue function based on the CES demand structure is still appropriate. Sorting these issues out more fully is an important area for future research likely requiring price and quantity data for both outputs and inputs. See Eslava and Haltiwanger (2020) for discussion of these issues.

Combining LBD employment (summed from the establishment to the firm level) with revenue measures in the Census Bureau's Business Register (BR) (aggregated across EIN reporting units to the firm level) yields an enhanced LBD that we refer to as the RE-LBD.²² Revenue data are available from 1996 to 2013 and are derived from business tax returns.²³ We construct annual firm employment growth rates on an "organic" basis to represent changes in establishment-level employment rather than artificial growth caused by mergers and acquisitions.²⁴

For firm-level exercises, we assign each firm a consistent "modal" industry code based on the NAICS industry in which it has the most employment over time. In exercises reported in an earlier working paper version (Decker et al. (2018)), we found our results are robust to an alternative approach in which we explicitly control for all industries in which firms have activity rather than assigning each firm a single industry code. We omit firms in the Finance, Insurance, and Real Estate sectors (NAICS 52-53) from all analysis due to the difficulty of measuring output and productivity in those sectors.

III. Empirical Approach and Results

A. "Shocks" hypothesis

We now study the dispersion of our various productivity measures. For this purpose, we use *within-industry* productivity: for any productivity measure z (which is in logs), we specify establishment- or firm-level productivity as $z_{jt} - \bar{z}_t$, where \bar{z}_t is the average for plant j's industry in year t. Figure 3a reports the standard deviation of our three (within-industry, log) measures for manufacturing —TFPS, TFPP, and TFPR—averaged for the 1980s, the 1990s, and the 2000s (up through 2013). Our preferred measure, TFPS, sees an increase from about 0.46 in the 1980s to 0.51 in the 2000s. The other measures also show widening dispersion. Figure 3b reports the

²² The Business Register is the main source dataset for a variety of Census Bureau products including the LBD, County Business Patterns, and Statistics of U.S. Businesses.

²³ See Appendix II for more details on revenue data construction. About 20 percent of LBD firm-year observations cannot be matched to BR revenue data because firms can report income under EINs that may fall outside of the set of EINs that Census considers part of that firm for employment purposes. We address potential match-driven selection bias by constructing inverse propensity score weights.

²⁴ The organic growth rate calculation is straightforward but requires highly specific definitions of firm-level employment. For a firm *j*, let E_{jt+1} be the sum of employment in March of year t + 1 among all establishments owned by firm *j* in year t + 1, and let E_{jt} be the sum of employment in March of year t among all establishments owned by firm *j* in March of years t + 1, and let $E_{jt+1} = (E_{jt+1} - E_{jt})/(0.5E_{jt} + 0.5E_{jt+1})$. See Haltiwanger, Jarmin, and Miranda (2013) for more discussion of organic firm growth.

standard deviation of (within-industry, log) revenue labor productivity (RLP) for the total U.S. economy (the first column). Since our RLP data cover a shorter time span than our TFP data, we show more time detail. As is apparent, RLP dispersion has risen over this time period for the whole economy, showing that rising productivity dispersion is not just a manufacturing phenomenon. The remaining bars in Figure 3b report RLP dispersion for manufacturing only; specifically, the second set of bars is RLP dispersion in the ASM-CM, and the third set of bars is RLP dispersion in manufacturing from the RE-LBD.

Figures 3a and 3b reveal several insights. First, consistent with previous literature (e.g., Syverson (2004, 2011)), within-industry dispersion in TFP is large; for example, a level of 0.51 (51 log points) for TFPS implies that an establishment one standard deviation above the mean for its industry is about $e^{.51} \approx 1.7$ times as productive as the mean. Within-industry RLP is even more dispersed—as may be expected given potential dispersion in non-labor production factors, especially capital. Second, the three TFP measures, while substantially different in construction, yield broadly similar dispersion trends. Third, the rise in revenue productivity dispersion observed in manufacturing survey data is confirmed by administrative data (compare the second and third sets of bars in Figure 3b). Bils, Klenow, and Ruane (2020) argue that rising revenue productivity dispersion observed in the ASM is due to increasing survey-based measurement error, but Figure 3b shows that the rise in various measures of productivity dispersion in the U.S. is evident in administrative data, apparently not an artifact of survey limitations.²⁵

Recall we assume TFP follows $a_{jt} = \rho_a a_{jt-1} + \eta_{jt}$. In Figure 3c we report the dispersion of revenue TFP *innovations* (η_{jt} ; bottom left panel of Figure 3), and Figure 3d reports the *persistence* of revenue TFP levels (ρ_a ; bottom right panel).²⁶ The dispersion of innovations has also risen, while the persistence of shocks has declined only modestly, in recent decades.

Figure 3 implies that shock dispersion has not declined, as might be expected from declining reallocation, but, if anything, has actually risen. In other words, the dispersion and volatility of shocks faced by businesses have not evolved in a way that could explain declining job reallocation. The business environment has not become less idiosyncratically turbulent; rather, it has become

²⁵ Recall that while the ASM measures of revenue and employment are from survey responses, the RE-LBD measures are from business tax returns (for revenue) and payroll tax records (for employment).

 $^{^{26}}$ The AR(1) estimates for TFPS and TFPP in Figure 3d are somewhat lower than those in the literature (e.g., Foster, Haltiwanger, and Syverson (2008), which use a narrow sample of products, and Cooper and Haltiwanger (2006), which use only large plants that are in existence continuously from 1972-1988).

more so. These findings for the standard deviation of TFPS and TFPP realizations (and innovations) are direct evidence of rising shock dispersion. The findings for TFPR and RLP are indirect evidence. All else equal, the data on shock dispersion should imply a *rising* pace of reallocation, while we observe the opposite. We therefore reject the "shocks" hypothesis.

B. "Responsiveness" hypothesis: Initial exploration

We next evaluate the "responsiveness" hypothesis for declining job reallocation—that is, the hypothesis that declining job reallocation is a result of dampened responsiveness of firms and establishments to their idiosyncratic productivity shocks. The evidence of rising revenue productivity dispersion we document above already is consistent with responsiveness weakening; as shown in our model discussion, rising revenue productivity dispersion may reflect either declining responsiveness or rising dispersion of fundamentals. We can directly test the responsiveness hypothesis by estimating responsiveness itself in the data.

We proceed in a manner analogous to our measurement of responsiveness in model-simulated data above; that is, we estimate an expanded version of equation (2):

(7)
$$g_{jt+1} = \beta_0 + \beta_1 a_{jt} + T(a_{jt}, t) + \beta_2 e_{jt} + T(e_{jt}, t) + X'_{jt}\Theta + \varepsilon_{jt+1}.$$

Equation (7) forms the core of our approach to measuring changes in responsiveness over time, so we will describe it in some detail. Individual establishments or firms are indexed by j, and time (in years) is indexed by t. Note carefully the naming and timing convention of variables in (7): The dependent variable, g_{jt+1} , is annual "DHS" employment growth between March of calendar year t and March of calendar year t + 1.²⁷ Productivity (a_{jt}) is measured for the calendar year t. Initial employment (e_{jt}) is log employment as of March of year t. The naming and timing conventions in (7) represent an empirical analogue to equation (2) given the timing of the measurement of growth and productivity in the data.²⁸

²⁷ DHS growth rates are commonly used in the literature and refer to Davis, Haltiwanger, and Schuh (1996). The DHS growth rate in equation (7) is $g_{jt+1} = (E_{t+1} - E_t)/(.5E_{jt} + .5E_{jt+1})$. It is measured from the LBD.

 $^{^{28}}$ At first glance it might appear that equation (7) has slightly different timing than (2). However, equation (2) can be interpreted as expressing the growth of employment from the beginning to the end of period t as a function of the realization of productivity in period t and initial (beginning of period t) employment. Equation (7) approximates this empirically by expressing growth of employment from March of calendar year t to March of calendar year t + 1 as a function of the realization of productivity during calendar year t and initial employment measured in March of calendar

In our baseline results, we measure productivity (a_{jt}) by the level of (log) TFP. We extend that baseline specification in a variety of ways, including the use of innovations to or changes in (rather than levels of) TFP. For our baseline specifications using the log of TFPS or TFPP (either realizations or innovations), β_1 estimates "responsiveness" (or the response of growth to productivity at the establishment or firm level) and corresponds to β_{1t} from equation (2), our responsiveness regression on model-simulated data. In extended analyses we obtain insights into changing responsiveness with respect to TFPR and RLP. For all of our measures of productivity, we permit this responsiveness to vary over time via $T(a_{it}, t)$ as described below.

Initial employment, another critical state variable in our model, is given by e_{jt} , which is measured as log establishment-level employment from the LBD in March of calendar year t. X'_{jt} includes detailed industry fixed effects (e.g., 6 digit NAICS) interacted with year effects, establishment size (in the case of specifications for manufacturing), firm size, state fixed effects, the change in state unemployment rates (to measure state-level business cycle effects), and interaction terms between the change in state unemployment rates and productivity; our liberal inclusion of cyclical indicators is intended in part to avoid result contamination from the Great Recession.²⁹ We estimate equation (7) on our manufacturing establishment sample (covering 1981-2013) for our TFPS, TFPP, and TFPR measures, and on our total economy firm sample (covering 1997-2013) in which a_{jt} is replaced with the log of revenue labor productivity (RLP).

As stated, equation (7) allows productivity responsiveness to vary over time via $T(a_{jt}, t)$, which we define variously as follows:

(8)
$$T(a_{jt}, t) \in \begin{cases} \delta a_{jt} Trend_t, \\ \gamma_{97} a_{jt} \mathbb{I}_{\{t \ge 1997\}}, \\ \lambda_{80s} a_{jt} \mathbb{I}_{\{t \in (1980, 1990)\}} + \lambda_{90s} a_{jt} \mathbb{I}_{\{t \in [1990, 2000)\}} \\ + \lambda_{00s} a_{jt} \mathbb{I}_{\{t \ge 2000\}} - \beta_1 a_{jt} \end{cases}$$

year t. We explore implications of timing assumptions further in Appendix I. For example, in our simulated models responsiveness to *lagged* realizations of productivity also declines as adjustment costs rise.

²⁹ In unreported exercises, we omit the cyclical controls in X'_{jt} and find very similar results. Moreover, to further ensure the Great Recession does not drive our results, in unreported exercises we end our sample in 2007 and still find similar results.

The first element of (8) defines the time function as a simple linear trend with coefficient δ . The second element uses a dummy variable to split the manufacturing sample roughly in half, such that overall responsiveness is equal to β_1 prior to 1997 and $\beta_1 + \gamma_{97}$ thereafter. The third element allows responsiveness to vary by decade, where the final "decade" is 2000-2013; by subtracting $\beta_1 a_{jt}$, we remove the main effect specified in (7) so the decade coefficients can be interpreted in a fully saturated manner. We also permit the effects of initial employment to vary over time in an analogous fashion.

We emphasize that the employment growth measure and the initial calendar-year t employment measure are from the LBD; this is important for two reasons. First, the LBD growth measure uses longitudinal linkages available for all establishments. This means we can track employment growth from March of year t to t + 1 for each establishment in the representative ASM-CM cross section for which we have TFP measures in calendar year t. When we use innovations to, or changes in, TFP we reduce the set of years available but, again, track the employment growth for all establishments for which we measure innovations in t. Second, the LBD's administrative employment measures we use to measure growth and initial employment are of high quality, minimizing concerns about possible division bias from measurement error in initial employment. For the manufacturing analysis, the employment measure used to construct the growth rates and initial employment differs from the source data for total hours used to construct the TFP measures. See Section V.E below for further discussion and robustness analysis.

Table 1 reports results from establishment- and firm-level regressions using annual DHS employment growth (inclusive of exit) for the dependent variable, as in equation (7). All regressions include the $X'_{jt}\Theta$ term from equation (7), but we do not report those coefficients.³⁰ We divide the table into four parts reflecting our four productivity concepts: Panel A includes regressions using TFPS (in which factor elasticities are revenue shares) and TFPP (in which factor elasticities are estimated by proxy method), while Panel B includes regressions using TFPR (in which factor elasticities are simply cost shares), and RLP (real revenue per worker).

Consider the first section of Panel A, under the heading "TFPS (revenue share based)". This section refers to establishment-level regressions in which the dependent variable is employment growth and the productivity variable a_{jt} is TFPS. The first column specifies changing

³⁰ We report only β_1 and the time function $T(a_{jt}, t)$ coefficients to satisfy data disclosure constraints.

responsiveness with the linear time trend described in (8). For TFPS, we estimate a base responsiveness coefficient of $\hat{\beta}_1 = 0.2965$, a significant positive number indicating strong selection early in the data time period, but we also find $\hat{\delta} = -0.0035$, which indicates responsiveness has weakened over time, as hypothesized. The regression reported in the next column uses the post-1997 responsiveness shifter from (8). Here we find a pre-1997 responsiveness coefficient of $\hat{\beta}_1 = 0.2905$, but after 1997 the responsiveness coefficient is equal to the base estimate plus the coefficient on the post-1997 interaction, $\hat{\gamma}_{97} - 0.0952$, for a total responsiveness coefficient in the post-1997 period of 0.1953—a number that is still consistent with positive responsiveness and productivity selection, but much weaker responsiveness than in the earlier period. The next column reports estimates from the fully saturated decade indicators (λ) from (8). Here we see the clear step down in productivity responsiveness, from about 0.29 in the 1980s to 0.20 in the 2000s, and the lower rows of this regression column report *p* values from *t* tests of equality between the various decade coefficients; in the case of TFPS, each decade coefficient is statistically different from the others.

The remainder of Table 1 proceeds analogously to the TFPS analysis, substituting the alternative productivity measures into otherwise identical regressions. Within manufacturing, while the quantitative results differ some between the alternative measures, and the exact timing of the decline in responsiveness varies somewhat, overall the qualitative results are strikingly similar and confirm a multi-decade decline in responsiveness.

Results for the whole economy using RLP (revenue labor productivity) as the productivity concept are in the second section (or right-hand side) of Panel B. Again, these data have a shorter time window. We estimate a clear decline in responsiveness as shown by the negative, significant value for $\hat{\delta}$ (the trend term). In other words, the weakening of responsiveness we observe in manufacturing has occurred across the economy generally.

Table 2 reports regressions that mimic those in Table 1 except that the dependent variable is now exit (firm shutdown, a binary indicator that is unity if the firm exits the data between years t and t + 1) rather than growth. While the DHS growth rate indicator used in Table 1 is inclusive of exit, it is useful to focus on the extensive margin in isolation.³¹ We focus on our preferred measure, TFPS. The second column of Panel A shows an exit coefficient that goes from -0.0801 in the pre-

³¹ The exit specifications eliminate any concerns about division bias.

1997 period to -0.0461 (i.e., -0.0801 + 0.0340) thereafter. The third column shows a substantial and statistically significant weakening of exit responsiveness from the 1980s to the 1990s along with some further modest (and marginally significant) weakening in the 2000s.

A useful way to quantify the magnitude of the Table 1 and Table 2 results—and of the decline in responsiveness—is to link them to the actual distribution of productivity. We compute the implied difference in employment growth (or exit) between the establishment (or firm) that is one standard deviation above its industry mean and the establishment (or firm) that is at the industry mean by multiplying each regression coefficient by the standard deviation of productivity. From Figure 3a we take the average of the standard deviation of TFPS across decades (0.48) to isolate the effect of changing responsiveness (i.e., avoid confounding the responsiveness change with changes in dispersion) and multiply it by the decade coefficients found in the TFPS regressions in Table 1 and Table 2.

The result is in Figure 4a, where we flip the sign of the exit coefficients for comparability. During the 1980s, an establishment that was one standard deviation above its industry in terms of TFPS grew its employment (over one year) by 14 percentage points more than the industry mean, a striking illustration of the intensity of productivity selection within industries.³² That same establishment also faced an exit risk 3.7 percentage points lower than its industry mean. In the 1990s, the growth rate differential fell to 12 percentage points while the exit risk differential narrowed to 2.8 percentage points. By the 2000s, the growth differential was 10 percentage points and the exit risk differential was 2.5 percentage points. While productivity selection is still clearly evident, the decline in responsiveness has weakened selection materially, substantially narrowing the growth and survival advantage of high-productivity establishments.

Figure 4b shows the same differentials for RLP economywide (using the standard deviation of RLP from Figure 3b); since we do not have decade dummy coefficients for RLP, here we use $\hat{\delta}$, the linear trend coefficient, combined with the base coefficient $\hat{\beta}_1$, to construct annual responsiveness coefficients, then we report multi-year averages at the beginning and end of the period. The result is qualitatively similar to the TFP-based manufacturing results: the growth and survival advantage of high-productivity firms (those firms whose revenue per worker is one standard deviation above their industry mean), while still evident, has deteriorated over time. The

³² Again, we obtain the 14 percentage point result by multiplying the 1980s coefficient from the third column of Table 1 (0.2859) by the average dispersion of TFPS from Figure 3a (0.48), that is, $0.2859 \times 0.48 = 0.14$.

growth differential between high- and average-productivity firms has fallen from 25 percentage points (1996-1999 average) to below 21 percentage points (2011-2013), while the exit probability differential has gone from 7.6 to 6.7 percentage points.

Tables 1 and 2 and Figure 4 strongly demonstrate that responsiveness has weakened among U.S. businesses. We observe weakening responsiveness to four independent measures of establishment- or firm-level productivity, and we see the decline in three different specifications: negative linear trend estimates, a negative and significant step down in the second half of the sample versus the first half (using the post-1997 indicator), and significantly different responsiveness coefficients in the 1980s, the 1990s, and the 2000-onward period. Employment growth responsiveness has weakened, as has the sensitivity of establishment or firm exit.

As discussed above, there should also be a decline in responsiveness to the *innovation* to, or the change in, productivity. As previously noted, our data are not ideally suited to measuring productivity changes; our ASM-CM sample is representative in any specific year but is not designed to be longitudinally representative. With that caveat in mind, we estimate our manufacturing regressions replacing the level of productivity (a_{jt}) with the innovation to productivity (given by $\eta_{jt} = a_{jt} - \rho_a a_{jt-1}$) and with the change in productivity $(\Delta a_{jt} = a_{jt} - a_{jt-1})$. We focus on our preferred productivity measure, TFPS, and report the results in Table 3.³³ Note that this exercise significantly reduces our sample size (from over 2 million establishment-year observations to fewer than 1 million). Regardless, we still observe a decline in responsiveness of employment growth to TFPS innovations or changes, with a particular step down between the 1990s and the 2000s. In other words, our main results are broadly robust to the use of productivity innovations or changes rather than productivity levels.

Taken together with the evidence of rising productivity dispersion, these results suggest that the costs or incentives to adjust employment in response to changing economic circumstances have changed over time. While declining responsiveness can arise from declining shock dispersion in certain theoretical environments (the model with non-convex adjustment costs above, though not the other specifications we describe), shock dispersion and the dispersion of revenue per worker have actually risen; the latter is consistent with model setups in which rising adjustment costs or,

³³ Tables 1 and 2 show that the various TFP measures deliver broadly similar results, so in all remaining empirical exercises we dispense with our TFPP and TFPR measures and report only TFPS results. As noted above, TFPS—the TFP measure based on elasticities from factor shares of revenue—is our preferred TFP measure, as it allows for endogenous prices while avoiding the imprecision of revenue function estimation. We also continue to report specifications based on RLP to gain perspective on non-manufacturing activity.

more generally, increasingly correlated wedges drive a decline in responsiveness. In robustness exercises described further below, we provide more evidence that changing responsiveness is not the result of changes in the distribution of shocks (whether in terms of dispersion or persistence).

C. "Responsiveness" hypothesis: Young versus mature firms

One critical change in the composition of U.S. businesses in recent decades may be affecting these results: the secular decline in young firm activity. If young firms are typically more responsive to shocks than are more mature firms, overall responsiveness would decline as young firm activity falls. The potential for firm age-based composition effects to affect our results is possibly a significant limitation of the exercises presented in Table 1, so we next study changing responsiveness within firm age groups by estimating the following:

(9)
$$g_{jt+1} = \left(\beta_{1}^{y}a_{jt} + T^{y}(a_{jt},t) + \beta_{2}^{y}e_{jt} + T^{y}(e_{jt},t)\right)\mathbb{I}_{\{y=1\}} + \left(\beta_{1}^{m}a_{jt} + T^{m}(a_{jt},t) + \beta_{2}^{m}e_{jt} + T^{m}(e_{jt},t)\right)\mathbb{I}_{\{m=1\}} + X_{jt}^{\prime}\Theta + \varepsilon_{jt+1},$$

where y indicates young firms (those with age less than five), m indicates mature firms (those with age five or greater), and each $\mathbb{I}_{\{.\}}$ is a corresponding age dummy indicator. We focus on *firm* age, even in establishment-level regressions; we assign each establishment the firm age of its parent firm. $T^{y}(a_{jt}, t)$ and $T^{m}(a_{jt}, t)$ (and the corresponding effects for e_{jt}) are defined as in (8) with the addition of firm age superscripts on all relevant coefficients. We also include interactions of the cyclical controls with firm age in X'_{it} .³⁴

Table 4 reports results from the regression in (9) for TFPS and RLP. As before, we report results with employment growth as the dependent variable (Panel A, on the left) and exit as the dependent variable (Panel B, on the right). As can be seen in all specifications—TFPS and RLP, with both growth and exit as dependent variables—young firms are indeed more responsive than mature firms (even within decades). In other words, young firms face more intense selection. As such, some portion of the decline in responsiveness reported in Table 1 does indeed reflect the changing

³⁴ Disclosure limitations preclude a more detailed age analysis.

age composition of firms. However, as the trend and decade coefficients demonstrate, responsiveness has declined over time within firm age groups.

Responsiveness has particularly declined among young firms, which have historically been more responsive. The 2000s growth coefficient for young firms, 0.25, is weaker than the initial 1980s coefficient for mature firms, 0.27. Following the exercise used for Figure 4, the growth differential for young firms with TFPS one standard deviation above their industry-year mean has declined from over 17 percentage points in the 1980s to just over 12 percentage points in the 2000s, while the exit risk differential has narrowed from 4.9 to 3.3 percentage points.

D. "Responsiveness" hypothesis: High-tech

As shown in Figure 1, patterns of reallocation in the high-tech sector have differed from the broader economy in recent decades. In particular, in high-tech the pace of reallocation rose during the 1980s and 1990s before declining in the 2000s. Given our shocks vs. responsiveness framework, the reallocation patterns lead us to expect productivity responsiveness to behave similarly; that is, we expect productivity responsiveness in the high-tech sector to strengthen during the 1980s and 1990s, then weaken thereafter.

We estimate equation (9) separately for high-tech and non-tech businesses (see the data discussion for details on industry classification). Again, we report results using the TFPS and RLP productivity concepts. Table 5 reports the results of these regressions, where we report only growth regressions (omitting exit regressions for brevity; recall that our DHS growth variable is inclusive of exit). We focus first on TFPS results, the first four columns of the table. While the results for non-tech establishments (the first two columns) are similar to those of the economy generally (shown in Table 4), responsiveness patterns in high-tech (the third and fourth columns) are different, in a manner consistent with aggregate reallocation patterns. This can be clearly seen in the decade-specific responsiveness coefficients: responsiveness rises between the 1980s and 1990s and steps back down in the 2000s. This rising and falling pattern is particularly evident among young high-tech firms. Figures 5a and 5b report growth differentials with the method from Figure 4; among young high-tech firms in manufacturing, the employment growth differential between high-productivity establishments and average establishments rose from 12 percentage

points in the 1980s to over 15 percentage points in the 1990s then fell to less than 8 percentage points in the 2000s.³⁵

The last two columns of Table 5 report regressions using RLP as the productivity concept. Our RLP data only begin in 1996, but the high-tech responsiveness pattern is evident even in the linear trend coefficients (δ^y and δ^m), particularly among young firms. Figures 5c and 5d report the coefficients in terms of growth rate differentials, averaged for 1996-1999 and 2011-2013; among young high-tech firms, the growth rate differential between high-productivity firms and their industry average has declined from 30 to 23 percentage points.

Table 5 and Figure 5 tell a rich story about productivity responsiveness and selection in recent decades. Consistent with patterns of job reallocation, responsiveness among non-tech businesses has declined steadily and significantly in recent decades, particularly among young firms, which have historically faced intense selection but increasingly behave more like mature firms. In high-tech, we observe rising responsiveness from the 1980s into the 1990s then falling responsiveness thereafter. These patterns are consistent with aggregate patterns of job reallocation. More broadly, Table 5 is consistent with the results of Tables 1-4, which tell a story of pervasive decline in productivity responsiveness by the end of the 2000s.

The high-tech productivity responsiveness pattern is also consistent with patterns of aggregate productivity growth during the 1990s and the 2000s. Aggregate productivity growth in the U.S. increased during the 1990s before stepping down in the early-to-mid 2000s, driven largely by industries that produce or heavily use ICT products (Fernald (2014)). The Table 5 patterns of responsiveness in high-tech have this same inverted u-shape pattern.

IV. Aggregate implications

A. Aggregate job reallocation

The motivating fact of the paper is the decline in job reallocation. While theory suggests a tight link between business-level responsiveness and aggregate job reallocation, a natural question is whether the responsiveness changes we document are sufficiently large to matter in the aggregate.

³⁵ Recall that this exercise compares high-productivity establishments—those with TFPS that is one standard deviation above their industryyear mean—to establishments at the industry-year mean.

In this section, we conduct a counterfactual exercise to relate our responsiveness estimates to changes in aggregate job reallocation.

Job reallocation from establishment-level (firm-level) data is equal to the employment-weighted sum of the absolute value of establishment-level (firm-level) employment growth rates. Establishment- and firm-level growth rates are the outcome variable for our responsiveness regressions above, so we can use our empirical responsiveness estimates to compute implied aggregate reallocation with and without estimated changes in establishment- or firm-level responsiveness. For this exercise, we focus on the specifications that permit changing responsiveness that varies by firm age group (reported in Table 4).

Consider the decade-specific regression coefficients reported in Table 4. For any given establishment's actual observed firm age, initial employment, TFPS productivity, and control variables, these regression coefficients provide a predicted employment growth rate \hat{g}_{t+1} that varies by decade. Moreover, these predicted growth rates also provide predicted employment for each establishment given by $\hat{E}_{jt+1} = (1 + \hat{g}_{jt+1})E_{jt}$, where E_{jt} is the establishment's actual observed initial employment in March of calendar year t. Given predicted employment, we can compute the establishment's predicted employment share of total employment as $\hat{\theta}_{jt+1} = (\hat{E}_{jt+1} + E_{jt})/\sum_j (\hat{E}_{jt+1} + E_{jt})$.³⁶ Thus, for any given year t + 1, we can predict aggregate job reallocation as $\widehat{JR}_{t+1} = \sum_j \hat{\theta}_{jt+1} |\hat{g}_{jt+1}|$.

Since we can construct predicted aggregate job reallocation for any given year, we can also construct an estimate of the effect of changing responsiveness on aggregate job reallocation using our establishment-level microdata. For any given year t, let $\widehat{JR}_{t+1}^{DR} = \sum_{j} \hat{\theta}_{jt+1}^{DR} |\hat{g}_{jt+1}^{DR}|$ be the predicted job reallocation rate given the actual decade-specific responsiveness coefficients described in Table 4 (where the "DR" superscripts refer to "declining responsiveness"). Separately, we construct a *constant-responsiveness* ("CR") version of predicted job reallocation, $\widehat{JR}_{t+1}^{CR} = \sum_{j} \widehat{\theta}_{jt+1}^{CR} |\hat{g}_{jt+1}^{CR}|$, where the 1980s coefficients are used for all years t.

The *effect* of declining responsiveness on aggregate job reallocation is given by the difference between the two predicted job reallocation objects:

³⁶ Here we measure employment shares using the Davis, Haltiwanger, and Schuh (1996) approach, which is necessary for correct aggregation. In implementing these counterfactuals, we are using the same timing conventions as in equation (7).

(10)
$$\Delta_{t+1}(JR) = \widehat{JR}_{t+1}^{DR} - \widehat{JR}_{t+1}^{CR} = \sum_{j} \widehat{\theta}_{jt+1}^{DR} |\widehat{g}_{jt+1}^{DR}| - \sum_{j} \widehat{\theta}_{jt+1}^{CR} |\widehat{g}_{jt+1}^{CR}|.$$

Equation (10) isolates the effect of declining responsiveness on aggregate job reallocation by comparing model-predicted reallocation under measured declining responsiveness trends to model-predicted reallocation when responsiveness is held constant at its 1980s strength (where by "model" we refer to the empirical regression model). While real-world reallocation may depend on additional factors beyond productivity responsiveness, these factors are controlled for in our counterfactual since both terms of equation (10) are based entirely on model-driven (i.e., responsiveness-driven) reallocation. If responsiveness coefficients do not change over time, then $\Delta_{t+1}(JR) = 0$, even if actual job reallocation has changed. Moreover, for each year *t*, equation (10) takes as given the actual empirical distribution of productivity and initial employment.

We calculate $\Delta_{t+1}(JR)$ in every year then take the average over the 2000s to obtain an average effect of -3.6 percentage points; that is, by the 2000s, declining responsiveness has reduced the aggregate job reallocation rate in manufacturing by 3.6 percentage points (relative to the higher responsiveness of the 1980s). Actual job reallocation in manufacturing, holding age composition constant, declined about 2.7 percentage points from the 1980s to the 2000s.³⁷ We seemingly over-account for the decline in the pace of job reallocation. However, the counterfactual exercise isolates the effect of declining responsiveness, taking the distribution of establishment-level productivity as given and abstracting from all other drivers of reallocation. Put differently, the counterfactual yields what reallocation would have been in the 2000s if responsiveness had been the same as it was in the 1980s, but with the 2000s distributions of productivity and employment.

We conduct an analogous exercise for the decline in job reallocation from the 1996-99 to 2011-13 for the entire U.S. private sector. That is, we use equation (10) where predicted employment shares and growth rates are based on responsiveness to revenue labor productivity as estimated in the third column of Table 4. The constant responsiveness (CR) scenario holds responsiveness at its initial value (i.e., setting the trend coefficients $\delta^y = \delta^m = 0$), while the declining responsiveness (DR) scenario allows responsiveness to vary according to the estimated trend coefficients (which effectively provide a specific responsiveness coefficient for any given year).

³⁷ For this purpose, we use the decline in Figure 1 adjusted for the changing firm age composition of manufacturing shown in Figure A5 of the appendix, since our counterfactual exercise relies on age-specific coefficients.

Notably, this exercise is subject to the (significant) limitation that, when studying the entire private sector, we cannot measure firm-level TFP but, rather, rely on firm-level real labor productivity. That said, the average effect for the 2011-2013 period is -1.4 percentage points, suggesting that declining responsiveness since the late 1990s has reduced economywide job reallocation rates by about 1.4 percentage points. The actual decline—again holding firm age composition constant—was about 3.0 percentage points.³⁸ Even with the limitation of using labor productivity instead of more precise "shock" measures, the decline in estimated responsiveness "accounts" for about half of the total decline in U.S. job reallocation.

B. Aggregate productivity

Our methodology for estimating aggregate reallocation effects can also be applied to estimating aggregate productivity implications of declining responsiveness. We first define an aggregate productivity index as the employment-weighted average of establishment (firm) productivity, $\sum_{j} \theta_{jt+1} a_{jt}$ (where θ_{jt+1} is an employment weight as specified previously). This index tracks the effect of changing the allocation of employment from *t* to *t+1* holding the productivity distribution in *t* constant. It is related to the type of weighted measure of productivity used by Olley and Pakes (1996) and Foster, Haltiwanger, and Krizan (2001)).³⁹ It is useful in our setting since it permits inference using a narrow counterfactual isolating the impact of declining responsiveness; in particular, we construct the effect of changing responsiveness on the aggregate productivity index for manufacturing as follows:

(11)
$$\Delta_{t+1}(Prod) = \sum_{j} \hat{\theta}_{jt+1}^{DR} a_{jt} - \sum_{j} \hat{\theta}_{jt+1}^{CR} a_{jt}$$

³⁸ Since this exercise relies on firm-level data and firm-level responsiveness coefficients, we cannot directly compare our estimated effect to the actual reallocation decline reported in Figure 1 (since that figure reports establishment-based reallocation rates, as is standard in the literature). Decker et al. (2016b) report the decline in the firm-based reallocation rate for the U.S. (see their Figure A1). To adjust this decline for changing age composition, we approximate the fraction of the firm-based reallocation decline that occurred within firm age classes using our Figure A5 (in the appendix), which provides this fraction for establishment-based reallocation.

 $^{39^{9}}$ The weighted average of establishment-level productivity tracks measured aggregate productivity well (see Figure A.1 of Decker et al. (2017)). However, conceptually it is only equal to aggregate productivity under the assumptions of constant returns to scale and perfect competition and using the composite input as weights.

where, as in our reallocation counterfactual, $\hat{\theta}_{jt+1}^{DR}$ refers to the predicted employment weight of establishment *j* given that establishment's initial employment E_{jt} , productivity a_{jt} , and actual estimated responsiveness coefficients (which decline over time, hence DR for "declining responsiveness") as reported in Table 4; while $\hat{\theta}_{jt+1}^{CR}$ refers to the predicted employment weight of that same establishment *j* if responsiveness stays constant over time at its 1980s value.

As in our reallocation counterfactual, our productivity counterfactual in (11) isolates the effect of declining responsiveness on the aggregate productivity index. $\Delta_{t+1}(Prod)$ measures the impact of changing responsiveness on the aggregate productivity index abstracting from other drivers of aggregate productivity dynamics; in Appendix I, we show that this difference-based counterfactual closely mimics the actual decline in aggregate productivity induced by rising adjustment costs in our theoretical model framework (see Figure A4).

As in our reallocation counterfactual, we construct this productivity effects measure for every year then average these effects over the 2000s. We obtain an average effect of about -2.3 log points; we interpret this as implying that, during the 2000s on average, the weighted index of establishment-level TFP is only 97.7 percent of what it would have been if responsiveness had stayed constant at its 1980s strength. To put this into perspective, according to the BLS, manufacturing TFP increased by about 30 percent from the 1980s to the 2000s; taking the difference-based counterfactual as an indicator of the actual drag on productivity, productivity would have risen by about 33 percent without the decline in responsiveness.

We also use the analogue of (11) for the U.S. private sector broadly during the 1996-99 to 2011-13 period. Again, this exercise suffers from the limitation of using revenue labor productivity for estimated responsiveness (as in our reallocation counterfactual) and, additionally, requires a productivity index based on labor productivity. With that in mind, we find an effect of -3.2 log points; this implies that for 2011-2013, the weighted index of firm-level revenue labor productivity is about 96.8 percent of what it would have been if responsiveness had stayed constant since the 1990s. According to the BLS, labor productivity increased by about 36 percent from the late 1990s to the 2011-13 period. Taking the difference-based counterfactual as an indicator of the actual drag on productivity, aggregate labor productivity would have risen by about 40 percentage points without the decline in responsiveness. We interpret these effects as nontrivial, though we emphasize that our productivity counterfactual methodology is an approximation; a more rigorous estimate would require a fully featured, carefully calibrated structural model.⁴⁰

Our reallocation and aggregate productivity exercises suggest that the decline in responsiveness we document has significant aggregate consequences. A large fraction of the decline in aggregate job reallocation can be explained by the decline in responsiveness, and the implications for aggregate productivity are nontrivial.

V. Robustness exercises

Tables 1-5 demonstrate robustness of our responsiveness results to alternative productivity measures, including innovations to productivity. We show that our changing responsiveness evidence holds up within firm age groups and is evident in a variety of time trend specifications. Our results typically have extremely high statistical significance. This robustness notwithstanding, we now briefly explore several other issues.

A. Investment responsiveness

One possible explanation for declining employment responsiveness is that businesses increasingly respond to idiosyncratic shocks by adjusting factors other than employment. This may be thought of as a capital/labor substitution mechanism. While our economywide firm dataset lacks information on factors other than labor, the manufacturing data are much richer. Here we focus on manufacturing establishments and study the responsiveness of equipment investment to productivity over time. We estimate specification (9) replacing the DHS employment growth rate with the investment rate I_{jt}/K_{jt} , where I_{jt} is establishment-level equipment investment throughout year *t*, and K_{jt} is the stock of capital equipment at the beginning of the year. Models of business dynamics with adjustment costs produce policy functions and intuition for capital investment that are similar to the policy functions and intuition for employment growth.⁴¹

⁴⁰ One approach would be to use a simulated method of moments estimation of competing driving forces in a structural model using our estimates of changes in the shock process and responsiveness as moments in the estimation. With estimated structural parameters in hand, a rigorous counterfactual could be constructed.

⁴¹ To make results comparable and minimize disclosure issues we use exactly the same specification as (9) except replacing the dependent variable. Unreported specifications where we also control for the initial capital stock each period (allowing such effects to vary over time as with other variables) produce similar results. Also in unreported exercises we find that adding the initial capital stock variable to the employment growth regressions in Table 5 does not materially affect our main results.

We estimate the investment responsiveness specification separately for high-tech and non-tech establishments, and for brevity we report only the regression in which the time functions $T^{y}(a_{jt}, t)$ and $T^{m}(a_{jt}, t)$ are specified as decade-specific coefficients (the last element in (8)). Table 6 reports the results based again on our preferred TFP measure, TFPS, standing in for a_{jt} . In the first column, among all industries we see rising investment responsiveness from the 1980s to the 1990s followed by a drop in the 2000s. This pattern is consistent with capital/labor substitution between the 1980s and 1990s while employment responsiveness was falling, after which investment joins employment in becoming less responsive in the 2000s. The non-tech column is similar.

The high-tech column is striking: while responsiveness among high-tech establishments was stronger than the rest of the manufacturing sector during the 1980s, by the 2000s it is far weaker than the rest of manufacturing and is in fact no longer statistically significant (though the difference between the 1990s and the 2000s is significant). We observe a similar pattern among both young and mature firms. In other words, productivity selection for growth via equipment investment appears completely absent in high-tech manufacturing during the 2000s. Similarly to our employment results, we can quantify investment responsiveness in terms of the differential between high-productivity establishments (those that are one standard deviation more productive than their industry mean) and average (within industry) establishments. Among high-tech young firm establishments, this differential in investment rates rose from 4.8 percentage points to 7.4 percentage points from the 1980s to the 1990s before falling close to zero in the 2000s.

In recent years researchers have given increasing focus to intangible capital (e.g., Corrado, Hulten, and Sichel (2009); Haskel and Westlake (2017)). Weakening responsiveness of employment and equipment investment could have been accompanied by changes in the responsiveness of intangible capital investment, which we cannot measure with our data. For example, high-productivity firms may be investing more in managerial capabilities (see, e.g., Bloom and Van Reenen (2007)). Investment in other forms of intangible capital, such as software or intellectual property, also may play a growing role (e.g., Crouzet and Eberly (2018)).

B. Worker hours

Another potential explanation for declining employment responsiveness is that businesses may increasingly adjust hours per worker in response to shocks, offsetting the decline in employment responsiveness. There are multiple factors that limit hours per worker as a margin of adjustment (see, e.g, the discussion in Cooper, Haltiwanger, and Willis, (2007) and citations therein). Still, we can study this adjustment margin for production worker hours in manufacturing.

For production workers in manufacturing, we have measures of both total hours and the number of production workers on a quarterly basis. This permits constructing a measure of production hours per worker.⁴² As noted above, the ASM is representative in the cross section for any given year, but it is not necessarily longitudinally representative, so the time series dimension of our measure of growth in production worker hours faces a similar limitation to our measures of TFP innovations and changes (whereas our employment growth exercises utilize the LBD universe data). With that caveat in mind, we estimate equation (7), replacing the dependent variable with the DHS growth rate of production hours per worker and limiting the sample to longitudinally covered establishments as in certain previous exercises.

The results are in Table 7, where we focus on decade-specific responsiveness coefficients and TFPS for parsimony. The responsiveness of hours per worker growth to productivity is small and does not change systematically between the 1980s and the 2000s. An establishment with a one standard deviation higher productivity than its industry mean has about a half percentage point higher growth in hours per worker in both the 1980s and the 2000s. During the 1990s, the relationship between hours per worker growth and productivity is not significant.

The change in the responsiveness of hours per worker to TFPS is much smaller than the analogous change in responsiveness of employment growth reported in Table 1. In Table 1, the responsiveness of employment growth declined by about 9 percentage points from the 1980s to the 2000s while there is no change in the responsiveness of hours per worker growth over the same period. For the change from the 1990s to the 2000s, responsiveness of employment growth declined by about 5 percentage points (Table 1), while responsiveness of hours per worker growth increased by about 1 percentage point. It is apparent that changing responsiveness on the hours per worker margin cannot mitigate the large decline in the responsiveness of employment growth.

⁴² As discussed in Appendix II, for our TFP estimates we construct a measure of total hours for all workers. However, for employment counts of all workers we only have a point-in-time measure for March. Dividing total hours by a point-in-time measure for employment would be a poor measure of hours per worker as it would reflect seasonal variation in employment. The ASM provides the number of *production* workers in March, June, August, and November; we divide annual total hours by the average of these quarterly numbers.

C. Shock mismeasurement and dispersion-dependent responsiveness

In section III.A, we presented evidence that TFP dispersion and revenue labor productivity dispersion are rising, contrary to the "shocks" hypothesis in which declining reallocation might be explained by declining shock dispersion. There are two potential questions about this result. First, some researchers have argued that the rise in productivity dispersion in manufacturing is, in some part, an artifact of increasing measurement error in the ASM establishment sample (Bils, Klenow, and Ruane (2020)). We have already shown evidence that this is not likely, since rising dispersion of revenue labor productivity is evident even in RE-LBD administrative data. However, there might still be concern that our finding of declining responsiveness reflects increasingly mismeasured productivity (e.g., from mismeasured inputs like capital). Second, in the case of non-convex adjustment costs, responsiveness varies with productivity dispersion in a manner that complicates the "shocks vs. responsiveness" dichotomy (see Appendix I).

These two issues are distinct, but both hinge on whether responsiveness is related to measured dispersion. We investigate this question using cross-industry variation in the time series of shock dispersion. In the ASM, we classify 4-digit NAICS industries into three groups based on terciles of the 1980s-2000s change in within-industry TFPS dispersion. The top tercile of industries saw a mean change of +0.11 in the standard deviation of TFPS; the middle tercile saw a mean change of +0.04, and the bottom tercile saw a mean change of -0.02. We interact dummy indicators of these terciles with our responsiveness coefficients in a regression otherwise similar to those reported in the first column of Table 1, focusing on the linear trend variable to avoid the proliferation of coefficients. These results are shown in the first column of Table 8. Focusing on the coefficients interacting productivity with tercile indicators and the trend, we do not find that the responsiveness trend is monotonic in the change in dispersion. The largest downward trend is seen in the middle tercile of industries, while the second-largest downward trend is in the bottom tercile of industries. The top tercile of industries-those with the largest gains in TFPS dispersion-see the smallest downward trend in responsiveness. We cannot reject the hypothesis that the trends in the top and bottom terciles are equal, but we (marginally) reject the hypothesis of equality between the top and middle terciles. These results provide evidence that the decline in responsiveness is not driven by changes in measured dispersion.

D. Changing shock persistence

Figure 3d shows the persistence of productivity shocks has declined from the 1980s to the 2000s. We view these changes as modest. However, in models of labor adjustment costs such as ours, a decline in shock persistence can, by itself, reduce responsiveness as firms are reluctant to pay a cost to respond to shocks that are likely to be transitory. We again exploit industry variation to explore whether declining shock persistence might explain declining responsiveness.⁴³ We divide 4-digit NAICS industries (within manufacturing) into terciles based on their change in persistence from the 1980s to the 2000s, where persistence is measured as the estimated AR(1) coefficient in the TFPS process. The top tercile has a mean change of +0.10, the middle tercile has a mean change of -0.02, and the bottom tercile has a mean change of -0.11. As in our dispersion exercises in the previous subsection, we interact tercile dummies with productivity and time trends to allow responsiveness trends to vary across the persistence terciles (and, again, we focus on the linear trend to avoid coefficient proliferation).

The results are reported in the second column of Table 8. At the industry level, the decline in responsiveness is not monotonic in the change in persistence. The largest decline in responsiveness is seen in the middle persistence tercile, followed by the bottom tercile then the top tercile. We cannot reject the hypothesis that the bottom and top persistence terciles see the same decline in responsiveness. Moreover, if declining persistence were to explain declining responsiveness, we would expect establishments in the top tercile, in which persistence actually rose, to see rising responsiveness over time, rather than the negative trend we observe. Changing shock persistence does not appear to be driving our findings on declining responsiveness.

E. Employment measurement error and division bias

Initial employment is a control in our responsiveness regressions but also enters the calculation of measured employment growth. This raises questions about division bias in the presence of measurement error in employment, which could introduce bias in our estimated responsiveness coefficients when using employment growth as the dependent variable.

⁴³ Here we note, yet again, that ASM data are not optimal for calculating longitudinal variables like persistence.

We first note that our analysis of *exit* uses a dependent variable equal to one if a plant or firm exits between t and t + 1 based on productivity and employment measures in t. Measurement error in initial employment does not cause division bias in this specification. Our results on exit show a systematic pattern of declining responsiveness.

For the specifications with the DHS growth rate as the dependent variable, we explore the potential consequences of measurement error in employment in period t by using lagged employment (employment in March of calendar year t - 1) as an instrument. This potentially mitigates measurement error division bias problems since employment in March of calendar year t - 1 does not appear in the DHS growth rate from March of calendar year t to t + 1. The potential division bias carries over to the specifications with RLP where the denominator of RLP in year t is employment in March of calendar year t from the LBD. (Recall, however, that the employment values used for constructing TFPS, TFPP, and TFPR are total hours from the ASM-CM, a different source of employment from the LBD). For the RLP regressions, we also instrument RLP in year t with lagged RLP (from year t - 1). We describe these results in Appendix III.B and Table A2. The IV results are broadly similar to our main regression results.

VI. Conclusion

Resource reallocation plays a critical role in productivity dynamics. The U.S. has seen a decline in the pace of job reallocation in recent decades that has proven difficult to understand. We study changing patterns of reallocation by drawing insight from canonical models of firm dynamics. In such models, a decline in the pace of job reallocation can arise from one of two sources: (1) the dispersion or volatility of idiosyncratic *shocks* faced by businesses—the business-specific conditions that drive hiring and downsizing decisions—could have declined; or (2) business-level *responsiveness* to idiosyncratic conditions may have weakened.

We show that shock dispersion has not declined but has actually risen, and business-level responsiveness to shocks—in terms of employment growth and survival—has weakened. Our finding of weakening responsiveness is robust to alternative productivity measures, three different time-trend specifications, and changing firm age composition. It holds for both the level of shocks and the innovations to shocks. Equipment investment has also become less responsive, at least in the last two decades, suggesting that capital/labor substitution is not a likely explanation, though

we cannot rule out substitution into intangible production factors. Hours per worker have not become more responsive. We find no evidence of a role for plausible sources of measurement error.

We further show that the decline in responsiveness can account for a large fraction of the observed decline in aggregate job reallocation: Within manufacturing, declining responsiveness can account for virtually the entire decline of reallocation since the 1980s. For the entire economy, the decline in responsiveness appears to account for about half of the decline in reallocation since the late 1990s. The decline in responsiveness we document is sufficient in magnitude to be a large driver of aggregate job reallocation trends.

Weakening business-level responsiveness to idiosyncratic productivity also has potentially significant implications for aggregate productivity since it implies weaker productivity selection. We find suggestive evidence that the decline in responsiveness has been a significant drag on aggregate productivity.

We view our results on declining responsiveness as independently significant and as substantial progress on questions about changing business dynamics in the U.S. Discovering the ultimate causes of declining responsiveness is important but beyond the scope of this paper. Our model framework suggests that changing firm-level behavior, rather than a more tranquil profitability environment, should be the focus of further research; for example, rising factor adjustment costs or changes in the relationship between revenue distortions and business-level productivity can readily produce both declining responsiveness and the rising dispersion of revenue-based productivity measures we observe. This could be the result of regulatory changes that affect the cost of hiring or downsizing; alternatively changes in the economic environment that are more interpretable in the "correlated wedges" framework could be to blame. As examples of the former, Davis and Haltiwanger (2014) focus on changes in employment-at-will doctrines in the U.S. judicial system, rising prevalence of occupational licensing, increasing use of non-competes even in sectors such as fast-food restaurants, and potential indirect factors (such as zoning) that impair Alternatively, the "correlated wedges" interpretation could geographic labor mobility. accommodate a role for rising product market power (as studied by De Loecker, Eeckhout, and Unger (2020)) or labor market power, as firms may absorb high productivity shocks through higher markups.
Our estimated negative effects of declining responsiveness on aggregate productivity are noteworthy. However, we also acknowledge that reallocation is costly. While job reallocation can facilitate job growth and job-to-job transitions—an important source of wage gains—job destruction imposes profound costs on workers, families, and communities. Optimal policies for labor and other markets must balance these considerations against the benefits of productivityenhancing reallocation for overall living standards. Identifying the source of declining responsiveness is a critical avenue for future research.

REFERENCES

- Andrews, Dan, Chiara Criscuolo, and Peter N. Gal. 2015. "Frontier Firms, Technology Diffusion and Public Policy: Micro Evidence from OECD Countries." OECD: *The Future of Productivity: Main Background Papers*.
- Autor, David, David Dorn, Lawrence Katz, Christina Patterson, and John Van Reenen. 2019. Forthcoming. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics*.
- Bachmann, Rüdiger and Christian Bayer. 2013. "Wait-and-See Business Cycles." *Journal* of Monetary Economics 60(6): 704-19.
- Baily, Martin Neil, Charles Hulten, and David Campbell. 1992. "Productivity Dynamics in Manufacturing Plants." In *Brookings Papers on Economic Activity: Microeconomics*, ed. Clifford Winston and Martin Neil Baily. Washington, DC: Brookings Institution Press.
- Barro, Robert J. 1972. "A Theory of Monopolistic Price Adjustment." *The Review of Economic Studies* 39(1): 17-26.
- Barth, Erling, Alex Bryson, James C. Davis, and Richard Freeman. 2016. "It's Where You Work: Increases in the Dispersion of Earnings Across Establishments and Individuals in the United States." *Journal of Labor Economics* 34 no. 2 pt. 2: S67-S97.
- Berger, David and Joseph Vavra. 2019. "Shocks vs. Responsiveness: What Drives Time-Varying Dispersion," *Journal of Political Economy* 127 no. 5:2104-2142.
- Bils, Mark, Peter J. Klenow, and Cian Ruane. 2020. "Misallocation or Mismeasurement?" NBER Working Paper no. 27611.
- Blackwood, Glenn, Lucia Foster, Cheryl Grim, John Haltiwanger and Zoltan Wolf. Forthcoming. "Macro and Micro Dynamics of Productivity: From Devilish Details to Insights." *American Economic Journal: Macroeconomics*.

- Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 no. 3:623-85.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica* 86 no. 3:1031-65.
- Bloom, Nicholas and John Van Reenen. 2007. "Measuring and Explaining Management Practices Across Firms and Countries." *The Quarterly Journal of Economics* 122 no. 4:1351-1408.
- Caballero, Ricardo, Eduard Engel and John Haltiwanger. 1997. "Aggregate Employment Dynamics: Building Up from Microeconomic Evidence" *American Economic Review*, 87 no. 1, 115-137.
- Census Bureau. 2015a. "Longitudinal Business Database (LBD)", 1976-2014." Federal Statistical Research Data Center.
- Census Bureau. 2015b. "Revenue-Enhanced Longitudinal Business Database (RE-LBD)", 1996-2014." Federal Statistical Research Data Center.
- Census Bureau. 2015c. "Census of Manufactures, 1972-2012" Federal Statistical Research Data Center.
- Census Bureau. 2015d. "Annual Survey of Manufactures, 1972-2014" Federal Statistical Research Data Center.
- 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.
- Corrado, Carol, Charles Hulten, and Daniel Sichel. 2009. "Intangible Capital and Economic Growth." *Review of Income and Wealth* 55 no. 3:661-685.
- Crouzet, Nicolas and Janice Eberly. 2018. "Intangibles, Investment, and Efficiency." American Economic Review Papers and Proceedings 108 no. 5:426-31.
- Davis, Steven J. and John Haltiwanger. 2014. "Labor Market Fluidity and Economic Performance." NBER Working Paper 20479 (published in the 2014 Federal Reserve Bank of Kansas City Jackson Hole Conference Symposium).
- Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda. 2007. "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms." Chap. 2 in *NBER Macroeconomics Annual 2006* edited by Daron

Acemoglu, Kenneth Rogoff, and Michael Woodford, vol. 21, pp. 107-80. Cambridge, MA: MIT Press.

- Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. Job Creation and Destruction. Cambridge, MA: MIT Press.
- De Loecker, Jan, Jan Eeckhout and Gabriel Unger. 2020. "The Rise of Market Power and the Macroeconomic Implications." *The Quarterly Journal of Economics* 135 no. 2:561-644.
- De Loecker, Jan, Pinelopi K. Goldberg, and Amit K. Khandelwal. 2016. "Prices, Markups, and Trade Reform." *Econometrica* 84 no. 2:445-510.
- 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. "Declining Business Dynamism: Implications for Productivity?" Hutchins Center Working Paper #23.
- 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).
- Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2018. "Changing Business Dynamism and Productivity: Shocks vs. Responsiveness." NBER Working Paper no. 24236.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2018. "How costly are markups?" *NBER Working Paper* no. 24800.
- Elsby, Michael W. L. and Ryan Michaels. 2013. "Marginal Jobs, Heterogeneous Firms, and Unemployment Flows." *American Economic Journal: Macroeconomics* 5 no. 1:1-48.
- Eslava, Marcela and John Haltiwanger. 2020. "The Life-Cycle Growth of Plants: The Role of Productivity, Demand and Wedges." NBER Working Paper No. 27184.
- 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* 34 no. S1 (Part 2, January 2016):S293-S331.
- 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 Chad Syverson, 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review*, 98(1): 394-425.
- Gopinath, Gita, Sebnem Kalemli-Ozcan, Loukas Karabarbounis and Carolina Villegas-Sanchez 2017. "Capital Allocation and Productivity in South Europe." *Quarterly Journal of Economics* 132 no. 4: 1915-1967.
- Guzman, Jorge and Scott Stern. 2016. "The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 15 US States, 1988-2014." *Mimeo*.
- Haltiwanger, John, Ian Hathaway, and Javier Miranda. 2014. "Declining Business Dynamism in the U.S. High-Technology Sector." The Kauffman Foundation.
- Haltiwanger, John, Ron S. Jarmin, Robert Kulick and Javier Miranda. 2017. "High Growth Young Firms: Contribution to Job Growth, Revenue Growth and Productivity" In *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges* edited by John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar. Chicago: University of Chicago Press.
- Haltiwanger, John, Ron S. Jarmin, and Javier Miranda. 2013. "Who Creates Jobs? Small versus Large versus Young." *The Review of Economics and Statistics* 95 no. 2: 347-361.
- Haskel, Jonathan and Stian Westlake. 2017. *Capitalism without Capital: The Rise of the Intangible Economy*. Princeton University Press.
- Hecker, Daniel. 2005. "High-technology employment: a NAICS-based update," *Monthly Labor Review* (U.S. Dept. of Labor and U.S. Bureau of Labor Statistics), Volume 128, Number 7, July 2005: 58.
- Hugo A. Hopenhayn. 1992. "Entry, Exit, and firm Dynamics in Long Run Equilibrium." *Econometrica*, Vol. 60(5): 1127-50.

- Hopenhayn, Hugo and Richard Rogerson. 1993. "Job Turnover and Policy Evaluation: A General Equilibrium Analysis." *Journal of Political Economy* 101(5): 915-38.
- Hottman, Collin, Stephen J. Redding and David E. Weinstein, 2016. "Quantifying the Sources of Firm Heterogeneity," *The Quarterly Journal of Economics* 131(3): 1291-1364.
- 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.
- Ilut, Cosmin L., Matthias Kehrig, and Martin Schneider. 2018. "Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News." *Journal of Political Economy* 126 no. 5:2011-71.
- Jovanovic, Boyan. 1982. "Selection and the evolution of industry." *Econometrica*, 50(3): 649–70.
- Karahan, Fatih, Benjamin Pugsley, and Aysegul Sahin. 2018. "Demographic Origins of the Startup Deficit." Mimeo.
- Kehrig, Matthias and Nicolas Vincent. 2017. "Growing productivity without growing wages: The micro-level anatomy of the aggregate labor share decline," Discussion Paper 244, Economic Research Initiatives at Duke (ERID) Working Paper.
- Kehrig, Matthias and Nicolas Vincent. 2020. "The micro-level anatomy of the labor share decline," NBER Working Paper No. 25275, March.
- 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.
- Pugsley, Benjamin, Petr Sedlacek, and Vincent Sterk. 2017. "The Nature of Firm Growth." Discussion Papers 1737, Centre for Macroeconomics (CFM).
- 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.
- Vavra, Joseph. 2014. "Inflation Dynamics and Time-Varying Volatility: New Evidence and SS Interpretation." *Quarterly Journal of Economics* 129(1): 215-58.

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



FIGURE 1. JOB REALLOCATION PATTERNS DIFFER BY SECTOR

Notes: HP trends using parameter set to 100. Industries defined on a consistent NAICS basis; high-tech is defined as in Hecker (2005). Data include all firms (new entrants, continuers, and exiters). Source: LBD.



FIGURE 2. THE SHOCKS AND RESPONSIVENESS HYPOTHESES, MODEL RESULTS (NON-CONVEX COST)

Notes: Panels c and d share same legend. Results relative to model baseline calibration (vertical purple line) with downward adjustment cost $F_{=0}$ and TFP dispersion $\sigma_A = 0.46$ (see Appendix I and Table A1 for model calibration details). "s.d. RLP" refers to the standard deviation of revenue labor productivity in model-simulated data.



FIGURE 3. WITHIN-INDUSTRY PRODUCTIVITY DISPERSION HAS RISEN

Notes: Dispersion measures refer to standard deviation of within-industry (log) productivity. Panels a, c, and d share same legend. Persistence measures refer to AR(1) parameter. Source: ASM-CM (panels a, c, and d); RE-LBD (panel b).



FIGURE 4. JOB GROWTH AND EXIT HAVE BECOME LESS RESPONSIVE TO PRODUCTIVITY

Notes: Compares employment growth rate or (inverse) exit probability of establishment (panel a) or firm (panel b) that is one standard deviation above its industry-year mean productivity, versus the mean. Source: ASM-CM (panel a); RE-LBD (panel b).



FIGURE 5. EMPLOYMENT GROWTH RESPONSIVENESS: YOUNG VS. MATURE FIRMS, HIGH-TECH VS. NON-TECH

Notes: Compares employment growth rate of establishment (panels a, b) or firm (panels c, d) that is one standard deviation above its industry-year mean productivity, versus the mean. Source: ASM-CM (panels a, b); RE-LBD (panels c, d).

	TABLE 1-	-BUSINESS-LEV	EL EMPLOYMEN	T GROWTH RESPO	NSIVENESS HAS	WEAKENED	
Panel A							
	TFPS	(revenue share	based)	TF	TFPP (proxy method)		
Productivity: β_1	0.2965	0.2905		0.2086	0.1876		
	(0.0097)	(0.0068)		(0.0086)	(0.0060)		
Prod*trend: δ	-0.0035			-0.0043			
	(0.0005)			(0.0004)			
Prod*post-97: γ ₉₇		-0.0952			-0.0958		
		(0.0084)			(0.0074)		
Prod*1980s: λ_{80s}			0.2859			0.2185	
			(0.0095)			(0.0086)	
Prod*1990s: λ _{90s}			0.2462			0.1053	
			(0.0060)			(0.0052)	
Prod*2000s: λ_{00s}			0.2001			0.0995	
			(0.0059)			(0.0051)	
p value: $\lambda_{80s} = \lambda_{90s}$			0.00			0.00	
p value: $\lambda_{80s} = \lambda_{00s}$			0.00			0.00	
p value: $\lambda_{90s} = \lambda_{00s}$			0.00			0.43	
Obs. (thousands)	2,375	2,375	2,375	2,375	2,375	2,375	
Panel R							
T unet D	TFPR (cost share based)		RLP (revenue per worker)		orker)		
Productivity: B_1	0.2040	0.1779		0.2762	()	
<i>J F</i> 1	(0.0094)	(0.0067)		(0.0003)			
Prod*trend: δ	-0.0046	(00000)		-0.0029			
	(0.0005)			(0.0000)			
Prod*post-97: Voz	(0.0000)	-0.0981		(0.000)			
Poor > / / 9/		(0.0081)					
Prod*1980s: loc		(0.0001)	0.1939				
1100 19000. 1805			(0.0094)				
Prod*1990s. Jos			0 1212				
1100 17703. Agus			(0.0058)				
Prod*2000s. Ja			0.0820				
1100 20005. A _{00s}			(0.0020)				
			(0.0034)				
p value: $\lambda_{80s} = \lambda_{90s}$			0.00				
p value: $\lambda_{80s} = \lambda_{00s}$			0.00				
p value: $\lambda_{90s} = \lambda_{00s}$			0.00				
Obs. (thousands)	2,375	2,375	2,375	58,700			

Notes: Dependent variable is annual employment growth. All coefficients are statistically significant with p < 0.01. TFPS, TFPP, and TFPR columns are establishment regressions in manufacturing for 1981-2013. RLP columns are economywide firm regressions for 1997-2013. All regressions include controls described in equation (7) and related text.

		TABLE 2—BUSI	NESS-LEVEL EXI	T RESPONSIVENES	SS HAS WEAKEN	ED
Panel A		<i>.</i> .	1 5			0
	TFPS	(revenue share	based)	TF	PP (proxy meth	nod)
Productivity: β_1	-0.0757	-0.0801		-0.0830	-0.0781	
	(0.0043)	(0.0030)		(0.0038)	(0.0027)	
Prod*trend: δ	0.0009			0.0014		
	(0.0002)			(0.0002)		
Prod*post-97: γ_{97}		0.0340			0.0352	
		(0.0037)			(0.0033)	
Prod*1980s: λ _{80s}			-0.0773			-0.0868
			(0.0042)			(0.0038)
Prod*1990s: λ _{90s}			-0.0586			-0.0473
			(0.0026)			(0.0022)
Prod*2000s: λ_{00s}			-0.0517			-0.0478
			(0.0026)			(0.0023)
p value: $\lambda_{80s} = \lambda_{90s}$			0.00			0.00
p value: $\lambda_{80s} = \lambda_{00s}$			0.00			0.00
p value: $\lambda_{90s} = \lambda_{00s}$			0.06			0.89
Obs. (thousands)	2,375	2,375	2,375	2,375	2,375	2,375
Panel B						
	TFP	R (cost share b	ased)	RLP	(revenue per w	orker)
Productivity: β_1	-0.0721	-0.0664		-0.0857	•	,
2 7 1	(0.0042)	(0.0030)		(0.0001)		
Prod*trend: δ	0.0014			0.0007		
	(0.0002)			(0.0000)		
Prod*post-97: Voz	(0.000-)	0.0330		(0.000)		
F / 9/		(0, 0036)				
Prod*1980s. Jac		(0.0020)	-0.0714			
1100 19005. Mg0s			(0, 0042)			
Prod*1000s. A.			-0.0430			
1100 17703. Ag _{0s}			(0.0025)			
Drod*2000a. 1			0.025)			
Prod 2000s: λ_{00s}			-0.0370			
			(0.0024)			
p value: $\lambda_{80s} = \lambda_{90s}$			0.00			
p value: $\lambda_{80s} = \lambda_{00s}$			0.00			
p value: $\lambda_{90s} = \lambda_{00s}$			0.09			
Obs. (thousands)	2,375	2,375	2,375	58,700		

Notes: Dependent variable is a binary establishment or firm exit indicator. All coefficients are statistically significant with p < 0.01. TFPS, TFPP, and TFPR columns are establishment regressions in manufacturing for 1981-2013. RLP columns are economywide firm regressions for 1997-2013. All regressions include controls described in equation (7) and related text.

	Innovation η_{jt}	Change (Δa_{jt})
Prod*1980s: λ_{80s}	0.3970	0.3318
003	(0.0279)	(0.0254)
Prod*1990s: λ_{90s}	0.3909	0.3334
	(0.0124)	(0.0117)
Prod*2000s: λ_{00s}	0.2999	0.2513
	(0.0126)	(0.0118)
p value: $\lambda_{80s} = \lambda_{90s}$	0.84	0.96
p value: $\lambda_{80s} = \lambda_{00s}$	0.00	0.00
p value: $\lambda_{90s} = \lambda_{00s}$	0.00	0.00
Observations (thousands)	854	854

Notes: First (second) column shows regression of employment growth on TFPS innovation (first difference) and controls as described in equation (7) and related text. All coefficients are statistically significant with p < 0.01.

Source: LBD, ASM-CM, and author calculations.

TABLE 4-GROWTH AND EXIT RESPONSIVENESS HAVE WEAKENED WITHIN FIRM AGE GROUPS

	Panel A—dependent variable:			Panel B—dependent variable: Exit between t and t ± 1		
		oymeni growi DS	$n y_{jt+1}$		Delween i and	
Prod*Voung: Ry	0 4069	13	0.3217	-0.1075	15	-0.1065
1100 10011g. p_1	(0.0137)		(0.0003)	(0.0060)		(0,0001)
Prod*Voung*trend· δ^{y}	-0.0054		-0.0034	0.0014		0.0009
fille found field.	(0.0006)		(0,0000)	(0,0003)		(0,0000)
Prod*Mature: Rm	0.2722		0.2493	-0.0690		-0.0747
From Wature: p_1	(0.0006)		(0.003)	(0, 0042)		(0,0001)
Prod*Matura*trand Sm	0.0020		0.0003	0.0072		0.0001
Prod Mature trend: 0	-0.0029		-0.0024	(0.0007)		(0.0003)
	(0.0005)		(0.0000)	(0.0002)		(0.0000)
Prod*Young*1980s: λ_{80s}^{y}		0.3666			-0.1020	
		(0.0136)			(0.0059)	
Prod*Young*1990s: λ_{pos}^{y}		0.3603			-0.0898	
0 905	(0.0092)			(0.0039)		
Prod*Young*2000s: λ_{000}^{y}		0.2542			-0.0689	
8 003		(0.0093)			(0.0039)	
Prod*Mature*1980s: λ_{soc}^{m}		0.2710			-0.0727	
003		(0.0094)			(0.0042)	
Prod*Mature*1990s: λ_{00c}^m		0.2185			-0.0529	
		(0, 0058)			(0, 0025)	
Prod*Mature*2000s: λ_{m}^{m}		0 1941			-0.0507	
1100 Mature 20005. 7005		(0.0059)			(0.0026)	
n value: $\lambda^{y} = \lambda^{y}$		0.67			0.06	
p value: $\lambda_{80s}^{y} = \lambda_{90s}^{y}$		0.00			0.00	
p value: $\lambda_{80s}^y = \lambda_{00s}^y$ p value: $\lambda_{90s}^y = \lambda_{00s}^y$		0.00			0.00	
p value: $\lambda_{80s}^m = \lambda_{90s}^m$		0.00			0.00	
p value: $\lambda_{B0s}^m = \lambda_{00s}^m$		0.00			0.00	
p value: $\lambda_{90s}^m = \lambda_{00s}^m$		0.00			0.55	
Observations (thousands)	2 3 7 5	2 375	58 700	2 3 7 5	2 375	58 700

Notes: All coefficients are statistically significant with p < 0.01. TFPS columns are establishment-level regressions in manufacturing data for 1981-2013. RLP columns are firm-level regressions on economywide data for 1997-2013. All regressions include controls described in equation (7) and related text. Young firms have age less than five.

	TFPS		R	RLP		
	Non	-tech	Hig	h-tech	Non-tech	High-tech
Prod*Young: β_1^{γ}	0.4171		0.3195		0.3207	0.3436
0 7 1	(0.0145)		(0.0387)		(0.0003)	(0.0015)
Prod*Young*trend: δ^{y}	-0.0054		-0.0058		-0.0033	-0.0054
	(0.0007)		(0.0020)		(0.0000)	(0.0001)
Prod*Mature: β_1^m	0.2786		0.1969		0.2481	0.2813
	(0.0101)		(0.0302)		(0.0003)	(0.0014)
Prod*Mature*trend: δ^m	-0.0030		-0.0022ª		-0.0023	-0.0044
	(0.0005)		(0.0016)		(0.0000)	(0.0001)
Prod*Young*1980s: λ_{80s}^{y}		0.3801		0.2488		
		(0.0144)		(0.0371)		
Prod*Young*1990s: λ_{90s}^{y}		0.3658		0.3190		
- ,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,,		(0.0097)		(0.0331)		
Prod*Young*2000s: λ_{00s}^{y}		0.2637		0.1575		
- 000		(0.0095)		(0.0332)		
Prod*Mature*1980s: λ_{80s}^{m}		0.2792		0.1657		
		(0.0099)		(0.0293)		
Prod*Mature*1990s: λ_{90s}^m		0.2212		0.1915		
		(0.0061)		(0.0176)		
Prod*Mature*2000s: λ_{00s}^{m}		0.1964		0.1340		
		(0.0060)		(0.0216)		
p value: $\lambda_{80s}^{y} = \lambda_{90s}^{y}$		0.36		0.14		
p value: $\lambda_{80s}^{y} = \lambda_{00s}^{y}$		0.00		0.06		
p value: $\lambda_{90s}^{y} = \lambda_{00s}^{y}$		0.00		0.00		
p value: $\lambda_{80s}^m = \lambda_{90s}^m$		0.00		0.45		
p value: $\lambda_{80s}^m = \lambda_{00s}^m$		0.00		0.38		
p value: $\lambda_{90s}^m = \lambda_{00s}^m$		0.00		0.04		
Observations (thousands)	2,239	2,239	136	136	55,870	2,826

Dependent variable: employment growth (g_{t+1}) from t to t + 1

Notes: High-tech industries defined as in Hecker (2005). All coefficients statistically significant with p < 0.01 unless otherwise noted. TFPS columns are establishment-level regressions in manufacturing data for 1981-2013. RLP columns are firm-level regressions on economywide data for 1997-2013. All regressions include controls described in equation (7) and related text. Young firms have age less than five.

a. Not statistically significant.

Source: LBD, ASM-CM, and author calculations.

TABLE 6-INVESTMENT RATE RESPONSIVENESS HAS ALSO WEAKENED (MANUFACTURING)

	Dependent variable: capital investment rate				
	All industries	Non-tech	High-tech		
Prod*Young*1980s: λ_{80s}^{y}	0.0670	0.0668	0.1005		
	(0.0098)	(0.0103)	(0.0330)		
Prod*Young*1990s: λ_{90s}^{y}	0.1768	0.1785	0.1550		
205	(0.0139)	(0.0147)	(0.0312)		
Prod*Young*2000s: λ_{00s}^{y}	0.1003	0.1048	0.0133ª		
	(0.0092)	(0.0086)	(0.0281)		
Prod*Mature*1980s: λ_{80s}^m	0.0414	0.0393	0.0746		
	(0.0068)	(0.0071)	(0.0283)		

Prod*Mature*1990s: λ_{90s}^m	0.1151	0.1144	0.1294
200	(0.0082)	(0.0087)	(0.0200)
Prod*Mature*2000s: λ_{00s}^m	0.0619	0.0658	-0.0020ª
	(0.0051)	(0.0053)	(0.0194)
p value: $\lambda_{\text{ROC}}^{y} = \lambda_{\text{ROC}}^{y}$	0.00	0.00	0.22
p value: $\lambda_{80s}^{y} = \lambda_{00s}^{y}$	0.00	0.00	0.04
p value: $\lambda_{90s}^{y} = \lambda_{00s}^{y}$	0.00	0.00	0.00
p value: $\lambda_{80s}^m = \lambda_{90s}^m$	0.00	0.00	0.12
p value: $\lambda_{80s}^m = \lambda_{00s}^m$	0.02	0.00	0.03
p value: $\lambda_{90s}^m = \lambda_{00s}^m$	0.00	0.00	0.00
Observations (thousands)	2.375	2.239	136

Notes: Equipment investment. Manufacturing only with TFPS productivity concept. High-tech industries defined as in Hecker (2005). All coefficients statistically significant with p < 0.05 (and almost always with p < 0.01) unless otherwise noted. All regressions include controls described in equation (7) and related text.

a. Not statistically significant.

Source: LBD, ASM-CM, and author calculations.

TABLE 7_	-RESPONSIVENESS OF HOURS PER WORKER	GROWTH IS SMALL AND LITTLE CHANGED
IADLL /	RESI ONSIVERESS OF HOORS FER WORREN	CORO W THE DEMALE AND ENTILE CHANGED

	TFPS
Prod*1980s: λ_{80s}	0.0110
	(0.0044)
Prod*1990s: λ_{90s}	0.0012ª
	(0.0035)
Prod*2000s: λ_{00s}	0.0102
	(0.0028)
p value: $\lambda_{80s} = \lambda_{90s}$	0.08
p value: $\lambda_{80s} = \lambda_{00s}$	0.87
p value: $\lambda_{90s} = \lambda_{00s}$	0.04
Observations (thousands)	854

Notes: Regression of hours per worker growth on TFPS and controls as described in section V.B. Coefficients are statistically significant with p < 0.01 unless otherwise noted.

a. Not statistically significant.

Source: LBD, ASM-CM, and author calculations.

 $TABLE\ 8---RESPONSIVENESS\ DOES\ NOT\ VARY\ SYSTEMATICALLY\ WITH\ INDUSTRY\ SHOCK\ DISPERSION\ OR\ PERSISTENCE$

	Responsiveness by dispersion tercile ^a	Responsiveness by persistence tercile ^b
Prod*BottomTercile	0.2792	0.2771
	(0.0120)	(0.0112)
Prod*MiddleTercile	0.3372	0.3338
	(0.0144)	(0.0139)
Prod*TopTercile	0.2782	0.2659
	(0.0160)	(0.0153)
Prod*BottomTercile*trend	-0.0036	-0.0032
	(0.0006)	(0.0006)
Prod*MiddleTercile*trend	-0.0045	-0.0045
	(0.0008)	(0.0007)

Prod*TopTercile*trend	-0.0026 (0.0007)	-0.0026 (0.0008)
p value: Prod*BottomTercile = Prod*MiddleTercile	0.00	0.00
p value: Prod*MiddleTercile = Prod*TopTercile	0.00	0.00
p value: Prod*BottomTercile = Prod*TopTercile	0.95	0.50
p value: Prod*Bottom*trend = Prod*Middle*trend	0.32	0.09
p value: Prod*Middle*trend = Prod*Top*trend	0.05	0.05
p value: Prod*Bottom*trend = Prod*Top*trend	0.25	0.47
Observations (thousands)	2,375	2.375

Notes: All regression coefficients are statistically significant with p < 0.01. Bottom, middle, and top terciles refer to the change in TFPS dispersion (first column) or persistence (second column), where top is the highest change and bottom is the lowest change.

a. The mean industry-level changes in dispersion are -0.02, +0.04, and +0.11 for the bottom, middle, and top terciles, respectively.

b. The mean industry-level changes in persistence are -0.11, -0.02, and +0.10 for the bottom, middle, and top terciles, respectively.

Online Appendix for

"Changing Business Dynamism and Productivity: Shocks vs. Responsiveness"1

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

Appendix I: Model framework

A. General formulation

Consider a class of models in which revenue of firm² *j* in time *t* is given by $A_{jt}E_{jt}^{\phi}$, where A_{jt} is a composite shock reflecting both technical efficiency and, potentially, demand shocks, E_{jt} is employment, and $\phi < 1$ reflects revenue function curvature arising from imperfect competition due to, for example, product differentiation (related arguments go through for decreasing returns to scale). Suppose the shock A_{jt} follows the process $\ln A_{jt} = \rho_a \ln A_{jt-1} + \eta_{jt}$. This setup is common to a wide range of models of firm dynamics and typically gives rise to an employment growth policy function given by:

$$g_{jt} = f_t \left(A_{jt}, E_{jt-1} \right) \tag{A1}$$

where g_{jt} is employment growth from t - 1 to t; this is the same as equation (1) in the main text. It is commonly the case that $\frac{\partial f}{\partial A} > 0$; that is, for any two firms with the same employment, the firm with higher A has higher growth. For empirical purposes, (A1) leads to the following log-linear approximation:

$$g_{jt} = \beta_0 + \beta_1 a_{jt} + \beta_2 e_{jt-1} + \varepsilon_{jt} \tag{A2}$$

While (A1) and its empirical counterpart (A2) are quite general, it is useful to illustrate the employment growth function using a special case of a simple model that is free of frictions or distortions (which we will add below). In this case, the firm's first-order condition, in logs (indicated by lowercase), is given by:

$$e_{jt} = \frac{1}{1 - \phi} \left(\ln \frac{\phi}{W_t} + a_{jt} \right) \tag{A3}$$

¹ All of the code used to produce the results in the paper can be found at openicpsr-120432.

² We use the term "firm" for expositional purposes; in model exercises we do not distinguish between firms and establishments. Our empirical exercises using TFP measures and manufacturing data rely on establishments, while our economywide exercises using RLP rely on firms.

where W_t is the industry wage. Taking time differences (indicated by Δ) and sweeping out year and industry effects yields the firm-level growth rate (measured as log first differences for convenience):

$$\Delta e_{jt} = \frac{1}{1 - \phi} \Delta a_{jt} \tag{A4}$$

Equation (A4) provides an employment growth function that is different from its expression in (A1); in particular, (A4) expresses employment growth as a function of the *change* in a_{jt} , which is intuitive in this frictionless environment (note also the importance of revenue function curvature parameter ϕ). However, (A4) can be transformed to express employment growth as a function of the productivity *level* instead. To see this, we start with (A3), consider it for t - 1, and invert it to express productivity in terms of employment:

$$a_{jt-1} = (1-\phi)e_{jt-1} - \ln\frac{\phi}{W_{t-1}}$$
(A5)

Substituting (A5) into (A4) (and, again, sweeping out industry and year effects) yields:

$$\Delta e_{jt} = \frac{1}{1 - \phi} a_{jt} - e_{jt-1}, \tag{A6}$$

That is, employment growth can be expressed as a function of the *level* of a_{jt} , as well as the level of e_{jt-1} , as in (A1) and (A2). This is useful for two reasons. First, as noted in the text, it is convenient to specify the growth function in terms of productivity levels for empirical purposes, since productivity data in manufacturing are constructed to be representative in the cross section but not necessarily longitudinally. Second, in models with labor adjustment costs (such as the one we will describe below), the productivity level is the relevant state variable arising from the firm value function.

We now turn to two illustrative special cases of the general model framework that can motivate (A1) and (A2): a model with labor adjustment costs, and a model with static distortionary wedges that are correlated with fundamentals. We explore these models to demonstrate how (A1) arises from firm optimization problems and how it is affected by model parameters and frictions or distortions on employment demand decisions.

B. Model with labor adjustment costs

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

$$V(E_{jt-1}, A_{jt}) = \max\{A_{jt}E_{jt}^{\phi} - W_t E_{jt} - C(H_{jt}, E_{jt-1}) + \beta \mathbb{E}_A(E_{jt}, A_{jt+1}\}$$
(A7)

with

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

where $\phi < 1$ due to product differentiation such that $A_{jt}E_{jt}^{\phi}$ is the revenue function for firm j, E_{jt} is employment for time t, H_{jt} is net hires made at the beginning of time t such that $H_{jt} = E_{jt} - E_{jt-1}$ (this can be positive or negative), W_t is the wage, and A_{jt} is a composite shock reflecting both technical efficiency and demand shocks. We interpret the revenue function curvature as reflecting product differentiation rather than decreasing returns to help draw out the potential relationship between revenue productivity and technical efficiency when prices are endogenous. That is, let firm-level price be given by $P_{jt} = D_{jt}Q_{jt}^{\phi-1}$, where $Q_{jt} = \tilde{A}_{jt}E_{jt}$ is firm-level output subject to a constant returns technology, with $A_{jt} = D_{jt}\tilde{A}_{jt}^{\phi}$. That is, A_{jt} is what we refer to as "TFP" in the main text and reflects both technical efficiency and demand shocks, both in the conceptual framework and empirical analysis. Since labor is the only production factor, TFPR and revenue labor productivity (RLP) are both given by $P_{jt}\tilde{A}_{jt}$. Note that in the alternative pricetaking version of the model (where $\phi = 1$), TFP, TFPR, and RLP are equivalent. We focus on the $\phi < 1$ case in our calibration. We also abstract from demand shocks for clarity of exposition (i.e., $D_{jt} = 1 \forall j, t$) in the remaining discussion. Our TFP shocks should be interpreted as reflecting the type of composite shocks we consider empirically.

This simple adjustment cost model is similar to Cooper, Haltiwanger, and Willis (2007, 2015), Elsby and Michaels (2013), and Bloom et al. (2018) and, in principle, accommodates both convex and non-convex adjustment costs. In particular, given the cost function $C(H_{jt})$, which depends upon E_{jt-1} , the policy rule for hiring depends on the initial *state* faced by the firm, which is summarized as (E_{jt-1}, A_{jt}) .

We view the model as primarily illustrative but seek a reasonable baseline calibration that matches key features of the data and the parameters of the existing literature. Appropriate caution is needed since we do not model entry or exit, and we do not have any lifecycle learning dynamics or frictions that make young firms different from more mature firms. We regard the calibration as providing guidance about the qualitative predictions for the key data moments we study but within a reasonable range of the parameter space.

Our main calibration exercise, described in detail below, implements "general equilibrium" in the sense that we fix the labor supply then find the wage that clears the labor market. Given a rigid labor supply, this may be thought of as an extreme scenario. However, in unreported exercises we consider the opposite extreme in which labor supply is perfectly elastic and the wage is fixed (i.e., partial equilibrium). A limitation of the partial equilibrium exercise is that when the wage is fixed, adjustment frictions can have large effects on average firm size and therefore productivity via channels that are unrelated to reallocation. However, our key results on how adjustment costs affect reallocation rates, firm-level productivity responsiveness, and the effect of changing responsiveness on aggregate productivity growth do not substantively depend on general versus partial equilibrium.

Our method for solving the model is as follows. We create a state space for employment, with 2,400 points (distributed more densely at lower values), and for TFP realizations, with 115 points. We specify firm-level TFP to follow an AR(1) process, $\ln A_{jt} = \rho_a \ln A_{jt-1} + \eta_{jt}$, and in practice we use a Tauchen (1986) method for generating TFP draws. Table A1 reports our calibration choices, some of which are standard in the literature and others of which are designed to target specific data moments. We describe two alternative adjustment cost specifications: kinked adjustment costs (as described in the main text) and convex adjustment costs. We start with the kinked adjustment cost case. Empirically determined calibration choices are intended to produce a model economy that resembles the U.S. manufacturing sector in the 1980s, the initial timing of our empirical exercises, but the qualitative model results in which we are interested are not sensitive to these specific calibration choices.

We obtain policy functions using value function iteration then simulate 2,000 firms for 1,000 periods, jumping off from the stationary distribution of productivity but discarding the first 100 periods. Given a fixed (inelastic) labor supply, we check market clearing then adjust the wage using simple bisection until the labor market clears. We estimate responsiveness regressions and construct other statistics described in the text by using the simulated data generated by the model when in equilibrium.

We perform several exercises on the model-simulated data with a focus on three key outcomes: aggregate job reallocation, the dispersion of revenue productivity (where in the model, revenue productivity is given by $A_{jt}E_{jt}^{\phi-1}$), and the responsiveness of growth to productivity as measured with the regression in equation (2) of the main text. In other words, we measure the standard deviation of labor productivity in the model economy, and we estimate the regression from (A2), that is,

$$g_{jt} = \beta_0 + \beta_1 a_{jt} + \beta_2 e_{jt-1} + \varepsilon_{jt}, \tag{A8}$$

where, as in the main text, g_{jt} is DHS employment growth from year t - 1 to year t, a_{jt} is productivity, and e_{jt-1} is (initial) employment. This is the same as equation (A2) and follows a timing convention that is analogous to our empirical work (though we confirm below that this timing convention is unimportant for the model's qualitative results). "Responsiveness" is measured by β_1 .

We study labor productivity dispersion and responsiveness under two model experiments starting from the model's baseline calibration. In our first experiment, we study the effects of declining responsiveness, in this case resulting from a rise in adjustment costs. In particular, starting with the baseline calibration (where upward adjustment has a cost parameter of F_+ = 1.03) we raise the cost of downward adjustments (F_-). Figure 2 in the main text shows the result of this experiment. Rising adjustment costs generate declining reallocation (Figure 2a) due to lower responsiveness (2c), with the additional result of wider labor productivity dispersion, each of which we observe in our empirical exercises. This experiment suggests that declining responsiveness, as generated by rising labor adjustment costs, can cause declining reallocation, with the additional symptom of rising labor productivity dispersion.

In our second experiment, we reduce the parameter governing TFP dispersion, starting from its baseline calibrated value of $\sigma_a = 0.46$. This is also reported in Figure 2 in the main text. As TFP dispersion falls, aggregate job reallocation declines (Figure 2b), labor productivity dispersion decreases, and responsiveness becomes weaker (Figure 2d; we discuss this more below). This summarizes the "shocks" hypothesis: the declining pace of job reallocation we observe empirically could be explained by declining dispersion of TFP realizations if we were to also observe declining labor productivity dispersion. As shown in the main text, however, we actually observe rising labor productivity dispersion in our empirical work.

We must make one side note here: As noted above and shown in Figure 2c, in the model with non-convex adjustment costs, when shock dispersion declines, so too does responsiveness. At first glance, this dispersion dependence of responsiveness in the non-convex costs model may

complicate the shocks vs. responsiveness hypothesis. However, three points are important to note. First, this is unique to the model with non-convex costs; as we will discuss below (and show in Figures A1 and A2), responsiveness is unaffected by changes in dispersion when adjustment costs are convex, or in the correlated wedges model without adjustment costs. Second, we can easily conclude that the declining responsiveness we observe in the data is not driven by declining shock dispersion because we also empirically find rising shock dispersion (apparent in our TFPS and TFPP productivity measures) and rising revenue productivity dispersion (apparent in our TFPR and revenue labor productivity measures). Third, as we show in the main text, using industry variation we find no monotonic relationship between changes in TFP dispersion and changes in responsiveness.

The model results are robust to a wide range of conditions. Figure A3a shows that responsiveness regressions using lagged (rather than current) TFP or current RLP make the same qualitative predictions as regressions using lag TFP, as do regressions using current TFP *innovations* or *differences* (in the main text, we also find that our empirical results are robust to using innovations or differences).

Figure A3b reports responsiveness coefficients from instrumental variables regressions performed on model-simulated data; these correspond with those we estimated on empirical data (described in the main text, Appendix III.B, and Table A2) and are motivated by concerns about division bias and measurement error in employment. Figure A3c addresses the measurement error issue more specifically by considering scenarios in which the econometrician observes firm employment, firm labor productivity, or both with error.³ Error in employment measurement is introduced with a multiplicative disturbance term drawn from an independent normal distribution with mean 1 and standard deviation 0.033 (such that employment disturbances of 10 percent map to three standard deviations from the mean);⁴ error in labor productivity measurement is generated by applying the employment disturbance term to the denominator in revenue per worker. As shown in Figure A3c, this source of measurement error does not dramatically affect responsiveness coefficients, such that the decline in responsiveness we observe empirically is unlikely to be caused by rising measurement error over time. That said, the best approach to

³ Recall that in our manufacturing regressions, the employment variable used to measure productivity, which comes from the ASM/CM, is independent of the employment variable used to measure employment levels and growth, which comes from the LBD.

⁴ This choice is arbitrary and does not have qualitative implications.

concerns about measurement error is our empirical investigation using cross-industry variation, covered in the main text and Table 8.

Figure A4a reports the effects of rising adjustment costs on aggregate productivity in the model with non-convex adjustment costs. The black solid line shows true (model) aggregate productivity. The dashed orange line replicates the productivity index exercise described in section IV.B of the main text; in that exercise, we empirically estimate the effects of declining responsiveness on aggregate productivity by constructing an aggregate productivity index that depends on estimated responsiveness coefficients (see that discussion for more detail).

The productivity index used in section IV.B is given by $\sum_{j} \theta_{jt} a_{jt}$, where θ_{jt} is the employment weight of firm *j* and a_{jt} is TFP; we can construct this index and related counterfactuals using model-simulated data to study the index's relationship with true productivity. For every adjustment cost scenario, we use simulated data and corresponding regression coefficients to construct $\sum_{j} \hat{\theta}_{jt}^{HC} a_{jt}$, the aggregate productivity index as predicted by the responsiveness regressions under that adjustment cost scenario (where "HC" stands for "high cost"). We then construct a counterfactual index using the same simulated data but applying the responsiveness coefficient from the low-cost baseline scenario, $\sum_{j} \hat{\theta}_{jt}^{B} a_{jt}$ (where "B" stands for baseline, referring to the use of the responsiveness coefficient from the low-cost baseline scenario). Then $\sum_{j} \hat{\theta}_{jt}^{HC} a_{jt} - \sum_{j} \hat{\theta}_{jt}^{B} a_{jt}$ is the effect of changing responsiveness on the aggregate productivity index, in the model-simulated data. The dashed orange line in Figure A4a shows this counterfactual productivity index, which tracks true aggregate productivity reasonably well, lending support to our empirical approach for estimating the effects of changing responsiveness on aggregate productivity.

Our shocks vs. responsiveness approach is also useful if changing responsiveness is generated by convex labor adjustment instead of non-convex. We construct an alternative baseline calibration of the model in which non-convex costs are set to zero ($F_- = F_+ = 0$), but γ =1.75 to again replicate a job reallocation rate of 0.18, leaving all other parameters unchanged relative to Table A1. (Recall from the model description that γ governs quadratic adjustment costs on employment). From this alternative convex cost baseline, we conduct both of our model experiments: (1) raise adjustment cost γ above its baseline value, and (2) reduce TFP dispersion σ_a . These results are in Figure A1. The qualitative results of the experiments for job reallocation, responsiveness, and revenue productivity dispersion are the same as those found in our non-convex cost experiments except that, as mentioned above, responsiveness is unaffected by changes in shock dispersion (providing an even cleaner shocks vs. responsiveness dichotomy). The productivity results for the convex cost case are reported in Figure A4b.

Finally, we note that declining responsiveness can also be derived from an increase in the curvature of the revenue function (generated by reducing ϕ). This is shown in Figure A3d; notably, while increased curvature reduces responsiveness in each of our example model frameworks, its implications for revenue productivity dispersion (not shown) are model dependent.

C. Alternative framework: Wedges

The shocks vs. responsiveness insight is more general than the specific adjustment costs models described above. As an example, here we show how a broader interpretation can be adopted, following the seminal work of Hsieh and Klenow (2009).

Hsieh and Klenow (2009) show how measured revenue productivity dispersion can exist in equilibrium if there are static distortions or "wedges" affecting firms' first-order conditions. This framework can be viewed as a reduced form way of capturing not only adjustment frictions (under certain specifications of the wedge process) but also a wide variety of other factors that distort first-order conditions.

Consider a simple one-factor (employment) model in the spirit of Hsieh and Klenow (2009). Firms maximize period t profits given by:

$$S_{jt}A_{jt}E_{jt}^{\phi} - W_t E_{jt} \tag{A9}$$

where $A_{jt}E_{jt}^{\phi}$ is revenue and S_{jt} is a firm-specific wedge, which can be thought of as a tax when $S_{jt} < 1$ or as a subsidy when $S_{jt} > 1$. Suppose the wedge S_{jt} follows the following process:

$$s_{jt} = -\kappa a_{jt} + \nu_{jt}, \tag{A10}$$

where lowercase indicates logs. Consistent with much of the recent literature, we assume $\kappa \in$ (0,1), and v_{jt} is independent of a_{jt} with $\mathbb{E}(v_{jt}) = 0.5$ Equation (A10) states that firms with more favorable fundamentals (e.g., higher TFP) face more substantial wedges (meaning, lower

⁵ By "consistent with the literature," we mean a common finding in the literature is that indirect measures of wedges (i.e., revenue productivity measures like TFPR) are positively correlated with measures of fundamentals (technical efficiency and demand shocks) and have lower variance than fundamentals. See Foster, Haltiwanger, and Syverson (2008) and Blackwood et al. (forthcoming).

 S_{jt}), but the variance of (log) wedges is lower than the variance of fundamentals. This relationship between S_{jt} and A_{jt} is critical for producing empirically plausible aggregate reallocation rates (under reasonable parameterizations of ϕ) in the absence of explicit adjustment frictions.

Given (A9) and (A10), the first-order condition, in logs (indicated by lowercase), is given by:

$$e_{jt} = \frac{1}{1-\phi} \left(\ln\left(\frac{\phi}{W_t}\right) + (1-\kappa)a_{jt} + \nu_{jt} \right).$$
(A11)

Taking time differences (indicated by Δ), sweeping out year and industry effects, and incorporating the transformation described in the first section of this appendix, we obtain an employment growth function (expressed in log differences):

$$\Delta e_{jt} = \frac{1}{1-\phi} \Big((1-\kappa)a_{jt} - (1-\phi)e_{jt-1} + \nu_{jt} \Big).$$
(A12)

Employment growth can be expressed as a function of the productivity level and lagged employment, along with the shock to the wedge and the model parameters. Equation (A12) shows that the relationship between employment growth and productivity depends not only on ϕ but also on κ , which determines the covariance between firm productivity and firm distortions. A higher value of κ results in weaker responsiveness of growth to productivity because stronger κ means that wedge shocks partially offset productivity shocks. In the text, we refer to a higher κ as reflecting a more positive correlation between fundamentals and distortions. By this we mean that the implicit tax on firms is increasing in fundamentals. In this case the implicit tax is larger the less positive is s_{jt} . Note also that aggregate job reallocation, which in this context can be thought of as the dispersion of employment growth rates, is decreasing in κ .

This framework also has implications for revenue productivity dispersion. Log revenue per worker is given by $\ln \frac{W_t}{\phi} + \kappa a_{jt} - v_{jt}$, such that the dispersion of revenue labor productivity is *increasing* in κ .

This model, albeit highly simplified, thus yields rich empirical predictions, which we report in a manner analogous to our simulations from the model with labor adjustment costs. That is, we calibrate the "wedges" model, using the wedge correlation parameter κ to target the empirical reallocation rate of the 1980s, then we conduct experiments varying κ and σ_A (the dispersion of TFP). These exercises are shown in Figure A2 in a manner comparable to Figures

2 and A1. Figures A2a and A2c report the results of raising κ from its baseline value; as discussed above, responsiveness and job reallocation fall while revenue labor productivity dispersion rises. Declining responsiveness through this mechanism, as in the other models, yields a decline in aggregate productivity, as shown in Figure A4c.

The wedge model also yields similar implications for changes in the variance of shocks, shown in Figure A2b. A decline in the variance of a_{jt} yields declining reallocation and revenue productivity dispersion but, as in the model with convex adjustment costs, does not affect responsiveness (thus, responsiveness depends on TFP dispersion only in the model with non-convex adjustment costs).

Finally, as in the models with adjustment costs, in the wedge model a decline in responsiveness can be generated through a decline in revenue function curvature, as shown in Figure A3d.

Appendix II: Data

A. Longitudinal Business Database

For longitudinal information we rely on the Longitudinal Business Database (LBD), which covers the universe of private nonfarm employer business establishments in the U.S. The LBD records establishment employment, payroll, detailed industry, and location annually (with employment corresponding to March 12). Establishments are linked over time by high-quality longitudinal identifiers, and firm identifiers link establishments of multi-establishment firms. See Jarmin and Miranda (2002) for a description of the LBD, which is constructed from the Census Bureau's Business Register. The LBD's high-quality longitudinal linkages make it ideal for studying growth and survival outcomes of businesses.

In our regression specifications we include several establishment characteristic controls derived from the LBD. Key among them is firm age. We follow the large LBD-based literature in defining firm age as follows. Upon the first appearance of a firm identifier in the LBD, we assign firm age as the age of the firm's oldest establishment, where an establishment has age 0 during the year in which it first reports positive payroll. Thereafter, the firm ages naturally (i.e., we add one year to the firm's age for each calendar year after the firm identifier's first observation). This allows us to abstract from spurious changes in firm identifiers. We also use firm identifiers to measure firm size, which is the sum of employment across all the firm's establishments. In our regressions we control for firm size based on four cutoffs: fewer than 250 employees, 250-499 employees, 500-999 employees, and 1,000 or more employees (these cutoffs follow Foster, Grim, and Haltiwanger (2016), hereafter FGH).

B. Revenue-enhanced LBD (RE-LBD)

While the LBD does not include revenue data, revenue information is available in the Business Register at the employer identification number (EIN) level starting in the mid-1990s. Importantly, EINs are not a straightforward firm or establishment identifier in that multiple establishments can have the same EIN, and some firms can have multiple EINs (e.g., splitting the firm by geography or separating tax functions from payroll functions). In the case of multiestablishment firms, in general revenue data are not broken out by establishment. Haltiwanger et al. (2017) deal with these various challenges and create firm-level revenue data by aggregating across EINs of the same firm.⁶ They then match these revenue figures to the LBD at the firm level, finding nominal revenue figures for about 80 percent of LBD firms. The resulting revenue dataset is roughly representative of the overall LBD in terms of observables like firm age, firm size, sector, multi- or single-establishment status, and patterns of firm growth. Nevertheless, Haltiwanger et al. (2017) construct propensity scores for the entire LBD using logistic regressions with dependent variable equal to 1 for firms with revenue data and 0 otherwise. These regressions are run separately for birth, deaths, and continuers, and they rely on observables including firm size, firm age, employment growth rate, industry, and multi-establishment status. We use the resulting propensity scores (in inverse) as sampling weights in all regressions. We deflate revenue with the GDP deflator, but this is unimportant as all empirical exercises will implicitly control for industry-level prices as we deviate firm productivity from industry-year means. More generally, we follow Haltiwanger et al. (2017) closely in our measurement approach using the RE-LBD.

C. Manufacturing data

We supplement the LBD with manufacturing data from the Census of Manufacturers (CM) and the Annual Survey of Manufacturers (ASM), a dataset we obtain from FGH and update through 2013. The CM surveys almost the universe of manufacturing establishments every five years (those ending in "2" and "7"); we use CM data from 1982 through 2012.⁷ The ASM, conducted in non-CM years, surveys roughly 50,000-70,000 establishments; we use ASM data from 1981 through 2013. The ASM is a series of five-year panels (starting in years ending in "4" and "9") with probability of panel selection being a function of industry and size.

We combine the CM and ASM into an annual manufacturing establishment dataset covering 1981-2013, and we link the combined ASM-CM with the LBD by establishment and year using internal Census Bureau establishment identifiers that are consistent across these datasets. We create a dummy variable equal to 1 for those establishments that appear in both the ASM-CM and the LBD and 0 for those establishments that appear only in the LBD. We then create propensity scores using a logistic regression to predict ASM-CM presence based on the following variables: whether the establishment is part of a multi-establishment firm, size

organization, which can affect the way in which revenue data are reported and the way EINs map to firms.

⁶ This is a complicated process involving careful attention to details including industry and legal form of

⁷ Very small establishments (those with fewer than five employees) are not surveyed by the CM; the Census Bureau fills in data for these with administrative records. We do not include these cases.

(employment), payroll, detailed industry, and firm age. We estimate these propensity scores separately for each year; we then use them (in inverse) as sampling weights in all regressions.

As discussed in the main text, we use the LBD to measure employment growth and survival for each plant-year observation for which we have the TFP measures. This implies we are using the LBD through 2014 for this purpose.

D. Output and production factors

We require measures of revenue and production factors to construct TFPS, TFPP, and TFPR. We calculate real establishment-level revenue (or, under TFPR assumptions, output) as $Q_{jt} = (TVS_{jt} + DF_{jt} + DW_{jt})/PISHIP_t$, where TVS_{jt} is total value of shipments, DF_{jt} is the change in (the value of) finished goods inventories, DW_{jt} is the change in (the value of) work-inprogress inventories, and $PISHIP_t$ is the *industry-level* shipments deflator, which varies by detailed industry (4-digit SIC prior to 1997 and 6-digit NAICS thereafter) and is taken from the NBER-CES Manufacturing Productivity Database. If the resulting Q_{jt} is not greater than zero, then we simply set $Q_{jt} = TVS_{jt}/PISHIP_t$.

For the purposes of TFP estimation, we construct labor from the ASM in terms of total hours (TH_{it}) as follows:

$$TH_{jt} = \begin{cases} PH_{jt} \frac{SW_{jt}}{WW_{jt}} & \text{if } SW_{jt} > 0 \text{ and } WW_{jt} > 0 \\ PH_{jt} & \text{otherwise} \end{cases}$$
(A13)

where PH_{jt} is production worker hours, SW_{jt} is total payroll, and WW_{jt} is the payroll of production workers.

We measure capital separately for structures and equipment using the perpetual inventory method: $K_{jt+1} = (1 - \delta_{t+1})K_{jt} + I_{jt+1}$ where *K* is the capital stock, δ is a year- (and industry-) specific depreciation rate, and *I* is investment. At the earliest year possible for a given establishment, we initialize the capital stock by multiplying the establishment's reported book value by a ratio of real capital to book value of capital derived from BEA data (where the ratio varies by 2-digit SIC or 3-digit NAICS). Thereafter, we observe annual capital expenditures and update the capital stock accordingly, where we deflate capital expenditures using BLS deflators.⁸

⁸ See FGH for more detail. In a small number of cases (less than 0.5 percent) we cannot initialize the capital stock as described; in such cases we follow Bloom et al. (2013) using I/K ratios.

We calculate materials as $M_{jt} = (CP_{jt} + CR_{jt} + CW_{jt})/PIMAT_t$, where *CP* is the cost of materials and parts, *CR* is the cost of resales, *CW* is the cost of work done for the establishment (by others) on the establishment's materials, and *PIMAT* is the industry materials deflator. We calculate energy costs as $N_{jt} = (EE_{jt} + CF_{jt})/PIEN_t$, where *EE* is the cost of purchased electricity, *CF* is the cost of purchased fuels consumed for heat, power, or electricity generation, and *PIEN* is the industry energy deflator.

We use the production factor and output measures described above for each of our three TFP measures (TFPS, TFPP, and TFPR).

E. Cost and revenue shares: TFPS and TFPR

TFPS and TFPR productivity estimates require industry-level factor expenditures as shares of revenue (for TFPS) or cost (for TFPR) to construct factor elasticity estimates. We obtain these shares at the detailed industry level (4-digit SIC prior to 1997, 6-digit NAICS thereafter) from the NBER-CES Manufacturing Productivity Database, which reports industrylevel figures for expenditures on equipment, structures, materials, energy, and labor. We average these cost shares across all of 1981-2013 to obtain time-invariant elasticities, though our results are robust to instead using time-varying elasticities as in FGH.

F. Proxy method: TFPP

Our TFPP productivity concept requires us to estimate factor elasticities using proxy methods. Given the challenge of identifying exogenous shocks to fundamentals, a long literature (e.g., Olley and Pakes (1996), Levinsohn and Petrin (2003)) proposes using a variable production factor as a "proxy" for identification. Blackwood et al. (forthcoming) compare multiple proxy-based TFP concepts with other concepts from the literature. Some literature achieves this using a two-step procedure (see Ackerberg, Caves, and Frazer (2015)), but we follow Wooldridge (2009) in implementing a single-step GMM approach using lagged values of capital and variable inputs as instruments. We refer the reader to the just-mentioned research for more detail on the general approach to proxy estimation of production functions. For our purposes, we estimate factor elasticities separately by 2- and 3-digit industries using energy as the proxy variable.

Appendix III: Additional empirical results

A. Reallocation has declined within firm age bins

As noted in the text, the aggregate decline in job reallocation is not simply a composition effect due to declining young firm activity. Rather, we also observe declining reallocation within firm age bins. To see this, we first create seven firm age groups (ages 0, 1, 2, 3, 4, 5 and 6+). We then study the change in aggregate (weighted average) job reallocation in year *t* relative to a base year *t*0 with the following shift-share decomposition:

$$R_{t} - R_{t0} = \sum_{a} \omega_{at0} (R_{at} - R_{at0}) + \sum_{a} R_{at0} (\omega_{at} - \omega_{at0}) + \sum_{a} (R_{at} - R_{at0}) (\omega_{at} - \omega_{at0})$$

where R_t is the aggregate (or, as we will specify it below, sector-level) job reallocation rate, a indexes age bins, ω_{at} is the employment share of age group a in time t, and R_{at} is the reallocation rate for age group a in time t. The first term is a within-age-group component based on the change in flows among firms of that age. The second term is a between-group component capturing the change in the age composition. The third term is a cross term. We focus on the overall component and the within component; the residual coming from composition shifts and cross terms reflects the extent to which composition effects account for the aggregate change.

To abstract from business cycle issues, we construct this counterfactual between the business cycle peaks of 1987-1989, 1997-1999, 2004-2006, and 2011-2013. We study the long differences in reallocation rates between these three periods. Figure A5 illustrates the results, showing both the overall change in reallocation for a sector and the change in the within-age-group term, indicated by the "Holding age constant" bars. As is evident, the decline in reallocation within age groups explains the bulk of the overall decline. In other words, the changing age composition of U.S. firms resulting from changing patterns of firm entry does not explain the patterns of reallocation that motivate this paper.

B. Instrumental variables: Empirical results

A particular challenge for our empirical approach is that our workhorse regressions given by equations (7) and (9) in the main text feature initial employment (E_{jt}) on the right-hand-side (as the state variable) and on the left-hand-side (in the DHS growth dependent variable). Additionally, in our economywide regressions using labor productivity, initial employment also appears in the denominator of the productivity indicator (which is real revenue per worker).⁹ In Appendix I, we explore this problem by running instrumental variables regressions on modelsimulated data. Regressions in which an employment lag is used to instrument for initial employment (i.e., use E_{jt-1} as an instrument for E_{jt}), and regressions in which we additionally instrument for productivity using a lag, find that responsiveness still declines as adjustment costs rise. This suggests that we can evaluate robustness of our main responsiveness results to the employment endogeneity issue using similar instrumental variables regressions in our empirical exercises.

For brevity, we focus on the time-trend regression specifications for studying changing responsiveness. Table A2 reports results of instrumental variables regressions. The first column reports establishment-based results for the manufacturing sector using our preferred productivity measure, TFPS, and instrumenting for initial employment. The second column reports economywide firm-based results instrumenting for initial employment, and the third column reports economywide firm-based results instrumenting for initial employment and for productivity. In each column, and for both young and mature firms, we observe declining responsiveness as indicated by the negative (and statistically significant) coefficient on the linear trend variables.

⁹ Initial employment is also used in TFP estimation in our manufacturing-only exercises; however, the employment variable used for TFP is independently constructed from our ASM-CM dataset (see Appendix II).

Appendix references

- Ackerberg, Daniel A., Kevin Caves, and Garth Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica* 83 no. 6:2411-2451.
- Blackwood, Glenn, Lucia Foster, Cheryl Grim, John Haltiwanger and Zoltan Wolf. Forthcoming. "Macro and Micro Dynamics of Productivity: From Devilish Details to Insights." *American Economic Journal: Macroeconomics*.
- Bloom, Nicholas, Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Itay Saporta-Eksten, and John Van Reenen. 2013. "Management in America." CES-WP-13-01, CES Discussion Paper Series, Washington, DC.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. 2018. "Really Uncertain Business Cycles." *Econometrica* 86 no. 3:1031-65.
- 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 no. 1:37-50.
- Elsby, Michael W. L. and Ryan Michaels. 2013. "Marginal Jobs, Heterogeneous Firms, and Unemployment Flows." *American Economic Journal: Macroeconomics* 5 no. 1:1-48.
- Foster, Lucia, Cheryl Grim, and John Haltiwanger. 2016. "Reallocation in the Great Recession: Cleansing or Not?" *Journal of Labor Economics* 34 no. S1 (Part 2, January 2016):S293-S331.
- Hsieh, Chang-Tai and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics*, 124(4): 1403-48.
- 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.
- Haltiwanger, John, Ron S. Jarmin, Robert Kulick and Javier Miranda. 2017. "High Growth Young Firms: Contribution to Job Growth, Revenue Growth and Productivity" In *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges* edited by John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar. Chicago: University of Chicago Press.

- 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.
- Levinsohn, James and Amil Petrin. 2003. "Estimating Production Functions Using Inputs to Control for Unobservables." *Review of Economic Studies* 70 no. 2:317-41.
- Olley, G. Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Industry." *Econometrica*, 64 no. 6:1263-98.
- Tauchen, George. 1986. "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions." *Economics Letters* 20 no. 2:177-181.
- Wooldridge, Jeffrey M. 2009. "On Estimating Firm-Level Productivity Functions Using Proxy Variables to Control for Unobservables." *Economics Letters* 104 no. 3:112-4.

	Description	Value	Calibration rationale
ϕ	Inverse demand elasticity parameter	0.67	Standard in literature
β	Discount factor	0.96	Standard in literature
$ ho_a$	(Log) TFP AR(1) coefficient	0.80	Estimated TFPS AR(1), 1980s average
σ_a	Standard deviation of (log) TFP	0.46	Estimated TFPS standard deviation, 1980s average
σ_{η}	Standard deviation of TFP innovation	0.28	Implied by ρ and σ_a
F_{+}	Upward kinked adjustment cost	1.03	Target job reallocation rate = 0.18 (1980s average)*
·			$F_{+} = 0$ in convex cost model.
<i>F_</i>	Downward kinked adjustment cost	0.00	(Rely on upward cost for baseline calibration)
γ	Convex adjustment cost parameter	0.00	No convex cost in non-convex cost model.
-			$\gamma = 1.75$ in convex cost model.
κ	Wedge/productivity correlation parameter	0.83	Wedge model only; target job reallocation rate 0.18*
			With $\sigma = 0.04$

Table A1: Baseline model calibration

Moment targets refer to U.S. manufacturing sector.

*1980s average reallocation rate among continuing establishments (Business Dynamics Statistics).

Table A2. Institumental variables regressions, employment growth responsiveness					
	TFPS:	RLP:	RLP:		
	IV for employment	IV for employment	IV for emp & RLP		
Prod*Young: β_1^{γ}	0.4358	0.3170	0.1499		
	(0.0177)	(0.0013)	(0.0016)		
Prod*Young*trend: δ^{γ}	-0.0042	-0.0033	-0.0015		
	(0.0009)	(0.0001)	(0.0001)		
Prod*Mature: β_1^m	0.3123	0.2581	0.1092		
	(0.0104)	(0.0010)	(0.0001)		
Prod*Mature*trend: δ^m	-0.0020	-0.0032	-0.0014		
	(0.0005)	(0.0001)	(0.0001)		
Observations (thousands)	2,179	4,909	4,909		

Table A2: Instrumental variables regressions, employment growth responsiveness

Note: All coefficients statistically significant with p < 0.01. All regressions include controls described in equation (7) and related text. RLP regressions use 10 percent random sample of RE-LBD.



Figure A1: The shocks and responsiveness hypotheses, model results (convex cost)

Note: Panels c and d share same legend. Results relative to model baseline calibration (vertical purple line) with downward adjustment cost γ =1.75 and TFP dispersion σ_A =0.46 (see Appendix I and Table A1 for model calibration details). "s.d. RLP" refers to the standard deviation of revenue labor productivity in model-simulated data.



Figure A2: The shocks and responsiveness hypotheses, model results (wedge model)

Note: Panels c and d share same legend. Results relative to model baseline calibration (vertical purple line) with TFP/ wedge correlation parameter κ =0.83 and TFP dispersion σ_A =0.46 (see Appendix I and Table A1 for model calibration details). "s.d. RLP" refers to the standard deviation of revenue labor productivity in model-simulated data.






Figure A4: Aggregate productivity and responsiveness

Note: Model simulations. 'Estimated effects' reflect productivity counterfactuals in simulated data mimicking empirical exercises.



Figure A5: Most variation in job reallocation is within firm age classes