

Cyclical Worker Flows: Cleansing vs. Sullyng*

John Haltiwanger[†]

Henry Hyatt[‡]

Erika McEntarfer[§]

Matthew Staiger[¶]

September 2022

Abstract

Do recessions speed up or impede productivity-enhancing reallocation? To investigate this question, we use U.S. linked employer-employee data to examine how worker flows contribute to productivity growth over the business cycle. We find that in expansions high-productivity firms grow faster primarily by hiring workers away from lower productivity firms. The rate at which job-to-job flows move workers up the productivity ladder is highly procyclical. Productivity growth slows during recessions when this job ladder collapses. In contrast, flows into nonemployment from low-productivity firms disproportionately increase in recessions, which leads to an increase in productivity growth. We thus find evidence of both sullyng and cleansing effects of recessions, but the timing of these effects differs. The cleansing effect dominates early in downturns but the sullyng effect lingers well into the economic recovery.

*We thank Paul Oyer, Kristin McCue, José Mustre-del-Río, and participants of the 2015 Society of Labor Economists Conference, the 2015 NBER Conference on the Role of Firms in Wage Inequality, the Spring 2017 Midwest Macro Conference, the 2017 Federal Reserve Board Conference on Labor Market Dynamics and the Macroeconomy, the 2021 ASSA Conference and seminar participants at Drexel University, Arizona State University and U.S. Census Bureau, for helpful comments and suggestions. Early results were circulated in a 2015 draft entitled “Do Workers Move Up the Firm Productivity Job Ladder?” John Haltiwanger was also a part-time Schedule A employee at the U.S. Bureau of the Census at the time of the writing of this paper. Matthew Staiger was a Pathways Intern at the U.S. Bureau of the Census at the time of the writing of this paper. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed (CBDRB-FY21-CED006-0002).

[†]Department of Economics, University of Maryland, College Park. E-mail: halt@umd.edu

[‡]Center for Economic Studies, U.S. Census Bureau. E-mail: henry.r.hyatt@census.gov

[§]Center for Economic Studies, U.S. Census Bureau. E-mail: erika.mcentarfer@census.gov

[¶]Opportunity Insights, Harvard University. E-mail: mstaiger@g.harvard.edu

1 Introduction

Economists have long sought to understand how business cycles affect the reallocation of resources. Do recessions promote economic efficiency by “cleansing” out less productive firms and redirecting labor to more productive uses? Or, does the decline in job mobility in recessions “sully” productivity-enhancing worker reallocation, leaving workers matched to mediocre firms? In this paper we use U.S. linked employer-employee data to decompose the employment growth of high- and low-productivity firms into two components: growth accounted for by job-to-job moves and growth accounted for by flows through nonemployment. We find that job-to-job flows move workers from less productive to more productive firms and the rate at which workers move up this job ladder is highly procyclical. In contrast, less productive firms rely heavily on hiring jobless individuals in expansions and are disproportionately more likely to displace workers back to nonemployment in contractions. In this way, worker flows through nonemployment shift workers away from low-productivity firms in contractions.

We thus find empirical evidence of both cleansing and sully effects of recessions, which feature in many models of the labor market. Much of the theoretical literature focuses on either cleansing or sully effects. Schumpeter (1939) originally proposed that recessions may be productivity enhancing, driving out less productive uses of capital and labor and freeing these resources for more efficient use. The notion of cleansing effects of recessions was later revived in Davis and Haltiwanger (1990), Caballero and Hammour (1994), and Mortensen and Pissarides (1994).¹ Barlevy (2002) notes that cleansing effects of recessions appear at odds with observed procyclical job quality. He proposes a model whereby declines in job-to-job moves cause a drag on productivity in recessions, a sully effect. Sully effects of recessions feature more recently in a set of papers by Moscarini and Postel-Vinay (2013, 2016) (MPV). The MPV framework in particular yields a rich set of predictions for the cyclical reallocation of workers across the firm productivity distribution that we largely confirm in our empirical analysis.² More recent research by Lise and Robin (2017) and Baley,

¹Unlike Schumpeter, the more recent literature does not argue that recessions are desirable but rather that, conditional on the adverse shock occurring, there may be an acceleration of the ongoing reallocation of resources to more productive uses in recessions.

²The MPV model builds on the job ladder framework of Burdett and Mortensen (1998) but importantly

Figueiredo, and Ulbricht (2022) use models with heterogeneity of both workers and firms that induce cyclical variation in sorting over the cycle. Within their frameworks, both of these papers suggest that the cleansing effect dominates the sullyng effect.³ Our empirical results suggest that the cleansing effects dominate at the start of a recession but that the sullyng effects continue long into the economic recovery.

In our empirical analysis, we classify firms into high- and low-productivity firms based upon the relative ranking of firms in the measured distribution of firm-level productivity. We find that high-productivity firms grow faster: the differential growth rate averages 0.20 percent of employment on a quarterly basis. Decomposing the growth rate differential into the components due to job-to-job moves versus worker flows through nonemployment, we find that the propensity of high-productivity firms to grow faster is driven by their advantage in poaching workers from less productive firms. This advantage is large: the difference in the growth rates due to job-to-job moves is 0.61 percent per quarter. Low-productivity firms actually lose workers on net through poaching, and so hire relatively more jobless workers to sustain growth: the differential employment growth rate from worker flows through nonemployment (relative to high-productivity firms) is 0.41 percent per quarter.

These patterns of worker reallocation between low- and high-productivity firms differ dramatically over the course of the business cycle. In expansions, high-productivity firms actively poach workers from less productive employers. During recessions, this job ladder collapses yielding a sullyng effect. The nonemployment margin also changes dramatically across the cycle. In expansions, low-productivity firms grow primarily through hiring jobless

for our purposes incorporates business cycle dynamics. In their framework, search frictions prevent workers from immediately moving into desirable job matches and workers move to better firm matches via job-to-job flows. High-productivity firms are able to offer higher wages and thus are able to grow faster in expansions than less-productive firms, who must rely on the pool of unemployed workers in filling vacancies. In recessions, this cyclical job ladder collapses, yielding a sullyng effect. Our empirical results are largely consistent with these predicted dynamics. Our results build on the findings in Haltiwanger, Hyatt, Kahn, and McEntarfer (2018) and Haltiwanger, Hyatt, and McEntarfer (2018). This earlier work provides support for a procyclical job ladder in terms of firm earnings and productivity. The current paper is distinguished by explicitly considering the sullyng and cleansing contributions of flows through the lens of the impact on the share of employment at high and low productivity firms. This earlier work does not explore decompositions of productivity growth into cleansing and sullyng components. Bertheau, Bunzel, and Vejlin (2020) report broadly similar patterns on poaching by high-paying and low-paying firms using Danish data.

³However, worker flows across firms ranked by firm productivity are not targeted directly when they estimate their models. While we do not explicitly consider worker heterogeneity in our empirical analysis, we do provide novel evidence on the patterns of worker flows across firms ranked by productivity that are relevant for these models of sorting.

workers. In contractions, worker separations to nonemployment from low-productivity firms exhibit a disproportionate increase. This cleansing effect peaks in the early stages of an economic downturn when the unemployment rate surges. The collapse of the job ladder begins later but lingers into the early stages of a recovery when the level of unemployment remains high.

What are the implications of these business cycle dynamics on aggregate productivity growth? We use an accounting decomposition to investigate the sources of aggregate productivity growth. On average, worker reallocation through job-to-job flows contributes 0.11 log points to the index of overall productivity growth each quarter. This is a substantial contribution to the average quarterly rate of aggregate productivity growth, which is 0.19 based on our data. However, during recessions there is clear evidence of a sullyng effect. This is particularly true during the Great Recession. Prior to the recession, in 2006:1, worker reallocation via job-to-job flows contributed 0.13 log points to quarterly aggregate productivity growth but this contribution declined to 0.02 by 2009:2. Acting against this is a cleansing effect that operates via worker flows through nonemployment. In 2006:1, worker reallocation via nonemployment contributed -0.10 log points to quarterly aggregate productivity growth but this increased to 0.08 in 2009:1. We show that during periods of rising unemployment the cleansing impact on productivity outweighs the sullyng impact. However, during periods when unemployment is above trend the sullyng impact on productivity outweighs the cleansing impact.

Our primary analysis uses revenue per worker to measure relative productivity within the firm’s industry. This has the advantage of allowing us to measure productivity within most industry sectors but has one key disadvantage.⁴ Specifically, revenue per worker reflects both “innate” firm productivity as well as sorting of workers across firms. We assess this issue using the AKM decomposition of earnings.⁵ The AKM firm fixed effect represents the pay-premium workers receive independent of worker quality and while not a direct measure of productivity, canonical models suggest that the relative ranking should be highly correlated with productivity differences within and across industries. We find that the cyclical patterns

⁴Total factor productivity measures are available for only handful of sectors that are a shrinking share of the U.S. economy.

⁵AKM refers to the decomposition developed by Abowd, Kramarz, and Margolis (1999).

of hires and separations via poaching and nonemployment are very similar whether we rank firms based on revenue per worker of the AKM firm fixed effect. Thus, our results do not appear to be driven by cyclical sorting of worker types across firms.⁶

Finally, we consider the implications of the worker flows across firms ranked by productivity for worker earnings. We find that worker movements from low- to high-productivity firms move workers into higher-paying firms, as measured by both firm average earnings and the AKM firm fixed effect. These earnings changes move strongly with their productivity analogues but are roughly half of the magnitude. Thus, workers obtain a substantial fraction of the gains from movements onto and up the firm productivity job ladder, but there is suggestive evidence of incomplete pass-through of gains to workers.

The paper proceeds as follows. Section 2 develops a methodology to decompose aggregate productivity growth into components attributable to worker reallocation from job-to-job flows and flows through nonemployment. Section 3 describes the data. Section 4 presents the results of our decomposition of aggregate productivity growth. Section 5 discusses implications for earnings. Section 6 concludes.

2 Decomposition Methodology

This section develops a methodology to decompose aggregate productivity growth into components attributable to worker reallocation through poaching and nonemployment flows. We begin with the perspective from the job ladder literature and show how worker flows, both poaching flows and flows through nonemployment, affect the share of workers at more productive firms. We then present the perspective from the literature on productivity growth and show how changes in the share of workers at more productive firms affect aggregate productivity growth. Lastly, we integrate these two perspectives to quantify how worker reallocation from less to more productive firms via poaching and nonemployment flows affects aggregate productivity growth.

⁶Haltiwanger, Hyatt, and McEntarfer (2018) and Crane, Hyatt, and Murray (2022) examine cyclical sorting of heterogeneous workers across heterogeneous firms. Both papers find evidence that the assortative matching of workers and firms is countercyclical: a greater share of low-productivity workers are able to match to better firms in expansions.

2.1 Job Ladders

The starting place for the job ladder literature is the following identity developed by Haltiwanger, Hyatt, and McEntarfer (2018) (HHM) and Haltiwanger, Hyatt, Kahn, and McEntarfer (2018) (HHKM):

$$\text{Net Job Flows} = H_t - S_t = \sum_{j \in \{p, n\}} (H_t^j - S_t^j) = \sum_{i \in \{l, h\}} \sum_{j \in \{p, n\}} (H_t^{ij} - S_t^{ij}), \quad (1)$$

where H_t is the number of hires and S_t is the number of separations in quarter t . The superscripts denote subsamples defined by the type of worker flow, where $j = p$ denotes poaching (job-to-job) flows and $j = n$ denotes flows through nonemployment, and the type of firm, where $i = h$ denotes high-productivity firms and $i = l$ denotes low-productivity firms. For example, H_t^p denotes poaching hires and H_t^{ph} denotes poaching hires at high productivity firms.⁷

The identity in equation 1 decomposes employment growth at high-productivity firms ($H_t^h - S_t^h$) into net growth due to two components: job-to-job moves of workers or poaching flows ($H_t^{ph} - S_t^{ph}$) and flows of workers through nonemployment ($H_t^{nh} - S_t^{nh}$).⁸ In the aggregate economy, employment growth is entirely attributable to net worker flows through nonemployment since poaching hires and poaching separations aggregated over both high- and low-productivity firms are equal. However, for any subset of firms in the economy, net poaching need not be zero, as some firms will be more successful poaching workers away from other employers. This “net poaching flows” component of growth captures the comparative growth advantage one group of firms has over another in their ability to attract workers away from other firms.

We extend equation 1 to decompose changes in the share of employment at high-productivity firms into components attributable to poaching and nonemployment flows. Let $\lambda_t^i = (H_t^{ip} - S_t^{ip})/E_{t-1}^i$ and $\delta_t^i = (H_t^{in} - S_t^{in})/E_{t-1}^i$ denote the employment growth rates at firm type $i \in \{l, h\}$ through net poaching flows and net nonemployment flows, respectively.

⁷Hires and separations characterize worker mobility between time $t - 1$ and t . For hires we count all worker flows into a firm in our sample at time t . For separations, we count all worker flows out of a firm in our sample at time $t - 1$.

⁸Correspondingly for low-productivity firms, $H_t^l - S_t^l = (H_t^{pl} - S_t^{pl}) + (H_t^{nl} - S_t^{nl})$.

Then,

$$\Delta\theta_t^h = \tilde{\lambda}_t^h + \tilde{\delta}_t^h, \quad (2)$$

where $\tilde{x}^h = (x^h - x^l)\theta_{t-1}^h\theta_{t-1}^l(E_{t-1}/E_t)$ for $x \in \lambda, \delta$ and $\Delta\theta_t^h$ is the change in the share of employment at high-productivity firms between quarter $t - 1$ and t . This expression shows that the sign of differential net poaching rate, $\lambda_t^h - \lambda_t^l$, determines whether poaching rates will increase or decrease the share of employment at high-productivity firms. The magnitude of this effect also depends on the share of workers at high productivity firms as well as the growth in overall employment.

2.2 Productivity Growth

We focus on a measure of aggregate productivity growth defined as $\sum_k \theta_{t-1}(k)\Delta R_t(k)$, where $\Delta R_t(k)$ is the change in the employment-weighted average of log revenue per worker of firms in industry k , and $\theta_{t-1}(k)$ is the employment share of industry k in $t - 1$. There are two things to note about this measure. First, our measure of industry-level productivity, $R_t(k)$, is the employment-weighted average of a firm-level measure of productivity, which is measured in logs. Such accounting indices of productivity have been widely used in the literature to quantify the contribution of reallocation effects to productivity (e.g., Olley and Pakes, 1996; Foster, Haltiwanger, and Krizan, 2001; and Melitz and Polanec, 2015).⁹ Second, our measure of aggregate productivity growth only captures productivity growth that occurs within-industries. This is common in the accounting reallocation literature discussed above as comparing levels of productivity across industries is difficult.

We isolate the component of productivity growth that is attributable to worker

⁹Conceptually, this aggregate index is consistent with aggregate productivity in a model with a single input (labor), constant returns to scale, and perfect competition in product markets. While these are strong assumptions (although not inconsistent with job ladder models), such indices track official statistics closely as we show below. Much of the literature that focuses on misallocation specifies curvature in the revenue function so that there is a well-defined size distribution even in the absence of distortions. Part of the reason for this is that this enables a measure of allocative efficiency that is relative to a frictionless/distortionless benchmark (e.g., Hsieh and Klenow, 2009; Bartelsman et al., 2013; and Blackwood et al., 2021). In principle, such curvature is not necessary in models with adjustment frictions such as search and matching frictions. Models with curvature in the revenue function have the property that it is not optimal to allocate all resources to the most productive firm. While this implies caution in using weighted average measures of firm-level productivity in quantitative analysis of models with such curvature, Decker et al. (2020) show that in models with adjustment frictions that this type of aggregate productivity index tracks structural measures of true productivity well even if there is curvature in the revenue function.

reallocation between more and less productive firms. Specifically,

$$\sum_k \theta_{t-1}(k) \Delta R_t(k) \approx (R_{t-1}^h - R_{t-1}^l) \Delta \theta_t^h + \sum_k \left[\theta_{t-1}(k) \sum_{i \in \{l, h\}} \theta_t^i(k) \Delta R_t^i(k) \right], \quad (3)$$

where the superscript i denotes high-productivity ($i = h$) and low-productivity ($i = l$) firms, $R_{t-1}^h - R_{t-1}^l$ is the employment-weighted average of the industry-level productivity differential between high- and low-productivity firms across industries, and $\Delta \theta_t^h$ is the change in the share of employment at high-productivity firms. The first term represents productivity growth that arises from worker reallocation that changes the share of workers at high-productivity firms. Intuitively, this term will capture the productivity gains from shifting activity from low productivity to high productivity firms. The second term represents productivity growth that arises from changes in the average productivity of high- and low-productivity firms. Equation 3 is an approximation because $\sum_k [\theta_{t-1}(k) (R_{t-1}^h(k) - R_{t-1}^l(k)) \Delta \theta_t^h(k)] \approx (R_{t-1}^h - R_{t-1}^l) \Delta \theta_t^h$. To show that the approximation performs well, we regress the true value, $\sum_k [\theta_{t-1}(k) (R_{t-1}^h(k) - R_{t-1}^l(k)) \Delta \theta_t^h(k)]$, on the approximate value, $(R_{t-1}^h - R_{t-1}^l) \Delta \theta_t^h$, which yields a coefficient of 0.85 and an R-squared of 0.91. See Appendix C for details. This approximation greatly simplifies the analysis as it allows us to focus on decomposing the change in the share of employment at all high-productivity firms (i.e., $\Delta \theta_t^h$) as opposed to decomposing the employment share within each industry (i.e., $\Delta \theta_t^h(k)$).

2.3 Integrating the Two Perspectives

Lastly, we combine the perspectives from the job ladder and productivity growth accounting literatures in order to decompose the component of aggregate productivity growth that is attributable to worker reallocation into components attributable to poaching flows and flows through nonemployment. Specifically, combining equations 2 and 3 yields,

$$\underbrace{(R_{t-1}^h - R_{t-1}^l) \Delta \theta_t^h}_{\text{Growth from reallocation}} = \underbrace{(R_{t-1}^h - R_{t-1}^l) \tilde{\lambda}_t^h}_{\text{Poaching}} + \underbrace{(R_{t-1}^h - R_{t-1}^l) \tilde{\delta}_t^h}_{\text{Nonemployment}} \quad (4)$$

An increase in the share of workers at high-productivity firms increases the index of aggregate productivity, and the magnitude of the effect is determined by the productivity differential

between high- and low-productivity firms.

Importantly, while our productivity measures reflect relative differences within industries, our poaching flow measures include workers flowing across industry boundaries. That is, we capture the change in employment at high (low) productivity firms from poaching flows regardless of the origin industry.

3 Data

A key contribution of our paper is the matching of U.S. Census Bureau linked employer-employee data to new productivity measures also developed at Census. We will first describe the linked employer-employee data, and how we use it to decompose firm growth via job-to-job moves versus flows through nonemployment. The Longitudinal Employer-Household Dynamics (LEHD) data contain quarterly earnings records collected by state unemployment insurance (UI) programs, linked to establishment-level data from the Quarterly Census of Employment and Wages (QCEW). LEHD employment coverage is quite broad, covering over 95 percent of private sector workers and almost all state and local government employment.¹⁰ State-level data availability varies by year, as states began sharing UI and QCEW data with the Census Bureau at different times. In this paper we use LEHD data for private-sector employers in 28 states from 1998-2015.¹¹ Our 28 states include many of the largest states so that our sample accounts for 65 percent of U.S. private sector employment.

The LEHD data allow us to decompose firm employment growth by worker hires and separations. Section 2 extends the decomposition developed in HHM and HHKM that yields an exact decomposition of firm employment growth due to workers switching jobs (what we call net job-to-job or net poaching flows) and growth due to flows between employment and nonemployment (what we call net nonemployment flows). A challenge for the identification of job-to-job flows in the LEHD data is that the data do not provide information on why a worker left one job and began another. We only have quarterly earnings, from which we infer

¹⁰For a full description of the LEHD data, see Abowd et al. (2009).

¹¹Our 28 states are CA, FL, GA, HI, ID, IL, IN, KS, ME, MD, MN, MO, MT, NC, NJ, ND, NM, NV, PA, OR, RI, SC, SD, TN, VA, WA, and WV. While we restrict our analysis to employers located in our 28-state sample, we use the complete set of available states to construct worker job histories. As described later in this section, our productivity measures reflect the labor productivity of the national firm.

approximately when workers left and began jobs. HHM and HHKM develop three alternative measures of job-to-job flows, and demonstrate that key findings on the nature of job ladders are robust to different approaches for identifying job-to-job moves in the LEHD data. We use the within/adjacent approach from HHM in this paper. This approach defines job-to-job transitions as those where the new job begins in the same or following quarter as the job separation. Based upon the robustness analysis in HHM, we are confident our main results are not sensitive to the specific rules we use amongst the set of rules they considered.¹²

To measure firm productivity, we use a relatively new firm-level database on productivity from Haltiwanger et al. (2017) based on the revenue and employment data from the Census Business Register and the Longitudinal Business Database (LBD). Since the underlying revenue and employment data are from the Census Business Register, this database offers much wider coverage of labor productivity at the firm level than earlier studies that focused on sectors like manufacturing or retail trade. These data allow us to measure the log of real revenue per employee on an annual basis for a wide coverage of the private, non-farm, for-profit firms. Revenue is deflated with the Gross Domestic Product price deflator. This measure of productivity is a standard gross output per worker measure of productivity that is commonly used to measure productivity at the micro and macro level but is a relatively crude measure compared to using total factor productivity (TFP). However, in the empirical literature, this revenue labor productivity measure has been shown to be highly correlated with TFP based measures of productivity across businesses within industries. That is, within detailed industry year cells, Foster, Haltiwanger, and Krizan (2001) and Foster, Haltiwanger, and Syverson (2008) find that the correlation between TFP and gross output (revenue) per worker across businesses is about 0.6 within industries in the manufacturing sector. This finding is consistent with the implications of models with labor market adjustment frictions which motivate our analysis.¹³ In our analysis below, we use this revenue labor productivity measure deviated from industry by year means.

¹²They also consider job-to-job flows restricted to those where the transition occurs within the same quarter and those with minimum disruptions in earnings. They find results that are very robust across these alternatives. Each of the different measures is highly correlated with the alternatives (pairwise correlations of about 0.98) and each of the LEHD based job-to-job flow series has a correlation of about 0.96 with CPS based job-to-job flows.

¹³See for example Decker et al. (2020). In their calibrated model of labor adjustment frictions, they obtain a correlation of TFP and revenue labor productivity of 0.90.

The gross output per worker data while offering much wider coverage than earlier studies has some limitations. The data only cover about 80 percent of firms in the Census LBD. The latter cover all firms with at least one paid employee in the private, non-farm sector. One reason is that the revenue data are not available for non-profits. For another, the revenue data derive from different administrative sources than the payroll tax data. Most of the matches between the payroll tax and revenue data are via Employer Identification Numbers (EINs) but firms can use different EINs for filing income taxes and filing quarterly payroll taxes.¹⁴ For such firms, name and address matching is required. Haltiwanger et al. (2017) also show that the missingness of revenue is only weakly related to observable firm characteristics such as industry, size, or age.¹⁵ We are able to construct measures of labor productivity at the firm (operational control) level given that the Census Business Register has a complete mapping of all EINs owned by any given parent firm. Even with these limitations, we have revenue per worker for more than 4 million firms in each calendar year which we integrate with the LEHD data infrastructure via EINs. For the remaining private-sector employers in the LEHD data for which we cannot match to our productivity data, we impute labor productivity using the size, age, and relative wages paid by the employer within their industry.¹⁶

3.1 Validating the Measure of Productivity Growth

Measures of aggregate productivity based on our data closely track the measures produced by the U.S. Bureau of Labor Statistics (BLS) and the U.S. Bureau of Economic Analysis (BEA). Figure 1 plots annual productivity growth from the BLS, the BEA, and the LBD. The

¹⁴Another source of mismatch is sole proprietors file income taxes on their individual income tax returns while payroll taxes are filed via their EIN. Administrative data are available that links the EINs to the filers via the SS-4 form (application for EINs). While this information is incorporated in the Census Business Register, it is imperfect.

¹⁵The productivity data explicitly excludes North American Industry Classification System (NAICS) 81 which is Other Services. This industry is very heterogeneous, including non-profits such as religious organizations where productivity is not well defined.

¹⁶The latest year for which we have firm productivity data is 2015, so we end our time series there although the LEHD data are more current. We investigated imputing post-2015 productivity using lagged productivity and other covariates but were not satisfied this 100 percent imputation was of sufficiently high quality. In unreported results, we have found that the patterns of worker flows are robust to excluding the imputed cases. Including the imputed cases facilitates our quantification of shares of employment at high and low productivity firms and in turn the productivity decomposition we use in the analysis.

measure produced by the BLS is an employment-weighted average of industry-level (4-digit NAICS) labor productivity growth rates.¹⁷ Using the LBD, we create a comparable measure using the full national sample to calculate an employment-weighted average of industry-level growth in log revenue per worker.¹⁸ The mean of the LBD measure is 2.0 log points while the mean of the BLS measure is 1.7 log points. The correlation of these two series is 0.86. The BEA measures the growth in value added per worker for the private, non-farm business sector, which is conceptually distinct from the LBD and BLS measures.¹⁹ Nevertheless, we see that the BEA measure exhibits similar patterns to the LBD and BLS measures. The fourth series is also based on the LBD but only uses the 28 states in our paper and presents estimates of the firm-level index of productivity growth described on the left-hand-side of equation 3. Despite differences in sample and methodology, this alternative measure from the LBD follows similar patterns. Taken together, Figure 1 gives us confidence to proceed with our measure of revenue labor productivity.

Our measure of firm productivity also exhibits a number of the key features that Syverson (2011) emphasized are common in the literature on firm productivity and dynamics. First, we find tremendous dispersion of revenue labor productivity within narrowly defined sectors. The within industry/year standard deviation of log real revenue per worker is about 0.80. This is in the range of labor productivity dispersion indices reported by Syverson (2004). Second, we find that log real revenue per worker is highly predictive of firm growth and survival, as shown in Table A.1.²⁰ We consider two dependent variables for all incumbents in period $t-1$. The first dependent variable is the Davis, Haltiwanger, and Schuh (1996) firm level growth rate of employment that is inclusive of firm exit from $t-1$ to t .²¹ The

¹⁷The industry-level growth rate is based on the growth rate in what BLS calls sectoral output per worker. Sectoral output is gross output less intrasectoral transactions. At the 4-digit level the adjustment for intrasectoral transactions is modest.

¹⁸Specifically, we calculate aggregate productivity growth in the following steps: (i) calculate total revenue and employment within each industry (defined at the four-digit NAICS) and year, (ii) calculate the difference in log revenue per worker within each industry and year, and (iii) take the weighted average across industries using the average level of employment in the current and previous years as weights.

¹⁹The BEA measure uses a conceptually different output measure and also reflect shifts in employment from low- to high-productivity industries, whereas the BLS and LBD measures only capture within-industry contributions to aggregate productivity growth.

²⁰For this analysis, we do not restrict the sample to those firms in our LEHD data sample. These regressions use all firm-year observations from the revenue-enhanced Census Business Register.

²¹This measure is given by $g_{it} = (E_{it} - E_{it-1}) / (0.5 * (E_{it} + E_{it-1}))$. It is a second order approximation to a log first difference that accommodates entry and exit.

second dependent variable is an exit indicator that takes on the value of one if the firm exits between $t-1$ and t and is zero otherwise. We use a linear probability model for this second specification. Firm exit and growth is organic growth and exit in the manner defined by Haltiwanger, Jarmin, and Miranda (2013) (i.e., it abstracts from changes in ownership or M&A activity). We regress these two outcomes on the deviation of log productivity from the industry average in $t-1$ and on log size in $t-1$ (log of firm employment in $t-1$). While these are simple reduced form specifications, these specifications are consistent with standard models of firm growth and survival since these are proxies for the two key state variables for the firm in making growth and survival decisions. The canonical model implies that holding initial size constant a firm with higher productivity is more likely to grow and less likely to exit. We find overwhelming evidence in support of these predictions in Table A.1. A one standard deviation increase in within-industry productivity yields a 20 percentage point increase in net employment growth and 5 percentage point decrease in the likelihood of exit.²²

Lastly, we show that intertemporal variation in factor utilization is unlikely to affect our main decomposition results. During a recession a firm may decide to cut back on production (possibly by reducing workers' hours or worker intensity), which could lead to a decline in log revenue per worker without any real changes in productivity. In Appendix B.1 we show that the productivity differential between high- and low-productivity firms ($R_{t-1}^h - R_{t-1}^l$) does not exhibit cyclical variation. Thus, our decomposition exercise does not appear to be affected by intertemporal variation in factor utilization.

3.2 Defining High- and Low-Productivity Firms

To help mitigate remaining concerns about measurement error, we construct robust rankings of firms by productivity. We first generate time-invariant measures of employer productivity, defined as the employment-weighted average of firm productivity over the life of the SEIN (the state tax identifier number, the key employer identifier in LEHD data). This approach is broadly consistent with the rank preserving equilibria assumption in Moscarini and Postel-

²²Decker et al. (2020) develop a simple model of firm dynamics with adjustment frictions that shows that the relationship between growth and survival from $t-1$ to t with realizations of labor productivity in period t is very similar as with TFP, holding firm size constant in $t-1$.

Vinay (2013). We then compute the employment-weighted and within-industry quintiles of the productivity distribution. Using these quintiles, we define high-productivity firms as those in the top two quintiles and low-productivity firms as those in the bottom three quintiles. In unreported analysis, we find that results are robust to permitting firms to change ranks over time. This is not surprising given the large differences between high- and low-productivity firms. For example, the within-industry differences in average gross output per worker between high- and low-productivity firms are typically in excess of 85 log points.

As a robustness check on our productivity measure, we also rank firms by the AKM firm fixed effect and average earnings. We estimate the AKM firm premiums by regressing log earnings on person fixed effects, firm fixed effects, and controls for time and worker age.²³ We estimate the model by implementing the iterative method proposed by Guimaraes and Portugal (2010). The AKM firm fixed effect abstracts from observable and unobservable individual characteristics and in canonical models, the firm specific pay premia should be closely related to productivity differences across firms. We also consider simple non-parametric measures of relative earnings by ranking firms based on the average log earnings of full-quarter workers. Just as we do for gross output per worker, we construct employment-weighted within-industry quintiles based on these measures and define high-ranked firms as those in the top two quintiles and low-ranked firms as those in the bottom three quintiles.

4 Worker Reallocation and Productivity Growth

Figure 2 shows our decomposition of net job flows for high- and low-productivity firms described in equation 1. As discussed previously, a key prediction of job ladder models is that job-to-job moves should reallocate workers away from less productive to more productive firms. Figure 2(a) shows that this prediction from the theory holds true in the data. The most productive firms have overall positive net employment growth on average and net

²³To control for time we include a set of year dummies that capture calendar year effects on earnings. To control for worker age, we follow the specification of Card, Cardoso, and Kline (2016). We center age around 40, include a quadratic and cubic transformation of worker age, but omit the linear term. Note that we estimate our AKM employer effects at the State Employer Identification Number level rather than the national level as in our productivity data. This method of estimating AKM effects follows the conventional implementation strategy when applied to LEHD data. This distinction should have a minimal effect on whether an employer is counted as low- vs. high-ranked.

poaching ($H_t^{ph} - S_t^{ph}$) is strongly positive. The average net employment growth of high-productivity firms is 0.33 percent per quarter with net poaching (the rate at which job-to-job moves reallocate workers to high-productivity firms) averaging 0.27 percent per quarter. In other words, during the 1998-2015 period, job-to-job moves of workers from less-productive employers account for most (80 percent) of the net employment growth of high-productivity firms.

The results of the decomposition are also striking for the less productive firms in the industry. In Figure 2(b), low-productivity firms grew at a rate of 0.14 percent per quarter on average from 1998-2015, which is slower than the high-productivity firms. Low-productivity firms lose -0.34 percent employment per quarter from workers “voting with their feet” and moving to firms ranked higher in firm productivity distribution. The positive growth rate for less productive firms is entirely due to strong hiring from nonemployment. In other words, in a typical quarter less productive firms recruit from the pool of unemployed individuals to replace workers moving to better firms. This is also consistent with job ladder models of the labor market. In job ladder models, it is the search and matching frictions that support the presence of low-productivity firms that primarily hire from nonemployment.

The patterns of hires and separations in Figure 2 are instructive for understanding the differences in the cyclical dynamics of job-to-job and nonemployment worker flows. Poaching hires and separations both decline for high-productivity firms in contractions with the decline in poaching hires larger so that net poaching declines significantly. Hires from nonemployment decline sharply for low-productivity firms in contractions accompanied by a surge in separations so that net employment growth declines sharply for low-productivity firms. There are similar qualitative patterns for hires from and separations to nonemployment for high-productivity firms but the magnitudes are smaller.²⁴

Figure 3 shows the time series of the main components of the decomposition in equation 2 and illustrates how worker reallocation affects the percent of employment at high-productivity firms.²⁵ Figure 3(a) presents the components of the change in percent of

²⁴Figure A.1 plots the differential growth rates from poaching and nonemployment flows between high- and low-productivity firms. The results both clearly illustrate the cyclical patterns of these differential flows and preview the robustness of our findings to using earnings-based measures of firm performance (i.e., the AKM firm effect).

²⁵To make the figure more readable, we present the results in percentage point terms.

employment at high-productivity firms ($\Delta\theta_t^h$) that are attributable to worker reallocation through poaching flows ($\tilde{\lambda}^h$) and nonemployment flows ($\tilde{\delta}^h$) and Figure 3(b) presents the combined effect.²⁶ During expansions, worker flows through nonemployment lead to a reduction in the share of workers at high-productivity firms whereas poaching flows lead to an increase in the share of workers at high-productivity firms. At the onset of a recession, the rate at which nonemployment flows contributes to the growth of employment share of low-productivity firms slows. Indeed, in the Great Recession this change was large enough such that nonemployment flows briefly contributed positively to the growth of the employment share at high-productivity firms. Throughout our sample period, poaching flows always contribute positively to the growth of employment at high productivity firms, but this largely collapses during recessions, particularly in the Great Recession. Consistent with our earlier analysis, the results highlight the staggered nature of the timing of these two effects. The cleansing effect begins at the onset of a recession where as the sullyng effect starts and peaks later on.

Figure 4(a) presents the decomposition of productivity growth into components attributable to poaching and nonemployment flows and shows clear evidence of the cleansing and sullyng effects of recessions. On average, worker reallocation through poaching flows contributes 0.11 log points to overall productivity growth each quarter (all statistics on productivity changes are quarterly and have not been annualized). This is a substantial contribution to the overall quarterly average rate of productivity growth of 0.19 when aggregating our micro data. However, during recessions there is clear evidence of a sullyng effect. In 2006:1 the poaching contribution is 0.13 log points but this declines to 0.02 by 2009:2. In contrast, worker reallocation through nonemployment tends to be a drag on productivity growth, on average, decreasing productivity by 0.06 log points each quarter.²⁷ However, during recessions there is evidence of a cleansing effect since during those times nonemployment flows yield declines in the employment share of low-productivity firms. In 2006:1, the nonemployment component is -0.1 log points but increases to 0.08 in

²⁶In our data equation 2 does not hold exactly. Figure B.3 presents evidence that this does not affect our main results by showing that the residual term is smaller in magnitude than the poaching and nonemployment flows and does not exhibit clear a cyclical pattern. See Appendix B.2 for details.

²⁷This drag on productivity is consistent with job ladder models with search and matching frictions. Dispersion in productivity is supported by such frictions in equilibrium.

2009:1. The figure illustrates the staggered nature of these effects in which the cleansing occurs at the outset of the recession—when unemployment rate is rising most rapidly—and the sully effect peaks relatively further on into the downturn—when the unemployment rate is highest. In addition, the sully effect lingers well into the recovery.

Figure 4(b) presents the combined effect of worker reallocation through poaching and nonemployment flows. During expansions, the total effect of worker reallocation through poaching and nonemployment flows contributes 0.05 log points to overall productivity growth each quarter. In recessions, this increases to 0.08. While this might suggest that cleansing effects dominate, these calculations neglect the fact that the sully effect lingers well into the recovery. Using an alternative indicator of the cycle, we find that the combined effect is 0.07 on average for quarters where the unemployment rate is below HP trend and 0.03 on average for quarters where unemployment is above the HP trend. This reversal is driven by much larger contributions of poaching flows during periods of low unemployment (0.13) compared to high unemployment (0.09).

Column 1 of Table 1 summarizes these patterns.²⁸ Here we show results from a set of regressions where different components of productivity growth are regressed on a cyclical indicator and a time-trend. We use two alternative cyclical indicators: the change in unemployment and unemployment deviated from a Hodrick-Prescott (HP) trend. Periods of rising unemployment correspond closely to NBER defined recessions. Panel B shows that productivity-enhancing reallocation through job-to-job moves is procyclical using both measures. Panel C shows that productivity-enhancing reallocation through worker flows through nonemployment is counter-cyclical using both measures. The net effect in Panel A, however, does depend on the cyclical indicator. Sully effects dominate when cyclicalities are measured using deviations from the unemployment rate: this is because it takes a while for productivity-enhancing reallocation from job-to-job moves to recover in expansions. In contrast, cleansing effects dominate when cyclicalities are measured using changes in the unemployment rate. Taken together, these results imply that the cleansing effect peaks earlier in a downturn compared to the collapse of the job ladder that lingers into the early stages of a recovery. These results also suggest that slow labor market recoveries will be

²⁸We discuss columns 2 and 3 of Table 1 in the next section.

generally more damaging to productivity growth than V-shaped recoveries as slow recoveries exhibit an accompanying slow recovery of job-to-job flows.

4.1 Robustness to Using the AKM Firm Premium

We assess the robustness of our results to using an alternative measure of firm performance: the AKM firm premium. While the revenue per worker measure of productivity has many strengths, it may partially reflect the sorting of workers across firms in addition to innate differences in firm productivity. The AKM firm premium is an alternative that mitigates this limitation. Furthermore, canonical models suggest that more productive firms offer higher wages and therefore the AKM firm premium ought to be closely related to innate differences in firm productivity.

Figure 5 decomposes growth in the aggregate firm premium attributable to worker reallocation across the firm premium ladder through poaching and nonemployment flows. Firms are re-ranked into high- and low-premium firms and the worker flows (i.e., $\tilde{\lambda}_t^h$ and $\tilde{\delta}_t^h$) and the differentials in the firm premium (i.e., $R_{t-1}^h - R_{t-1}^l$) are re-calculated to reflect the new rankings and alternative measure of firm performance. Figure 5(a) follows the methodology described in equation 4 but uses the AKM firm premium as the measure of firm performance instead of revenue per worker. For comparison, Figure 5(b) presents based on the same methodology but instead of using the AKM firm effect, we use average earnings. In both cases, firms are ranked within 4-digit NAICS industry codes.²⁹

Regardless of whether we rank firms by the AKM firm effect or average earnings, both series in Figure 5 present clear evidence of the cleansing and sully effects that were apparent in worker reallocation across the firm productivity ladder using revenue per worker. When ranking firms based on the AKM firm premium, we find that, on average, poaching flows to higher-paying firms lead to an increase in firm premium by 0.09 log points per

²⁹Unlike the gross output per worker measure, the AKM firm premium is directly comparable across industries. Figure A.2 presents results in which firms are ranked by the AKM firm pay premium across all industries, effectively treating all firms as being part of the same industry. These results capture the effects of worker reallocation both within and across industries. The figure illustrates that poaching flows contribute 0.3 log points to the growth in the AKM firm premium each quarter whereas nonemployment flows lead to a decline by 0.4 log points per quarter. The estimates are about three times as large as those for the within industry based tabulations. While much larger in magnitude, the qualitative patterns are very similar.

quarter whereas nonemployment flows lead to a 0.09 log point decrease in the firm premium per quarter. These estimates are quite similar qualitatively and quantitatively to those found in Figure 4(a), which uses revenue per worker to measure productivity differences within industries. We also find the cyclical patterns of these components are very similar to those using the direct measure of productivity.

Columns 2 and 3 of Table 1 quantify the cyclicity of the contributions to the AKM firm premium and average earnings. As in the column 1, the cleansing contribution outweighs the sullyng contribution in response to an increase in the unemployment rate while the opposite is true in response to a positive deviation of the unemployment rate from trend. These findings hold regardless of whether we rank firms by the AKM firm premium (column 2) or the average earnings at the firm (column 3). Since the AKM firm premium abstracts from worker heterogeneity, this suggests our main results by firm productivity are not being driven by variation in the patterns of sorting of heterogeneous workers across heterogeneous firms over the cycle.

Some caution is needed in interpreting the magnitudes of the impact of worker reallocation on the average firm premium in terms of implications for productivity. The AKM firm premium should be correlated with true productivity but variation may reflect additional factors (e.g., the proportion of rents shared with workers). Viewed from this perspective, the AKM firm premium results are of interest in their own right and provide guidance regarding the sources of variation in this important component of earnings.

5 Implications for Earnings

Whereas Section 4 focused on implications for productivity growth, the current section asks if moving up the firm productivity ladder benefits workers in the form of higher earnings. We use the same methodology described in equation 4, but we replace the productivity differentials, $R_{t-1}^h - R_{t-1}^l$, with earnings differentials. We measure firm-level earnings in two ways: (i) the AKM firm fixed effect, and (ii) average log earnings of all workers. The earnings differential is the difference between the employment-weighted average of earnings at high- and low-productivity firms. Note that the exercise is distinct from the analysis discussed in

Section 4.1, since we are investigating the earnings implications of worker flows across firms ranked by the revenue per worker measure of productivity (as opposed to firms ranked by the earnings-based measures).

Figure 6 shows that worker reallocation up the firm productivity ladder through poaching and nonemployment flows has meaningful implications for the earnings of workers. Figures 6(a) and 6(b) present results in which the earnings differentials are measured with the AKM firm premium and average log earnings, respectively. The results in the two figures are quantitatively and qualitatively similar. Job-to-job transitions add an average of 0.05 log points to average earnings in any given quarter. The earnings contribution of job-to-job flows is lower during recessions. During the 2007-2009 recession, the earnings contribution of job-to-job flows fell to a series low of 0.01. The contribution of nonemployment transitions to earnings growth is in the opposite direction and similar in magnitude to the contribution of job-to-job flows.³⁰ In the average quarter, worker movements into and from nonemployment subtract an average of 0.03 log points from earnings. This negative effect of nonemployment transitions is less present during recessions. During the 2001 recession, the contribution of nonemployment transitions is close to zero. During the 2007-2009 recession, the contribution of nonemployment transitions to earnings growth is briefly positive.

Comparing these results to Figure 4, we can draw conclusions about the extent to which the gains from productivity-enhancing reallocation are realized by workers. The earnings and productivity implications of employment transitions have similar signs, with job-to-job transitions providing gains, while nonemployment transitions subtract from each. The magnitudes, however, differ. The proportionate changes in earnings are roughly half the magnitude of the analogous changes in productivity.³¹ This is suggestive of incomplete pass-through of the gains in revenue productivity from worker flows into earnings.

³⁰Hahn, Hyatt, and Janicki (2021) consider the implications of employment transitions for earnings growth without considering productivity. They report that job-to-job and nonemployment transitions move in opposite directions and are roughly similar in magnitude, and follow opposite cyclical patterns.

³¹If revenue per worker and the AKM firm premium produced the exact same ranking then we would expect Figure 5(a) and Figure 6(a) to be identical, which they are not. The difference between the two results is partly attributable to the fact that the difference in the average AKM firm premium between high- and low-premium firms (as used in Figure 5(a)) is greater than the difference in the average AKM firm premium between high- and low-productivity firms (Figure 6(a)).

6 Conclusion

Consistent with the existing literature on firm heterogeneity, we find evidence of large differences in productivity across firms within the same industry. We also find that more productive firms in the same industry are more likely to grow and less productive firms more likely to contract and exit. The dispersion of productivity across firms is large in magnitude contributing to a high pace of reallocation of workers across firms. Using a decomposition of net job flows into those accounted for by job-to-job flows and those accounted for net flows from nonemployment, we find that much of the overall reallocation of employment from less productive to more productive firms is accounted for by job-to-job flows. The pace at which workers move up the productivity job ladder is highly procyclical. The collapse of the productivity job ladder is consistent with a sullyng effect of recessions. In recessions, we find that the reallocation of workers away from less productive firms via nonemployment flows increases. This occurs through a spike in separations to nonemployment along with a decline in hires from nonemployment at low productivity firms. Thus, we also find evidence that this component of reallocation is consistent with a cleansing effect of recessions.

The timing of the cleansing and sullyng effects differs across stages of the cycle. The cleansing effect peaks relatively early in a downturn coincident with the relatively early spike in separations. The sullyng effect peaks later in a downturn but lingers into the early stages of a recovery when unemployment is falling but remains well above trend.

Our findings are robust to using a direct measure of productivity based on relative differences in revenue per worker across firms within the same industry and an alternative measure of firm performance based on using the AKM firm premium. Since the AKM firm premium abstracts from worker heterogeneity, this robustness suggests our results are not being driven by variation in the patterns of sorting of heterogeneous workers across heterogeneous firms over the cycle. This is not to suggest that the latter is unimportant but rather that there may be additional effects of cleansing and sullyng from sorting above and beyond those we have quantified. We recognize that any conclusions about the role of sorting over the cycle for productivity fluctuations are tentative at best. Further theoretical and empirical work is needed in this area.


References

- [1] Abowd, John, Bryce Stephens, Lars Vilhuber, Fredrik Andersson, Kevin McKinney, Marc Roemer, and Simon Woodcock. 2009. “The LEHD Infrastructure Files and the Creation of the Quarterly Workforce Indicators.” In *Producer Dynamics: New Evidence from Micro Data*, 68, Studies in Income and Wealth, ed. Timothy Dunne, J. Bradford Jensen and Mark J. Roberts, 149-230. Chicago: University of Chicago Press.
- [2] Abowd, John, Francis Kramarz, and David Margolis. 1999. “High Wage Workers and High Wage Firms.” *Econometrica* 67(2): 251-333.
- [3] Baley, Isaac, Ana Figueiredo, and Robert Ulbricht. 2022. “Mismatch Cycles.” Forthcoming, *Journal of Political Economy*.
- [4] Barlevy, Gadi. 2002. “The Sullyng Effect of Recessions.” *Review of Economic Studies* 69(1): 65-96.
- [5] Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta. 2013. “Cross Country Differences in Productivity: The Role of Allocative Efficiency.” *American Economic Review* 103(1): 305-334.
- [6] Bertheau, Antoine, Henning Bunzel, Rune Vejlin. 2020. “Employment Reallocation over the Business Cycle: Evidence from Danish Data.” IZA Discussion Paper # 13681.
- [7] Blackwood, Glenn, Lucia Foster, Cheryl Grim, John Haltiwanger, and Zoltan Wolf. 2021 “Macro and Micro Dynamics of Productivity: From Devilish Details to Insights.” *American Economic Journal: Macroeconomics* 13(3): 142–172.
- [8] Burdett, Kenneth, and Dale Mortensen. 1998. “Wage Differentials, Employer Size, and Unemployment.” *International Economic Review* 39(2): 257-273.
- [9] Caballero, Ricardo J., and Mohamad L. Hammour. 1994. “The cleansing effect of recessions.” *The American Economic Review*, pp.1350-1368.
- [10] Card, David, Ana Rute Cardoso, and Patrick Kline. 2016. “Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women.” *The Quarterly Journal of Economics*, 131(2), pp.633-686.
- [11] Crane, Leland, Henry Hyatt, and Seth Murray. 2022. “Cyclical Labor Market Sorting.” Forthcoming, *Journal of Econometrics*.
- [12] Davis, Steven and John Haltiwanger. 1990. “Gross job creation and destruction: Microeconomic evidence and macroeconomic implications.” *NBER Macroeconomics Annual*, 5, pp.123-168.
- [13] Davis, Steven, John Haltiwanger, and Scott Schuh (1996) *Job Creation and Destruction*, Cambridge: MIT Press.
- [14] Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2017. “Declining dynamism, allocative efficiency, and the productivity slowdown.” *American Economic Review*, 107(5), pp.322-26.
- [15] Decker, Ryan, John Haltiwanger, Javier Miranda, and Ron Jarmin. 2020. “Changes in Business Dynamism and Productivity: Shocks vs. Responsiveness” *American Economic Review* 110(12): 3952-90.
- [16] Fernald, John. 2014. “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity.” Federal Reserve Bank of San Francisco Working Paper 2012-19.

- [17] Foster, Lucia, John Haltiwanger, and C. J. Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*, 63, Studies in Income and Wealth, ed. Charles Hulten, Edwin Dean, and Michael Harper, 303-372. Chicago: University of Chicago Press.
- [18] 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.
- [19] Guimarães, Paulo, and Pedro Portugal. 2010. "A simple feasible procedure to fit models with high-dimensional fixed effects." *The Stata Journal*, 10(4): 628-649.
- [20] Hahn, Joyce, Henry Hyatt, and Hubert Janicki. 2021. "Job Ladders and Growth in Earnings, Hours, and Wages." *European Economic Review* 133: 103654.
- [21] Haltiwanger, John, Henry Hyatt, Lisa Kahn, and Erika McEntarfer. 2018. "Cyclical Job Ladders by Firm Size and Firm Wage." *American Economic Journal: Macroeconomics* 10(2): 52-85.
- [22] Haltiwanger, John, Henry Hyatt, Erika McEntarfer. 2018. "Who Moves Up the Job Ladder?" *Journal of Labor Economics* 36(S1): S301-S336.
- [23] Haltiwanger, John, Ron Jarmin, Robert Kulick, and Javier Miranda. 2017. "High Growth Young Firms: Contribution to Job, Output, and Productivity Growth." In: *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, 75, Studies in Income and Wealth, ed. John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar, Chicago: University of Chicago Press: 11-62.
- [24] Haltiwanger, John, Javier Miranda, and Ron Jarmin. 2013. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95(2): 347-361.
- [25] Hsieh, Chang-Tai and Peter Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124(4): 1403-1448.
- [26] Lise, Jeremy, and Jean-Marc Robin. 2017. "The Macrodynamics of Sorting between Workers and Firms." *American Economic Review* 107(4): 1104-1135.
- [27] Melitz, Marc, and Sašo Polanec. 2015. "Dynamic Olley-Pakes productivity decomposition with entry and exit." *The RAND Journal of Economics* 46(2): 362-375.
- [28] Moscarini, Giuseppe, and Fabien Postel-Vinay. 2013. "Stochastic Search Equilibrium." *Review of Economic Studies* 80(4): 1545-1581.
- [29] Moscarini, Giuseppe, and Fabien Postel-Vinay. 2016. "Did the job ladder fail after the Great Recession?" *Journal of Labor Economics*, 34(S1), pp.S55-S93.
- [30] Mortensen, Dale, and Christopher Pissarides. 1994. "Job Creation and Destruction in the Theory of Unemployment." *Review of Economic Studies* 61: 397-415.
- [31] Olley, Steven and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64(6): 1263-1297.
- [32] Schumpeter, Peter. 1939. *Business Cycles* New York: McGraw-Hill.
- [33] Syverson, Chad. 2004. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy* 112(6): 1181-1222.
- [34] Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49(2): 326-365.


7 Tables and Figures

Figure 1: Aggregate Productivity Growth



Notes: This figure presents annual aggregate productivity growth based data from the U.S. Bureau of Labor Statistics (BLS), U.S. Bureau of Economic Analysis (BEA), and the Longitudinal Business Database (LBD). The two series based on the LBD are employment-weighted averages of industry-level measures of productivity growth. To calculate industry-level productivity growth we either calculate the change in total revenue and employment within an industry or an employment-weighted average of firm-level measures of log revenue per worker. The former closely follows the methodology used by the BLS while the latter is the productivity index described in our decomposition.


Figure 2: Poaching and Nonemployment Flows by Firm Productivity




Notes: High-productivity indicates that the firm is in the top two quintiles of the within-industry productivity distribution. Low-productivity indicates the firm is the bottom three quintiles of the within-industry productivity distribution. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

Figure 3: Changes in Percent of Employment at High-Productivity Firms

(a) Poaching and Nonemployment Flows

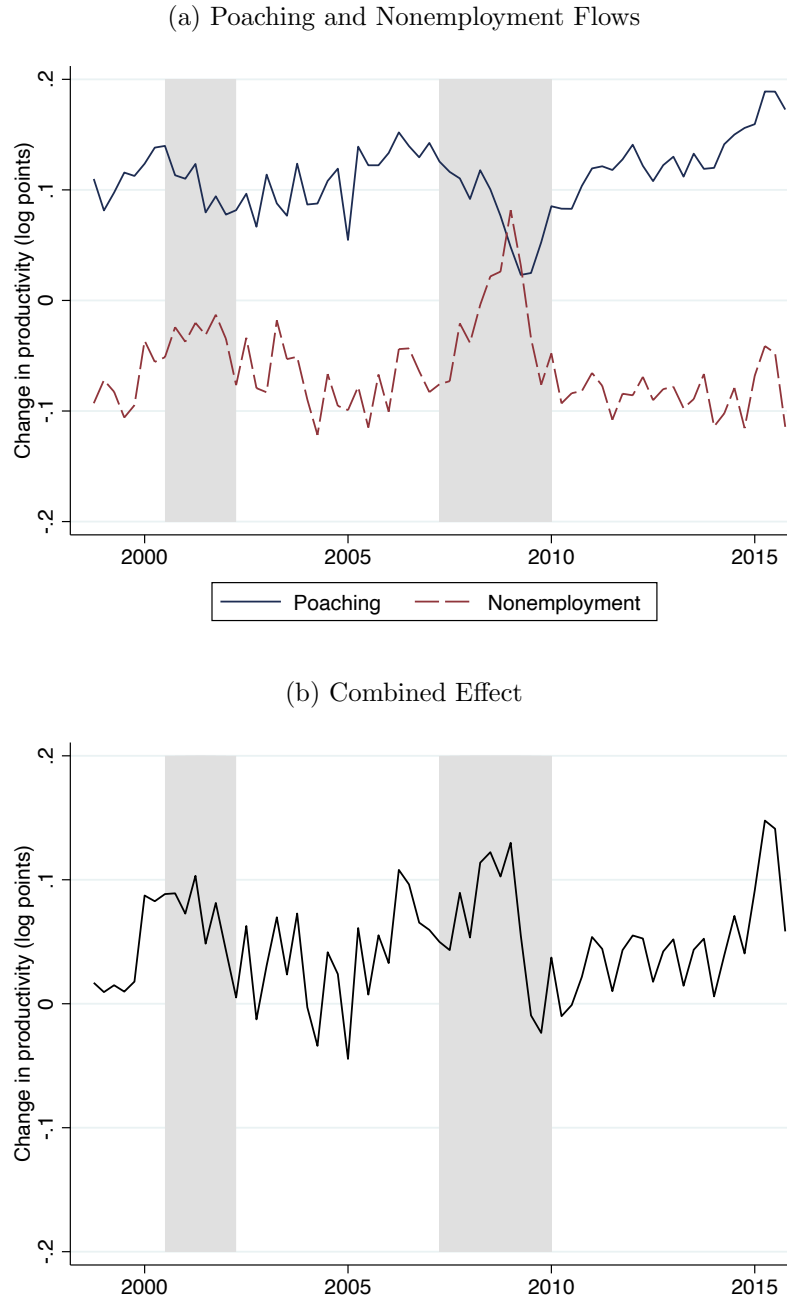


(b) Combined Effect



Notes: Panel (a) presents the change in the percent of employment at high-productivity firms that is attributable to worker reallocation through poaching and nonemployment flows and Panel (b) presents the sum of these two components. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

Figure 4: Decomposition of Growth in Productivity Over the Cycle



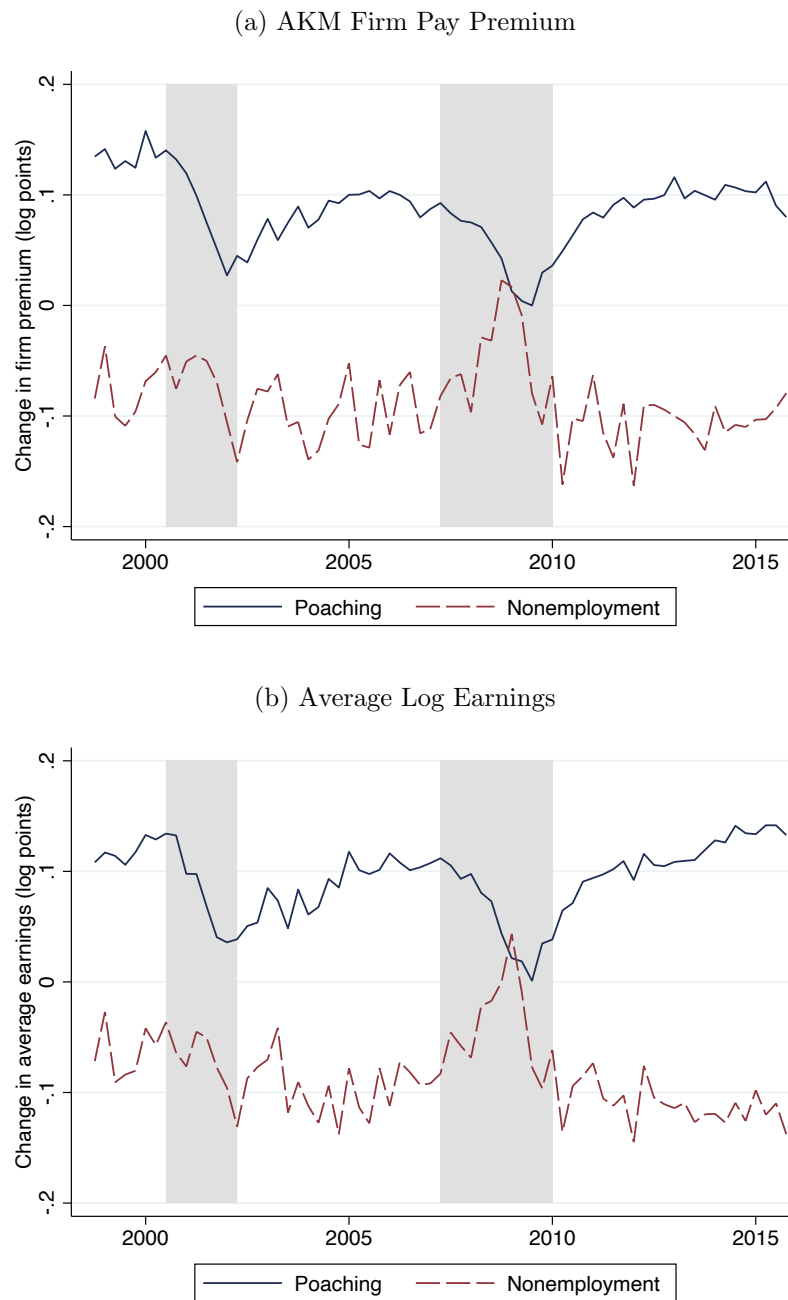
Notes: Panel (a) presents the components of quarterly productivity growth that are attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows and Panel (b) presents the sum of these two components. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

Table 1: Productivity Growth from Job Flows Over the Cycle

	Log Revenue Per Worker (1)	AKM Firm Pay Premium (2)	Log Average Earnings (3)
A. Net Job Flows			
Change in unemployment rate	0.040 (0.015)	0.010 (0.016)	0.014 (0.016)
Deviated unemployment rate	-0.024 (0.005)	-0.026 (0.006)	-0.029 (0.005)
B. Poaching Job Flows			
Change in unemployment rate	-0.052 (0.010)	-0.062 (0.009)	-0.061 (0.010)
Deviated unemployment rate	-0.024 (0.004)	-0.021 (0.004)	-0.026 (0.004)
C. Nonemployment Job Flows			
Change in unemployment rate	0.092 (0.009)	0.072 (0.010)	0.075 (0.009)
Deviated unemployment rate	0.000 (0.006)	-0.005 (0.005)	-0.003 (0.005)

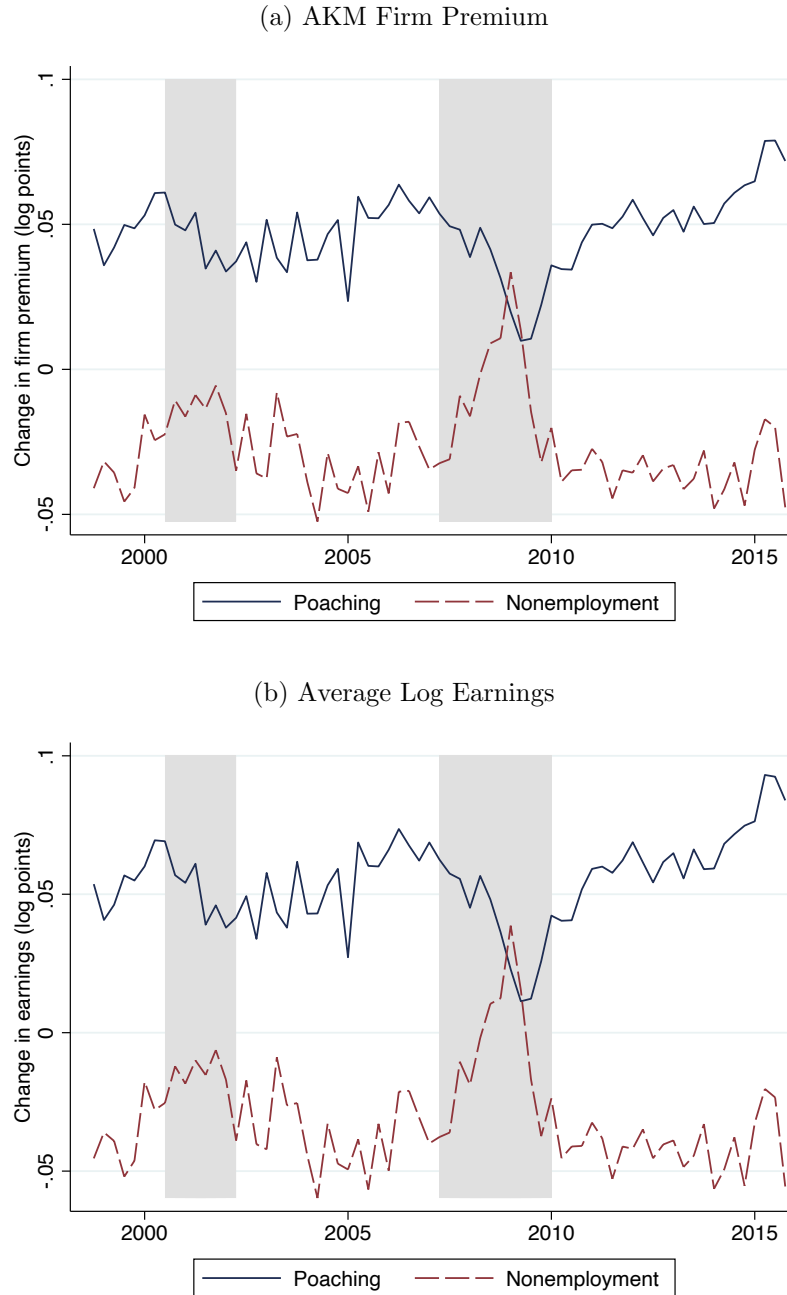
Notes: Each cell presents results from a separate regression estimated on national quarterly data. The dependent variable is the growth in log revenue per worker, AKM firm pay premium, or average log earnings. The independent variables in each regression include a cyclical indicator as well as a linear time-trend and a constant, which are not reported. The cyclical indicators considered include the change in the unemployment rate and the deviations of unemployment from the Hodrick-Prescott trend. Aside from the linear trend, all dependent and independent variables are measured in percentage point units. Standard errors are presented in parentheses.

Figure 5: Robustness to Ranking Firms Using Earnings-Based Measures of Firm Performance



Notes: This figure presents the main decomposition results using earnings-based measures of firm performance instead of log revenue per worker. Panels (a) and (b) measure firm performance using the AKM firm pay premium and average log earnings, respectively. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

Figure 6: Earnings Growth from Worker Reallocation up the Firm Productivity Ladder



Notes: The figure shows the components of earnings growth that attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows. Panel (a) uses the differences between the AKM firm premium at high- and low-productivity firms in order to quantify the implications of changes in the share of workers at high-productivity firms whereas panel (b) uses average log earnings of workers. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.


Appendix A Additional Results

Table A.1: Productivity, Employment Growth, and Firm Death

	Employment Growth Rate (1)	Firm Death (2)
Productivity	0.216 (0.00011)	-0.066 (0.00005)
Log of firm size	0.056 (0.00006)	-0.045 (0.00002)


Notes: Each column presents estimates from a separate regression. The dependent variable in column 1 is the employment growth rate between the current and subsequent year and the dependent variable in column 2 is an indicator equal to one if the firm dies in the subsequent year. The independent variables in each regression are the log of firm size and productivity, which is defined as the log revenue per worker deviated from the industry (defined by the 4-digit NAICS code) average. Standard errors are presented in parentheses.

Figure A.1: Differential Flows between High- and Low-Ranked Firms



Notes: Differential growth rates are the difference in quarterly employment growth rates between firms in the high category (top two quintiles) and those in the low category (bottom three quintiles). The different series present results in which the high and low categories are defined by productivity, average log earnings, and the AKM firm premium. For all measures, the quintiles are calculated within NAICS 4-digit industry codes. Results are presented separately for poaching and nonemployment flows. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

Figure A.2: Ranking Firms Across Industries using the AKM Firm Pay Premium



Notes: This figure presents the main decomposition results using the AKM firm pay premium measure of firm performance instead of log revenue per worker. Firms are ranked in the pooled sample across industries instead of within industries. This approach uses the same methodology described in the text but effectively treats all firms as being part of the same industry. Data are seasonally adjusted using X-12. The shaded regions mark quarters in which there was a recession.

Appendix B Assessing Measurement Issues

B.1 Productivity

Intertemporal variation in factor utilization does not appear to meaningfully affect the decomposition results. This is because, to the extent that log revenue per worker is subject to this concern, it affects high- and low-productivity firms to an equal extent. Figure B.1(a) plots, $R_{t-1}^h - R_{t-1}^l$, which is the productivity differential between high- and low-productivity firms. The productivity differential is orders of magnitudes larger than the short-term variation, which may be driven by intertemporal variation in factor utilization. To illustrate that the short-term variation in these differentials does not affect the decomposition exercise, we construct a smoothed series by fitting a linear time trend to the productivity differentials. We then use this smoothed series to implement the productivity growth decomposition. Figure B.1(b) presents the results and shows that the decomposition using the actual and smoothed productivity series yield essentially the same results. To quantify this we regress the component of productivity growth attributable to poaching flows using the observed productivity differentials on the series using the smoothed differentials. The R-squared is 0.99. The analogous R-squared for the nonemployment flows is 0.999.

Variation in factor utilization over the business cycle is a greater issue for our measure of aggregate productivity growth. Fernald (2014) produces a series of growth in business sector TFP that adjusts for factor utilization. Figure B.2(a) presents both the adjusted and unadjusted growth rates from these data. The unadjusted series exhibit a larger decline in TFP during recessions relative to the adjusted series. To get a sense of how sensitive our measure of productivity growth is to changes in factor utilization, we use the difference between the unadjusted and adjusted series from Fernald (2004) to adjust for variation in factor utilization.³² Figure B.2(b) compares the productivity growth attributable to worker reallocation to the adjusted and unadjusted measures of aggregate productivity growth.

B.2 Worker Flows


There are two reasons why the equality described in equation 2 does not hold exactly in our data (i.e., why $\Delta\theta_t^h$ does not exactly equal the sum of $\tilde{\lambda}_t^h$ and $\tilde{\delta}_t^h$). First, some workers may move to or from an employer located in a state outside of our 28-state sample. Because we aim to implement an exact decomposition, the counts of hires and separations only include worker flows where both the origin and destination employers are in one of the 28 states in our sample. Second, the administrative code that identifies the employer, the SEIN, can change over time and create a spurious flow of workers between the old and new SEIN. We are able to flag when these changes occur and omit these flows from the poaching and nonemployment flows. However, there is no straightforward way to account for this issue when measuring productivity. Thus, a change in an SEIN could lead to a change in the share of workers at high productivity firms but have no corresponding flow of workers. In

³²Specifically, we calculate the difference between the adjusted and unadjusted growth rates in 2001 and 2008. We adjust our growth rate by adding these differences to the growth rates in 2001 and 2008. Because our data also exhibit a large decline in 2009, we also add the 2008 difference from the Fernald (2014) series to the 2009 growth rate.

unreported results, we directly measure flows of workers in and out of the states in our sample and show that the residual term is primarily attributable to changes in the SEIN over time, not migration in and out of the sample.

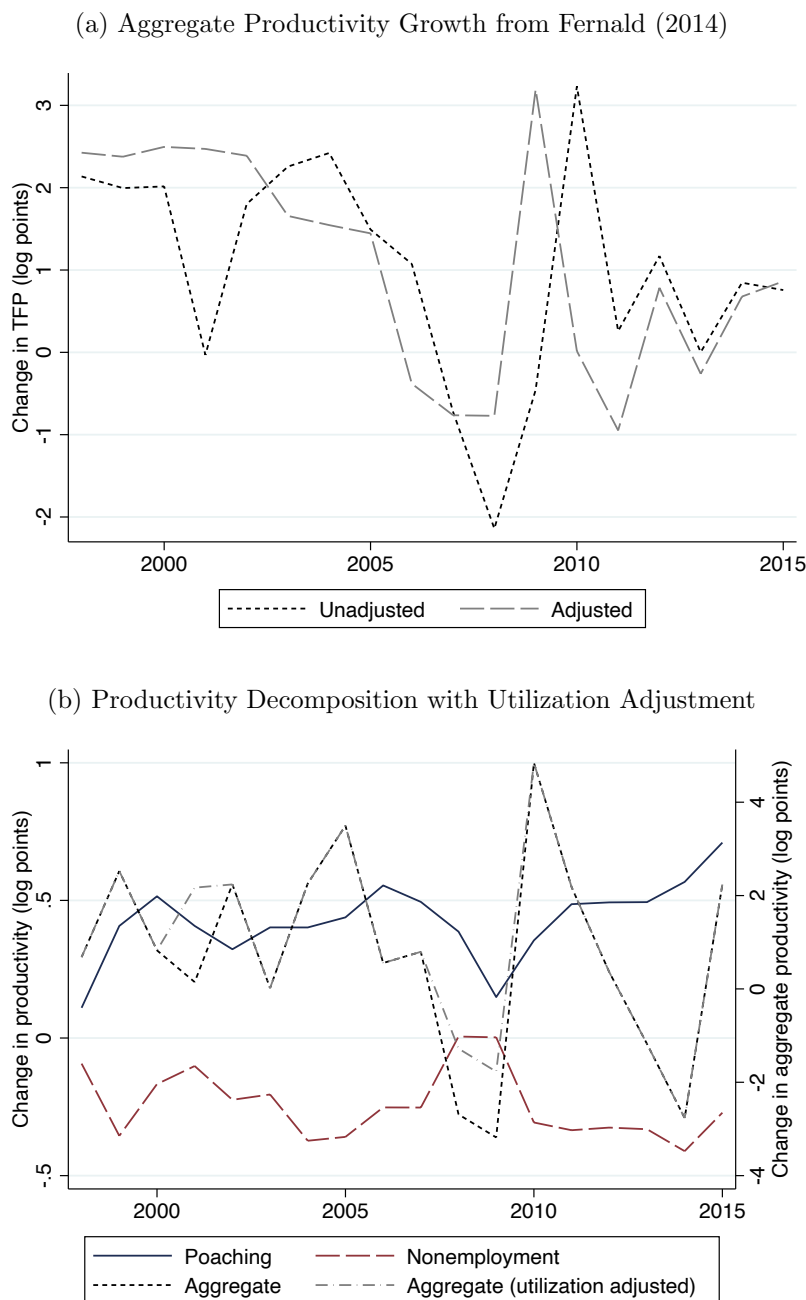
Regardless of the source, these residuals flows are both small in magnitude and do not exhibit a clear pattern across the business cycle. Specifically, the average size of the differential net poaching and nonemployment growth rates are three and five times as large as the differential residual flows, respectively. Figure B.3 in Appendix B presents a version of Figure A.1 that contains the residual flows as well as the poaching and nonemployment flows constructed with and without the restriction that both origin and destination employers are in the 28-state sample. The results indicate that the residual flows do not exhibit any notable cyclicity across the business cycle and the differential net poaching and nonemployment growth rates that exclude workers moving in and out of our 28-state sample are very similar (in levels and movements across time) to the results from Figure A.1. We infer that this residual term is not important for the main results.

Figure B.1: Decomposition and Intertemporal Variation in Factor Utilization



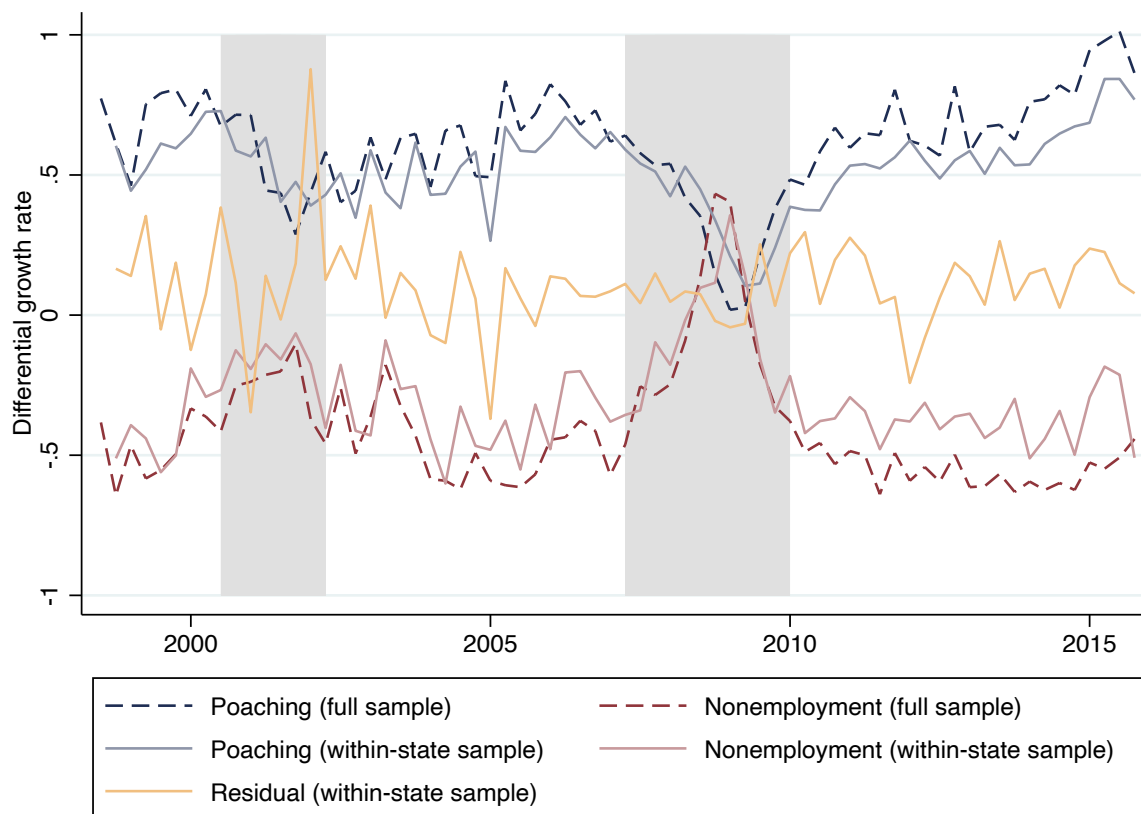
Notes: Panel (a) plots the difference between the average productivity at high- and low-productivity firms. Panel (b) presents the components of productivity growth that attributable to worker reallocation between high- and low-productivity firms through poaching and nonemployment flows. The results in Panel (b) implement the decomposition using both the observed productivity differentials (depicted in Panel (a)) as well as the productivity differentials from a smoothed series generated by fitting the observed productivity differentials with a linear time trend. Data are seasonally adjusted using X-12.

Figure B.2: Aggregate Productivity Growth with Factor Utilization Adjustment



Notes: Panel (a) presents data from Fernald (2014) on the growth in business sector total factor productivity (TFP) as well as a measure that implements an adjustment for variation in factor utilization. Panel (b) presents the annual productivity growth from worker reallocation through poaching and nonemployment flows as well as aggregate productivity growth. In addition, Panel (b) includes a series that adjusts aggregate productivity growth using the difference between the unadjusted and adjusted measures of growth from Fernald (2014).

Figure B.3: Differential Flows between High- and Low-Productivity Firms with Residual



Notes: Differential growth rates are the difference in quarterly employment growth rates between high- and low-productivity firms. Results are presented separately for poaching, nonemployment, and residual flows. The poaching and nonemployment flows depicted by the dashed lines are based on the full sample. The poaching and nonemployment flows depicted by the solid lines are based on a sample that is limited to worker flows in which both the origin and destination employers are in one of the 28 states in our sample. The residual growth rate is the difference between the observed changes in the share of employment at high- and low-productivity firms and what is predicted by the poaching and nonemployment flows. Data are seasonally adjusted using X-12.

Appendix C Decomposition Methodology

Productivity growth for a given industry can be decomposed into two components,

$$\Delta R_t(k) = (R_{t-1}^h(k) - R_{t-1}^l(k))\Delta\theta_t^h(k) + \sum_{i \in \{l, h\}} \theta_t^i(k)\Delta R_t^i(k), \quad (\text{C.1})$$

where the first and second components are productivity growth attributable to worker reallocation between high- and low-productivity firms and productivity growth within the high- and low-productivity firm groups, respectively. Aggregate productivity growth can then be written as the employment weighted average across industries,

$$\sum_k \theta_{t-1}(k)\Delta R_t(k) = \sum_k [\theta_{t-1}(k)(R_{t-1}^h(k) - R_{t-1}^l(k))\Delta\theta_t^h(k)] + \sum_k \left[\theta_{t-1}(k) \sum_{i \in \{l, h\}} \theta_t^i(k)\Delta R_t^i(k) \right]. \quad (\text{C.2})$$


Equation C.2 is an exact decomposition whereas equation 3 is an approximation. The two equations differ only with respect to the first term where,

$$\sum_k [\theta_{t-1}(k)(R_{t-1}^h(k) - R_{t-1}^l(k))\Delta\theta_t^h(k)] \approx (R_{t-1}^h - R_{t-1}^l)\Delta\theta_t^h \quad (\text{C.3})$$

This approximation suggests two ways of quantifying how worker reallocation affects productivity growth. The left-hand-side suggests we could decompose changes in the share of workers at high-productivity firms within each industry, and then aggregate results have. The right-hand-side suggests that we can simply decompose the aggregate share of workers at high-productivity firms. Given that there are 311 distinct 4-digit industry codes, the latter is preferable to the extent that this approximation is accurate.

We directly calculate both sides of equation C.3 show that the approximation is highly accurate. To show that the approximation performs well, we regress the true value, $\sum_k [\theta_{t-1}(k)(R_{t-1}^h(k) - R_{t-1}^l(k))\Delta\theta_t^h(k)]$, on the approximate value, $(R_{t-1}^h - R_{t-1}^l)\Delta\theta_t^h$, which yields a coefficient of 0.85 and an R-squared of 0.91. Figure C.1 plots these two series.

Figure C.1: True Value and Approximation



Notes: The figure presents the exact values of $\sum_k [\theta_{t-1}(k)(R_{t-1}^h(k) - R_{t-1}^l(k))\Delta\theta_t^h(k)]$ and its approximation $(R_{t-1}^h - R_{t-1}^l)\Delta\theta_t^h$.