Reallocation in the Great Recession: Cleansing or Not?

Lucia Foster, *US Census Bureau*

Cheryl Grim, *US Census Bureau*

John Haltiwanger, *University of Maryland and National Bureau of Economic Research*

The high pace of reallocation across producers is pervasive in the US economy. Evidence shows that this high pace of reallocation is closely linked to productivity. While these patterns hold on average, the extent to which the reallocation dynamics in recessions are "cleansing" is an open question. We find that downturns prior to the Great Recession are periods of accelerated reallocation even more productivity enhancing than reallocation in normal times. In the Great Recession, we find that the intensity of reallocation fell rather than rose and that the reallocation that did occur was less productivity enhancing than in prior recessions.

I. Introduction

The Great Recession was unusually severe and persistent relative to post-WWII recessions; we explore whether its impact on productivity-enhancing reallocation was also unusual. A pervasive feature of the US

Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the US Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Haltiwanger thanks the Kauffman Foundation for support. We thank Robert Kulick for his superb research assistance. We thank David Card, Erik Hurst, Ron Jarmin, Alex Mas, Javier Miranda, and participants at the COMP_NET conference in Dublin and the NBER Conference on "The Labor Market in the Aftermath of the Great

[*Journal of Labor Economics*, 2016, vol. 34, no. 1, pt. 2]
 0734-306X/2016/34S1-0010$10.00
Submitted September 9, 2013; Accepted March 2, 2015; Electronically published December 16, 2015

economy is a high pace of output and input reallocation across producers.[1] The annual average job creation rate for the US private sector over the past 30 years is close to 18%, while the analogous job destruction rate is 16%. Evidence shows that this high pace of reallocation is closely linked to productivity dynamics: resources are shifted away from low-productivity producers toward high-productivity producers. An open question is whether recessions have a "cleansing" impact by accelerating this productivity-enhancing reallocation. Theory suggests that the nature and extent of productivity-enhancing reallocation could be fundamentally altered by the nature of the downturn. Using micro-level data, we examine how the pattern of reallocation differed in the Great Recession in terms of both intensity and the extent to which it was productivity enhancing.

Whether recessions are a period of productive winnowing or of counterproductive destruction has been the subject of a long ongoing debate. Economists trace the genesis of the debate back to Schumpeter's (1939, 1942) discussion of creative destruction. The cleansing hypothesis is that a recession is a time of accelerated productivity-enhancing reallocation because it is a relatively low-cost time for reallocation.[2] Alternative hypotheses highlight the potential distortions to reallocation dynamics in recessions. Such distortions could arise from many factors. For example, if credit markets are distorted in a recession, reallocation may be driven more by credit constraints and less by market fundamentals such as productivity, demand, and costs. The close connection between the financial crisis and the Great Recession suggests that this hypothesis might be especially relevant in the recent period.[3]

Prior research suggests that the recession in the early 1980s is consistent with the cleansing hypothesis.[4] Of particular relevance, Davis and

---

Recession" for their helpful comments. Information concerning access to the data used in this article is available as supplementary material online in a zip file. Contact the corresponding author John C. Haltiwanger at haltiwan@econ.umd.edu.

[1] We use productivity differences across producers as a placeholder more generally for differences across producers in terms of technical efficiency, demand, and costs. All of these factors contribute to our empirical measure of establishment-level productivity, as we discuss in Sec. II.

[2] It is important to emphasize that the modern view of the "cleansing" hypothesis does not mean that recessions are necessarily welfare enhancing. The social planner may prefer to avoid cyclical variation in consumption along with the loss of activity from unemployment. But, conditional on the cycle occurring, it may be optimal to increase the pace of productivity-enhancing reallocation given that the opportunity cost of time is low.

[3] Greenstone and Mas (2012) and Chodorow-Reich (2014) find that a relationship between the Great Recession credit market shock and subsequent employment declines.

[4] Foster, Haltiwanger, and Krizan (2001) find that the 1977–82 and 1982–87 periods were times of especially intense productivity-enhancing reallocation. Davis, Haltiwanger, and Schuh (1996) highlight the increased intensity of reallocation

Haltiwanger (1990, 1992, 1999) show that job reallocation activity increased during recessions in the manufacturing sector from the late 1940s through the 1990s.[5] Extending the analysis to the entire private sector, Davis, Faberman, and Haltiwanger (2006, 2012) find that these patterns for manufacturing also hold for the entire private sector for the recessions in the 1990s and 2000s prior to the Great Recession.

The empirical finding that job destruction is more cyclical than job creation is directly related to the cleansing hypothesis, as shown in a series of models developed in the early 1990s. In these models, the marginal cost of creating jobs is lower in recessions, so while creation falls in recessions, it falls less than the rise in destruction. Possible reasons for this include that the opportunity cost of time is low in recessions (Davis and Haltiwanger 1990), the sunk cost of job creation is increasing in the level of aggregate activity (Caballero and Hammour 1994), or the search and matching framework is such that the marginal cost of creating a job is lower in recessions because it is easier for firms to fill jobs in slack labor markets (Mortensen and Pissarides 1994). In all of these models, reallocation is productivity enhancing. These models provide rationales for why the intensity of reallocation may increase in recessions.

Caballero and Hammour (1996) highlight distortions, such as hold up problems and bargaining problems, which may have an impact on incentives for job creation and job destruction over the cycle. In particular, they note that if the marginal cost of creating jobs is lower in recessions, then the social planner would have job destruction rise first, followed quickly by an increase in job creation in recessions. They emphasize that recessions with a rise in job destruction, a decline in job creation, and only a very slow recovery in job creation are a sign of inefficiency.

Beyond the distortions emphasized by Caballero and Hammour (1996), numerous mechanisms can yield "sullying" or "scarring" effects of recessions. Barlevy (2003) develops a model building on the credit market imperfections of Bernanke and Gertler (1989). In his model, recessions are cleansing in the absence of financial constraints, but the cleansing effect can

during the 1982–83 recession. Recent research by Collard-Wexler and De Loecker (2013) shows that this type of reallocation was responsible for much of the productivity growth in the US steel industry over the past several decades.

[5] Blanchard and Diamond (1990), Davis and Haltiwanger (1990, 1999), and Caballero and Hammour (2005) use vector autoregression analysis to conduct a more nuanced and sophisticated analysis of the behavior of job reallocation over the cycle. As they emphasize, exploring the cumulative impulse response functions of job creation, destruction, and—in turn—reallocation in response to an econometric specification that explicitly identifies the aggregate shocks provides a more comprehensive analysis than simple descriptive statistics of the cyclical patterns of job creation, destruction, and reallocation. Given that we do not conduct such analysis here, our characterization of the cyclical dynamics of creation and destruction here should be viewed as suggestive.

be reversed when financial constraints are present. He shows that it is possible that recessions are times of increased, but noncleansing, reallocation.[6] In contrast, Osotimehin and Pappadà (2013) develop an alternative related model where credit frictions have a distortionary effect on the selection of exiting firms but do not reverse the cleansing effect of recessions. The difference in these models is the interaction of productivity and credit constraints. Barlevy argues that the most productive businesses are likely to be more subject to credit constraints. Osotimehin and Pappadà argue that the most productive firms face more forgiving net worth exit thresholds and are more likely to face better draws of idiosyncratic productivity shocks.[7]

With these remarks as background and motivation, this paper addresses five empirical questions concerning the potential cleansing effects of the Great Recession. First, do patterns of reallocation over the business cycle change in the Great Recession? Second, is reallocation productivity enhancing? Third, does the nature of the relationship between productivity and reallocation change over the business cycle? Fourth, is the relationship between productivity and reallocation we see in earlier recessions different in the Great Recession? Fifth, what are the aggregate implications of changes in these micro-level relationships?

In the first part of our empirical analysis, we find a significant change in the responsiveness of job creation and destruction to cyclical contractions in the Great Recession relative to prior recessions. In earlier cyclical downturns, periods of economic contraction exhibit a sharp increase in job destruction and a mild decrease in job creation, consistent with the earlier literature. However, in the Great Recession, job creation fell by as much or more than the increase in job destruction. In this respect, the Great Recession was not a time of increased reallocation (whether productivity enhancing or not).

The second part of our analysis investigates the relationship between productivity and reallocation. We find that low-productivity establishments are more likely to exit, whereas high-productivity establishments are more likely to grow. In turn, we find that the marginal impact of productivity on exit and growth changes over the cycle. For recessions before the Great Recession, the marginal impact of productivity on exit and growth increases with the magnitude of the contraction. However, this is reversed in the Great Recession. More productive establishments still have lower exit rates and higher growth rates in the Great Recession,

[6] Using Colombian establishment-level data, Eslava et al. (2010) present empirical evidence that the exit margin is distorted in times of financial constraints in a manner consistent with the model of Barlevy (2003). Barlevy (2002) focuses on worker matching issues as other possible reasons recessions can have sullying effects.

[7] Other mechanisms that work against cleansing have been proposed. For example, Ouyang (2009) argues that recessions stifle learning opportunities important for the development of young firms.

but the difference in the exit and growth rates between high- and low-productivity establishments declines with sharp contractions. We also find that these patterns are primarily driven by establishments of young firms.[8] Finally, we find that changes in these micro-level relationships have aggregate implications. In sum, we find that the cleansing impact of earlier recessions attenuates in the Great Recession. If the cleansing effect of a recession is its "silver lining," we find that this silver lining is tarnished in the Great Recession.

The paper proceeds as follows. The next section describes the data and measurement issues. In Section III, we analyze job reallocation over the business cycle. We bring together reallocation and productivity measures in Section IV to address our central question about the cleansing effect of the Great Recession. Section V concludes and offers ideas for future related areas of research.

II. Data and Measurement Issues

We describe our measures of reallocation and productivity in this section. We rely heavily on the growing existing literature on measuring these concepts using micro-level data. Our primary data sources are administrative, census, and survey establishment-level data from the US Census Bureau. These annual data cover the period from about the mid-1970s to 2011, thus enabling us to compare the Great Recession to earlier recessions.[9] We are able to examine reallocation for the entire US economy, but for reasons of data availability, we are constrained to the manufacturing sector when analyzing productivity. We begin by describing how we measure reallocation over the business cycle (this relates to the analysis in Sec. III). We then describe how we measure productivity and reallocation in an integrated manner (this relates to the analysis in Sec. IV). Details about data sources and measuring productivity, weights, and job flows are given in the appendix, available online (in Secs. A–D, respectively).

A. Reallocation

Our annual job reallocation measures for the entire US economy and the manufacturing sector are from the Business Dynamics Statistics series,

[8] Fort et al. (2013) find that young and small firms are hit especially hard in the Great Recession. They find that the decline in housing prices is important in that context.

[9] We use a variety of data sources, some of which cover different periods. The public domain Business Dynamics Statistics (BDS) has job flows from 1977 to 2011. The internal version of the Longitudinal Business Database (LBD), on which the BDS is based, is available from 1976 to 2011. The Annual Survey of Manufactures (ASM) and the Census of Manufactures (CM) data that we use to measure productivity are available from 1972 to 2010. We integrate these data with the LBD so that we can examine outcomes in the LBD from t to $t + 1$ (starting in 1981 and looking at outcomes through 2011) using productivity through 2010. (See Sec. A of the online appendix).

which is a public use data set derived from the Longitudinal Business Database.[10] The LBD is a longitudinally linked version of the Census Bureau's business register. As such, the LBD covers all establishments with paid employees in the nonagricultural private sectors of the US economy (see Jarmin and Miranda 2002).

Measures of job flows in the BDS are consistent with the methodology from Davis et al. (1996). Davis et al. (1996) measure job creation as the employment gains from all expanding establishments including startups and job destruction as the employment losses from all contracting establishments including shutdowns. The job reallocation rate is the sum of the job creation and job destruction rates (see appendix D online).

Measures of reallocation can be calculated for various groups of establishments, including establishment and firm age groups, establishment and firm size groups, establishment location (region, state) groups, and establishment industry groups.[11] In addition, the measures of reallocation can be disaggregated into intensive and extensive margins. Establishment births are those establishments that did not exist in time $t - 1$ but exist in time t; analogously, establishment deaths are those establishments that existed in time $t - 1$ but do not exist in time t. All designations of births and deaths rely upon the complete universe of information from the LBD.[12]

While most of our analysis of job flows relies on the BDS, we supplement this analysis with an alternative public domain source of jobs flows. The Business Employment Dynamics (BED) is a longitudinal version of Bureau of Labor Statistics' (BLS) Quarterly Census of Employment and Wages. The BED covers the private economy and thus provides a quarterly analog to the annual data provided by the BDS (although coverage and measurement issues make comparability complicated).[13] The methodology for measuring job flows in the BED is essentially the same as that for the BDS.[14]

[10] BDS data are available at http://www.census.gov/ces/dataproducts/bds/.

[11] We follow Haltiwanger, Jarmin, and Miranda (2013) in our measurement and definitions of establishment and firm size and age. Age of a firm is based on the age of the oldest establishment at the time of the new firm's inception. After that, a firm ages naturally regardless of changes in composition. See Haltiwanger et al. (2013) for more on the distinction between establishments and firms in the LBD.

[12] The establishment links in the LBD are of high quality given the comprehensive administrative data underlying the LBD. Davis et al. (1996) rely upon the ASM and the CM to create measures of job creation and destruction using US Census micro-level data. Using the ASM, with its rotating panels of establishments, introduces measurement complexities we avoid by using the LBD.

[13] See Davis et al. (2012) for a discussion of the cyclical dynamics of job flows in the BED. See Haltiwanger, Jarmin, and Miranda (2011) for a discussion of the cyclical dynamics of job flows in the BDS.

[14] We use BED statistics from Davis et al. (2012) that have been extended back to 1990:2.

B. Connection between Productivity and Reallocation

To explore the connection between productivity and reallocation, we use establishment-level data from the US Census Bureau. We integrate the establishment-level LBD with establishment-level data from the ASM and the CM.

Our analysis of the relationship between reallocation and productivity dynamics over the cycle is restricted to the US manufacturing sector. We find that job creation and destruction dynamics for manufacturing largely mimic the patterns for the whole economy. While there are some differences between the overall private and manufacturing sectors in terms of their cyclical dynamics of job flows in the Great Recession, we believe our analysis of the connection between productivity and establishment survival and growth in the manufacturing sector should be of relevance more broadly as well.

We begin by identifying all manufacturing establishments in the LBD from 1976 to 2011. We compute measures of growth and survival using Davis et al.'s (1996) methodology for these establishments. Specifically, let E_{it} be employment at establishment i in year t, that is, the number of workers on the payroll in the pay period covering March 12. The employment growth rate is $g_{it} = (E_{it} - E_{it-1})/X_{it}$, where $X_{it} = .5(E_{it} + E_{it-1})$.[15] In turn, we generate indicators of the components of growth—Davis et al. (1996) growth rates for continuing establishments and indicators of establishment entry and exit. All of these measures are based on the full LBD and do not require any information from the ASM/CM data. Our measures of firm size and firm age are also derived from the full LBD and are not dependent on the ASM/CM data. We adopt the timing convention that the growth rate from March of year t to March of year $t + 1$ represents the t to $t + 1$ growth rate (e.g., a 2010 outcome reflects the change from March 2010 to March 2011). Thus, our analysis of the connection between productivity and reallocation reflects outcomes from t to $t + 1$ as a function of establishment-level total factor productivity (TFP) and other measures (e.g., firm size and firm age) in period t. We now turn to how we construct establishment-level measures of TFP in year t.

To construct a measure of TFP to integrate with these LBD measures, we rely on the subsample of establishments present each year in either the ASM or the CM from 1972 to 2010 to create an analytic data set for 1981–2010.[16] The CM is, in principle, the universe of establishments, but data

[15] This growth rate measure has become standard in analyses of establishment and firm dynamics because it shares some useful properties of log differences while also accommodating entry and exit. See Tornqvist, Vartia, and Vartia (1985) and Davis et al. (1996) for discussion.

[16] While we use data back to 1972 to get the best possible capital stock measures, our analysis uses data from 1981 to 2010. We focus on this period since we are

are collected only from those establishments' mailed forms. Very small establishments (where the size threshold varies by industry) have their data imputed from administrative data, so we exclude those cases. The CM is collected every 5 years in years ending in 2 and 7.

The ASM is collected in all years where a CM is not collected and is a sample of roughly 50,000–70,000 manufacturing establishments. Probability of selection into the ASM sample is a function of industry and size. Thus, in both ASM and CM years, we have a subset of establishments of the comprehensive universe from the LBD. To deal with this issue, we estimate propensity score weights for each establishment-year observation in the LBD. The weights are based on the probability that an establishment is in the ASM or CM (nonadministrative record cases) in a specific year. As we show in online appendix C, using such propensity score weights enables our weighted sample to replicate the size, age, and industry distributions in the LBD as well as the overall patterns of employment in the LBD. Note that we estimate the propensity score models separately for each year, which enables us to take into account the changing nature of our samples (e.g., CM vs. ASM years). For all of our statistical analysis using the matched ASM/CM/LBD data, we use these propensity score weights.[17]

We measure TFP at the establishment level by constructing an index in a manner similar to that used in Baily, Hulten, and Campbell (1992) and a series of papers that built on that work.[18] The index is given by

$$\ln \mathrm{TFP}_{et} = \ln Q_{et} - \alpha_K \ln K_{et} - \alpha_L \ln L_{et} - \alpha_M \ln M_{et}, \tag{1}$$

where Q is real output, K is real capital, L is labor input, M is real materials, α denotes factor elasticities, the subscript e denotes individual establishments, and the subscript t denotes time. Details on measurement of TFP are in appendix B online, so here we focus on the most relevant features of how these various components are measured. Operationally, we define nominal output as total shipments plus the change in inventories. Output is deflated using an industry-level measure from the NBER-CES Manufacturing Industry Database. Capital is measured separately for

interested in classifying establishments based on the age of their parent firm. Our firm age measure is left-censored for firms born in or before 1976. As such, in 1981 and beyond, we can consistently classify firms into age classes of less than 5 and 5 or more years old.

[17] The ASM has sample weights, which could in principle be used instead. However, the sample-weighted ASM is not designed to match published totals, as discussed in Davis et al. (1996). Moreover, our method implies that we are capturing the patterns of the universe LBD data. Finally, our method facilitates using the CM and ASM records in a consistent manner.

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

structures and equipment using a perpetual inventory method. Labor is measured as total hours of production and nonproduction workers. Materials are measured separately for physical materials and energy and where each are deflated by an industry level deflator. Outputs and inputs are measured in constant 1997 dollars.

We measure the factor elasticities using industry-level total factor cost shares. We could measure these factor elasticities at the establishment level; however, arguments against using an establishment-level approach can be made when factor adjustment costs exist (see Syverson 2011). Moreover, for related reasons, Syverson (2011) notes that some time averaging may be warranted at the industry level. Accordingly, for an establishment in a given industry in period t, we use industry-level measures of cost shares for period t based on the average of the t and $t - 1$ cost share for the factor elasticity.[19]

Given the large differences in output measures across industries (e.g., steel vs. food), our TFP measures need to control for industry differences in any comparison over industries. We do this by creating measures of (log) TFP that are deviations from the industry-by-year average. We refer to this as TFP in the remainder of the paper, but it should be interpreted as the deviation of establishment-level TFP from the industry-by-year average.

As noted above, our measure of productivity is a revenue measure. This means that differences in establishment-level prices are embedded in our measure of productivity. Unfortunately, the Census Bureau does not collect establishment-level prices. However, as Foster, Haltiwanger, and Syverson (2008) have shown, it is possible to back out the establishment-level price effects for a limited set of products in Economic Census years. Foster et al. (2008) create a physical quantity measure of TFP removing the establishment-level price for establishments producing a set of 11 homogeneous goods (e.g., white pan bread). The within-industry correlation between revenue and physical productivity measures in Foster et al. (2008) is high (about 0.75). However, they also find that there is an inverse relationship between physical productivity and prices consistent with establishments facing a differentiated product environment. In addition, they find that establishment-level prices are positively related to idiosyncratic

[19] As discussed in Syverson (2011), there are numerous alternative ways to measure factor elasticities (e.g., estimation methods using either IV or proxy methods to address endogenous factors), but these alternative methods tend to produce similar establishment-level TFP measures (even if they produce somewhat different factor elasticities). We also consider industry-level cost shares averaged over our entire sample and obtain very similar results. Our approach is related to the Divisia/Tornqvist index number approach, but this latter approach is focused on an index of TFP growth over time. Our focus is on generating a relative productivity measure across establishments within years.

Table 1
Descriptive Statistics, ASM/CM/LBD Matched Sample

	Mean	SD
Overall growth rate (continuers + exit)	−.17	.65
Young	−.26	.85
Mature	−.15	.59
Establishment exit	.08	.27
Young	.15	.35
Mature	.07	.25
Conditional growth rate (continuers only)	−.01	.38
Young	.04	.49
Mature	−.02	.35
Establishment entry	.07	.25
TFP	.000	.360
Young	−.011	.353
Mature	.003	.362
Cycle	.0004	.0107
Young	.19	.39
GR	.09	.28
Years	1981–2010	
N (millions)	2.2	

Source.—Authors' calculations on the Annual Survey of Manufactures, the Census of Manufactures, and the Longitudinal Business Database.

Note.—Statistics use propensity score weights to make the sample representative of the LBD. Statistics are not activity weighted. Rates are in fractions (not percents). Employment growth and exit are measured from period t to period $t + 1$. TFP is the deviation of establishment-level log TFP from its industry-year mean in year t, so the mean is, by construction, equal to zero. Cycle is the state-year change in the unemployment rate from t to $t + 1$. Young is a dummy variable equal to one for establishments that belong to firms less than 5 years old, and GR is a dummy variable equal to one for years from 2007 to 2009.

demand shocks. As such, our measure of establishment-level productivity should be interpreted as reflecting both idiosyncratic technical efficiency and demand factors. However, we only capture the idiosyncratic demand factors as they translate into establishment-level prices. It is important to emphasize that it is only idiosyncratic (not aggregate or industry-level) demand shocks that are potentially captured by our TFP measure given that our measure deviates from industry by year means.

Summary statistics of our integrated establishment-level sample are provided in table 1. We have roughly 2.2 million establishment-year observations from 1981 to 2010. We measure growth rates and survival rates for all of these establishments based on the LBD from t to $t + 1$. In table 1, the growth rate for incumbent establishments is negative.[20] By design, this growth rate does not include the contribution of entry. The growth rate for continuing establishments is about −1%, and the slightly higher exit rate (8%) compared to entry rate (7%) implies that the overall growth

[20] These statistics use the propensity score weights to adjust the sample, but they are not activity weighted.

rate inclusive of entry (not reported) is about −2%. Variable TFP represents the deviation from industry-year means, so by construction it has a mean of zero. The within-industry-by-year dispersion in TFP is similar to that reported in Syverson (2004). The cyclical variable that we focus on (called "Cycle," the change in the state-level unemployment rate from the Current Population Survey [CPS]) has a mean around zero but with substantial variation.[21] It is not uncommon for individual states to experience changes in unemployment of 0.03 in a given year in the Great Recession. About 20% of establishments belong to young firms, and the Great Recession dummy applies to fewer than 10% of our establishment-year observations.

We also show summary statistics with establishments classified into young and mature (based upon the age of the firm). We find that growth rates for young businesses (excluding startups) are lower than for mature businesses, but this reflects a substantially higher growth rate for continuing young businesses and a substantially higher exit rate for young businesses.

III. Did Reallocation Dynamics Change in the Great Recession?

In this section, we present results of our analysis of the patterns of job creation and job destruction over the cycle. We start by examining job flows using data from the BDS series, which provides annual job flow statistics for the entire US private sector. Panel *A* of figure 1 shows the job creation and job destruction rates for the US economy from 1981 to 2011. The figure also includes the change in the unemployment rate.[22] It is apparent that job destruction tends to rise and job creation tends to fall during periods of increasing unemployment. Interestingly, it appears that this pattern changed in the Great Recession. Job destruction did rise sharply in the 2008–9 period, but what is more striking is the sharp fall in

[21] We use this measure because it allows for variation at national and state levels and is highly correlated with other measures indicative of the business cycle. Correlations between the change in the national unemployment rate and other cyclical indicators are as follows: GDP growth (−0.92), net employment growth (0.93), and change in employment-to-population (over 16 years old) ratio (−0.95). We prefer measures of the cycle that correspond to measures of change and growth, as opposed to measures that capture deviations of levels from trends, because the change and growth measures are much more highly correlated with our outcomes of interest (i.e., employment growth). For example, at the national level, the correlation between the Hodrick-Prescott filtered unemployment rate and net employment growth is only −0.23.

[22] The change in the unemployment rate is the March-to-March change to match the timing of our job flows series. All measures of growth and change (e.g., job flows and unemployment rate) are measured as percents in this section, while they are measured as fractions in other parts of the paper. We use rates in percents in this section since it facilitates discussion of trends.

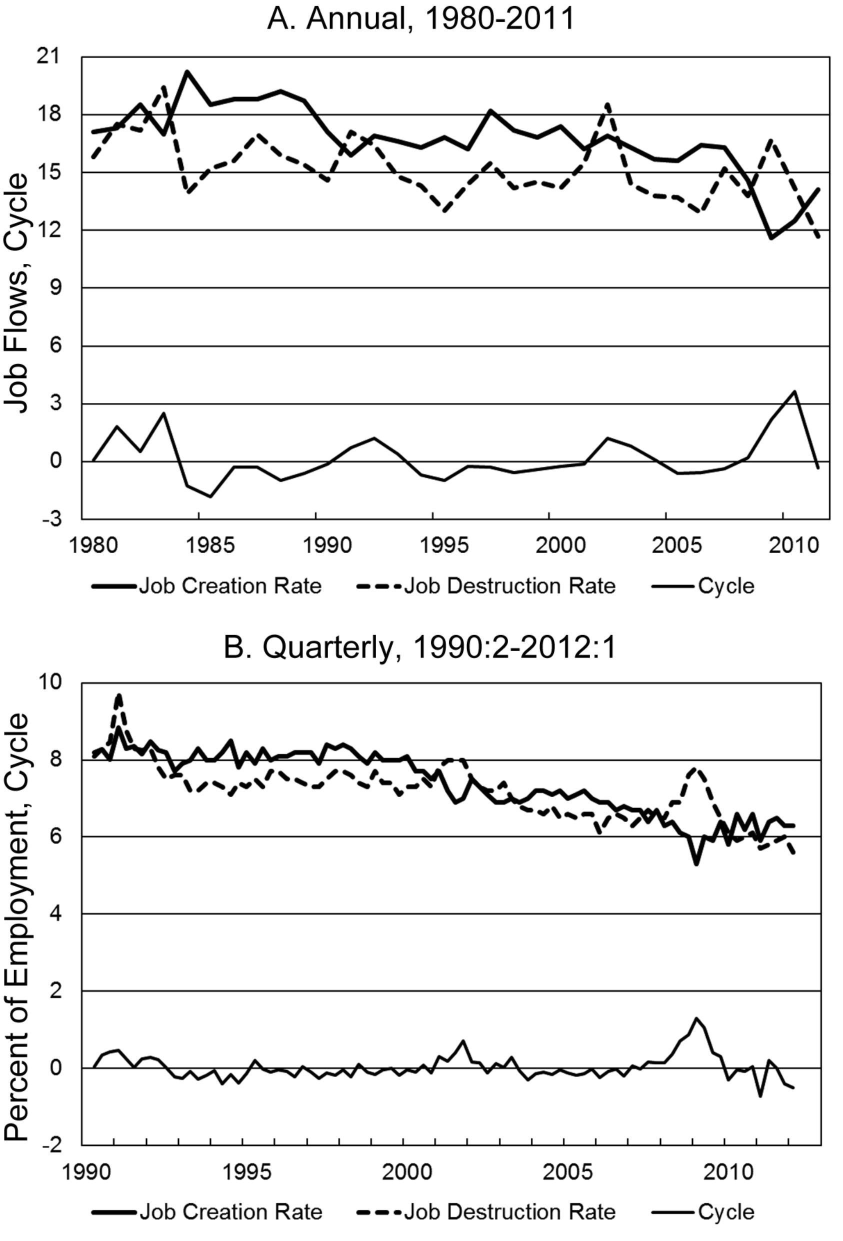

FIG. 1.—Job flows and the business cycle. Authors' calculations using Business Dynamics Statistics (annual), Business Employment Dynamics (quarterly), and the Current Population Survey. Cycle is the change in the unemployment rate.

job creation that starts in 2007 and persists through 2010. We also note that job flows exhibit a downward trend—a point we return to below.

As both a cross check and to explore higher frequency data, we use job creation and destruction series from the BED statistics, which also cover the US private sector. Panel *B* of figure 1 shows quarterly job creation and job destruction rates with the change in the unemployment rate for the period 1990:2–2012:1. The quarterly numbers reinforce the message from the annual data that recessions are periods in which job destruction rises and job creation falls. Again, however, job creation falls sharply in 2007 and stays low. The downward trend in job flows is even more pronounced in the BED. An advantage of the BED is that it is more current: panel *B* of figure 1 shows that the slow recovery from the Great Recession through the first quarter of 2012 is due to anemic job creation rather than from job destruction staying persistently high. Other related data sources (e.g., the Job Openings and Labor Turnover Survey [JOLTS] confirm that this pattern has continued past the first quarter of 2012.

Job creation is as low during the Great Recession as during any period in the past 30 years (see fig. 1). Moreover, job reallocation (creation plus destruction) is at its lowest point in 30 years during the Great Recession and its immediate aftermath. Comparing the Great Recession to the early 1980s recession, job reallocation is 28% in 2009 in contrast to 35% in 1983 (both time periods are when job destruction peaked and are measured using March-to-March BDS data). These patterns are driven in part by the substantial downward trends in job flows evident in both the BDS and the BED.[23] It is well beyond the scope of this paper to explore the determinants of the declining trends in job flows; however, it is clear that downward trends are important, so we take them into account in our analysis.[24]

To assess the changing pattern of job creation during cyclical downturns, we begin with a simple calculation quantifying the fraction of the changes in net employment accounted for by changes in job creation during periods of net contraction. For each episode of net contraction lasting for one or more periods, we cumulate the net employment losses during the episode (in percentage terms) starting from the beginning of each episode. We also cumulate the change (typically a reduction) in job creation over the same episode. These cumulative changes permit computing the fraction of net employment contraction accounted for by the reduction in job creation.[25] A simple example helps illustrate the calculation.

[23] Figure E1 in the online appendix depicts the Hodrick-Prescott trends in the job flows that clearly depict the downward trends.

[24] See Davis et al. (2007), Hyatt and Spletzer (2013), Davis and Haltiwanger (2014), and Decker et al. (2014a, 2014b) for discussions of determinants of declining trends in job flows.

[25] By construction, overall net contraction is accounted for by the cumulative reductions in job creation and the cumulative increases in job destruction.

Suppose that over four consecutive periods net growth is {0, −4, −6, 0}, job creation is {15, 14, 13, 15}, and job destruction is {15, 18, 19, 15}. There is a net contraction during periods 2 and 3. The cumulative net employment decline in periods 2 and 3 is −10 and the cumulative decline in job creation is −3, so the fraction is 0.3.[26]

We sum up these cumulative changes from each cyclical contraction for two subperiods (pre– and post–Great Recession) and compute the fraction for each of these changes.[27] Using this cumulative change per episode largely mitigates concerns about trends, since the cumulative changes are from the start of each cyclical episode.[28] One limitation of this approach when using the national BDS and BED series is the relatively small number of periods over which to make these calculations. To overcome this limitation, we also compute this using state-level job flow series. We then take the average of these fractions across all states.

Table 2 shows the share of the decline in net employment accounted for by declines in job creation during net contractions. The share is substantially below 0.5 for net contractions prior to the Great Recession. Thus, most of the net decline during contractionary periods prior to the Great Recession is accounted for by a rise in job destruction rather than a fall in job creation. In contrast, this share rises substantially above 0.5 in the post-2007 period; during the Great Recession, most of the net decline is accounted for by a decline in job creation.

We shed further light on these patterns by taking advantage of state-level variation in the covariance between cyclical indicators and the job flows. We consider simple descriptive regressions relating job flows to a cyclical indicator and a dummy variable for the Great Recession period interacted with the cyclical variable. For this purpose, we use state-level changes in the unemployment rate.[29] Since we see a negative trend in job flows, we include a linear trend in our specifications.[30] The results are shown in table 3. The specifications have a main effect of the cyclical indicator and an interaction effect. As such, the overall effect for the Great

[26] Notice it is the cumulative decline in job creation from just prior to the start of the current contraction (i.e., the job creation is −1 in period 2 and −2 in period 3 relative to job creation just prior to the start of the current contraction).

[27] This is equivalent to taking the weighted average of the per episode fractions, where the weight is the cumulative net change for the episode.

[28] We cumulate first differences in net employment and job flows so we are effectively detrending.

[29] We consider other cyclical indicators, such as the change in the employment–to-population ratio (population over age 16) and obtain very similar results.

[30] In unreported results, we find similar patterns using the national sample in spite of the relatively sparse number of observations. We also find that the patterns are robust to using alternative detrending methods.

Table 2
Share of Change in Net Employment Growth Due to Change in Job Creation in Periods of Net Contraction

	National		State
Period	BDS (Annual)	BED (Quarterly)	BDS (Annual)
Pre–Great Recession	.21	.28	.39
Post-2007	.61	.59	.65

SOURCE.—Authors' calculations on Business Dynamics Statistics (BDS) and Business Employment Dynamics (BED).

NOTE.—The calculations take advantage of the identity that Net = Job Creation – Job Destruction. For periods of net contraction lasting one or more periods, the cumulative change in net employment growth and cumulative change in job creation are calculated over the entire consecutive period of net contraction. In turn, these cumulative changes are cumulated further within the periods in the table. The share is the fraction of the overall cumulative change in net employment growth over the specified period accounted for by the overall change in job creation over the specified period. For BDS, pre–Great Recession is 1981–2007, post-2007 is 2008–11. For the BED, pre–Great Recession is 1990:2–2007:3, post-2007 is 2007:4–2012:1. As noted, these statistics are only calculated for periods with net employment growth less than zero. For example, this is 2007:4–2010:1 for the BED. For the BDS National Annual there are only 6 years of net contraction with only 2 years in the post 2007 period. For the BED Quarterly, there are 22 quarters of net contraction with 9 quarters in the post-2008 period. For the BDS State Annual, there are 393 state-year observations with net contraction with 112 state-year observations with net contractions in the post-2007 period.

Recession is the sum of the main and interaction effects. During the Great Recession, the relationship between the change in unemployment and job creation becomes more negative, it becomes less positive with job destruction, and its relationship with the reallocation rate shifts from its usual positive relationship to a negative one.

We also explore the extent to which earlier recessions are different from each other (see table E1 in the online appendix). In particular, we estimate specifications equivalent to table 3 where we include a dummy for the 1981–83 recession interacted with the cyclical indicator as well as the GR

Table 3
Job Flows and Change in the Unemployment Rate at the State-Level (Annual), 1981–2011

	Job Creation Rate	Job Destruction Rate	Reallocation Rate
Cycle	−.631***	1.194***	.563***
	(.046)	(.053)	(.068)
GR × cycle	−.371***	−.421***	−.793***
	(.079)	(.079)	(.128)
Trend	−.168***	−.136***	−.304***
	(.010)	(.011)	(.020)

SOURCE.—Authors' calculations on Business Dynamics Statistics.

NOTE.—N = 1,581. GR is a dummy variable equal to one for years from 2008 to 2010 (job flows from March 2007 to March 2010). Cycle is the state-year change in the unemployment rate. All specifications include state fixed effects. Standard errors in parentheses are clustered at the state level.

*** $p < .01$.

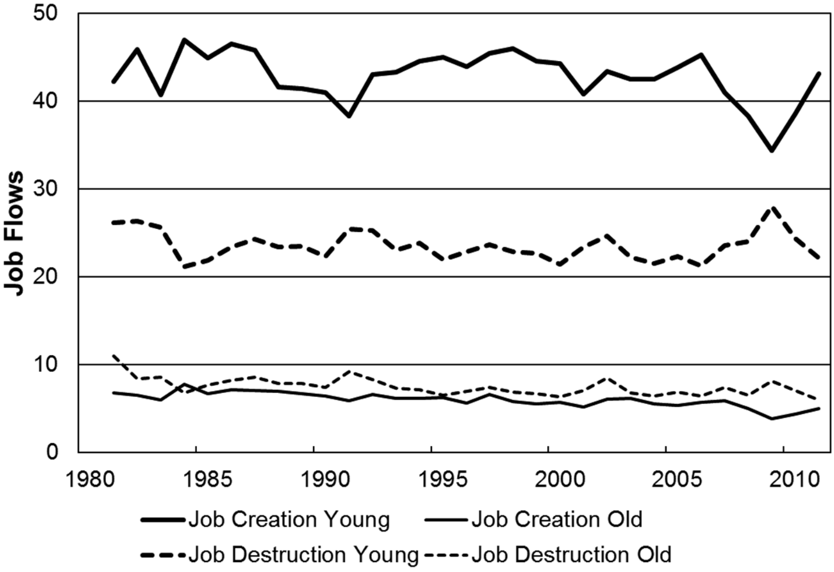

FIG. 2.—Job flows by age, 1981–2011. Authors' calculations on Business Dynamics Statistics. Young is for establishments owned by firms less than 5 years old. Mature is for establishments owned by firms 5 or more years old. Job flows are establishment based and are classified by firm age characteristics.

dummy interacted with the cyclical indicator as in table 3.[31] We find no evidence that the 1981–83 recession differs from other recessions prior to the Great Recession in terms of the nature of the covariance between job flows and the cycle. In contrast, the Great Recession is different. Reallocation fell rather than increased in the Great Recession.

Earlier studies emphasize that the large decline in job creation in the Great Recession is driven by a decline in job creation for young businesses (see Fort et al. 2013). Defining young firms as those less than 5 years old, figure 2 shows patterns of job creation and destruction at the establishment level by firm age class (young and mature).[32] Job creation fell substantially, especially among the very young businesses.[33]

Overall, our evidence points toward the cyclical covariance of job creation and destruction exhibiting different patterns in the Great Recession.

[31] These are simple specifications with main effects and interaction effects, so the overall effect for the early 1980s recession is the main effect plus the interaction effect for the early 1980s recession. The same remarks apply to the Great Recession.

[32] This analysis is based on establishments classified by the characteristics of the parent firm.

[33] We repeat the same type of simple descriptive regressions as in table 3 by these age categories and find that young businesses have greater sensitivity to the cyclical indicator in terms of both job creation and job destruction. We also find that job creation for young businesses fell more with the increase in unemployment in the Great Recession than in prior recessions (see table E2 in the online appendix).

Prior to the Great Recession, destruction is more cyclically sensitive and reallocation rises in cyclical downturns. These patterns are consistent with the reallocation timing and cleansing models of Davis and Haltiwanger (1990), Caballero and Hammour (1994), and Mortensen and Pissarides (1994). However, in the Great Recession these patterns changed. Job creation fell much more substantially and job destruction rose less, resulting in little, if any, increase in reallocation (the BDS estimates actually yield a decline in reallocation in the Great Recession). The trend decline in job flows also plays a role in these dynamics. The low job creation and reallocation rates in the Great Recession and its aftermath are driven by both trend and cyclical factors.

These patterns do not provide direct information about whether the greater intensity of reallocation in prior recessions is actually productivity enhancing or whether the slowdown in reallocation in the Great Recession also exhibited changes in the nature of reallocation. To address these questions, we need to explore the relationship between productivity and reallocation.

As a final point for this section, we also find that the patterns for the private sector tend to hold for the manufacturing sector (shown in online appendix fig. E2). This is relevant since our analysis of the cyclical relationship between productivity and reallocation is confined to the manufacturing sector, where we can measure TFP at the micro level. The different patterns of recessions are especially apparent in comparing the 2001 downturn and the Great Recession. During the 2001 downturn, there was a sharp rise in job destruction, with relatively little response of job creation in the manufacturing sector. In contrast, in the Great Recession, while job destruction also exhibits a substantial increase, there is a much more notable decline in job creation. When we conduct the same type of exercise as in table 2 for manufacturing, we find that the share of cumulative net losses during contractions accounted for by job creation is equal to 0.13 in contractions prior to the Great Recession and equal to 0.28 post-2007.[34] In manufacturing, variation in job destruction still dominates, but variation in job creation plays a substantially larger role in the Great Recession.

IV. Did Cleansing Effects Change in the Great Recession?

Existing research shows a tight connection between reallocation and productivity dynamics: exit is much more likely for low-productivity establishments, while establishment growth is increasing in productivity. A large fraction of industry-level productivity growth is accounted for by this reallocation of outputs and inputs from low-productivity to high-

[34] We calculate these fractions using periods of net contraction for the overall economy.

productivity businesses.[35] Canonical models of firm dynamics by Jovanovic (1982), Hopenhayn (1992), Hopenhayn and Rogerson (1993), and Ericson and Pakes (1995) provide a structure for heterogeneous firm dynamics models, where firms are subject to idiosyncratic productivity, demand, and cost shocks, which have an impact on their growth and survival.

In the analysis that follows, we use empirical specifications consistent with these models to examine whether there is a connection between productivity-enhancing reallocation and the business cycle. We use a simple regression model linking the growth and survival dynamics of incumbent establishments to productivity. We examine entry indirectly by focusing on young versus mature businesses. Complementing our analysis of the dynamics of young firms, we provide some descriptive analysis of where entrants fall in the productivity distribution at the point of entry.

A. Growth and Survival of Incumbents

Our core specification is given by:

$$\begin{aligned} Y_{es,t+1} = \lambda_s + \lambda_{t+1} + \beta(\mathrm{TFP}_{est}) + \gamma(\mathrm{Cycle}_{s,t+1}) \\ + \delta(\mathrm{TFP}_{est} \times \mathrm{Cycle}_{s,t+1}) + \mathrm{X}'_{est}\Theta + \varepsilon_{es,t+1}, \end{aligned} \tag{2}$$

where e is establishment, s is state, Y is a set of outcomes, TFP is total factor productivity deviations from industry by year means, and Cycle is the change in the relevant state unemployment rate from t to $t + 1$.[36] There are three outcomes (all measured from t to $t + 1$): "Overall Growth" (continuers + exit), "Exit," and "Conditional Growth" (conditional on survival, i.e., continuers).[37]

In considering the specification, timing is important. We explore the determinants of growth and survival from t to $t + 1$ based on the pro-

[35] See, e.g., Baily et al. (1992), Campbell (1998), Bartelsman and Doms (2000), Baily, Bartelsman, and Haltiwanger (2001), Foster et al. (2001, 2006), and Syverson (2011).

[36] We report in the online appendix three robustness checks for the cyclical indicator, all of which produce results very similar to those reported in the main text. First, we use specifications without year effects so that variation in the national cycle is used (table E3). Second, we use specifications with year effects but without state effects (table E4). Third, we use the change in the employment-to-population (age 16+) ratio as the cyclical indicator (table E5).

[37] One potential limitation of our approach in using outcomes for manufacturing establishments is that they may be less sensitive to local business cycle conditions than establishments in other sectors. We find that there is a strong relationship between the outcomes of manufacturing establishments and local business conditions. Note: Syverson (2004) finds that many manufactured goods are shipped less than 500 miles. In future work, it would be interesting to consider how the patterns vary by sector (and, in turn, the local nature of the market for the goods).

ductivity of the establishment in period t and the business cycle conditions from t to $t + 1$ (the change in state level unemployment from the CPS).

We estimate this specification for 1981–2010 pooling all years with year effects and controlling for establishment characteristics (including establishment size, firm size, and state effects).[38] The inclusion of year effects implies that we are exploiting state-specific variation in the cycle and that we have abstracted from any of the trend issues (at least national trends) discussed in the previous section. While this is a reduced-form specification, it is broadly consistent with the specifications of selection and growth dynamics from the literature. A common prediction from the models discussed above is that low-productivity plants are more likely to exit. Similarly, the adjustment cost literature for employment dynamics predicts that, conditional on initial size, plants with positive productivity shocks are more likely to grow (e.g., Cooper, Haltiwanger, and Willis 2007). In terms of the empirical literature, there is already much evidence that high-productivity establishments are more likely to survive and grow (e.g., Syverson 2011). Our innovation is to consider how these effects vary over the cycle and in turn across different cycles.

The unit of observation is the establishment in a given state and year. Since some key right-hand-side variables vary only at the state-year level, standard errors are clustered at the state level. We focus on results using clustering at the state level, since Arellano (1987) and Angrist and Pischke (2009) suggest clustering at the state level given potential serial correlation in the state-level regressors. Clustering errors at the state-year level or the state level yields similar results.

To examine the impact of the Great Recession, we expand equation (2) to include effects of the Great Recession:

$$\begin{aligned} Y_{es,t+1} = {} & \lambda_s + \lambda_{t+1} + \beta(\mathrm{TFP}_{est}) + \gamma(\mathrm{Cycle}_{s,t+1}) + \delta(\mathrm{TFP}_{est} \times \mathrm{Cycle}_{s,t+1}) \\ & + \chi(\mathrm{GR}_{t+1} \times \mathrm{TFP}_{est}) + \mu(\mathrm{GR}_{t+1} \times \mathrm{Cycle}_{s,t+1}) \\ & + \phi(\mathrm{GR}_{t+1} \times \mathrm{Cycle}_{s,t+1} \times \mathrm{TFP}_{est}) + X'_{est}\Theta + \varepsilon_{es,t+1}, \end{aligned} \tag{3}$$

where GR is a dummy for the Great Recession taking on values of 1 in years 2007–9.[39]

Results of these regressions are shown in table 4. We first consider specifications without interactions with the Great Recession (cols. 1, 3, and 5). In these specifications, the cross-sectional impact of productivity

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

[39] The dummy GR indicates outcomes from March 2007 to March 2010.

Table 4
Reallocation and Productivity over the Business Cycle

	Overall Growth Rate (Continuers + Exiters)		Exit		Conditional Growth Rate (Continuers Only)	
	(1)	(2)	(3)	(4)	(5)	(6)
TFP	.157***	.159***	−.060***	−.060***	.041***	.042***
	(.006)	(.006)	(.003)	(.003)	(.003)	(.003)
Cycle	−3.307***	−2.961***	.671***	.497***	−2.143***	−2.128***
	(.459)	(.483)	(.176)	(.179)	(.247)	(.286)
TFP × cycle	1.542**	2.182**	−.655***	−.927***	.494	.534
	(.643)	(.862)	(.226)	(.265)	(.412)	(.567)
GR × TFP		.030		−.018*		−.005
		(.023)		(.011)		(.011)
GR × cycle		−3.116**		1.581***		−.126
		(1.349)		(.523)		(.770)
GR × TFP × cycle		−2.961*		1.466**		.066
		(1.619)		(.684)		(.764)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm size class FE	Yes	Yes	Yes	Yes	Yes	Yes
N (millions)	2.2	2.2	2.2	2.2	2.1	2.1

Source.—Authors' calculations on the Annual Survey of Manufactures, the Census of Manufactures, and the Longitudinal Business Database.

Note.—Regressions are weighted by propensity score weights. Weight calculation is described in the online appendix. Standard errors (in parentheses) are clustered at the state level. Employment growth and exit are measured from period t to period $t + 1$. Regression for exit is a linear probability model where exit = 1 if the establishment has positive activity in period t but no activity in period $t + 1$. TFP is the deviation of establishment-level log TFP from its' industry-year mean in year t. GR is a dummy variable equal to one for years from 2007 to 2009 (reflecting outcomes from March 2007 to March 2010). Cycle is the state-year change in the unemployment rate from t to $t + 1$. Establishment size (log employment in t) is included as a control.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

on growth and survival (when the change in the unemployment rate is zero) is given by the first row of columns 1, 3, and 5. Consistent with earlier studies, we find that establishment-level productivity is positively related to growth and negatively related to exit in the cross section. All of these effects are statistically significant.

To assess quantitative significance, figure 3 depicts the implied differences in growth and survival between an establishment one standard deviation below the within industry–year mean and an establishment one standard deviation above the industry-year mean for the main TFP effect (independent of the cycle, so Cycle = 0). For now, we focus on the bars in figure 3 labeled "All." The difference in overall growth rates between an establishment one standard deviation below and above the mean is about 11 percentage points, the analogous difference in exit rates is 4 percentage points, and the difference in the growth of survivors is 3 percentage points. Comparing the magnitudes of the difference for overall growth with the difference for conditional growth, it is evident that the predicted

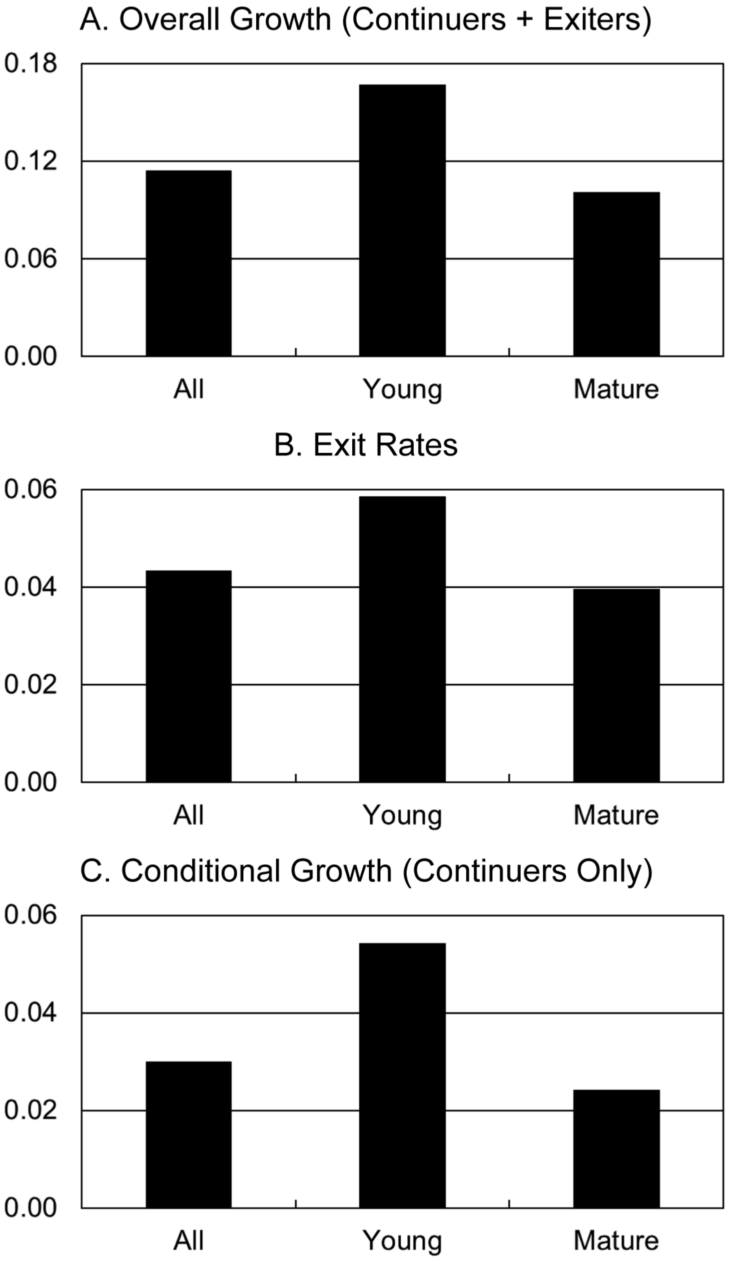

FIG. 3.—Differences in growth rates between high-productivity and low-productivity establishments, normal times. Authors' calculations on Annual Survey of Manufactures, Census of Manufactures, and Longitudinal Business Database. Depicted is the predicted difference in growth rates (panels *A* and *C*, high minus low) and the predicted difference in probability of exit (panel *B*, low minus high) between an establishment one standard deviation above industry-by-year mean productivity and an establishment one standard deviation below industry-by-year mean productivity. Normal is zero change in state-level unemployment.

difference in overall growth rates is largely accounted for by the predicted difference in exit rates.[40]

Returning to table 4, we also find that growth and survival of manufacturing establishments are related to local business cycle conditions. Increases in the state-level unemployment rate are associated with declines in growth and increases in exit. All of these effects are statistically significant and large in magnitude.

Of primary interest, we find that the relationship between productivity and reallocation is enhanced in business cycle contractions. The positive impact of productivity on overall growth and the negative impact of productivity on exit are both increased in magnitude during periods with increases in state-level unemployment. Both of these effects are large in magnitude and statistically significant. We find that the point estimate for this interaction effect is positive for the growth of continuing establishments but not statistically significant at conventional levels. As we discuss below, this is sensitive to permitting effects to vary with firm age.

Did these patterns change in the Great Recession? Columns 2, 4, and 6 of table 4 speak to this question. We are particularly interested in the interaction effect of TFP and the cycle. First, we find the magnitude of the estimated interaction effect between TFP and the cycle is larger for the period prior to the Great Recession than what we find in columns 1, 3, and 5 when we pool all recessions together. This pattern is especially notable for the overall growth and exit specifications. Driving this is the estimated three-way interaction between TFP, the cycle, and the Great Recession, which is reported in the last row of columns 2, 4, and 6. For overall growth, we find the three-way estimated effect is negative and statistically significant. Observe as well that the magnitude of the overall interaction between TFP and the cycle is negative in the Great Recession (adding 2.182 and −2.961). Thus, instead of the cycle enhancing the impact of TFP on overall growth, it tends to diminish it on the margin in the Great Recession. A similar pattern is observed for exit. The estimated three-way interaction effect is positive and larger in magnitude than the two-way interaction effect of TFP and the cycle. Instead of the cycle enhancing the impact of TFP on exit, it tends to diminish it on the margin in the Great Recession. For growth of continuing establishments, we find less systematic patterns. It appears the three-way interaction for overall growth is being driven mostly by the exit margin.

There are other estimated interactions of interest in columns 2, 4, and 6. In particular, we find that the impact of the cycle is even more severe in

[40] Note that the "exit" outcome in table 4 and fig. 3 is from a linear probability model, so there is no simple aggregation of the survival growth and exit outcomes to obtain the overall growth outcome. This requires translating exit into job destruction from exit. The difference between the overall growth and survival growth yields an estimate of the latter (appropriately weighting survival growth for the share of continuing establishments).

terms of its impact on growth and survival in the Great Recession. We also find that the main effects of TFP (independent of the cycle) on growth and survival are slightly enhanced in the Great Recession (although only statistically significant for exit).

We use the same type of exercise as in figure 3 to quantify how the relationship between productivity, growth, and survival changes with the cycle. Figure 4 depicts such exercises for the overall growth and exit outcomes. We focus on overall growth and survival, since we obtain statistically significant effects for the interaction between the effects of TFP and the cycle for these outcomes.[41] The left-most bar, labeled "Normal" (zero change in unemployment), is taken from figure 3. The remaining bars of each figure show how these outcomes vary with the cycle. A "Mild" contraction is a 1 percentage point increase in state-level unemployment, a "Sharp" contraction is a 3 percentage point increase in state-level unemployment, and "GR" is for the period 2007–9 (reflecting outcomes from March 2007 to March 2010). The "Mid" and "Sharp" GR can be thought of as the effect of the Great Recession across different states; some states contract more than others.

We find the difference in overall growth between high-productivity and low-productivity establishments increases substantially when unemployment rises in periods before the Great Recession. In a sharp contraction, the difference in overall growth rates exceeds 15 percentage points (see panel *A* of fig. 4). The Great Recession modifies these patterns. The difference in growth rates between high-productivity and low-productivity establishments is still large in the Great Recession but rather than increasing with unemployment, it falls with increases in unemployment. In a mild contraction in the Great Recession, the difference in growth rates between high-productivity and low-productivity establishments is about 13 percentage points. For a sharp contraction, this falls to about 12 percentage points.

Closely related patterns are exhibited in panel *B* of figure 4 for the exit margin. In cyclical contractions before the Great Recession, the difference in exit rates between low-productivity and high-productivity establishments rises with larger increases in unemployment (note that in panel *B* of fig. 4 we use the difference in exit rates between low-productivity and high-productivity establishments). However, in the Great Recession, this pattern reverses. While there is still a substantially higher probability of exit of low-productivity businesses during the Great Recession, this difference declines with larger increases in unemployment.[42]

[41] For completeness, we show the results for continuing establishments in appendix fig. E3.

[42] In online appendix table E6, we show that the results in table 4 are broadly similar if we exclude the 1981–83 recession, suggesting that our results are not simply driven by differences between the 1981–83 recession and the Great Recession.

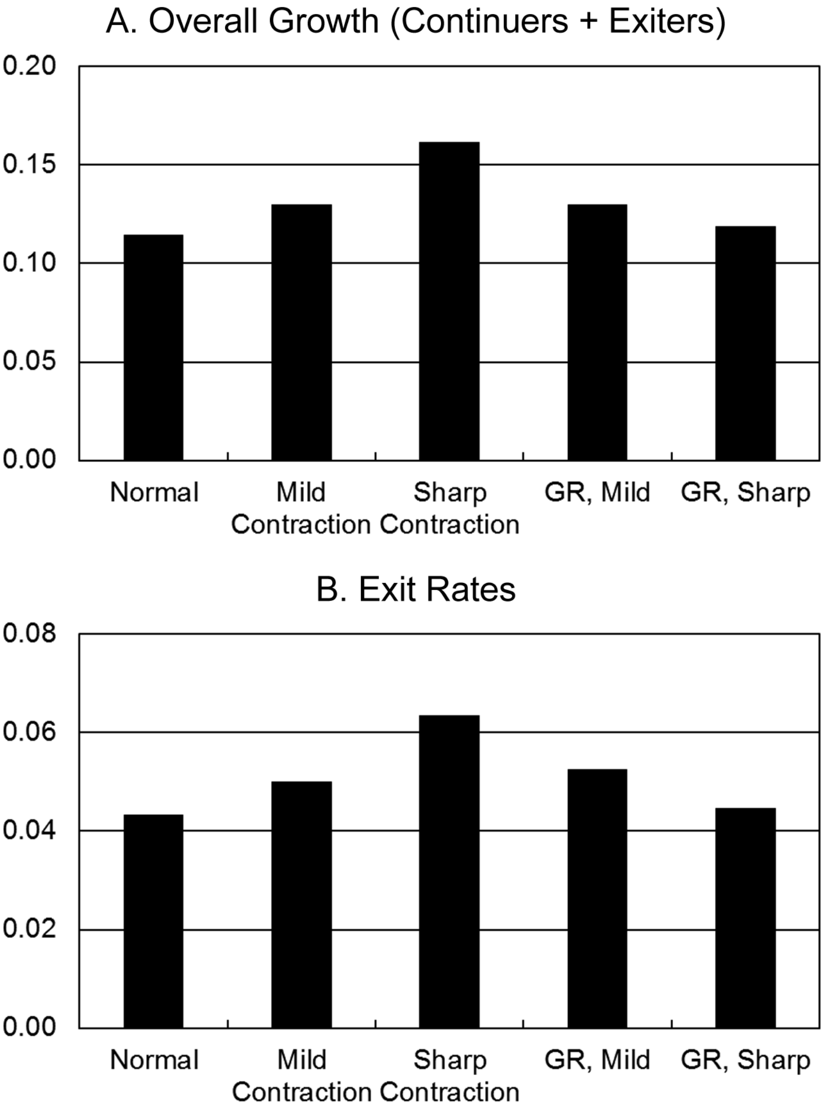

FIG. 4.—Differences in growth and exit rates between high-productivity and low-productivity establishments over the business cycle. Authors' calculations on Annual Survey of Manufactures, Census of Manufactures, and Longitudinal Business Database. Depicted is the predicted difference in growth rates (panel *A*, high minus low) and the predicted difference in probability of exit (panel *B*, low minus high) between an establishment one standard deviation above industry-by-year mean productivity and an establishment one standard deviation below industry-by-year mean productivity. Normal is zero change in state-level unemployment, mild contraction is 1 percentage point increase in state level unemployment, sharp contraction is 3 percentage point increase in state-level unemployment, and GR is for the period 2007–9.

We now turn to exploring whether these patterns vary by firm age. As before, we denote as "Young" establishments that are part of young firms and call the remaining establishments of mature firms "Mature."[43] The results of these regressions are shown in table 5. We find that the general patterns for the full sample hold for both "Young" and "Mature" (compare cols. 1, 3, and 5 in table 5 to those in table 4). However, the quantitative magnitudes are substantially larger for the establishments of young firms. To see this, we start by returning to figure 3. For young businesses, we find that the difference in growth rates for an establishment one standard deviation below and above mean productivity is about 17 percentage points. In contrast, the analogous difference for mature establishments is about 10 percentage points. The exit and growth rates of continuers for young establishments are also substantially more sensitive to productivity than those for mature establishments.

Table 5 shows that establishments of young firms are also more sensitive to the cycle and that the interaction effect of the cycle and productivity is larger in magnitude for establishments of young firms and is statistically significant for all three outcomes. The significance of the estimated two-way interaction between TFP and the cycle for the growth of continuing young establishments is especially notable since it contrasts with the results of table 4, where we could not detect a statistically significant relationship. Table 5 helps account for this, as we find that the two-way interaction effect between TFP and the cycle is actually negative (although not significant) for mature continuing establishments. Apparently cleansing effects on this margin (growth of continuing establishments) are present only for young businesses.

Did these patterns change in the Great Recession? For the three-way interaction effect of interest between TFP, the cycle, and the Great Recession, we find point estimates largely consistent with those for the full sample but with less systematic statistical significance. Part of the challenge here is that the number of establishments from young firms is only about 20% of the overall sample and the three-way interactions are focusing on a specific 3-year period (2007–9). Based on the point estimates for establishments of young firms, we find that the three-way interaction effect between TFP, the cycle, and the Great Recession tends to offset the two-way interaction effect between TFP and the cycle. However, while the patterns are systematic, they are not precisely estimated. For mature establishments, we find smaller three-way interaction estimates, but they still tend to systematically offset the two-way interaction of TFP and the cycle. We know from table 4 that, when pooled, we obtain large, statistically significant effects for the three-way interaction that offset the

[43] Results are similar when we use measures of "Young" that rely on establishment age.

Table 5
Reallocation and Productivity over the Business Cycle by Firm Age

	Overall Growth Rate (Continuers + Exiters)		Exit		Conditional Growth Rate (Continuers Only)	
	(1)	(2)	(3)	(4)	(5)	(6)
Young	−.059***	−.054***	.050***	.049***	.047***	.050***
	(.005)	(.005)	(.002)	(.002)	(.003)	(.003)
TFP × mature	.138***	.139***	−.054***	−.055***	.034***	.033***
	(.007)	(.007)	(.003)	(.003)	(.003)	(.003)
TFP × young	.237***	.236***	−.085***	−.083***	.075***	.077***
	(.013)	(.015)	(.006)	(.006)	(.006)	(.007)
Cycle × mature	−2.590***	−2.487***	.345**	.230	−2.047***	−2.183***
	(.401)	(.402)	(.143)	(.141)	(.232)	(.270)
Cycle × young	−6.626***	−5.274***	2.196***	1.775***	−2.578***	−1.878***
	(.988)	(1.152)	(.407)	(.463)	(.412)	(.455)
TFP × cycle × mature	.674	1.112	−.429*	−.720***	−.031	−.193
	(.620)	(.733)	(.234)	(.235)	(.354)	(.538)
TFP × cycle × young	3.886**	4.336**	−1.147*	−1.088	2.476**	2.811**
	(1.568)	(2.016)	(.649)	(.759)	(1.030)	(1.110)
GR × TFP × mature		.015		−.010		−.005
		(.029)		(.014)		(.007)

GR × TFP × young		.046		−.030		−.004
		(.076)		(.034)		(.042)
GR × cycle × mature		−1.685		1.234**		.769
		(1.268)		(.505)		(.758)
GR × cycle × young		−8.627***		2.940***		−4.152***
		(2.318)		(.934)		(.907)
GR × TFP × cycle × mature		−1.708		1.162		.686
		(1.691)		(.759)		(.657)
GR × TFP × cycle × young		−3.566		.965		−.999
		(4.585)		(1.934)		(2.115)
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm size class FE	Yes	Yes	Yes	Yes	Yes	Yes
N (millions)	2.2	2.2	2.2	2.2	2.1	2.1

SOURCE.—Authors' calculations on the Annual Survey of Manufactures, Census of Manufactures, and Longitudinal Business Database.
NOTE.—See note to table 4. Young (mature) is establishments that belong to firms less than (greater than or equal to) 5 years old.
* $p < .10$.
** $p < .05$.
*** $p < .01$.

two-way interaction. We are pushing the data hard in seeking to identify differential three-way interaction effects by firm age—especially given that the group that is more sensitive to the cycle (young) has relatively small samples for the 2007–9 period.

We illustrate the predictions from table 5 in figure 5 in the same manner as figure 4. We focus on overall growth for the sake of brevity.[44] Figure 5 shows that the differences in growth rates between high-productivity and low-productivity establishments are much larger for establishments of young as opposed to mature firms. For example, the difference in growth rates between high-productivity and low-productivity establishments of young firms is over 15 percentage points, while the difference for establishments of mature firms is generally around 10 percentage points. This differential grows for both young and mature, but especially for young during periods of rising unemployment prior to the Great Recession. For young, it grows to over 25% in a sharp contraction. During the Great Recession, this differential is only at 21% for a sharp contraction. While appropriate caution is needed for the latter interaction with the Great Recession given the lack of statistical precision, it suggests that the result that the Great Recession is less productivity enhancing is being driven disproportionately by young establishments.

A possible concern about our results is that we have made no adjustments for cyclical variation in capacity utilization in our measures of TFP. It is well known that capacity utilization is procyclical (see Basu and Fernald 2001), likely due to capacity utilization increasing in times of higher demand. Thus, standard aggregate measures of TFP that are not adjusted for time-varying capacity utilization are spuriously procyclical. This concern is substantially mitigated in our setting, because our measures of TFP are deviations from industry-by-year means. If specific years, or even specific industries within years, are hit especially hard in a recession by demand shocks, our measure of TFP abstracts from any common time variation in capacity utilization at the industry by year level. Still, it may be that when a specific industry is hit especially hard in a downturn, not all plants in the industry are equally affected, leading to idiosyncratic variation in capacity utilization over the cycle. We address this issue in a sensitivity analysis and find that these concerns are not driving our results. In this analysis, we include as extra controls the energy to capital ratio at the plant level both separately and interacted with all variables in the same way as TFP. Using the energy to capital ratio at the plant level is a common way to capture variation in capacity utilization

[44] The results for exit and growth of continuing establishments are shown in online appendix figs. E4 and E5, respectively. The much larger response of young continuing establishments to TFP and to the interaction of TFP and the cycle is evident in fig. E5.

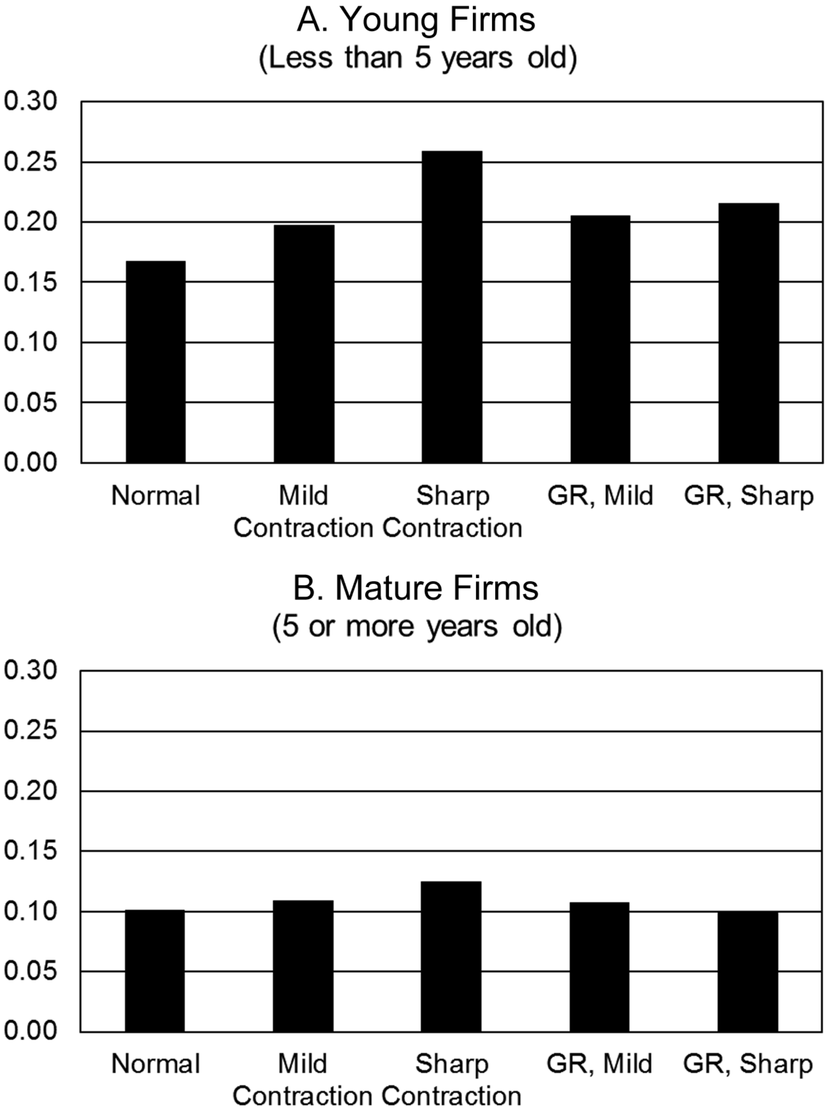

FIG. 5.—Differences in overall growth rates (continuing + exiting establishments) between high-productivity and low-productivity establishments over the business cycle by firm age. Authors' calculations on Annual Survey of Manufactures, Census of Manufactures, and Longitudinal Business Database. Depicted is the predicted difference in growth rates (high minus low) between an establishment one standard deviation above industry-by-year mean productivity and an establishment one standard deviation below industry-by-year mean productivity. Normal is zero change in state-level unemployment, mild contraction is 1 percentage point increase in state level unemployment, sharp contraction is 3 percentage point increase in state-level unemployment, and GR is for the period 2007–9.

(see Burnside, Eichenbaum, and Rebelo 1995). The results show that our findings on the marginal impact of productivity over the cycle on growth and survival are robust to including these additional controls (see online appendix table E7). We find that high energy-to-capital ratio plants are less likely to exit and more likely to grow, consistent with predictions, but this does not change our results on productivity.

B. Where Do Entrants Fit In?

The specifications of growth and survival we use in the prior section, while not derived explicitly from a structural model, are consistent with theoretical models of firm dynamics in the literature. An equivalent specification for entry would require capturing the decision rules of potential entrants, which is well beyond the scope of the current paper. Instead, we conduct a simple descriptive analysis of where entrants fit in the productivity distribution relative to incumbents and how this changes over the cycle. For this purpose, we estimate a simple descriptive linear probability specification based upon classifying establishments in any given year into two groups: new entrants (establishments in the first year of operation) and existing establishments (establishments with activity in prior years).

The specification has as the left-hand-side variable *entry*, equal to one if the establishment is a new entrant and equal to zero otherwise. On the right-hand side, we include TFP in the current year, a measure of the Cycle (in this case from $t-1$ to t since the designation of entry is for establishments that entered between $t-1$ and t), and the interaction. We also include a specification where we permit these relationships to differ in the Great Recession using a GR dummy (again being careful to treat the timing differently since this outcome is between $t-1$ and t).

We report results for this descriptive regression in table 6. We find that higher productivity establishments are slightly less likely to be entrants. The estimated effect is statistically significant given our sample size but is quantitatively small. Moving from one standard deviation below the (within-industry) mean to one standard deviation above the mean implies a difference in the likelihood of being an entrant of less than half a percent. Thus, entrants have slightly lower productivity than incumbents. This finding is consistent with findings in Foster, Haltiwanger, and Krizan (2001) and Foster et al. (2008). In terms of Foster et al. (2008), recall that this pattern may reflect lower prices for entrants compared to incumbents (given that our TFP measure is revenue based rather than a physical quantity measure of TFP).

Not surprisingly, the likelihood that an establishment is an entrant is lower in times of rising unemployment in the state. In terms of the interaction between TFP and the cycle, we find a positive and significant

Table 6
Entry and Productivity over the Business Cycle

	Establishment Entry	
	(1)	(2)
TFP	−.006***	−.006***
	(.002)	(.002)
Cycle	−.388***	−.376***
	(.136)	(.142)
TFP × cycle	.274***	.239**
	(.075)	(.103)
GR × TFP		.006*
		(.004)
GR × cycle		−.176
		(.504)
GR × TFP × cycle		−.088
		(.199)
Year FE	Yes	Yes
State FE	Yes	Yes
Firm size class FE	Yes	Yes
N (millions)	2.2	2.2

SOURCE.—Authors' calculations on the Annual Survey of Manufactures, the Census of Manufactures, and the Longitudinal Business Database.

NOTE.—Regressions are weighted by propensity score weights. Weight calculation is described in the online appendix. Standard errors (in parentheses) are clustered at the state level. Entry is measured from $t - 1$ to t. Regression is linear probability model with entry = 1 if this is first year of operation of establishment. TFP is the deviation of establishment-level log TFP from its industry-year mean in year t. GR is a dummy variable equal to one for years from 2008 to 2010 (given $t - 1$ to t). Cycle is the state-year change in the unemployment rate from $t - 1$ to t. Establishment size (log employment in t) is included as a control.

* $p < .10$.
** $p < .05$.
*** $p < .01$.

point estimate, suggesting that entrants in contractions are relatively more productive than in expansions. Again, however, this effect is relatively small. For an increase in unemployment of 3 percentage points, the probability that a high productivity establishment is an entrant is positive but very small. Moving from one standard deviation below mean productivity to one standard deviation above mean productivity yields a one-tenth of 1% higher probability that an establishment is an entrant during a period of a sharp contraction. We find little evidence that these patterns changed substantially in the Great Recession. We know from earlier work that job creation from entry fell substantially in the Great Recession (e.g., Fort et al. 2013). This is consistent with the patterns here given the large negative coefficient on the cyclical variable. It is a bit surprising that the interaction between GR and the cycle is not statistically significant (although it is negative, consistent with earlier work).

C. Aggregate Implications

The analysis of the relationship between productivity and reallocation above is based on the relationship between growth, survival, and productivity at the establishment level. A strength of this approach is the rich set of controls we are able to use while focusing on within-state variation in the cycle over time to identify the effects of interest. A limitation of the analysis is that it is difficult to draw inferences about aggregate consequences for productivity. A full analysis of the latter is beyond the scope of this paper, but in this section we conduct a counterfactual exercise to shed light on the aggregate consequences.

Much of the literature on the aggregate relationship between productivity and reallocation revolves around the extent to which resources are shifted away from less productive to more productive establishments (see Syverson [2011] for a recent survey). Our micro analysis is very much about such shifts, a fact that we now exploit in a simple counterfactual exercise to provide some perspective on aggregate implications. In each year we first compute the following base year index using the actual data:

$$P_t = \sum_i \theta_{it} P_{it},$$

where θ_{it} is the employment weight for plant i in period t and P_{it} is plant-level productivity (deviated from the industry-year mean). Then we use the model to generate a counterfactual index, given by

$$P^C_{t+1} = \sum_i \theta^C_{it+1} P_{it},$$

where θ^C_{it+1} is the predicted employment share for plant i in period t based upon the estimated model. We compute the predicted employment share using base year employment levels and the predicted growth rates in employment from the model.[45] We measure the gains from reallocation as $P^C_{t+1} - P_t$. We conduct this exercise in each year and then take time averages of these differences depending on different assumptions for the counterfactual (where the assumptions differ in terms of the assumed state of the cycle).[46]

[45] For this purpose, we use all of the terms in the model involving TFP, the cycle, and the GR dummies.

[46] The index of productivity is an employment-weighted average of establishment-level productivity. In this respect, it is related to the indices used in Baily et al. (1992), Griliches and Regev (1995), Olley and Pakes (1996), and Foster et al. (2001). Much of the work using activity-weighted averages of establishment-level TFP uses either output or composite input weights. We do not have that information for our counterfactual (we use the LBD to generate outcome measures for the counterfactual), so we are restricted to using the activity measures in our outcome measures, namely, employment. Foster et al. (2001) show that these activity-weighted indices are similar using output, input or employment weights.

Given that we use plant-level productivity measured as deviations from within-industry-by-year means, this calculation yields an estimate of the implied increase in within-industry productivity from reallocation effects alone. Moreover, since the plant-level distribution of productivity is held constant (in each year) for this exercise, the change only reflects the interaction of the predicted changes in the distribution of employment with where plants sit in the productivity distribution.

Figure 6 shows the results of this exercise. The bar labeled "Normal" implies that in a year with no change in the unemployment rate, the average increase in productivity from reallocation effects from one year to the next across incumbents is about 2.1 log points. During mild and sharp contractions prior to the Great Recession, this contribution increases to 2.4 and 2.9 log points, respectively. However, during mild and sharp contractions in the Great Recession (which can be thought of as the effect across different states), the reallocation contribution is 2.3 and 2.1 log points, respectively. Consistent with our micro evidence, the contribution of reallocation to this aggregate index of establishment-level productivity decreases in the Great Recession.

These estimates of the contribution of reallocation are large relative to those in the literature. In accounting decompositions, such as those in Foster et al. (2001, 2008), reallocation effects account for up to half of industry-level productivity growth using similar activity-weighted establishment-level productivity as indices of industry-level productivity.

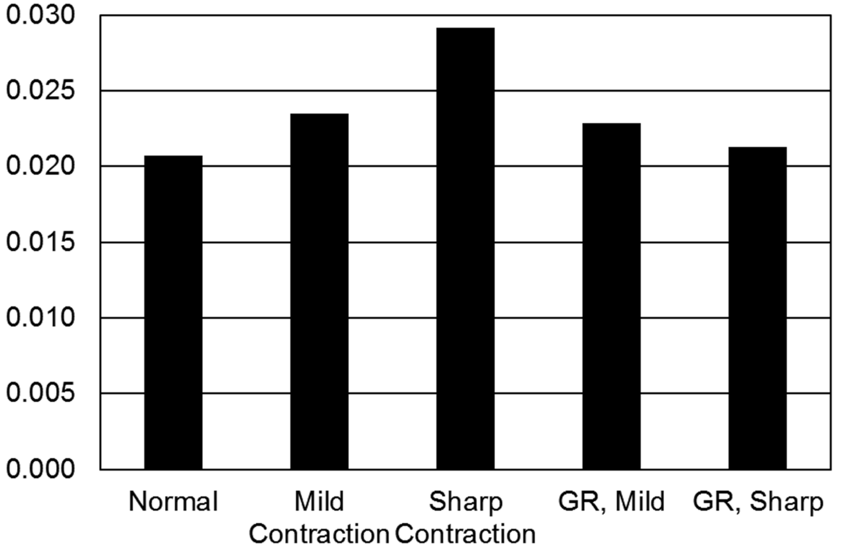

FIG. 6.—Predicted contribution of reallocation to aggregate (industry-level) productivity. Authors' calculations from estimated models.

In these papers, this type of average industry index grows by about 1 log point per year, so that half of this is substantially below the greater than 2 log point effects we are capturing. However, a strength of our current approach relative to this existing literature is that our counterfactual exercise focuses on the reallocation effects induced by productivity differences. That is, in this earlier literature, the accounting decompositions capture the contribution of reallocation of activity across establishments regardless of the source of that reallocation.[47] The work of Foster et al. (2008, 2013) emphasizes that much reallocation is induced by demand-side effects as opposed to productivity effects. Our revenue-based measure of TFP captures some but hardly all of the demand side effects identified in this recent work. Instead, our counterfactual exercise is based on the reallocation that is directly linked to productivity differences. Taking our results at face value yields a substantial contribution to productivity growth from productivity difference–induced reallocation.[48]

V. Conclusions and Future Work

We address the question "Was the Great Recession a cleansing recession?" by building up five related facts. First, we show that reallocation in the Great Recession differs markedly from that of earlier recessions. Job creation falls much more substantially than in prior recessions, and job destruction rises less than in prior recessions—taken together they yield less of an increase (or even a decline) in the intensity of reallocation. Second, we find that reallocation is productivity enhancing. Less productive establishments are more likely to exit, while more productive establishments are more likely to grow. Third, we show that these patterns are enhanced in recessions prior to the Great Recession. Fourth, we show that reallocation is less productivity enhancing in the Great Recession as contractions become more severe. The gap in growth rates and exit rates between high-productivity and low-productivity businesses decreases rather than increases with larger increases in unemployment in the Great Recession. Fifth, we find that the implied increases in aggregate (industry-level) productivity indices from productivity-induced reallocation are substantial,

[47] In terms of the above exercise, this is equivalent to using actual θ_{it+1} in calculating the gains from reallocation.

[48] Our counterfactual exercise cannot provide a full accounting of overall industry-level productivity growth. The ASM is not well suited for capturing the within-establishment productivity growth that is a critical part of the overall growth. High-frequency ASM data can measure the cross-sectional distribution of TFP within industries in a given year, but they do not provide a high-quality measure of within-establishment productivity growth given the ASM's sample limitations. The ASM is not well suited for longitudinal analysis of plants, and thus our longitudinal outcomes are derived from the LBD.

with even larger effects in sharp contractions prior to the Great Recession and smaller effects in sharp contractions in the Great Recession.

Our analysis is mostly descriptive—evaluating how the patterns and nature of reallocation change over the cycle and how they differ in the Great Recession. We do not directly address why the Great Recession is different. As such, our contribution is much more about what happened than why it happened. The obvious next step is to explore why the patterns are different. A clear candidate is the role of the financial collapse. Our finding that the patterns change more for young businesses is at least suggestive that the financial collapse (which arguably hit young firms much harder) is relevant. But to provide convincing evidence, we need to find ways to integrate direct measures of the financial collapse at the firm, or at least the regional level, into the type of analysis we have conducted here.[49]

This paper raises questions that bear looking into in future research. One interesting question concerns the heterogeneity of recessions in general. In comparing the Great Recession to earlier recessions in our productivity analysis, we group all of the earlier recessions for which we have data into one category in our main analysis. Much of the thinking about cleansing recessions was motivated by the patterns seen in the 1981–83 recession. The 1981–83 recession has a big surge in destruction and exits of low-productivity establishments, followed by a big surge in creation as early as 1984. That recession is very different from the relatively mild recessions of 1991 and 2001.[50] We do sensitivity analysis that suggests that our results are not driven by the differences between the 1981–83 recession and the Great Recession, but there is much room for further research in this area. In particular, investigating differences across recessions taking into account the different driving forces of recessions would be a promising area for future research. This would be one way to help understand why the Great Recession looks different in terms of its reallocation dynamics.

Another interesting area for future research is to explore the implications of the declining trend in job flows exhibited in the United States over the past few decades for productivity growth. Both the BDS and the BED show pronounced downward trends in job flows and thus the pace of

[49] Fort et al. (2013) present evidence that the fall in housing prices is important for understanding the especially large decline of young businesses in the Great Recession.

[50] Our descriptive analysis in Sec. III shows that these shallower recessions did not differ much from the early 1980s recession in terms of the covariance between job flows and the cycle. The 1991 and 2001 recessions differ in terms of the severity of the recessions, but the covariance between job flows and changes in unemployment are similar across the 1981–83, 1991, and 2001 recessions.

reallocation. Since we find that reallocation is productivity enhancing in general (ignoring the cycle), the obvious question is whether this has implications for long-run trend productivity growth in the United States.

Finally, we note that a core limitation of our current analysis is that we study the relationship between productivity and reallocation only for the manufacturing sector. While manufacturing is interesting and important, much of the changing patterns of job reallocation in terms of trends and the cycle are driven by other sectors. Our focus on manufacturing is driven by data limitations. There are sources that can be used for measuring productivity (even TFP) for establishments and firms in other sectors—but this will require addressing a variety of challenges in terms of measurement and methodology. The high pace of reallocation in nonmanufacturing sectors and the changing patterns of reallocation suggest that addressing such challenges would have substantial payoffs.

References

Angrist, Joshua, and Jorn-Steffen Pischke. 2009. *Mostly harmless econometrics*. Princeton, NJ: Princeton University Press.

Arellano, Manuel. 1987. Computing robust standard errors for within-group estimators. *Oxford Bulletin of Economics and Statistics* 49, no. 4: 431–34.

Baily, Martin Neil, Eric J. Bartelsman, and John Haltiwanger. 2001. Labor productivity: Structural change and cyclical dynamics. *Review of Economics and Statistics* 83, no. 3:420–33.

Baily, Martin Neil, Charles Hulten, and David Campbell. 1992. Productivity dynamics in manufacturing plants. In *Brookings papers on economic activity: Microeconomics*, vol. 4, ed. Clifford Winston and Martin Neil Baily, 187–267. Washington, DC: Brookings Institution Press.

Barlevy, Gadi. 2002. The sullying effect of recessions. *Review of Economic Studies* 69, no. 1:65–96.

———. 2003. Credit market frictions and the allocation of resources over the business cycle. *Journal of Monetary Economics* 50, no. 8:1795–1818.

Bartelsman, Eric J., and Mark Doms. 2000. Understanding productivity: Lessons from longitudinal microdata. *Journal of Economic Literature* 38, no. 3:569–95.

Basu, Susanto, and John Fernald. 2001. Why is productivity pro-cyclical? Why do we care? In *New developments in productivity analysis*, ed. Charles R. Hulten, Edward R. Dean, and Michael J. Harper, 225–302. Chicago: University of Chicago Press for National Bureau of Economic Research.

Bernanke, Ben, and Mark Gertler. 1989. Agency costs, net worth, and business fluctuations. *American Economic Review* 79, no. 1:14–31.

Blanchard, Oliver Jean, and Peter Diamond. 1990. The cyclical behavior of the gross flows of U.S. workers. In *Brookings papers on economic activity*, no. 2, ed. William C. Brainard and George L. Perry, 85–155. Washington, DC: Brookings Institution Press.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 1995. Capital utilization and returns to scale. In *NBER macroeconomics annual 1995*, vol. 10, ed. Ben S. Bernanke and Julio J. Rotemberg, 67–124. Cambridge, MA: National Bureau of Economic Research.

Caballero, Ricardo J., and Mohamad L. Hammour. 1994. The cleansing effect of recessions. *American Economic Review* 84, no. 5:1350–68.

———. 1996. On the timing and efficiency of creative destruction. *Quarterly Journal of Economics* 111, no. 3:805–32.

———. 2005. The cost of recessions revisited: A reverse-liquidationist view. *Review of Economic Studies* 72, no. 2:313–41.

Campbell, Jeffrey R. 1998. Entry, exit, embodied technology, and business cycles. *Review of Economic Dynamics* 1, no. 2:371–408.

Chodorow-Reich, Gabriel. 2014. The employment effects of credit market disruptions: Firm-level evidence from the 2008–09 financial crisis. *Quarterly Journal of Economics* 129, no. 1:1–59.

Collard-Wexler, Alan, and Jan De Loecker. 2013. Reallocation and technology: Evidence from the U.S. steel industry. NBER Working Paper no. 18739, National Bureau of Economic Research, Cambridge, MA.

Cooper, Russell, John Haltiwanger, and Jonathan Willis. 2007. Search frictions: Matching aggregate and establishment observations. *Journal of Monetary Economics* 54, suppl.: 56–78.

Davis, Steven J., R. Jason Faberman, and John Haltiwanger. 2006. The flow approach to labor markets: New data sources and micro-macro links. *Journal of Economic Perspectives* 20, no. 3:3–26.

———. 2012. Labor market flows in the cross section and over time. *Journal of Monetary Economics* 59, no. 1:1–18.

Davis, Steven J., and John Haltiwanger. 1990. Gross job creation and destruction: Microeconomic evidence and macroeconomic implications. In *NBER macroeconomics annual 1990*, vol. 5, ed. Olivier Jean Blanchard and Stanley Fischer, 123–86. Cambridge, MA: MIT Press.

———. 1992. Gross job creation, gross job destruction, and employment reallocation. *Quarterly Journal of Economics* 107, no. 3:819–63.

———. 1999. On the driving forces behind cyclical movements in employment and job reallocation. *American Economic Review* 89, no. 5:1234–58.

———. 2014. Labor market fluidity and economic performance. Paper presented at Federal Reserve Bank of Kansas City Conference, Jackson Hole, WY, August.

Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda. 2007. Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. In *NBER macroeconomics annual 2006*, vol. 21, ed. Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 107–56. Cambridge, MA: MIT Press.

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

Decker, Ryan, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2014a. The role of entrepreneurship in US job creation and economic dynamism. *Journal of Economic Perspectives* 28, no. 3:3–24.

———. 2014b. The secular decline of business dynamism in the United States. Working paper, University of Maryland (June).

Ericson, Richard, and Ariel Pakes. 1995. Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies* 62, no. 1:53–82.

Eslava, Marcela, Arturo Galindo, Marc Hofstetter, and Alejandro Izquierdo. 2010. Scarring recessions and credit constraints: Evidence from Colombian firm dynamics. CEDE Working Paper no. 2010-07, Universidad de los Andes, Bogota.

Fort, Teresa, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2013. How firms respond to business cycles: The role of firm age and firm size. *IMF Economic Review* 61, no. 3:1–40.

Foster, Lucia, John Haltiwanger, and C. J. Krizan. 2001. Aggregate productivity growth: Lessons from microeconomic evidence. In *New developments in productivity analysis*, ed. Charles R. Hulten, Edward R. Dean, and Michael J. Harper, 303–72. Chicago: University of Chicago Press for National Bureau of Economic Research.

———. 2006. Market selection, reallocation and restructuring in the U.S. retail trade sector in the 1990s. *Review of Economics and Statistics* 88, no. 4:748–58.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review* 98, no. 1:394–425.

———. 2013. The slow growth of new plants: Learning about demand. NBER Working Paper no. 17853, National Bureau of Economic Research, Cambridge, MA.

Greenstone, Michael, and Alexandre Mas. 2012. Do credit market shocks affect the real economy? Quasi-experimental evidence from the Great Recession and "normal" economic times. Working Paper no. 12–27, Department of Economics, Massachusetts Institute of Technology, Cambridge, MA.

Griliches, Zvi, and Haim Regev. 1995. Productivity and firm turnover in Israeli industry, 1979–1988. *Journal of Econometrics* 65, no. 1:175–203.

Haltiwanger, John, Ron Jarmin, and Javier Miranda. 2011. Historically large decline in job creation from startups and existing firms in the

2008–09 recession. Brief no. 5, Kauffman Foundation Business Dynamics Statistics Briefing, Kansas City.

———. 2013. Who creates jobs? Small versus large versus young. *Review of Economic and Statistics* 95, no. 2:347–61.

Hopenhayn, Hugo. 1992. Entry, exit, and firm dynamics in long run equilibrium. *Econometrica* 60, no. 5:1127–50.

Hopenhayn, Hugo, and Richard Rogerson. 1993. Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy* 101, no. 5:915–38.

Hyatt, Henry, and James Spletzer. 2013. The recent decline in employment dynamics. Working Paper no. 13–03, Center for Economic Studies, Washington, DC.

Jarmin, Ron S., and Javier Miranda. 2002. The longitudinal business database. Working Paper no. 02–17, Center for Economic Studies, Washington, DC.

Jovanovic, Boyan. 1982. Selection and the evolution of industry. *Econometrica* 50, no. 3:649–70.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. Job creation and job destruction and the theory of unemployment. *Review of Economic Studies* 61, no. 3:397–415.

Olley, Steven, and Ariel Pakes. 1996. The dynamics of productivity in the telecommunications equipment industry. *Econometrica* 64, no. 6:1263–1310.

Osotimehin, Sophie, and Francesco Pappadà. 2013. Credit frictions and the cleansing effect of recessions. Unpublished manuscript, University of Virginia, Charlottesville.

Ouyang, Min. 2009. The scarring effect of recessions. *Journal of Monetary Economics* 56, no. 2:184–99.

Schumpeter, Joseph A. 1939. *Business cycles: A theoretical, historical and statistical analysis of the capitalist process*. 2 vols. New York: McGraw Hill.

———. 1942. *Capitalism, socialism and democracy*. New York: Harper.

Syverson, Chad. 2004. Product substitutability and productivity dispersion. *Review of Economics and Statistics* 86, no. 2:534–50.

———. 2011. What determines productivity? *Journal of Economic Literature* 49, no. 2:326–65.

Tornqvist, Leo, Pentti Vartia, and Yrjo Vartia. 1985. How should relative changes be measured? *American Statistician* 39, no. 1:43–46.