

NBER WORKING PAPER SERIES

IDIOSYNCRATIC RISK AND
AGGREGATE EMPLOYMENT DYNAMICS

Jeffrey R. Campbell
Jonas D.M. Fisher

Working Paper 7936
<http://www.nber.org/papers/w7936>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2000

We thank John Leahy for valuable input. Campbell gratefully acknowledges research support from the National Science Foundation through grant SBR-9730442. The views expressed are those of the authors and not necessarily those of the National Bureau of Economic Research, the Federal Reserve Bank of Chicago or of the Federal Reserve System.

© 2000 by Jeffrey R. Campbell and Jonas D.M. Fisher. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Idiosyncratic Risk and Aggregate Employment Dynamics
Jeffrey R. Campbell and Jonas D.M. Fisher
NBER Working Paper No. 7936
October 2000
JEL No. E00, L00, J00

ABSTRACT

This paper studies how producers' idiosyncratic risks affect an industry's aggregate dynamics in an environment where certainty equivalence fails. In the model, producers can place workers in two types of jobs, organized and temporary. Workers are less productive in temporary jobs, but creating an organized job requires an irreversible investment of managerial resources. Increasing productivity risk raises the value of an unexercised option to create an organized job. Losing this option is one cost of immediate organized job creation, so an increase in its value induces substitution towards cheaper temporary jobs. Because they are costless to create and destroy, a producer using temporary jobs can be more flexible, responding more to both idiosyncratic and aggregate shocks. If all of an industry's producers adapt to heightened idiosyncratic risk in this way, the industry as a whole can respond more to a given aggregate shock. This insight is used to better understand the observation from the U.S. manufacturing sector that groups of plants displaying high idiosyncratic variability also have large aggregate fluctuations.

Jeffrey R. Campbell
Department of Economics
University of Chicago
1126 East 59th Street
Chicago, IL 60637
and NBER
jcampbe@midway.uchicago.org

Jonas D.M. Fisher
Research Department
Federal Reserve Bank of Chicago
230 South LaSalle Street
Chicago, IL 60604
jfisher@frbchi.org

1. Introduction

This paper analyzes how idiosyncratic risk affects individual and aggregate employment dynamics when producers can choose the flexibility of their input use. Consider the problem of a producer that can place workers in two types of jobs, organized and temporary. Workers are less productive in temporary jobs, but creating an organized job requires an irreversible investment of managerial resources. Increasing productivity risk raises the value of an unexercised option to create an organized job. Losing this option is one cost of immediate organized job creation. An increase in option value induces substitution towards cheaper temporary jobs. Because they are costless to create and destroy, a producer using temporary jobs can be more responsive to both idiosyncratic and aggregate shocks. Therefore, increased idiosyncratic risk can raise a producer's observed flexibility. If all of an industry's producers respond to heightened idiosyncratic risk in this way, the industry as a whole can display greater flexibility, responding more to a given *aggregate* shock.

We develop this intuition analytically in a simplified setting with risk lasting one period. To assess its robustness and quantitative importance for aggregate dynamics, we calibrate, solve, and simulate an industry equilibrium model. We use Davis, Haltiwanger, and Schuh's (1996) job creation and destruction data for U.S. manufacturing plants of different ages and in different four-digit industries to calibrate our model and evaluate its behavior. These observations are useful for our purposes because both idiosyncratic risk and aggregate dynamics vary significantly across plants in different industries and at different stages of the life cycle.

Our analysis builds on Campbell and Fisher's (1999) model of a competitive industry of plants subject to proportional costs of job creation and destruction and idiosyncratic productivity risk. Transitory movements in the real product wage drive aggregate fluctuations. A distinguishing feature of that model is its ability to reproduce the U.S. manufacturing sector's greater use of job destruction relative to job creation to accommodate transitory employment changes. To that framework, we incorporate temporary jobs, as in Bentolila and Saint-Paul (1992) and Cabrales and Hopenhayn (1997), and a simple life-cycle in which plants exogenously exit and are replaced at a constant rate. New plants begin their lives with high productivity risk, and as a plant ages

its productivity risk declines. These assumptions are consistent with the higher observed average rates of job creation and destruction among young plants. As detailed in Campbell and Fisher (2000), our industry equilibrium framework can be interpreted as a model of general equilibrium.

A consistent pattern in the U.S. manufacturing sector's gross job flow data is that employment aggregates become more variable when producers face greater idiosyncratic risk. This is true whether we compare plants of different ages or look across four-digit industries. The time-series variances of young plants' aggregate job creation, destruction, and growth rates are greater than those of their older counterparts. Job creation, destruction, and growth all vary more in four-digit industries with high idiosyncratic risk, even after controlling for the industry's parent two-digit industry. Our model economy replicates these findings, because greater idiosyncratic risk induces plants to use more temporary jobs, which increases both individual and aggregate flexibility.

For the case of a single investment project which is costly to start and shut down, Dixit (1989) demonstrated that increasing risk induces delay in both decisions by increasing the value of the option to adjust investment in the future. This paper shows how these option value considerations induce changes in producers' flexibility and impact aggregate dynamics. Increasing the incentive for an individual producer to delay costly job creation and destruction need not dampen aggregate fluctuations. In fact, our results point to the conclusion that increasing idiosyncratic risk increases producers' flexibility and *amplifies* their responses to aggregate shocks. Our application of these results focuses on aggregate employment dynamics, but the same approach can be used to study producers' other input choices, including capital and inventories. For example, if producers can substitute between general purpose machinery and illiquid special purpose machinery, increasing idiosyncratic risk can increase the variance of aggregate investment by changing its composition. In general, our results suggest that the increase in microeconomic turbulence described by Sargent and Ljungqvist (1998), Comín (2000) and Campbell, Lettau, Malkiel and Xu (2000) will amplify macroeconomic responses to a given aggregate shock.

Our characterization of the differences between young and old plants' employment dynamics is related to earlier research by Gertler and Gilchrist (1994) and Kashyap, Lamont, and Stein

(1994) which documents small and large firms' distinct inventory adjustment patterns following contractionary monetary policy. Gertler and Gilchrist find that small firms' inventories contract substantially following a monetary tightening, while large firms' inventories expand mildly. They interpret this as reflecting the relative inability of small firms to take on debt as their balance sheets deteriorate in a recession. For large firms with virtually unimpeded access to credit markets, these balance sheet effects are assumed to be negligible. Kashyap, *et al* find similar results for firms distinguished by bond-market rating status. Firms with no bond rating tend to be small, so these results reinforce those of Gertler and Gilchrist.

To the extent that small firms tend to operate young plants, this financial theory offers another potential explanation for young and old plants' different aggregate employment dynamics. However, this theory does not explain the complete dynamics of young and old plants' employment data. Theoretically, the asymmetry between small and large firms should only be important during aggregate downturns, when balance sheets of small firms are already weak. Gertler and Gilchrist and Kashyap *et al* verify this prediction in their data. In particular, small and large firms' inventories and short-term borrowing do not respond differently to monetary disturbances if the economy is already expanding. However, in the employment data, young plants' greater volatility is most evident during aggregate *expansions*. This highlights the need for a more complete explanation for young and old producers' different dynamic behavior.

The remainder of this paper is organized as follows. In the next section, we empirically characterize the relationship between plants' idiosyncratic risks and aggregate employment dynamics. In Section 3, we develop analytic solutions to a producer's job creation problem in a simplified setting. We use this to develop economic intuition regarding the impact of idiosyncratic risk on producers' responses to aggregate shocks. In Section 4 we present a more general model useful for quantitative analysis. In Section 5 we study a calibrated version of this model and compare its observable behavior to the data. Section 6 concludes.

2. Evidence on Idiosyncratic Risk and Aggregate Dynamics

This section provides an empirical basis for evaluating our theory of idiosyncratic risk, producer flexibility, and aggregate employment dynamics. Our data source is Davis, Haltiwanger, and Schuh's (1996) (hereafter DHS) observations of job creation and destruction rates for the U.S. manufacturing sector. For a given population of plants, the rate of job creation in a period is defined as the total number of jobs added since the previous period at plants that increased employment, divided by the average of total employment in the current and previous periods. The rate of job destruction is similarly defined in terms of employment losses at shrinking plants. The difference between job creation and destruction is the rate of job growth. DHS propose using the sum of job creation and destruction rates as a measure of job reallocation across plants. We follow their terminology and refer to this sum as the job reallocation rate.

A well known characteristic of these data is that industries' rates of job creation and destruction exceed those necessary to implement observed net job growth. This implies that jobs are continually being reallocated across plants within the same industry. DHS document that this is the case even when looking at very narrowly defined (four-digit) industries within specific geographic regions, so this reallocation does not reflect movements between industries within a broadly defined sector. Rather, the high rates of job creation and destruction seem to reflect plants' idiosyncratic risks and the resulting heterogeneity in their labor demand within narrowly defined industries. For this reason, these data are particularly useful for examining how idiosyncratic risk is related to aggregate dynamics.

We begin by comparing the aggregate employment dynamics of young and old manufacturing plants, because previous empirical work indicates that young plants face more risk than their older counterparts. We then characterize the relationship between idiosyncratic risk and aggregate employment dynamics across four-digit manufacturing industries. For this analysis, we measure idiosyncratic risk faced by an industry's plants using its job creation and destruction data.

2.1. Employment Dynamics over the Plant Life–Cycle

Dunne, Roberts, and Samuelson (1987) used plant-level employment observations to examine how the employment growth process changes as a plant ages. A robust finding of this study is that the cross-sectional variance of plant employment growth rates declines significantly with age. Jovanovic’s (1982) industry dynamics model replicates this finding because plants face less productivity risk as they age and learn about their production technology. This model suggests interpreting the decline in the variance of employment growth with plant age as a consequence of falling idiosyncratic risk. This risk need not be directly connected to the production process. A decline in risk regarding demand for a plant’s specific product could also generate the observed variance decline. Nevertheless, not all models generate a direct connection between idiosyncratic risk and the variance of employment growth. In the presence of labor adjustment costs, both the risks the plant faces and its adjustment technology influence realized employment variability. If a plants’ adjustment costs exogenously increase as it ages, then the variance of its employment growth will fall even if the risks it faces do not change. Subject to this caveat, we identify young plants as “high-risk” in comparison to their older counterparts for our empirical analysis.

For our comparison of young and old plants’ employment dynamics, we use DHS’s quarterly job creation and destruction series for plants in three different age categories. DHS (p. 225) recommend aggregating the two categories that include the youngest establishments, and we do this here. We refer to this combination as ‘young’ plants. These plants are typically less than 10 years old and account for 22.5% of total employment on average over the sample period.¹ We refer to the remaining plants as ‘old’ plants.² To seasonally adjust the time series, we first regressed each series against a series of quarterly dummies. Our seasonally adjusted series is then the sum of the raw series’ average value and the residuals from this regression.

Figure 1 plots the rates of job growth for young and old plants separately over our sample period, 1972:II to 1988:IV. The vertical lines are the business cycle peaks and troughs, as defined

¹We aggregate the two youngest categories by adding their job creation and destruction rates after weighting them by their total employment shares. We then divide the results by the sum of the two groups’ employment shares.

²Because of the sample design of the underlying data source, the threshold between young and old changes over time. The minimum age of an old establishment is at least 9 years and at most 13 years. These changes only occur at the end of a year.

by the NBER.³ This figure illustrates that employment growth at young plants is typically higher than at old plants and visibly more variable, especially in the periods between recessions. This excess variability of young plants between recessions was highlighted in the introduction as suggesting that factors other than the balance sheet effects emphasized by Gertler and Gilchrist (1994) and Kashyap, Lamont and Stein (1994) may be important for understanding the dynamics of young and old plants.

The top panel of Table 1 reports the sample means of the job growth, creation, destruction, and reallocation rates for young and old plants separately, as well as for the U.S. manufacturing sector as a whole. The bottom panel reports sample standard deviations and some selected sample correlations. As Figure 1 suggests, young plants' average job growth rate is substantially higher than that for old plants, 0.90% versus -0.70%. More importantly the young plants' job growth rate has a higher sample standard deviation than that of old plants, 2.32% versus 1.81%. A Wald test (not reported) rejects the hypothesis that the underlying variances are the same for young and old plants at a very high level of confidence. Thus, we observe that job growth at young plants as a group displays significantly more time series variability. Consideration of the other statistics in Table 1 reinforces this impression of greater variability at young plants. Young plants' average reallocation rate is much higher than that of old plants, 14.1% versus 9.8% (also statistically significant), as we would expect if young plants faced more idiosyncratic risk. Job creation, destruction, and reallocation all vary more over time for young plants. This is particularly the case for job creation. Its standard deviation is 1.66% for young plants and only 0.71% for old plants.

Another feature of this data, focused on by DHS, is the relative variability of job creation and destruction. For the U.S. manufacturing sector as a whole, the variance ratio of job destruction and job creation equals 2.89, so job destruction fluctuates much more than job creation. This statistic has been the subject of considerable theoretical attention, because models in which heterogeneous plants can be aggregated into a well-defined representative producer typically

³The NBER's dates of business cycle peaks and troughs are expressed as monthly dates. The usual transformation to quarterly dates is somewhat problematic given the non-standard quarterly timing of the job creation and destruction data as described in the technical appendix of DHS. The only change to the standard dating convention is that the March 1975 trough is assigned to 1975, second quarter.

imply that job creation and destruction vary equally. The variance ratio is also considerably greater than one for old plants (see Table 1), which is as expected because old plants account for the majority of manufacturing employment. A more interesting finding, first noted by DHS, is that young plants' job creation and destruction rates have approximately equal variances. Theoretical approaches to the empirical failure of the representative producer model have focused on job search and matching frictions (Mortensen and Pissarides (1994)), asymmetries in aggregate driving processes (Caballero (1992), Foote (1998)), and employment adjustment costs (Campbell and Fisher (2000)). The failure of young plants' dynamics to reproduce a variance ratio significantly above one challenges these theories much as the high variance ratio for manufacturing as a whole originally cast doubt on the representative producer paradigm. Examining this aspect of young and old plants' different aggregate dynamics is therefore a useful way to evaluate our theory.

The job creation and destruction series for young and old plants that we analyze here do not distinguish between employment changes due to plant births and deaths and the expansion and contraction of plants which simply continue operation, but DHS report job creation and destruction statistics separately for entering and exiting plants. The theoretical analysis we outlined in the introduction does not necessarily have strong implications for producers' use of entry and exit to adjust employment, but it may be of interest to determine whether or not young and old plants systematically use these margins differently. The data seem ambiguous about whether or not the asymmetry between young and old plants primarily represents the asymmetric behavior of plant births (which by construction are young) and plant deaths, depending on the particular variable in question and the data's periodicity. Using quarterly data, the job growth rates for young and old plants look much more symmetric if births and deaths are excluded, but the analogous annual job growth data (not presented here) displays strong asymmetries between young and old plants. Regardless of the data's periodicity, the job creation rate of continuing young plants has a larger variance than that of continuing old plants. However, the job destruction rates at young and old continuing plants have approximately equal variances in both quarterly and annual data. We conclude that the greater variance of young plants' job flows does not exclusively represent their different use of entry and exit.

2.2. Cross-Industry Evidence on Idiosyncratic Risk and Aggregate Dynamics

The pace of job reallocation varies greatly across four-digit manufacturing industries. To the extent that these differences reflect the varying degree of idiosyncratic risk these industries' producers face, a comparison across industries can provide more evidence regarding the impact of idiosyncratic risk on aggregate employment dynamics. Of course, industries differ along more than just one dimension. Their plants also operate different production and adjustment technologies. Most importantly, industries clearly vary in their exposure to aggregate disturbances such as oil-price shocks (for energy intensive industries) and monetary policy shocks (for durable goods industries). It is possible to control for exposure to macroeconomic disturbances, but these other industry characteristics are difficult to measure. Nevertheless, if the impact of idiosyncratic risk on producer flexibility and aggregate dynamics is quantitatively important, it should manifest itself in a cross-industry comparison, in spite of these additional complicating factors.

To carry out this comparison, we use DHS's job creation and destruction series for 446 four-digit industries. With these data, we estimate linear regressions in which the sample standard deviations of job growth, creation, destruction, and reallocation are explained using a measure of the industry's idiosyncratic risk and dummy variables for all two-digit industries. We intend these dummy variables to control for each industry's exposure to macroeconomic disturbances.

Unlike in the comparison of young and old plants, we have little guidance from the literature regarding which industries' plants are exposed to more idiosyncratic risk. We base our measure of risk on the available job creation and destruction data. This poses difficulties since it is possible to generate the result we expect (*i.e.*, time-series variability is positively related to idiosyncratic risk) spuriously using job reallocation to measure idiosyncratic risk. To see why, suppose counterfactually that the industries in our data were each populated by a large number of identical representative producers. These plants would face no idiosyncratic risk, so the industries' job reallocation rates would all equal the absolute values of the representative producers' employment growth rates. Therefore, an industry with high employment growth variance must also have a high average value of job reallocation. Any regression across indus-

tries of employment growth’s variance against the average job reallocation rate would find a positive association between the two variables, although the interpretation that idiosyncratic risk increases aggregate variance would not be valid. When producers are heterogeneous within industries, such spurious regression results need not arise, but the possibility certainly exists.

To mitigate this problem, we use the sum of an industry’s job creation and destruction rates’ *median* values in our sample to measure the idiosyncratic risk faced by that industry’s plants. To further minimize the impact of aggregate volatility on our measure of idiosyncratic risk, we use DHS’s annual job creation and destruction data instead of their quarterly series. This eliminates seasonal fluctuations from our analysis. The data used to construct each industry’s sample statistics cover the years 1973 through 1986, the year before an extensive redefinition of the four-digit industry codes.

Figure 2 illustrates this cross-industry comparison by plotting the job growth rate’s sample standard deviation versus our measure of idiosyncratic risk for all 446 industries in our sample. The data were first regressed against a full set of two-digit industry dummies, and the residuals from this regression were used to construct Figure 2. The figure’s solid line is the estimated linear regression curve for this data. Its slope equals the estimated coefficient on our idiosyncratic risk measure from the multivariate regression. Although there is considerable variation of the data around the estimated regression line, the plot clearly indicates a positive relationship between our measure of idiosyncratic risk and the job growth rate’s standard deviation. The estimated regression coefficient equals 0.17 (with heteroskedasticity-consistent standard error 0.012.) To help gauge this estimate’s economic significance, Table 2 reports summary statistics for our regressions’ independent and dependent variables. The cross-industry standard deviation for our measure of idiosyncratic risk is 5.90%. Our estimated coefficient then implies that a one standard deviation increase in our measure of idiosyncratic risk raises the standard deviation of an industry’s job growth rate by approximately one percentage point.

Table 3 reports the coefficients on our measure of idiosyncratic risk from all four regressions. The first column reports ordinary least squares estimates. Figure 1 clearly indicates the presence of outliers in our data. The second column reports least absolute deviations estimates to verify that our results do not depend on a few outlying observations. The typical four-digit indus-

try accounts for only a small fraction of total manufacturing employment, and our estimation procedure weights each industry equally, regardless of its size. The final column reports OLS estimates after deleting the smallest one quarter of our sample’s industries, so we can verify that small and unimportant industries are not essential for our results. Below each estimate is its heteroskedasticity-consistent standard error.

The estimated coefficients for all of the OLS regressions are positive and statistically significant. The least absolute deviations estimates and the OLS estimates with the restricted sample are also positive and statistically significant. Furthermore, they differ little from the baseline OLS estimates. Apparently, our measure of idiosyncratic risk is associated with higher aggregate employment variance, regardless of the estimation method, sample, or measure of aggregate volatility. The baseline estimates imply that raising our measure of idiosyncratic risk by one standard deviation increases the standard deviations of job creation, destruction, and reallocation by 0.77, 0.55, and 0.73 percentage points. One interesting feature of the OLS estimates is that job creation’s standard deviation increases more with idiosyncratic risk than does job destruction’s, suggesting that idiosyncratic risk reduces an industry’s variance ratio. Indeed, when we repeat our analysis using the variance ratio’s square root as the dependent variable, we find that idiosyncratic risk reduces the variance ratio. Recall that the variance ratio for young plants is much smaller than that for old plants, so this effect of idiosyncratic risk is consistent across our two ways of grouping plants.

Whether we look across industries or across plants of different ages, we find that groups of plants with high idiosyncratic risk also display high *aggregate* time-series variances. Before considering our model of this relationship, it is worth noting one important caveat to our results. Our measure of a four-digit industry’s risk is based on idiosyncratic employment volatility, and we argued that young plants face more risk than their older counterparts because they also display high idiosyncratic variances of employment growth. Because our risk measures are indirect, they could reflect differences across plants that have nothing to do with productivity or demand risk, most notably differences in plants’ adjustment costs. In the model of Campbell and Fisher (2000), an industry with high adjustment costs will display less job reallocation and respond less to aggregate shocks simply because doing so is expensive. Such an industry will

also have a relatively high variance ratio, because adjustment costs create asymmetries in plants' responses to transitory real wage shocks that generate a variance ratio above one. Therefore, it is possible to explain our results as only reflecting plants' different adjustment costs. This explanation may have some appeal when considering our cross-industry evidence, because it is natural to suppose that adjustment technologies vary across industries. However, it can explain our results over the plant life-cycle only if plants' adjustment costs exogenously increase as they age. In the remainder of this paper, we provide a model in which plants' adjustment costs endogenously rise with age in response to exogenously decreasing idiosyncratic risk.

3. Microeconomic Effects of Idiosyncratic Risk

This section develops intuition about the impact of idiosyncratic risk on an employer's responses to aggregate shocks, by recasting the two-period investment and option pricing model of Dixit and Pindyck (1994) as a job creation decision. We extend that framework by allowing the holder of a job creation opportunity to fill it temporarily, which requires no sunk cost.

Consider the problem of a producer with an opportunity to create a single job. At discrete time intervals, the producer can choose between three options. It can irreversibly expend managerial resources to create an organized job, it can create a temporary job, or it can leave the opportunity unexploited. The cost of the managerial resources used to create an organized job is τ_c , which we refer to as the job creation cost. After creating an organized job, the producer incurs the wage cost W per period to fill it with one worker. If it is created, the organized job's output of the numeraire good in the initial period is Z_0 . For simplicity, we assume that the decision to create an organized job is irreversible and that the job cannot be abandoned to avoid the wage cost. Creating a temporary job does not require paying a job creation cost, it is perfectly reversible in future periods, and it does not foreclose the option to create an organized job later. The cost of filling a temporary job with one worker is also W per period, and its output in the initial period is θZ_0 , where θ is strictly between zero and one.

After the initial period, the job's productivity is subject to a shock. With probability $1/2$, it grows to $(1 + \sigma) Z_0$, and with equal probability it shrinks to $(1 - \sigma) Z_0$. Thereafter, the job's

productivity remains unchanged. So that productivity is always positive, we proceed under the assumption that σ is less than one. Denote the organized job's productivity in the second and all subsequent periods with Z_1 . The output of a temporary job is θZ_1 following the productivity shock.

Suppose that the producer is risk-neutral and discounts future profits at the constant net interest rate r . Furthermore, assume that

$$\tau_c < (1 - \theta) \left(\frac{1 + r}{r} \right) \frac{W}{\theta}. \quad (1)$$

This upper bound on the job creation cost is sufficient to guarantee that temporary jobs are never used after the initial period.⁴ Given this restriction, the producer's optimal employment policy takes a simple form. If Z_0 falls below $\left(\frac{r}{1+r} \times \tau_c + W \right) / (1 + \sigma)$, then the producer will never create an organized job, even following a favorable productivity shock. In this case, creating a temporary job is also not profitable under any circumstances, so the job creation opportunity will go unexploited.⁵ If Z_0 is greater than a second threshold, which we denote with Z_0^* , the producer should immediately create an organized job. If Z_0 lies between these two thresholds, the producer's optimal policy is to delay creating an organized job until Z_1 is known. The organized job should be created in the second period if and only if the productivity shock is favorable.

A producer waiting to learn Z_1 should create a temporary job only if doing so is profitable, which will be the case if Z_0 is greater than W/θ . This is only possible if Z_0^* is greater than W/θ . Otherwise, all job creation opportunities that can be profitable temporary jobs are immediately created into organized jobs. Thus, the producer's use of temporary jobs depends on its policy for creating organized jobs.

We wish to demonstrate that increasing σ , which increases the producer's idiosyncratic risk, increases Z_0^* and thereby increases the producer's propensity to use temporary jobs. It is well

⁴To see this, note that creating an organized job following a favorable productivity shock is more profitable than perpetually maintaining a temporary job if and only if $Z_0(1 + \sigma) - \frac{r}{1+r}\tau_c > \theta Z_0(1 + \sigma)$. The analysis of the case following an unfavorable shock is identical and straightforward.

⁵To see that creating a temporary job is not optimal in this case, assume that the contrary is true and use the inequality in (1) to show that Z_0 must be greater than $\left(\frac{r}{1+r} \times \tau_c + W \right) / (1 + \sigma)$.

known that standard methods for the pricing and optimal exercise of financial call options can be used to characterize Z_0^* , and we follow this approach here. To do so, let $V_0(Z_0)$ denote the *fundamental value* of organized job creation, the expected present discounted value of profits from immediately creating an organized job, gross of the job creation cost, and let $F_0(Z_0)$ denote the value of an unexercised option to create an organized job in the following period. In this case, the expressions for $V_0(Z_0)$ and $F_0(Z_0)$ are

$$V_0(Z_0) = \frac{1+r}{r} \times (Z_0 - W),$$

$$F_0(Z_0) = \begin{cases} 0 & \text{if } Z_0 < \frac{\frac{r}{1+r} \times \tau_c + W}{1+\sigma} \\ \frac{1}{2} \times \frac{1}{1+r} \left(\frac{1+r}{r} \times ((1+\sigma)Z_0 - W) - \tau_c \right) & \text{if } \frac{\frac{r}{1+r} \times \tau_c + W}{1+\sigma} \leq Z_0 < \frac{\frac{r}{1+r} \times \tau_c + W}{1-\sigma} \\ \frac{1}{1+r} \left(\frac{1+r}{r} \times (Z_0 - W) - \tau_c \right) & \text{otherwise.} \end{cases}$$

Immediate creation of an organized job involves three costs, the direct cost of job creation, τ_c , the value of the option to create the job tomorrow, $F_0(Z_0)$, and the profits from operating the job opportunity as a temporary job. Immediate job creation is only optimal if the value of the job created is greater than its total cost. The value of Z_0 which equates the benefit and cost of immediately creating an organized job is the desired threshold. That is

$$V_0(Z_0^*) = \tau_c + F_0(Z_0^*) + \max\{\theta Z_0^* - W, 0\}. \quad (2)$$

Figure 3's upper panel illustrates the determination of Z_0^* by plotting the left and right hand sides of Equation (2) as functions of Z_0 , which we refer to as the fundamental value curve and the total cost curve. The total cost curve is kinked at $Z_0 = W/\theta$. If the kink occurs to the right of Z_0^* , then the firm will never create a temporary job. Otherwise, job opportunities with productivity between W/θ and Z_0^* will be filled with temporary jobs in the initial period. In this panel of the figure, Z_0^* is less than W/θ , so the producer will not use a temporary job.

Figure 3's lower panel repeats the analysis of the upper panel using a higher value of σ to show how idiosyncratic risk affects the use of temporary jobs. As is familiar from the study of irreversible investment, increasing idiosyncratic risk increases the value of the unexercised job creation option, rotating the total cost curve upward and increasing Z_0^* . In the figure, this

increase pushes Z_0^* past W/θ , making the creation of a temporary job optimal for some values of Z_0 . With the higher value of σ , creating a temporary job is optimal if Z_0 lies between W/θ and Z_0^* . With the lower value of σ in Figure 3's upper panel, all jobs created are organized. In this sense, increasing idiosyncratic risk increases producers' use of temporary jobs.

To examine how idiosyncratic productivity risk influences aggregate employment dynamics in this setting, consider a group of plants, within an industry, say, or alternatively, of a certain age classification. These plants consist of job creation opportunities with different values of Z_0 . The industry's employment growth in the initial period will equal the number of jobs created, which will depend on both the distribution of productivity across job creation opportunities and the optimal employment policy. If a job creation opportunity's productivity is greater than either Z_0^* or W/θ , it will be filled by some worker. We refer to the minimum of these two productivity thresholds as the effective job creation threshold. The number of jobs created in the initial period will equal the number of job creation opportunities with Z_0 greater than the effective job creation threshold.

Consider the industry's employment response to a very simple aggregate productivity shock which simply shifts the distribution of Z_0 to the right. This productivity shock leaves producers employment policies unchanged and affects aggregate employment only by increasing the total number of job creation opportunities with productivity above the effective job creation threshold. In this case, whether or not the industry's marginal job is temporary or organized is irrelevant for the industry's employment in the initial period. The only industry characteristic that matters for the response of employment to this shock is the number of job creation opportunities with Z_0 near the effective job creation threshold. This illustrates that producers' use of temporary jobs need not matter for aggregate dynamics.⁶

Our conclusions regarding temporary jobs and aggregate employment are less ambiguous when we instead consider an aggregate shock that necessarily changes producers' employment policies. Suppose now that an aggregate shock changes the wage in the initial period to W_0 , but the wages in all subsequent periods are unaffected. Both W_0 and W are known to the producer

⁶In this setting, producers' use of temporary jobs will matter for job destruction in the next period, because all temporary jobs that do not receive favorable productivity shocks are not replaced with permanent jobs. However, this is a consequence of the assumption that producers' productivity risk lasts only one period.

when it makes its initial job creation decision. If no producer uses temporary jobs in the initial period, then the derivative of the effective job creation threshold with respect to W_0 is

$$\frac{dZ_0^*}{dW_0} = \frac{r}{1 - \frac{1}{2} \frac{1+\sigma}{1+r}} < 1 + r.$$

The upper bound on dZ_0^*/dW_0 follows from the assumption that σ is less than one. When some producers use temporary jobs, the effective job creation threshold's derivative with respect to W_0 equals $1/\theta$, which is greater than $1 + r$ if θ is sufficiently below one. The intuition for this difference is that a temporary wage change has a relatively small effect on the total cost of organized job creation when compared to its benefit. Therefore, any use of temporary jobs induced by high idiosyncratic risk will increase the response of the effective job creation threshold to the transitory wage shock. In this sense, idiosyncratic risk makes individual producers' employment policies more flexible.

Whether or not the flexibility of producers' policies has a significant impact on the industry's employment growth ultimately depends on the distribution of Z_0 . If increasing σ moves the effective job threshold from a point in the distribution of Z_0 with many job creation opportunities to one with very few, then a high-risk industry's employment will respond very little to a wage change. On the other hand, the effective job creation threshold's location matters little for aggregate dynamics if the distribution of Z_0 is relatively flat. In this case, greater flexibility of producer's job creation policies will directly produce a larger response in aggregate employment to a given wage shock. In the current example, the distribution of Z_0 must be exogenously specified and unrelated to producers' productivity risk, so it is impossible to determine the importance of these distributional considerations for aggregate dynamics. For this reason, we proceed to consider a richer model in which the idiosyncratic risk an industry's producers face determines its productivity distribution. Nevertheless, this simple framework suggests that the effects of idiosyncratic risk on the flexibility of producers' employment policies may explain the empirical evidence relating idiosyncratic risk and aggregate employment dynamics.

4. The Model

The remainder of this paper quantitatively evaluates the impact of idiosyncratic risk and temporary jobs on aggregate employment dynamics. For this purpose we generalize the decision problem studied in the previous section and place the producer within the context of an industry. We consider an industry of plants that produce a single good for sale in a competitive product market. Plants use a single factor of production, labor, that they hire in a competitive labor market. Labor comes in fixed shift lengths. We refer to an employee working one shift as a job. As in the previous section, a worker can be used in one of two kinds of jobs, organized and temporary, which differ in productivity and adjustment costs.

Let n_t and l_t denote the number of organized and temporary jobs filled by a particular plant in period t . The plant's output in that period is $z_t^{1-\alpha} (n_t + \theta l_t)^\alpha$, where both α and θ are between zero and one. We refer to z_t as the plant's productivity, but it can also be interpreted as its stock of an endowed and untradable input, such as managerial skill. Productivity evolves as

$$\ln z_{t+1} = \ln z_t + \varepsilon_{t+1}.$$

We assume that ε_{t+1} has a Gaussian distribution that is truncated above and below by very large (in absolute value) numbers, a technical condition used in the analysis of the plant's dynamic programming problem.

The per period cost of a job measured in units of the constant output price, which we call the wage, is denoted by W_t , so the plant's total wage bill in period t is $W_t (n_t + l_t)$. The wage follows a two-state Markov chain between low and high values, W^l and W^h , with transition matrix Π . As described in Campbell and Fisher (2000), variations in the wage can be interpreted within a general equilibrium framework as arising from shocks to the supply of labor to the industry and shocks to the demand for the industry's product. In the baseline version of our model, changes in W_t are the only source of aggregate fluctuations and ε_t is purely idiosyncratic. However, we also consider specifications of our model in which plants' productivity is the product of idiosyncratic and aggregate components, where the latter component is the source of aggregate fluctuations.

The plant incurs employment adjustment costs, but only for changes in n_t . When the plant

increases n_t , it incurs the job creation cost of τ_c units of lost output for every organized job added. If the plant reduces n_t , it incurs the job destruction cost of τ_d units of lost output for every organized job destroyed. So the adjustment costs incurred by the plant when it changes employment from n_{t-1} to n_t are

$$\tau(n_{t-1}, n_t) = (\tau_c I\{n_t > n_{t-1}\} - \tau_d I\{n_t < n_{t-1}\})(n_t - n_{t-1}), \quad (3)$$

where $I\{\cdot\}$ is the indicator function. The job creation and destruction costs are unrelated to the identity of the worker filling the job, so they can be interpreted as *net adjustment costs*, as defined by Hammermesh and Pfann (1996). In a model similar to ours but without aggregate fluctuations, Aguirregabiria and Alonso-Borrego (1999) interpret τ_c as the cost of hiring a “permanent” worker, and τ_d is the cost of dismissing that worker. In this case it is costly to replace one worker with another, so there is surplus value to be split between a worker and her employer, as in Mortensen and Pissarides (1994). Our analysis goes through under this interpretation if we assume that W_t is the return a worker gets to operating a home production technology, and that the employer has all bargaining power when negotiating wages.

To examine how groups of plants facing different degrees of idiosyncratic risk respond to the same aggregate shock, we incorporate a plant life-cycle into the model. An exogenous flow of new plants begins production every period. So that the mass of plants operating is constant over time, all incumbents face an exogenous probability of exit as well. Because we assume exit is exogenous, we are unable to characterize the effect of idiosyncratic productivity risk on producers’ exit decisions. However, this assumption preserves the technical simplicity of the plant’s optimization problem. At any given plant, the process governing exit and the evolution of productivity depends on an exogenous state variable, s_t , which can equal either s^n , for “new,” or s^m , for “mature.” Plants begin life with $s_t = s^n$ and after every period there is a constant probability, ϕ , of transiting to $s_t = s^m$, which is an absorbing state. A plant’s probability of exit after period t is $\delta(s_t)$. The mean and variance of the plant’s productivity innovation, $\mu(s_t)$ and $\sigma^2(s_t)$, may also depend on the state variable. Note that plants for which $s_t = s^n$ will tend to be younger than average. Hence, while not important from a technical standpoint, it is natural

to assume that $\sigma(s^n) > \sigma(s^m)$ and $\delta(s^n) > \delta(s^m)$, so that young plants are subject to more productivity and exit risk than their older counterparts.

Finally, we assume that risk-neutral managers who discount future profits with the constant discount factor $\beta \in (0, 1)$ choose employment to maximize the expected present discounted value of the plants' profits. We now describe the decision problems of mature and new plants.

4.1. Mature Plants

The manager of a mature plant faces the problem of choosing the number of organized and temporary jobs to fill every period in order to maximize the present discounted value of profits. Conditional on a plant's choice of organized jobs, its use of temporary jobs solves a static profit maximization problem. Denote this problem's attained value with $f(z, n, W)$,

$$f(z, n, W) = \max_{l \geq 0} z^{1-\alpha} (n + \theta l)^\alpha - W(n + l). \quad (4)$$

The solution to this problem is straightforward. If n/z is greater than $(W/(\alpha\theta))^{1/(1-\alpha)}$, then the marginal product of the first temporary job is below the wage and the plant uses no temporary jobs. Otherwise, the plant's choice of l satisfies

$$\theta\alpha z^{1-\alpha} (n + \theta l)^{\alpha-1} = W.$$

To characterize the choice of organized jobs, we cast the relevant profit maximization problem as a dynamic program. There are three state variables, the plant's productivity, the number of organized jobs it used in the previous period, and the wage. Using m to denote the number of organized jobs in the previous period and dropping time subscripts, the dynamic program is

$$g(z, m, W) = \max_{n \geq 0} f(z, n, W) - \tau(m, n) + \beta(1 - \delta(s^m)) \mathbf{E}_m[g(z', n, W') | z, W].$$

Here, $\mathbf{E}_m[\cdot | \cdot]$ is the conditional expectations operator associated with the transition matrix for W and the evolution of z for mature plants, and we use the "'' notation in the usual fashion. It is straightforward to show, as in Campbell and Fisher (2000), that $g(z, m, W)$ can be written

as $g(z, m, W) = z \times v(m/z, W)$, where $v(m/z, W) = g(1, m/z, W)$. Define the scaled numbers of organized jobs, $x = m/z$ and $y = n/z$. Substituting the scaled version of the value function into the Bellman equation and utilizing the fact that $\ln z$ follows a random walk, we can define a Bellman equation for the scaled value function, $v(x, W)$,

$$v(x, W) = \max_{y \geq 0} [\pi(x, y, W) + \beta(1 - \delta(s^m)) \mathbf{E}_m[uv(y/u, W') | W]],$$

where $u = z'/z = \exp(\varepsilon)$,

$$\pi(x, y, W) = \begin{cases} \tilde{y}(W)^\alpha - W \left(y + \frac{\tilde{y}(W) - y}{\theta} \right) - \tau(x, y) & \text{if } y < \tilde{y}(W) \\ y^\alpha - Wy - \tau(x, y) & \text{if } y \geq \tilde{y}(W). \end{cases}$$

and

$$\tilde{y}(W) = \left(\frac{W}{\theta\alpha} \right)^{1/(\alpha-1)}.$$

This final quantity, $\tilde{y}(W)$, equals the scaled value of optimal employment at a plant that faces no adjustment costs and a wage equal to W/θ . For this reason, we refer to it as the plant's frictionless labor demand schedule.

In general, the optimal choice of y can involve either increasing the number of organized jobs, decreasing it, or leaving it unchanged. If the plant increases the number of organized jobs, then the optimal value of y must satisfy

$$\begin{aligned} W \left(\frac{1}{\theta} - 1 \right) - \tau_c + \beta(1 - \delta(s^m)) \mathbf{E}_m[v'(y/u, W') | W] &= 0, \text{ if } y < \tilde{y}(W) \\ \alpha y^{\alpha-1} - W - \tau_c + \beta(1 - \delta(s^m)) \mathbf{E}_m[v'(y/u, W') | W] &= 0, \text{ if } y \geq \tilde{y}(W), \end{aligned} \quad (5)$$

where $v'(x, W)$ is the derivative of $v(x, W)$ with respect to its first argument, which exists almost everywhere due to that function's concavity. If the plant decreases its stock of organized jobs, the first-order condition for y is

$$\alpha y^{\alpha-1} - W + \tau_d + \beta(1 - \delta(s^m)) \mathbf{E}_m[v'(y/u, W') | W] = 0. \quad (6)$$

To characterize the optimal policy, let $\underline{y}^m(W)$ and $\bar{y}^m(W)$ be the unique values of y that satisfy (5) and (6) with equality. The optimal policy is to set $y = \underline{y}^m(W)$ if $x \leq \underline{y}^m(W)$, to set $y = \bar{y}^m(W)$ if $x \geq \bar{y}^m(W)$, and to set $y = x$ if $\underline{y}^m(W) < x < \bar{y}^m(W)$. To describe the decision rule for n , define the job creation and destruction schedules $\underline{n}^m(z, W) = z \times \underline{y}^m(W)$ and $\bar{n}^m(z, W) = z \times \bar{y}^m(W)$, respectively. Then, the optimal policy is to set n equal to the job creation schedule if m falls below it and to set n equal to the job destruction schedule if m falls above it. If m falls between these two schedules, the optimal policy is to leave the stock of organized jobs unchanged. If the optimal choice of n is less than $z \times \tilde{y}(W)$, then the optimal number of temporary jobs is

$$l = \frac{(z \times \tilde{y}(W) - n)}{\theta}.$$

Otherwise, the optimal number of temporary jobs equals zero.

Figure 1 illustrates the optimal employment policy for a mature plant for a given value of W . Here, we have assumed $\tilde{y}(W) < \underline{y}^m(W)$, so that no mature plant uses temporary jobs at this wage. The job creation and destruction schedules are both log linear in z with intercepts $\ln \underline{y}^m(W)$ and $\ln \bar{y}^m(W)$ and a common slope of one. The figure illustrates the employment decisions of three plants, labeled A , B , and C , with the same value of lagged employment but different current productivity realizations. Plant A lies above the job destruction schedule, so it reduces employment. Plant B lies between the two schedules in the region of inaction, so it leaves employment unchanged. Plant C lies below the job creation schedule, so it increases employment. The logarithmic employment growth rates of plants A and C equal the vertical distances from their initial locations to the job creation and destruction schedules.

4.2. New Plants

The problem faced by a new plant is similar to that of mature plants. The only differences are that the probability of exit and the distribution of the idiosyncratic productivity innovations correspond to those for new plants, and there is a constant probability of transiting from the new to the mature state. Just as with mature plants, job creation and destruction schedules characterize the optimal choices of organized jobs for new plants. The Bellman equation characterizing

a new plant's value function, $r(z, m, W)$, is

$$r(z, m, W) = \max_{n \geq 0} \{f(z, n, W) - \tau(m, n) + \beta(1 - \delta(s^n))(\mathbf{E}_n[\phi g(z', n, W') + (1 - \phi)r(z', n, W') | z, W])\}.$$

The expectations operator $\mathbf{E}_n[\cdot]$ is that associated with the transition matrix for W and the evolution of z for new plants. Just as with the mature plant's problem, the unique solution to the new plant's problem can be written as $r(z, m, W) = z \times r(1, m/z, W)$, so that there exists a scaled version of the Bellman equation with a unique, concave, solution. The optimal employment policies for new plants therefore have the same form as those for old plants, but they may take on different values. That is, there are job creation and destruction schedules, $\underline{n}^n(z, W) = z \times \underline{y}^n(W)$ and $\bar{n}^n(z, W) = z \times \bar{y}^n(W)$. The optimal policy is to leave the plant's stock of organized jobs unchanged if the plant falls between the two schedules. Otherwise, the plant should adjust employment in organized jobs to the closest schedule. Optimal employment in temporary jobs can then be determined exactly as it is for mature plants.

Figure 2 illustrates the optimal employment policy for new plants, assuming $\tilde{y}(W) > \underline{y}^n(W)$, so that a new plant might use temporary jobs at this wage. The figure plots job creation and destruction schedules as well as the frictionless labor demand schedule a plant would use if it faced no adjustment cost and a wage equal to W/θ . Adding the possibility of using temporary jobs complicates plants' employment dynamics, but it also enlarges the scope of observable behavior that the model can generate. To be concrete, consider the employment decisions of six plants, labeled A through F in the figure. All of these plants have the same lagged value of employment in organized jobs, n_{t-1} , but different realizations of current productivity. Suppose that the wage does not change between the previous and current period, so that these plants' employment changes reflect only idiosyncratic shocks.

As in Figure 1, plant A lies above the job destruction schedule. Its optimal employment policy is to lower n down to that schedule and to set its employment of temporary jobs equal to zero. Therefore, we know that this plant destroys jobs. Unlike in the case without temporary jobs, the extent to which plant A destroys jobs depends not only on the vertical distance from

A to the job destruction schedule but also on that plant's use of temporary jobs in the previous period. Plants *B* and *C* both lie between the job creation and destruction schedules, so they will leave their stock of organized jobs unchanged. Because these plants also lie above the frictionless labor demand schedule, they will not use temporary jobs this period. Even though their stock of organized jobs will be unchanged, we may observe them destroying jobs if they used temporary jobs in the previous period. For example, *B* will destroy temporary jobs if it was at a point like *D* in the previous period, but it will not if it was at a point like *C*. Plants *D* and *E* are also between the job creation and destruction schedules, but they both lie below the frictionless labor demand schedule. Therefore, these plants may either create or destroy temporary jobs. For example, *D* will create jobs if it ended the previous period at *C*, and it will destroy jobs if it transited from *E*. Finally, plant *F* lies below the job creation schedule. It will increase its stock of organized jobs and its employment of temporary jobs.⁷

The preceding discussion illustrates that not all plants increasing or decreasing employment lie on the job creation and destruction schedules. Those plants that fall in the region of inaction between the two schedules may change their use of temporary jobs. However, the model displays an empirically relevant feature of the simpler models without temporary jobs, namely that a positive fraction of plants leave their employment unchanged between periods. The variety of plants' responses to a change in W_t is correspondingly wide. If W_t rises from W^l to W^h , we expect those plants on the job destruction schedule to adjust their stock of organized jobs further downward. Plants in the region of inaction above the frictionless labor demand schedule will not respond at all to the shock, while those below it will respond to the shock only by changing their use of temporary jobs. Finally, those plants on the job creation schedule will change both their organized and temporary jobs in response to the wage increase.

⁷When W_t changes between periods, a number of new possibilities arise. For example, a plant like *F* could be observed to reduce employment if the frictionless labor demand schedule falls so much that the plant reduces its use of temporary jobs in spite of its favorable idiosyncratic productivity shock.

5. Aggregate Implications of Idiosyncratic Risk

This section uses the industry model to assess the quantitative impact of idiosyncratic risk on aggregate employment dynamics. These dynamics are derived by aggregating the individual decisions of the population of plants in the model.⁸ We are interested in two key issues left unresolved by studying the example in Section 3. First, do plants' employment decisions behave as suggested by the example? Second, how do these microeconomic policies affect aggregate variables when the distribution of plants in the state space is endogenous? Since the industry model is more general than the example, we are also interested in documenting other effects of idiosyncratic risk not suggested by the example.

To guide our analysis we compare the employment dynamics emerging from various parameterizations of the model to data on job flows in the U.S. manufacturing sector compiled by DHS. We find that the microeconomic policies in the industry model behave as predicted by the example and that because of this the model shares many of its features of gross job flows with the data. Idiosyncratic risk's impact on the use of temporary jobs is quantitatively the most important mechanism in the model.

The analysis proceeds in four steps. First, we discuss how parameter values were selected. Second, we document numerically the impact of idiosyncratic uncertainty on a plant's optimal employment policy. The model's implications for aggregate dynamics depend on plant-level decisions, but they also depend on how the distribution of employment and productivity across plants evolves. As is familiar from Caballero (1992), endogenous changes in the plants' distribution can make the evolution of an aggregate variable quite unlike that of its counterpart at any particular plant. The third step, then, is to verify that these considerations do not completely offset the effects of the microeconomic policies. To do so, we display the dynamics of employment growth of the populations of young and old plants for several special cases of our model. Finally, we compare the statistics on gross job flows resulting from simulations of the model to

⁸The solution to the decision problems of individual plants is quite straightforward and can be computed almost exactly. See Campbell and Fisher (1999) for the method. Computing the aggregate dynamics in this model is less straightforward. For a description of the relevant transition equations for the distribution of employment across plants and the expressions for gross job flows derived from the time series of this distribution, see the technical appendix of this paper, available on the World Wide Web at <http://home.uchicago.edu/~jcampbe>.

the evidence from the U.S. manufacturing sector.

5.1. Parameter Values

To simulate our model we need to specify the following parameters:

Plant-level parameters: $\alpha, \beta, \mu_n, \sigma_n, \mu_m, \sigma_m, \delta_n, \delta_m, \tau_c, \tau_d, \theta, \phi$

Aggregate Uncertainty: Π, W^h, W^l .

We proceed by finding parameters so that the model comes close to mimicking the gross job flows in a “representative” manufacturing industry when temporary jobs are used sometimes. To quantify the various mechanisms operating in the model, we study perturbations of the model around this baseline calibration. We use the gross job flow data by plant age for the manufacturing sector as a whole to choose parameters. Data by plant age are not available by industry.

The elasticity of production with respect to labor input, α , is approximately equal to the share of output paid to labor.⁹ We set α to 0.66. The discount factor, β , is set to $1.05^{-1/4}$ so that the annual risk free real interest rate is 5%, and a period in the model corresponds to one quarter of a year.

Holding fixed all other model parameters, the job reallocation rate for young plants (consistent with the DHS data we study below, age 40 quarters or less) and old plants (all other plants) in the model is increasing in the amount of idiosyncratic uncertainty. Therefore we choose the two standard deviations, σ_n and σ_m , by matching the steady state of the model to the average job reallocation rates (the sum of the rates of job creation and destruction) for young and old plants in the U.S. manufacturing sector, 14.1% and 9.8%, respectively.¹⁰ The expected value of productivity growth for a plant of type $x = n, m$ is given by $\exp(\mu_x + \sigma_x^2/2)$. We use this formula to fix μ_n and μ_m so that productivity growth is zero on average for both new and mature plants.

⁹This will not be exact because of the adjustment costs.

¹⁰Rates of job creation and destruction in the model are the population counterparts of the measures defined as in DHS. For a particular set of plants, e.g. young plants, the rate of job creation at date t is the sum of all jobs created between dates t and date $t - 1$ divided by the average of total employment for the group of plants in dates t and $t - 1$. The rate of job destruction is defined similarly.

This ensures that total industry employment does not have a trend.

Increases in the exogenous exit rates, δ_n and δ_m , increase the fraction of employment at plants which exit (*NEGD* in the notation of DHS). Therefore, we choose these parameters so that the fraction of employment at young and old plants which exit in the steady state of the model corresponds to averages for these variables in the U.S. manufacturing sector. The target averages are 1.23% for young plants and 0.47% for old plants.

We follow Campbell and Fisher (2000) in selecting baseline values for values τ_c and τ_d , the costs of creating and destroying organized jobs. These net adjustment costs involve disruptions to production and all other costs that are not related to the identity of a worker but depend solely on changing the number of employees. We choose the process for the wage so that on average it equals unity in the model. Consequently, values for the adjustment cost parameters can be interpreted as fractions of a quarterly wage. Campbell and Fisher (2000) survey the available evidence and determine that $\tau_c = 1.5$ and $\tau_d = 0.5$ is a reasonable benchmark. Note that this is consistent with the generally held view that the most relevant case for the U.S. labor market is that creation costs exceed destruction costs. However, our qualitative conclusions also hold if the cost of job destruction is larger than that for job creation.

The final two plant-level parameters are ϕ , the probability of a new plant transiting to mature status, and θ , the relative productivity of temporary jobs. We fix ϕ so that the average time before transiting to mature status following birth is 5 years. This corresponds to $\phi = 0.05$. We assume $\theta = 0.7$. With this choice temporary jobs are used by new plants only.

We consider several specifications of the process governing the wage. In each case $\mathbf{\Pi}$ has the basic structure

$$\mathbf{\Pi} = \begin{bmatrix} \frac{1+\rho}{2} & \frac{1-\rho}{2} \\ \frac{1-\rho}{2} & \frac{1+\rho}{2} \end{bmatrix}, \quad (7)$$

and the two wage states are $W^h = \exp(\sigma_W)$ and $W^l = \exp(-\sigma_W)$. Here ρ and σ_W are the first order autocorrelation coefficient and the standard deviation of $\ln(W_t)$, respectively. In our baseline calibration we fix $\rho = 0$ so that the wage is *i.i.d.* and select σ_W to match the standard deviation of employment growth for the manufacturing sector, 1.90%. We do consider other values for ρ .

The parameters underlying our baseline calibration are summarized in Table 1. Notice that the standard deviation of productivity and the exit rate at new plants are greater than at mature plants. Since plants begin life as new-type plants and only later become mature-type plants, this corresponds to younger plants in the model being faced with greater idiosyncratic risk than old plants, consistent with the evidence reported by Dunne, Roberts and Samuelson (1989).¹¹ Hence our baseline calibration corresponds to the conventional understanding of how idiosyncratic risk varies by plant (or firm) age.

The implied employment policy is shown in the first row of Table 5. This table lists the intercepts of the creation and destruction schedules of new and mature plants as well as the intercepts of the frictionless employment schedule. As this table shows, $\tilde{y}(W^l) = 0.117 > 0.098 = \underline{y}^n(W^l)$, which means temporary jobs are used by new plants in the low wage state. However, $\underline{y}^n(W^h)$ is greater than $\tilde{y}(W^h)$, so new plants do not use temporary jobs in the high wage state. Mature plants never use temporary jobs. Since idiosyncratic risk is greater for new plants than mature plants in the calibrated model, this is consistent with the example.

5.2. Effects of Idiosyncratic Risk on Optimal Employment Policies

In this section we consider the impact of idiosyncratic risk on the optimal employment policies of mature plants. We fix all other parameters at their calibrated values except those noted below. Consider first the impact of idiosyncratic risk on the decision to use temporary jobs. Unlike the example in Section 3, the industry model has exit risk in addition to productivity risk, and this too may influence the use of temporary jobs because it affects the rate at which future adjustment costs are discounted. Figure 6 shows the effects of increasing productivity (top graph) and exit risk (bottom graph) on the log level of the intercepts of the organized creation schedule in the steady state. To help evaluate the effects, note that the calibrated values of productivity risk and exit risk at new plants are near the midpoints of the horizontal axes. The solid line in each plot corresponds to the calibrated values for the adjustment costs, the short dashed lines correspond to a case in which τ_d is raised to 2.0 and τ_c is lowered to

¹¹Note that the magnitudes of σ_n and σ_m overstate the degree of productivity risk faced by plants since productivity is raised to the power $1 - \alpha$ in the production function.

zero, and the long dashed lines correspond to a case where τ_c equals 2.0 and τ_d equals zero. The horizontal lines indicate the threshold determining whether temporary jobs are ever used. When the intercept of the log creation schedule lies below this threshold, temporary jobs are used. The extent to which a plant relies on temporary jobs is directly related to the vertical distance between the threshold and the intercept.

Consistent with the example developed in Section 3, the top graph of Figure 6 shows that increasing productivity risk does indeed increase the use of temporary jobs in the industry model. Moreover, according to the graph, this fact does not depend on the configuration of adjustment costs. Another result is that the intensity of temporary job use is higher when adjustment costs fall more heavily on the creation margin. When adjustment costs are configured in this way the total cost of creating a job rises because future job destruction costs, another cost of creating a job, are discounted and job creation costs are not. Hence there is greater substitution away from organized jobs and toward temporary jobs.

The bottom graph of Figure 6 shows that increasing exit risk, δ_m , can increase the use of temporary jobs as well, but that this depends on the configuration of adjustment costs. If the adjustment costs fall more heavily on the destruction margin, then the exit rate does not have this effect; otherwise, it does. Increasing the exit rate increases the rate at which future costs of adding an organized job are discounted. Since the future destruction costs associated with an organized job are a cost of adding such a job, increasing exit risk *reduces* the total cost of the marginal job. With diminishing returns, this implies that more organized jobs must be added. Increasing exit risk lowers the value of future creation costs avoided by creating an organized job immediately, and so when adjustment costs fall more heavily on this margin the total cost of organized job creation *rises*.

5.3. Effects of Idiosyncratic Risk on Employment Growth Aggregates

Here we study the dynamics of employment growth of young and old plants implied by the optimal employment policies of new and mature plants for two polar cases for the wage transition matrix. The first case we consider involves specifying $\rho = -1$ so that the off-diagonal elements of Π equal unity. This implies that the wage follows a two-period cycle. We consider two

parameterizations. In one, the remaining parameters are fixed at their calibrated values, in particular $\theta = 0.7$. With a two period cycle this implies new plants use temporary jobs in both wage states, but mature plants never use temporary jobs. The other parameterization is identical except $\theta = 0$ so that only organized jobs are used by both new and mature plants.

Demeaned simulated employment growth for young plants (solid lines/circles) and old plants (dash lines/squares) implied by the two parameterizations are illustrated in Figure 7, where the top plot corresponds to the model in which temporary jobs are used by new plants and the bottom plot corresponds to the parameterization with only organized jobs. The top plot shows that when temporary jobs are used by new plants only, they can have a substantial impact on employment growth, making young plants much more variable than old plants. The bottom plot shows that differences in the employment policies of young and old plants translate only slightly into differences in their aggregate dynamics when temporary jobs are not used. Nevertheless, young plants are more variable than old plants, as is the case in the U.S. manufacturing sector. Consistent with the example, overall volatility is clearly lower than in the case without temporary jobs.

The polar opposite of the two-period cycle for the wage is to suppose that when computing their optimal employment policy, plant managers expect the current wage to prevail forever. This is captured by assuming $\rho = 1$ in (7). With $\theta = 0.7$ and $\rho = 1$ temporary jobs are only used by new plants in the low wage state. Figure 8 plots the dynamics following a wage change from W^h to W^l for the two parameterizations, $\theta = 0.7$ and $\theta = 0$. The change in the wage occurs after first simulating the model for several periods at the initial wage until employment and the distribution of employment across plants converge. Employment growth rates are plotted relative to their values before the change. Generally, the findings are similar to the two-period cycle cases. With temporary jobs used by new plants, young plants clearly respond more to the wage change than old plants. Without temporary jobs, the difference is less pronounced. Note that in both parameterizations there is some persistence in the response of employment growth. The key here is the region of inaction in the state space of the employment policy induced by proportional adjustment costs. This feature of the model provides for an endogenous source of persistence since individual plants are slow to respond to wage changes.

5.4. Quantifying the Effects of Idiosyncratic Risk

We now study the statistical properties of the calibrated model and its perturbations. The perturbations are designed to shed light on the quantitative impact of (i) the use of temporary jobs, (ii) increased idiosyncratic risk at new plants, (iii) higher adjustment costs, (iv) persistence in the wage process and (v) aggregate technology shocks that leave plants' scaled employment policies unchanged. These perturbations and their implications for the employment policies are summarized in the remaining rows of Table 5. In Tables 6 and 7 we display statistics on gross job flows for the U.S. manufacturing sector reproduced from Table 1 and statistics based on the various versions of the model. Table 6 presents steady state statistics and Table 7 presents cyclical features of the data based on simulating the model.¹² In the tables, ' G ' denotes employment growth, ' C ' denotes the rate of job creation, ' D ' denotes the rate of job destruction and ' R ' denotes the rate of job reallocation (recall that by definition $R = C + D$.)

5.4.1. The Calibrated Model

As already mentioned, the policy intercepts in Table 5 indicate that in the calibrated model new plants choose to use temporary jobs only in the low wage state, and mature plants never use them. Table 5 also shows that the intercept for the organized job creation schedule of new plants is essentially the same across wage states, unlike with mature plants whose organized creation intercept is clearly higher in the low wage state.

The implications of these employment policies for gross job flows are summarized in Panel B of Tables 6 and 7. Consider first the steady state values (which are close but not identical to the means calculated from simulating the model with aggregate uncertainty). The parameters were chosen to match the empirical reallocation rates for young and old plants, yet the model does reasonably well at accounting for the total reallocation rate. This is a result of the fact that the share of employment at young plants, 17%, is not far off the value of this share in the data, 22.5%. Abstracting from the assumed absence of an employment trend in the model, the model does quite well at accounting for the remaining qualitative features of the data. For example,

¹²All statistics for the model are based on simulating the model for 1050 periods, starting at the steady state distribution of employment across plants and then discarding the first 50 observations.

average growth at young plants is above zero and at old plants is less than zero. This difference results from the fact that mature plants are larger than new plants and that a large fraction of employment growth at young plants is accounted for by the switch from new to mature status. At old plants exit contributes to the net job loss on average.

Now consider the cyclical features of the data summarized in Panel B of Table 7. This shows that idiosyncratic risk impacts aggregate employment growth in the manner anticipated by the example in Section 3. The model was calibrated to the aggregate growth rate's variance, but *not* the variances of the age-specific growth rates. In the data the standard deviation of employment growth, σ_G , at young plants is 2.32%, and in the model (where these plants rely on temporary jobs) it is 3.69%. In the model old plants rely exclusively on the less flexible organized jobs and so σ_G for these plants is much lower, 1.55%, compared to 1.81% in the data. If anything, then, the calibrated model goes too far in reproducing the observed differences in the aggregate growth rates.

The calibrated model also does well at capturing the differences in the volatility of creation, destruction and reallocation at young and old plants. In the data these variables fluctuate more at young plants than at old, and this is the case in the model as well. Moreover, the model reproduces the qualitative result that the destruction-creation variance ratio is much lower for young plants than for old plants. Due to the way job reallocation and job growth are defined in terms of job creation and destruction (one is the sum and the other is the difference of these variables), qualitative success in terms of the variance ratio translates into a qualitative success with respect to the correlation between job reallocation and job growth: job reallocation is strongly countercyclical at old plants and less so at young plants.

To understand this success, first note that, as in Campbell and Fisher (2000), proportional adjustment costs induce plants to move the organized job destruction schedule more than the organized job creation schedule following a wage change. This asymmetric policy response creates asymmetries in gross job flows. Because young plants use temporary jobs as well as organized jobs, the asymmetric fluctuations of the policies for creating and destroying organized jobs have a smaller impact on their variance ratios. The evolution of the distribution of plants across the state space also plays a crucial role. Since young plants begin life creating jobs, a

relatively large fraction of the employment of young plants is bunched near the organized job creation margin. This amplifies the variation in job creation and contributes further to the low variance asymmetry of young plants.

There are two main drawbacks of this calibration. Ultimately, these stem from the simplicity of the model. First, the reallocation rates for young and old plants, and hence the whole population of plants, are much less volatile in the model than in the data. The reason for this is that, with only one source of aggregate variation in the model, job creation and destruction are essentially perfectly negatively correlated. Since job reallocation is just the sum of job creation and destruction, this dampens variation in job reallocation. The second drawback of the model is that it lacks persistence. We consider this below.

5.4.2. Assessing the Role of Temporary Jobs

To assess the role of temporary jobs in generating the asymmetries in gross job flows across young and old plants, consider the results for the $\theta = 0$ and $\theta = 0.75$ perturbations of the model. Obviously, when temporary jobs are not productive, no temporary jobs are used at all. In this case we expect gross job flows to be less variable for all plants, and especially young plants, compared to the calibrated model. When the productivity of temporary jobs is increased above its calibrated value, Table 5 shows that new plants use temporary jobs in both wage states and mature plants use them in the low wage state. Therefore, we expect gross job flows to be more variable than in the calibrated model. Finally, note from Table 5 that with $\theta = 0.75$ the intercept of the organized creation margin for new plants increases with the wage so that the use of organized jobs is countercyclical at these plants.

Panels C and D of Tables 6 and 7 show the aggregate implications of these employment policies. Notice that with organized jobs only, $\theta = 0$, the model generates an asymmetry in the growth rates of young and old plants, but it is much smaller than in the calibrated model with temporary jobs. This is consistent with the simulation results from the model with a two-period cycle plotted Figure 7. The asymmetry is an implication of the fact that the creation and destruction margins respond more in percentage terms at new plants than at mature plants. From Table 5 we see that the creation margin changes by 7% more for new plants than for mature

plants when the wage changes, and the destruction margin changes by 16% more. The fact that the aggregate growth rates do not fully reflect these differences is due to the fact that when the wage changes not all plants change employment — some are in the region of inaction. On the other hand, the fact that higher variance induces plants to respond more frequently implies that among new plants a greater fraction of plants are at or near their adjustment thresholds compared to mature plants. This amplifies the effects of the policy differences.

The fact that increasing productivity risk makes plants which use only organized jobs more responsive to a wage shock can be understood as follows. Note first that when a plant is expected to spend a long time in the region of inaction, a change in the current wage has a small percentage effect on the total expected wage bill for the last organized job created or retained. Increasing productivity risk tends to increase the frequency of job changes and reduce the time spent in the region of action. This implies that the effect of a wage change on the total expected wage bill for the last organized job created or retained is larger, and we expect there to be a correspondingly larger impact on employment. Due to its effects on discounting the future, increasing exit risk has a similar qualitative effect.

Since employment dynamics in the $\theta = 0$ case are not influenced by the use of temporary jobs, the Campbell and Fisher (2000) asymmetry effect is strong. Therefore, the model without temporary jobs is relatively successful in accounting for the variance ratio. This version of the model does less well at capturing other differences between young and old plants, as creation, destruction and reallocation are all less volatile at young plants in this version of the model.

With $\theta = 0.75$ temporary jobs are used by a larger number of plants and are used more intensively at plants that were already relying on them, compared to the calibrated model. This means that a greater fraction of employment dynamics is governed by plants operating on the more flexible margin. So, increasing θ above its baseline value generally raises overall variability near the calibrated parameter values. However, this parameter change reduces relative variability in terms of ratios of gross flow measures. Of course if only temporary jobs are used, then the model acts like a frictionless representative producer and all asymmetries disappear.

Several other features of this simulation are worth noting. First, the variance ratio uniformly drops below one. Also notice that having more temporary jobs moves the correlations between

creation and destruction towards zero. This is a direct result of the fact that the use of organized jobs is countercyclical for this parameterization. Finally, since plants' use of temporary jobs reduces the size of the region of inaction, there is less endogenous persistence in this version of the model.

We now briefly discuss some implications of using other values of the relative productivity parameter, θ . It turns out that a plant's use of temporary jobs is quite sensitive to changes in the magnitude of θ . This can be seen by noting that, with all other parameters kept at their baseline values, if $\theta < 0.66$, then no temporary jobs would be used at all, and if $\theta > 0.77$, then both new and mature plants would use temporary jobs in both wage states. Since intensity of temporary job use changes smoothly with θ , the impact of this sensitivity is mitigated in the implied aggregate employment dynamics. As θ is increased so that both new and mature plants use temporary jobs, the differences between young and old plants begin to disappear. Higher productivity uncertainty for new plants lessens the sensitivity of temporary job use to θ . For example, if σ_n is twice the magnitude of σ_m , then the threshold below which no temporary jobs are used by new plants is 0.59.

5.4.3. Increasing Idiosyncratic Risk

The next two perturbations of the calibrated versions of the model involve increasing productivity risk and exit risk at new plants. Panels E and F in Tables 6 and 7 cover these cases. Comparing the policy intercepts indicated in Table 5 for these cases to the calibrated case, we see that increasing idiosyncratic risk in these ways has the qualitative effects anticipated above. For example, when either source of risk is increased, the use of temporary jobs by new plants increases noticeably relative to the calibrated model. The volatility of gross job flows for both young and old plants is increased on all dimensions.

Given the cross-industry empirical evidence discussed in Section 2, it is also interesting to examine how industry-wide job growth depends on overall productivity risk. We can do this by scaling the variances, σ_n and σ_m , relative to the calibrated model to generate data from model industries facing different amounts of risk. Plotting the standard deviation of total employment growth versus the sum of the job creation and destruction rates' sample medians produces a

model analogue of Figure 2. This is done in Figure 9. To generate the data for this figure, we used scale factors ranging from 0.8 to 2.4. A circle marks the data from the model's baseline calibration. Because we used annual data to construct Figure 2 while our model produces quarterly data, the two figures are not strictly comparable. However, Figure 9 does show that the model economy can replicate the sort of cross-industry observations from Section 2 by only changing the model's idiosyncratic risk.

Increasing idiosyncratic risk raises both the standard deviation of employment growth and the median job creation and destruction rates, so Figure 9 displays a monotone increasing relationship. The observed relationship is nonlinear. Around the calibrated model, the slope is approximately 0.25, a number surprisingly close to the slope of the OLS regression line from Figure 2, 0.17. When idiosyncratic risk is larger and more plants use temporary jobs, the slope of the risk-variance relationship increases. In the middle of the graph, it is approximately 1.3. Finally, as idiosyncratic risk increases even further, plants use temporary jobs so intensively that the industry begins to resemble a representative producer operating a Cobb-Douglas technology. Microeconomic uncertainty is irrelevant for such an industry, so the relationship between our risk measure and employment volatility flattens. Overall, the figure displays a considerable impact of idiosyncratic risk on aggregate uncertainty. Doubling the calibrated value of idiosyncratic risk increases the standard deviation of aggregate employment growth by more than a factor of four.

5.4.4. Higher Adjustment Costs and Persistent Wages

When adjustment costs are increased compared to the baseline calibration, Table 5 shows that the intensity of temporary job usage rises substantially. Now both new and mature plants use temporary jobs in both wage states. Since the organized creation intercepts are small compared to the frictionless ones, plants use a lot of temporary jobs. Aggregate dynamics for this case are summarized in Panel G of Tables 6 and 7. Due to their effect on the use of temporary jobs, higher adjustment costs amplify fluctuations but reduce relative differences in gross job flows between young and old plants. The impact on overall variability contrasts sharply with the usual effect of increasing adjustment costs, which is to reduce volatility (this is what happens

in a version of the calibrated model in which $\theta = 0$.)

A persistent wage process (here the persistence parameter is the calibrated value from Campbell and Fisher (2000), $\rho = 0.47$) does not change the predictions of the model relative to the baseline very much, except that it increases the persistence of job growth. Job growth is similarly strongly positively autocorrelated at young and old plants in the data. In the baseline parameterization of the model it is identically negatively autocorrelated at young and old plants. In a model without adjustment costs and with an *i.i.d.* wage process, the autocorrelation of job growth would be -0.5 . Yet in the baseline case which has this wage process, the magnitude is -0.34 . The region of inaction is the source of the endogenous persistence indicated here. It makes plants slow to respond to wage changes on average. When wage changes are made more persistent (but still negatively autocorrelated in growth rates), employment growth becomes positively autocorrelated. The best the model can do along this dimension is to assume ρ is near unity so that wage growth is close to *i.i.d.* In this case the autocorrelation of employment growth is about 0.5. Additional persistence could be generated by the model if we assumed a driving process which implied that wage growth was positively autocorrelated, such as a deterministic sine wave.

5.4.5. Aggregate Technology Shocks

In the simple framework of Section 3, the use of temporary jobs does not necessarily change aggregate employment dynamics if all shocks simply shift the distribution of productivity and leave job creation policies unchanged. In the industry model, idiosyncratic risk endogenously determines both the number of plants which use temporary jobs and the intensity with which they use them. This provides a natural channel through which idiosyncratic risk can influence the industry's dynamics, even if plants' employment policies are invariant to aggregate shocks.

To see this, consider a version of the model in which an aggregate technology shock drives all aggregate employment fluctuations. The shock follows a random walk in logarithms and multiplies z_t in the plant's production function. The aggregate technology innovations have Gaussian distributions with mean μ_A and variance σ_A^2 . In this case the model resembles a standard (S, s) framework. The scaled employment policies are invariant to the shock, and aggregate

dynamics only reflect changes in the mean of the aggregate productivity distribution. Substitution toward temporary jobs induced by higher idiosyncratic risk can affect aggregate dynamics because plants adjust employment more frequently. The higher adjustment frequency implies that a greater fraction of the industry is ready to respond to a given aggregate shock when it arrives.

We assess this effect in the last perturbation of the calibrated model, Panel I of Tables 6 and 7, where a multiplicative aggregate technology shock drives employment dynamics. We set the variance of this shock, σ_A^2 , to the calibrated variance of the wage shock and constrain its mean so that there is no drift. When we conduct this experiment using the parameter values from the baseline calibration, we find almost no difference between the dynamics of young and old plants. This is because young plants fill very few temporary jobs in the absence of transitory wage shocks. To demonstrate that the use of temporary jobs can influence aggregate dynamics by changing the frequency of employment adjustment, we increased σ_n to the value used to construct Panel E of Tables 6 and 7, 0.831. This change induces new plants to use more temporary jobs.

This version of the model does roughly as well as or better than the calibrated model, replicating all the data's second moments except the variance ratio and the closely related growth-reallocation correlation. The fall in variance ratios is due to the absence of the microeconomic policy asymmetries discussed in Campbell and Fisher (2000). As expected, this version of the model vastly overpredicts the rate of job reallocation among new plants. Still, the results are encouraging since they suggest that many of the model's desirable characteristics do not depend on the source of the aggregate shock.

6. Conclusion

Are our results robust to sensible alternative adjustment cost specifications? Our analysis emphasized the value of the options to create and destroy jobs in the future that kinked adjustment costs generate. Plant level data from the U.S. show that inaction in employment adjustment is common, so adjustment cost specifications that generate this feature are the empirically relevant

case. Fixed costs of increasing or decreasing employment, as in Bertolla and Caballero (1990), are the other prominent adjustment cost specification that generates inaction. We anticipate that increasing idiosyncratic risk will induce plants to substitute more flexible factors for those that require fixed costs of adjustment. The question of how such policy differences impact aggregate dynamics in that framework is an open and interesting question.

Macroeconomic interest in models of infrequent employment adjustment centers around their ability to directly connect microeconomic observations with aggregate fluctuations. Nevertheless, current models are themselves too stylized to directly confront microeconomic data. For instance, they counterfactually predict that job creation very rarely follows job destruction. Empirical characterization of these models with plant-level data has relied on estimation and simulation using reduced form microeconomic policies, called “adjustment hazard functions” by Caballero, Engel, and Haltiwanger (1997), without directly characterizing structural parameters or optimal policies. Our model with temporary jobs generates a wide variety of employment paths, some of which are very unlike those in other models of infrequent employment adjustment. For example, plants in our model can quickly reverse job creation and destruction if they are using temporary jobs. Therefore, the addition of temporary jobs can make infrequent adjustment models suitable for structural estimation. Indeed, Aguirregabiria and Alonso-Borego (1999) have estimated a model similar to ours using Spanish social security data on plants’ use of workers with temporary and permanent contracts. A potentially fruitful extension of their work is the incorporation of aggregate uncertainty. Estimation of such a model using U.S. data would provide more evidence on the connection between idiosyncratic risk, plants’ employment policies, and aggregate fluctuations.

Our emphasis on producers’ substitution between rigid and flexible factors places this paper in a literature which seeks to understand organizational choices and their impact on flexibility. Jones and Ostroy (1984) define an action as increasing flexibility if it lowers the cost of any future adjustment. Similarly, Athey and Schmutzler (1995) model organizational flexibility as a reduction in adjustment costs resulting from a direct investment. We believe that a more natural definition of organizational flexibility focuses directly on producers’ observed dynamics: an organization is more flexible if its observed input choices respond more to a shock of a given size.

The distinction between these two approaches can be quite stark. If the job destruction cost in our model is set equal to zero, then increasing a plant's stock of organized jobs reduces the cost of attaining any desired level and configuration of employment in the future, so it increases flexibility according to Jones and Ostroy (1984). However, the accumulation of organized jobs may move the plant into the region of inaction, where it will choose to behave quite rigidly, not responding to either idiosyncratic or aggregate shocks. For this reason, a definition of organizational flexibility that focuses directly on adjustment costs might be unappealing. Consideration of producers' substitution between similar factors with different costs of adjustment provides an alternative approach to understanding the organization of production and its implications for observed flexibility.

In our industry model, increasing producers' idiosyncratic risks increases the aggregate employment response to a given aggregate shock. In the light of this result, an interesting conjecture is that the increase in microeconomic turbulence that we alluded to in the introduction has caused the U.S. economy to respond more to aggregate disturbances, such as shocks to monetary policy. To formally evaluate this conjecture, a general equilibrium analysis is clearly desirable. Campbell and Fisher (2000) demonstrate that our analysis can be recast in a general equilibrium environment, but that model makes the strong simplifying assumption that consumers are risk-neutral. This is an important caveat, because Veracierto (1998) and Thomas (2000) show that results from models of irreversible and lumpy investment similar to ours that also abstract from endogenous price determination are not robust. Thomas (2000) emphasizes the role of consumption smoothing in determining aggregate investment almost independently of the microeconomic investment technology. Consumption smoothing behavior has no necessary implications for the elasticity of labor supply, so we expect these considerations to be relatively unimportant when considering the effects of idiosyncratic uncertainty on economy-wide employment dynamics. Nevertheless, further development of our analysis in general equilibrium settings is an important task for future research.

References

- Aguirregabiria, Víctor and César Alonso-Borrego. 1999. Labor Contracts and Flexibility: Evidence from a Labor Market Reform in Spain. Working Paper, University of Chicago.
- Athey, Susan and Armin Schmutzler. 1995. Product and Process Flexibility in an Innovative Environment. *RAND Journal of Economics*, **26**(4), 557–574.
- Bentolila, Samuel and Gilles Saint-Paul. 1992. The macroeconomic impact of flexible labor contracts, with an application to Spain. *European Economic Review*, **36**(3), 1013–1053.
- Bertola, Giuseppe and Ricardo J. Caballero. 1990. Kinked Adjustment Costs and Aggregate Dynamics, pp. 237–288 in Olivier Jean Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual*. Cambridge, MA: MIT Press.
- Caballero, Ricardo J. 1992. A Fallacy of Composition. *American Economic Review*, **82**(5), 1279–1292.
- Caballero, Ricardo J., Eduardo M.R.A. Engel, and John Haltiwanger. 1997. Aggregate Employment Dynamics: Building from Microeconomic Evidence. *American Economic Review*, **87**(1), 115–137.
- Cabrales, Antonio and Hugo Hopenhayn. 1997. Labor-market flexibility and aggregate employment volatility. *Carnegie–Rochester Conference Series on Public Policy*, **46**, 189–228.
- Campbell, Jeffrey R. and Jonas D. M. Fisher. 1998. Organizational Flexibility and Employment Dynamics at Young and Old Plants. NBER Working Paper No. 6809.
- Campbell, Jeffrey R. and Jonas D. M. Fisher. 1999. Technical Appendix to “Aggregate Employment Fluctuations with Microeconomic Asymmetries.” Working Paper, University of Chicago.
- Campbell, Jeffrey R. and Jonas D. M. Fisher. 2000. Aggregate Employment Fluctuations with Microeconomic Asymmetries. Forthcoming *American Economic Review*.
- Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu. 2000. Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk. NBER Working Paper No. 7590.
- Comín, Diego. 2000. An Uncertainty-driven Theory of the Productivity Slowdown: Manufacturing. Working Paper, Harvard University.
- Davis, Steven J., John Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*, MIT Press, Cambridge, Mass.
- Dixit, Avinash. 1989. Entry and Exit Decisions under Uncertainty. *Journal of Political Economy*, **97**(3), 620–638.
- Dixit, Avinash K. and Robert S. Pindyck. 1994. *Investment under Uncertainty*, Princeton University Press, Princeton, New Jersey.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson. 1989. The Growth and Failure of U.S. Manufacturing Plants. *Quarterly Journal of Economics*, **104**(4), 671–698.
- Foote, Christopher L. 1998. Trend Employment Growth and the Bunching of Job Creation and Destruction. *Quarterly Journal of Economics*, **113**(3), 809–834.

- Gertler, Mark and Simon Gilchrist. 1994. Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms. *Quarterly Journal of Economics*, **109**(2), 311–340.
- Hamermesh, Daniel S. and Gerard A. Pfann. 1996. Adjustment Costs in Factor Demand. *Journal of Economic Literature*, **34**(3) 1264–1292.
- Jones, Robert A. and Joseph M. Ostroy. 1984. Flexibility and Uncertainty. *Review of Economic Studies*, **51**(1), 13–32.
- Jovanovic, Boyan. 1982. Selection and the Evolution of Industry. *Econometrica*, **50**(3), 649–670.
- Kashyap, Anil K., Owen A. Lamont, and Jeremy C. Stein. 1994. Credit Conditions and the Cyclical Behavior of Inventories. *Quarterly Journal of Economics*, **109**(3), 565–592.
- Mortensen, Dale T. and Christopher A. Pissarides. 1994. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies*, **61**(3), 397–415.
- Sargent, Thomas J and Lars Ljungqvist. 1998. The European Unemployment Dilemma. *Journal of Political Economy*, **106**(3), 514–550.
- Thomas, Julia K. 2000. Is Lumpy Investment Relevant for the Business Cycle? Working Paper, Carnegie Mellon University.
- Veracierto, Marcelo. 1998. Plant Level Irreversible Investment and Equilibrium Business Cycles. Federal Reserve Bank of Chicago Working Paper 98-1.

Table 1: Gross Job Flows in the U.S. Manufacturing Sector

A. Means ^{(i),(ii)}								
Plant Type	$E(G)$	$E(C)$	$E(D)$	$E(R)$				
All	-0.33 (0.36)	5.39 (0.16)	5.72 (0.28)	11.1 (0.28)				
Young	0.95 (0.46)	7.52 (0.33)	6.56 (0.32)	14.1 (0.45)				
Old	-0.70 (0.33)	4.54 (0.12)	5.24 (0.27)	9.8 (0.26)				
B. Second Moments ^{(i),(iii)}								
Plant Type	σ_G	σ_C	σ_D	σ_R	$[\sigma_D/\sigma_C]^2$	$\rho(C,D)$	$\rho(R,G)$	$\rho(G_t,G_{t-1})$
All	1.90 (0.23)	0.85 (0.06)	1.45 (0.21)	1.42 (0.16)	2.89 (0.87)	-0.33 (0.10)	-0.51 (0.12)	0.64 (0.08)
Young	2.32 (0.25)	1.66 (0.19)	1.65 (0.20)	2.26 (0.20)	0.99 (0.37)	-0.07 (0.13)	-0.002 (0.19)	0.66 (0.08)
Old	1.81 (0.24)	0.71 (0.06)	1.43 (0.20)	1.35 (0.15)	4.06 (1.12)	-0.36 (0.09)	-0.63 (0.09)	0.60 (0.10)

Notes: (i) Here G denotes job growth, C denotes the rate of job creation, D denotes the rate of job destruction and R denotes the rate of job reallocation parameterization. These job flow measures are defined as in Davis, Haltiwanger and Schuh (1996). Numbers in parenthesis are standard errors computed using a GMM procedure. For estimation of the relevant zero-frequency spectral density, a Bartlett window truncated at lag three was used. (ii) $E(x)$ denotes the expected value of variable $x = G, C, D, R$. (iii) $\sigma(x)$ denotes the standard deviation of variable x and $\rho(x, y)$ denotes the correlation between variables x and y .

Table 2: Summary Statistics for Employment Fluctuations and Idiosyncratic Risk in Four-Digit Industries

Industry	Statistic ^{(ii),(iii)}	Summary Statistic Across Industries ⁽ⁱ⁾		
		Average	Standard Deviation	Correlation with \hat{R}_0
	$\hat{\sigma}_G$	8.95	3.78	0.33
	$\hat{\sigma}_C$	4.24	1.83	0.50
	$\hat{\sigma}_D$	6.17	3.00	0.22
	$\hat{\sigma}_R$	5.77	3.00	0.27
	\hat{R}_0	18.41	5.90	1.00

Note: (i) For each of the 446 four-digit manufacturing industries, the sample statistics in the table's leftmost column were calculated using annual observations from 1973 through 1986. The reported numbers are the averages, standard deviations, and correlations of these industry sample statistics across all of the industries. (ii) The sample standard deviation of variable X is $\hat{\sigma}_X$, and G , C , D , and R , refer to the industry's job growth rate, job creation rate, job destruction rate, and job reallocation rate. (iii) The sum of job creation and destruction's sample medians is \hat{R}_0 .

Table 3: Estimated Coefficients from Cross-Industry Regressions

Dependent Variable ⁽ⁱⁱⁱ⁾	Estimation ⁽ⁱ⁾⁽ⁱⁱ⁾		
	OLS ^(iv)	LAD ^(iv)	OLS, No Small Industries ^(v)
$\hat{\sigma}_G$	0.170 (0.032)	0.154 (0.043)	0.161 (0.036)
$\hat{\sigma}_C$	0.131 (0.017)	0.094 (0.015)	0.112 (0.016)
$\hat{\sigma}_D$	0.094 (0.028)	0.094 (0.021)	0.084 (0.031)
$\hat{\sigma}_R$	0.124 (0.036)	0.107 (0.023)	0.095 (0.032)

Note: (i) The reported coefficients are from regressions of the indicated dependent variable on the sum of job creation and destruction's sample medians and a full set of two-digit industry dummy variables. The number below each reported coefficient is its heteroskedasticity consistent standard error. (ii) The industry sample statistics used for the regressions were calculated using annual observations from 1973 through 1986. (iii) The sample standard deviation of variable X is $\hat{\sigma}_X$, and G , C , D , and R , refer to the industry's job growth rate, job creation rate, job destruction rate, and job reallocation rate. (iv) The 446 four-digit manufacturing industries comprised the sample for these regressions. (v) The 334 largest industries comprised the sample for this regression.

Table 4: Baseline Calibration

Parameter	Value	Parameter	Value
α	0.66	σ_n	0.518
β	$1.05^{-1/4}$	σ_m	0.415
τ_c	1.5	μ_n	-0.134
τ_d	0.5	μ_m	-0.086
θ	0.7	ϕ	0.05
δ_n	0.0245	σ_W	0.044
δ_m	0.0046	ρ	0

Table 5: Policy Intercepts for the Model Parameterizations Considered

Parameterization	Organized Intercepts								Temporary Intercepts	
	New Plants				Mature Plants				$\tilde{y}(W^l)$	$\tilde{y}(W^h)$
	$\underline{y}^n(W^l)$	$\overline{y}^n(W^l)$	$\underline{y}^n(W^h)$	$\overline{y}^n(W^h)$	$\underline{y}^m(W^l)$	$\overline{y}^m(W^l)$	$\underline{y}^m(W^h)$	$\overline{y}^m(W^h)$		
Calibrated	0.098	0.519	0.098	0.457	0.123	0.498	0.115	0.445	0.117	0.091
$\theta = 0$	0.104	0.520	0.098	0.458	0.123	0.498	0.115	0.445	0	0
$\theta = 0.75$	0.082	0.517	0.086	0.455	0.124	0.505	0.114	0.440	0.144	0.111
$\sigma_n = 0.831$	0.043	0.534	0.046	0.456	0.123	0.498	0.115	0.445	0.117	0.091
$\delta_n = 0.049$	0.092	0.527	0.094	0.462	0.123	0.498	0.115	0.445	0.117	0.091
$\tau_c = \tau_d = 3.0$	0.026	0.774	0.027	0.685	0.048	0.648	0.049	0.587	0.117	0.091
$\rho = 0.47$	0.100	0.528	0.096	0.450	0.125	0.506	0.114	0.439	0.117	0.091
$\sigma_n = 0.831, \sigma_W = 0,$ $\sigma_A = 0.044$	0.045	0.492	0.045	0.492	0.119	0.470	0.119	0.470	0.103	0.103

Note: The calibrated parameters are listed in Table 4. For the other parameterizations all parameters are left at their calibrated values except as indicated, with one exception. The exception is the ‘ $\sigma_n = 0.831$ ’ case in which $\mu_n = -0.345$ so that productivity growth continues to be zero among new plants in expected value.

Table 6: Steady State Gross Job Flows in U.S. Data and in the Model


Plant Type	Statistics ⁽ⁱ⁾			
	$E(G)$	$E(C)$	$E(D)$	$E(R)$
A. U.S Manufacturing Sector				
All	-0.33	5.39	5.72	11.1
Young	0.95	7.52	6.56	14.1
Old	-0.70	4.54	5.24	9.8
B. Baseline Calibration ⁽ⁱⁱ⁾				
All	0.00	5.25	5.25	10.5
Young	2.23	8.18	5.90	14.1
Old	-0.47	4.65	5.12	9.8
C. Temporary Jobs Never Used				
All	0.00	5.21	5.21	10.4
Young	2.28	7.97	5.69	13.7
Old	-0.47	4.65	5.12	9.77
D. Temporary Jobs More Productive				
All	0.00	6.71	6.71	13.4
Young	2.24	11.4	9.18	20.6
Old	-0.47	5.73	6.20	11.9
E. Higher Productivity Risk at New Plants				
All	0.00	6.42	6.42	12.8
Young	2.39	15.3	12.9	28.2
Old	-0.46	4.72	5.18	9.9
F. Higher Exit Risk at New Plants				
All	0.00	5.45	5.45	10.9
Young	2.10	9.09	6.09	16.1
Old	-0.47	4.65	5.11	9.8
G. Higher Adjustment Costs				
All	0.00	11.7	11.7	23.3
Young	2.25	17.9	15.7	33.6
Old	-0.46	10.4	10.9	21.2
H. Persistent Wage Process				
All	0.00	5.25	5.25	10.5
Young	2.28	8.18	5.90	14.1
Old	-0.47	4.65	5.12	9.8
I. Aggregate Technology Shock				
All	0.00	6.42	6.42	12.8
Young	2.39	15.3	12.9	28.2
Old	-0.46	4.72	5.18	9.9


Note: (i) $E(x)$ denotes the expected value of variable $x = G, C, D, R$. Here G denotes job growth, C denotes the rate of job creation, D denotes the rate of job destruction and R denotes the rate of job reallocation parameterization. The numbers for the US data are taken from Table 1. (ii) Panels B-I correspond to the parameterizations of the model listed in Table 5.



Table 7: Variation of Gross Job Flows in U.S. Data and in the Model
Statistics⁽ⁱ⁾


Plant Type	σ_G	σ_C	σ_D	σ_R	$[\sigma_D/\sigma_C]^2$	$\rho(C, D)$	$\rho(R, G)$	$\rho(G_t, G_{t-1})$
A. U.S Manufacturing Sector								
All	1.90	0.85	1.45	1.42	2.89	-0.33	-0.51	0.64
Young	2.32	1.66	1.65	2.26	0.99	-0.07	-0.002	0.66
Old	1.81	0.71	1.43	1.35	4.06	-0.36	-0.63	0.60
B. Baseline Calibration ⁽ⁱⁱ⁾								
All	1.90	0.87	1.03	0.23	1.40	-0.99	-0.70	-0.39
Young	3.69	2.07	1.66	1.04	0.64	-0.87	0.41	-0.47
Old	1.55	0.63	0.94	0.32	2.22	-0.99	-0.97	-0.35
C. Temporary Jobs Never Used								
All	1.57	0.65	0.92	0.28	2.01	-0.99	-0.96	-0.34
Young	1.70	0.81	0.86	0.11	1.15	-0.99	-0.52	-0.34
Old	1.54	0.62	0.93	0.33	2.28	-0.99	-0.97	-0.34
D. Temporary Jobs More Productive								
All	6.65	3.84	3.18	2.36	0.69	-0.79	0.29	-0.50
Young	9.30	5.25	4.25	3.05	0.66	-0.81	0.34	-0.50
Old	6.12	3.54	2.96	2.20	0.70	-0.79	0.28	-0.50
E. Higher Productivity Risk at New Plants								
All	2.43	1.16	1.29	0.33	1.24	-0.97	-0.40	-0.43
Young	7.09	3.77	3.33	1.70	0.78	-0.89	0.27	-0.49
Old	1.86	0.65	1.13	0.32	2.16	-0.99	-0.97	-0.35
F. Higher Exit Risk at New Plants								
All	2.04	0.97	1.09	0.26	1.26	-0.98	-0.47	-0.41
Young	4.43	2.52	2.01	1.42	0.63	-0.83	0.38	-0.48
Old	1.55	0.62	0.93	0.32	2.25	-0.99	-0.97	-0.34
G. Higher Adjustment Costs								
All	12.4	6.56	5.98	2.67	0.83	-0.91	0.23	-0.50
Young	15.9	8.17	7.44	2.83	0.83	-0.94	0.26	-0.50
Old	11.6	6.22	5.65	2.61	0.83	-0.91	0.22	-0.50
H. Persistent Wage Process								
All	1.89	0.90	0.99	0.22	1.21	-0.98	-0.41	0.02
Young	3.02	1.78	1.34	1.11	0.57	-0.78	0.42	-0.12
Old	1.67	0.73	0.96	0.28	1.77	-0.98	-0.85	0.08
I. Aggregate Technology Shock								
All	1.81	1.15	0.96	1.10	0.69	-0.46	0.20	0.57
Young	2.23	1.49	1.04	1.38	0.49	-0.45	0.38	0.24
Old	1.78	0.94	0.85	0.22	0.82	-0.97	0.41	0.61

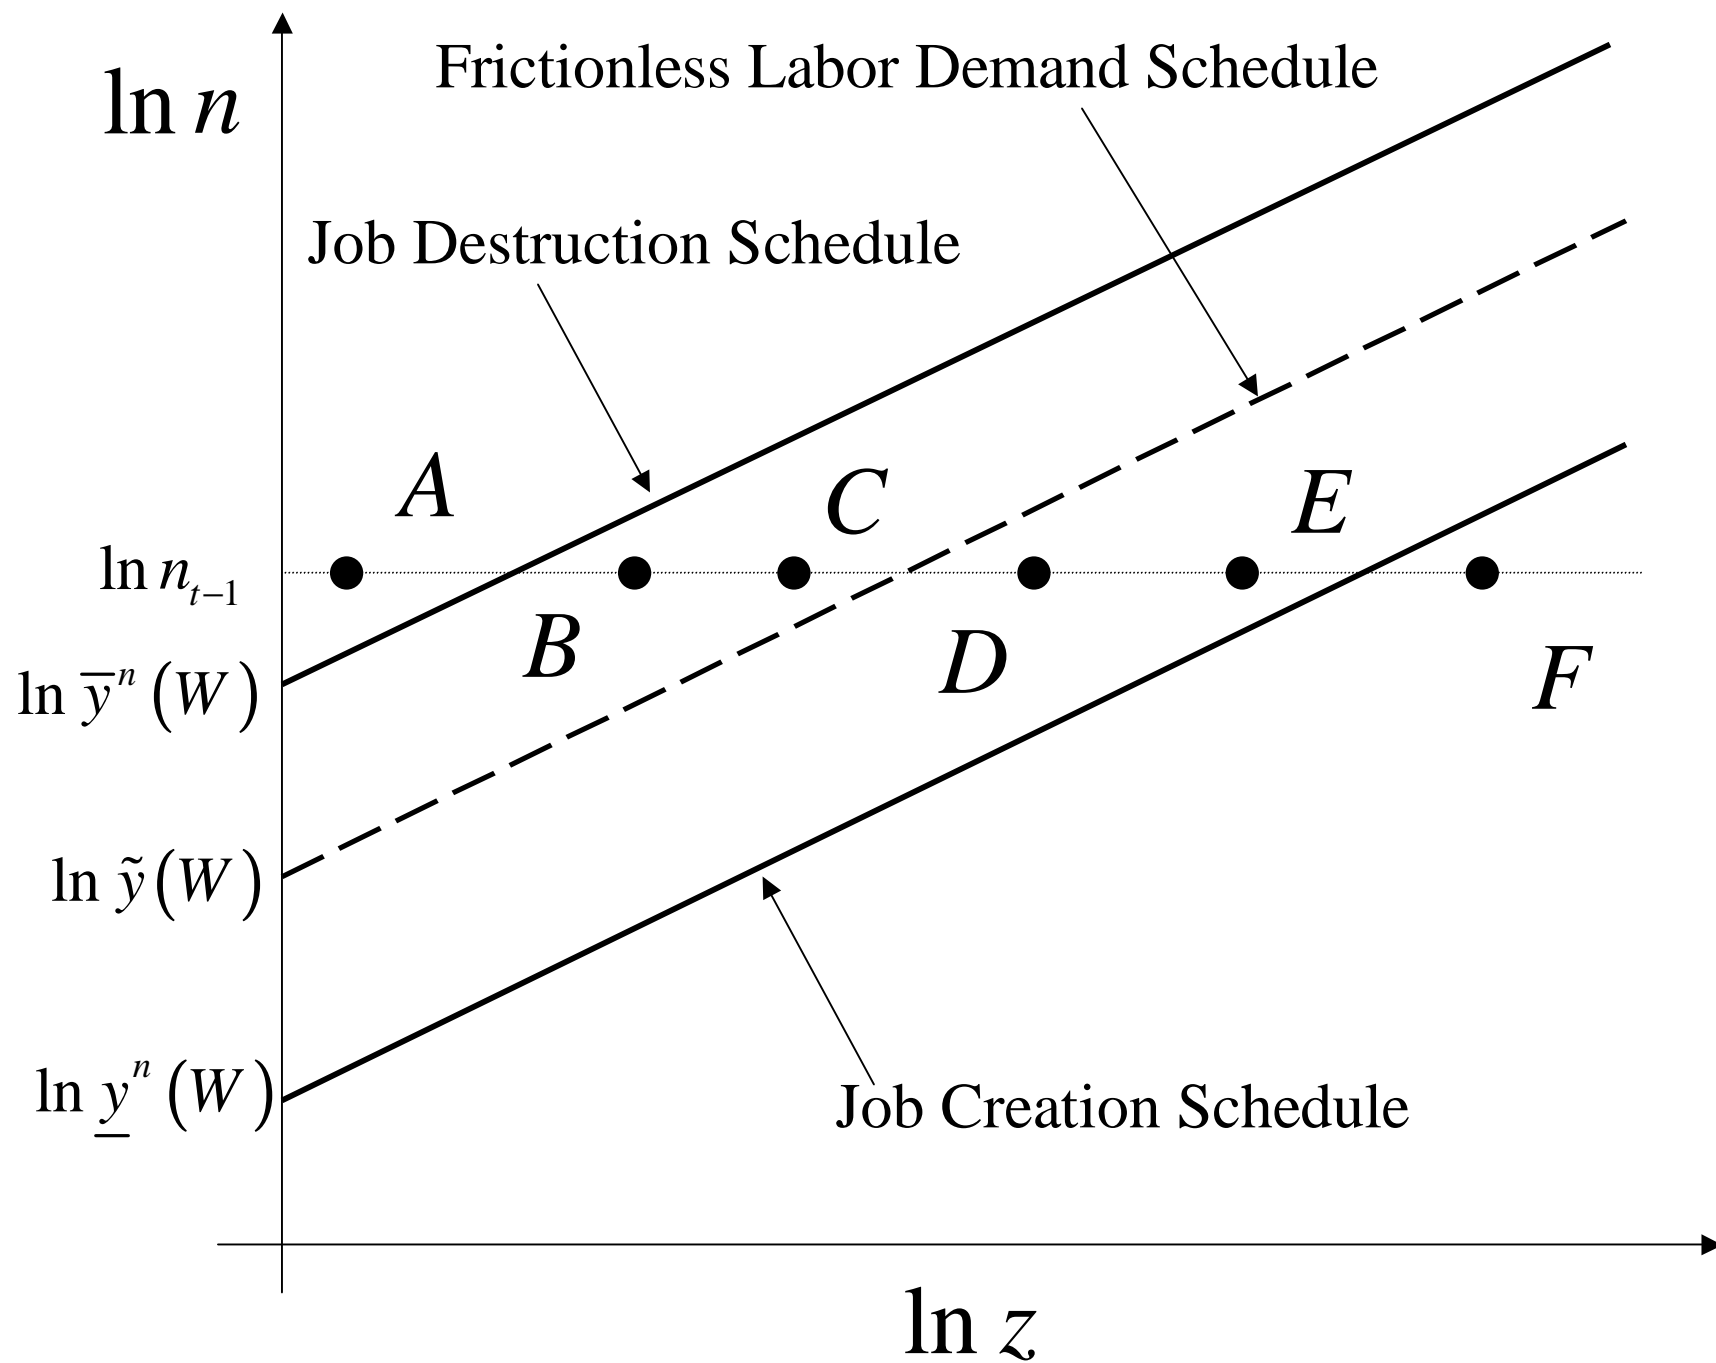
Note: (i) $\sigma(x)$ denotes the standard deviation of variable x and $\rho(x, y)$ denotes the correlation between variables x and y . For the definitions of the variables, see Table 6. The numbers for the U.S. data are from Table 1. (ii) Panels B-I correspond to the parameterizations of the model listed in Table 5.














Impact of σ_m on Organized Creation Margin




Impact of δ_m on Organized Creation Margin




Organized and Temporary Jobs Used



Organized Jobs Only



Organized and Temporary Jobs Used



Organized Jobs Only

